



FortiOS™ Handbook v3

for FortiOS 4.0 MR3



FortiOS™ Handbook

v3

13 March 2012

01-435-99686-20120313

for FortiOS 4.0 MR3

Copyright© 2012 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, and FortiGuard®, are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance metrics contained herein were attained in internal lab tests under ideal conditions, and performance may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to the performance metrics herein. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any guarantees. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

Visit these links for more information and documentation for your Fortinet products:

Fortinet Knowledge Base - <http://kb.fortinet.com>

Technical Documentation - <http://docs.fortinet.com>

Training Services - <http://campus.training.fortinet.com>

Technical Support - <http://support.fortinet.com>

You can report errors or omissions in this or any Fortinet technical document to techdoc@fortinet.com.



Contents Quick Look

	Best Practices	9
Chapter 1	What's New	97
	Upgrading to FortiOS 4.0 MR3	99
	FortiOS 4.0 MR3 New Feature Highlights	103
	Logging and reporting enhancements	127
	FortiOS 4.0 MR3 Usability improvements	139
	More New Features	155
Chapter 2	Firewall	181
	Understanding the FortiGate firewall	183
	Working with NAT in FortiOS	189
	Firewall components	195
	Security policies	217
	Monitoring firewall traffic	229
	Internet Protocol version 6 (IPv6)	235
	Advanced FortiGate firewall concepts	265
Chapter 3	System Administration	287
	Using the web-based manager	289
	Using the CLI	315
	Basic setup	335
	Interfaces	375
	Central management	395
	Best practices	401
	FortiGuard	407
	Monitoring	421
	Multicast forwarding	471
	Virtual LANs	503
	PPTP and L2TP	537
	Session helpers	551
	Advanced concepts	561

Chapter 4	Logging and Reporting	623
	Logging overview	625
	The SQLite log database	633
	Log devices	639
	Logging FortiGate activity	655
	Log message usage	673
	Reports	677
Chapter 5	Troubleshooting	697
	Life of a Packet	699
	Troubleshooting process	713
	Troubleshooting tools	719
	Technical Support Organization Overview	751
	Troubleshooting common issues	763
	Troubleshooting advanced	789
	Troubleshooting 'get' commands	805
	Troubleshooting bootup and FSSO	871
Chapter 6	UTM Guide	877
	UTM overview	879
	Network defense	883
	AntiVirus	895
	Email filter	913
	Intrusion protection	943
	Web filter	997
	FortiGuard Web Filter	1023
	Data leak prevention	1035
	Application control	1061
	DoS policy	1081
	Endpoint Control and monitoring	1089
	Vulnerability Scan	1113
	Sniffer policy	1127
	Other UTM considerations	1137
Chapter 7	User Authentication	1161

Introduction to authentication	1163
Authentication and User in the web-based manager	1171
Authentication servers	1201
Users and user groups	1223
Configuring authenticated access	1243
Certificate-based authentication	1263
FSSO integration with Windows AD or Novell	1283
Dynamic profiles and end points	1325
Monitoring authenticated users	1359
Examples and Troubleshooting	1361
Chapter 8 IPsec VPNs	1379
IPsec VPN concepts	1381
IPsec VPN Overview	1389
IPsec VPN in the web-based manager	1393
Auto Key phase 1 parameters	1407
Phase 2 parameters	1425
Defining VPN security policies	1431
Gateway-to-gateway configurations	1437
Hub-and-spoke configurations	1453
Dynamic DNS configuration	1469
FortiClient dialup-client configurations	1483
FortiGate dialup-client configurations	1501
Supporting IKE Mode config clients	1509
Internet-browsing configuration	1515
Redundant VPN configurations	1519
Transparent mode VPNs	1543
Manual-key configurations	1551
IPv6 IPsec VPNs	1555
L2TP and IPsec (Microsoft VPN)	1567
GRE over IPsec (Cisco VPN) configurations	1579
Protecting OSPF with IPsec	1589
Hardware offloading and acceleration	1597
Monitoring and troubleshooting	1603
Chapter 9 SSL VPN	1611

Introduction to SSL VPN	1613
Basic Configuration	1617
The SSL VPN client	1651
Setup examples	1655
Chapter 10 Advanced Routing	1669
Advanced Static routing	1671
Dynamic Routing Overview	1699
Routing Information Protocol (RIP)	1715
Border Gateway Protocol (BGP)	1751
Open Shortest Path First (OSPF)	1787
Intermediate System To Intermediate System Protocol (IS-IS)	1827
Router Reference	1841
Chapter 11 Virtual Domains	1873
Virtual Domains	1875
Virtual Domains in NAT/Route mode	1903
Virtual Domains in Transparent mode	1921
Inter-VDOM routing	1941
Troubleshooting Virtual Domains	1977
Chapter 12 High Availability	1983
Solving the High Availability problem	1987
An introduction to the FortiGate Clustering Protocol (FGCP)	1991
Configuring and connecting HA clusters	2021
Configuring and connecting virtual clusters	2089
Configuring and operating FortiGate full mesh HA	2111
Operating a cluster	2127
HA and failover protection	2167
HA and load balancing	2217
HA with third-party products	2233
VRRP	2237
TCP session synchronization	2243
Chapter 13 Traffic Shaping	2249
The purpose of traffic shaping	2251

Traffic shaping methods	2261
Examples	2277
Troubleshooting	2285
Chapter 14 FortiOS Carrier	2289
Overview of FortiOS Carrier features	2291
Carrier web-based manager settings	2315
MMS UTM features	2353
Message flood protection	2373
Duplicate message protection	2385
MMS Replacement messages	2393
Configuring GTP on FortiOS Carrier	2401
GTP message type filtering	2409
GTP identity filtering	2417
Troubleshooting	2425
Chapter 15 Deploying Wireless Networks	2433
Introduction to wireless networking	2435
Configuring a WiFi LAN	2445
Access point deployment	2459
Wireless network monitoring	2469
Configuring wireless network clients	2475
Wireless network examples	2487
Using a FortiWiFi unit as a client	2501
WiFi Reference	2503
WiFi Controller Reference	2505
Chapter 16 VoIP Solutions: SIP & FortiGate Voice	2517
FortiGate VoIP solutions: SIP	2519
Example FortiGate Voice branch office configuration	2603
FortiGate Voice web-based manager configuration reference	2621
Using the PBX user web portal	2649
FortiGate Voice VoIP, PBX, and PSTN CLI Reference	2653

Chapter 17	WAN Optimization, Web Cache, Explicit Proxy, and WCCP	2671
	WAN optimization, web cache, explicit proxy, and WCCP concepts	2673
	WAN optimization and Web cache storage	2693
	WAN optimization peers and authentication groups	2697
	Configuring WAN optimization rules	2705
	WAN optimization configuration examples	2715
	Web caching	2735
	Advanced configuration example	2757
	SSL offloading for WAN optimization and web caching	2781
	FortiClient WAN optimization	2805
	The FortiGate explicit web proxy	2807
	The FortiGate explicit FTP proxy	2831
	FortiGate WCCP	2845
	WAN optimization, web cache, explicit proxy and WCCP get and diagnose commands	2859
Chapter 18	Load Balancing	2867
	Configuring load balancing	2869
	Load balancing configuration examples	2899
Chapter 19	Hardware	2917
	FortiGate installation	2919
	AMC module configuration	2929
	FortiGate hardware accelerated processing	2933
	Configuring RAID	2961
	FortiBridge installation and operation	2967
Chapter 20	Certifications and Compliances	2999
	FIPS-CC operation of FortiGate units	3001
	Configuring FortiGate units for PCI DSS compliance	3023
	Appendix	3037
	Index	3043



Best Practices

Administrator password best practices	312
Hardware best practices	401
Shutting down best practices	403
Performance best practices	403
Firewall best practices	404
Intrusion protection best practices	404
Antivirus best practices	404
Web Filtering best practices	405
Email filtering best practices	405
Security best practices	405
Log management best practices	631
Troubleshooting best practices	713
Password policy best practices	1244
Password best practices	1245
HA best practices	2013
Encapsulating IP traffic filtering best practices	2403
Deep SIP message inspection best practices	2588
WAN Optimization best practices	2692



Contents

Best Practices	9
How this Handbook is organized	95
Chapter 1 What's New	97
Upgrading to FortiOS 4.0 MR3	99
General firmware upgrade steps	99
Backing up and restoring your FortiGate configuration file	100
Temporarily installing FortiOS 4.0 MR3	100
FortiOS 4.0 MR3 New Feature Highlights	103
Flow-based UTM Extensions.	103
UTM Configuration and Inspection Enhancements	104
UTM profile and sensor configuration improvements	104
Archive inspection for antivirus profiles	105
Improved IPS default block rate	105
Web Filter profiles	105
Web Filtering Overrides	106
Application Control Sensors and filters	107
Geography-based filtering for firewall addresses	107
DLP document fingerprinting	108
Internet Content Adaptation Protocol (ICAP)	109
Profile Group	110
Modem interface Improvements	111
WiFi Extensions.	111
WiFi controller redesign	112
Captive portal enhancements	112
Rogue AP detection and reporting.	113
Custom AP profiles.	113
Distributed ARRP (Automatic Radio Resource Provisioning)	113
WiFi monitor	114
New WiFi commands.	114

Strong Authentication Enhancements	115
FortiToken support	115
Two-factor authentication	116
Enabling two-factor authentication for administrators	117
Multiple authentication group enforcement	118
Dynamic Profiles	118
Hard-timeout enhancement	119
PKI certificate authentication enhancement	119
NTLM authentication enhancements.	120
New PCI Compliance Features.	120
Feature Improvements to extend IPv6 support	122
Top Session dashboard widget IPv6 support	122
OSPFv3 NSSA extension	122
Explicit proxy and web caching improvements	123
Explicit FTP proxy	123
Form-based user authentication for explicit web proxy	125
Web caching in security policies	125
Logging and reporting enhancements	127
The FortiGate UTM Weekly Activity Report	127
Viewing the current and historical reports	129
Creating custom reports from the CLI	130
Log Access Improvements	130
Viewing log messages	131
Filtering log messages	131
Downloading log messages	131
New Unified UTM Log Access	132
SQL logging enabled by default	132
Sending DLP archives to multiple FortiAnalyzer units	133
Remote logging configuration enhancements	133
Log and Report Monitoring.	134
Logging Monitor	134
Log Message Enhancements.	134
Event logs	134
Traffic logs	135
Chat message log support for MSNP21	135
SSL connection encryption level option over OFTP	135
Uploading logs to a FTP server in text format	136
Example for uploading logs to a FTP server in text format	136
Deleting all local logs, archives and user-configured report templates	137

FortiGuard Analysis and Management Service (FAMS)	137
FortiAnalyzer with FAMS support	137
FAMS enhancements	137
FortiOS 4.0 MR3 Usability improvements	139
High-level web-based manager menu changes	139
New FortiGate Setup Wizard	140
FortiExplorer enhancements	140
Dashboard Widgets	140
Traffic History	141
System Resources	141
Network Protocol Usage	142
Chart display improvements	143
Monitoring Improvements	143
DHCP Monitor	143
Modem Monitor	144
Session Monitor	144
Policy Monitor	144
Load Balance Monitor	144
Traffic Shaper Monitor	144
AV Monitor	145
Intrusion Monitor	145
Web Monitor	145
Email Monitor.	145
Archive & Data Leak Monitor.	145
Application Monitor	146
IPsec Monitor	146
SSL-VPN Monitor	146
Web Cache Monitor	146
WAN optimization Peer Monitor	146
WAN optimization web cache monitor.	146
Filtering web-based manager lists	147
Reference count column (object usage visibility).	148
Configuration object tagging and coloring	150
Adding tags to configuration objects	151
Example of how to find a security policy using Tag Management.	151
Adding tags to predefined signatures and applications	152
Security configuration object icons	153
Access to online help.	153
Backing up and restoring configuration files per-VDOM.	153
More New Features	155
New features for FortiOS 4.0 MR3 Patch 5	156

New features for FortiOS 4.0 MR3 Patch 4	156
New features for FortiOS 4.0 MR3 Patch 3	156
New features for FortiOS 4.0 MR3 Patch 2	156
New features for FortiOS 4.0 MR3 Patch 1	158
Login grace timer for SSH connections	159
FortiManager automatic authorization	159
Dynamic DNS commands	159
New diagnose commands	160
Real-time session, traffic shaper bandwidth and CP6 statistics.	160
diag sys session filter proto-state	160
diag log-stats show	160
New get commands	160
IPsec get commands.	160
Traffic shaper and per-IP shaper.	161
Management checksum configuration information for FortiManager	161
MTU configuration support on non-IPsec tunnel interfaces	162
Customizing maximum number of invalid firewall authentication attempts	162
Controlling the connection between a FortiManager unit and a FortiGate unit	162
Bringing up or down IPsec tunnels.	163
Configuring active CPUs	163
Formatting multiple disk partitions	163
Transparent mode port pairs	164
DNS server changes	165
DHCP Server changes	165
DHCP IP Reservation	165
Installing firmware on a partition without a reboot	166
Example of installing a firmware on a partition without rebooting	166
SNMP enhancements	167
WAN optimization, Web Cache and Explicit proxy MIBs	167
SNMPv3	167
Replacement message changes	167
Archive replacement messages and FTP proxy replacement message	168
Successful firewall authentication replacement message	168
Web filtering disclaimer replacement message	168
Video chat block replacement message	168
Replacement message images	168
VDOM and global privileges for access profiles	169
Example of incorporating the new access profile to existing administrator accounts.	169

HA dynamic weighted load balancing	170
Configuring weighted-round-robin weights	170
Dynamic weighted load balancing	172
VRRP virtual MAC address support	174
FGCP HA subsecond failover	175
Static Route enhancements	175
Monitoring ISIS from the Routing Monitor page	176
Security Policy and Firewall Object Enhancements	176
Source IP addresses for FortiGate-originating traffic	176
Local-in security policies	177
Protocol Options	177
FTPS support	177
Virtual IP source address filter support.	178
Virtual IP port forwarding enhancements.	178
Load balancing HTTP host connections	178
Web Proxy Service and Web Proxy Service Group	178
SSL renegotiation for SSL offloading provides allow/deny client renegotiation	179
SSL VPN Port forwarding support	179
IKE negotiation	179
SHA-384 and SHA-512 support for IKE	180
FortiOS Carrier URL extraction feature.	180

Chapter 2 Firewall 181

Understanding the FortiGate firewall 183

What is the FortiGate firewall?	183
FortiGate firewall components	183
How the firewall components create a FortiGate firewall and help in protecting your network	184
Understanding how a packet travels through the FortiGate unit.	185
How packets flow in and out of the FortiGate unit.	186

Working with NAT in FortiOS 189

NAT in FortiOS	189
NAT/Route mode.	189
Route mode	190
Transparent mode	191
Types of NAT in FortiOS	191
Static NAT (SNAT)	192
Dynamic NAT (DNAT)	192

Combining types of NAT	193
Firewall components	195
Using Interfaces and zones in the FortiGate firewall	195
How to apply VLANs and zones and to a security policy	195
Understanding the firewall address component	196
IP addresses for self-originated traffic	197
IP pools.	198
IP Pools for security policies that use fixed ports	199
Source IP address and IP pool address matching.	199
Geography-based addressing	201
Wildcard addresses	202
Fully Qualified Domain Name addresses.	204
Address groups	204
Virtual IP addresses	205
Services.	206
Service groups	212
Firewall schedules	213
Schedule groups	213
Schedule expiry	213
UTM profiles	214
How to use UTM profiles to monitor and protect your network	214
Security policies	217
Security policy overview	217
Security policy list details	218
Viewing security policies	219
Policy order.	219
How to arrange policies	221
Security policies	221
Identity-based policies	222
SSL VPN policies.	223
IPsec policies.	224
Accept policies.	224
Deny policies	224
IPv6 policies	225
Security policy 0	225
Local-in policies	226
Creating a basic security policy	226
How to create a basic security policy for Internet access	227
How to test the basic security policy	227
How to verify if traffic is hitting the basic security policy	228

Monitoring firewall traffic	229
Session tables	229
Viewing session tables in the web-based manager	229
Sessions Monitor.	229
Viewing session tables in the CLI	230
Monitoring security policy traffic activity	232
Internet Protocol version 6 (IPv6)	235
What is IPv6?	235
IPv6 in FortiOS	236
Dual stack routing configuration	236
IPv4 tunneling configuration	237
Remotely connecting to an IPv6 network over the Internet	237
IPv6 overview.	237
Differences between IPv4 and IPv6	237
IPv6 MTU.	238
IPv6 address format	238
IP address notation	239
Netmasks.	240
Address scopes	240
Address types	240
IPv6 neighbor discovery	244
Transition from IPv4 to IPv6	245
Configuring FortiOS to connect to an IPv6 tunnel provider	246
Create a SIT-tunnel interface.	247
Create a static IPv6 route into the tunnel-Interface	247
Assign your IPv6 network to your FortiGate	247
Create a security policy to allow traffic from port1 to the tunnel interface	248
Test the connection	248
FortiGate IPv6 configuration	248
Displaying IPv6 options on the web-based manager	249
UTM protection for IPv6 networks	249
Configuring IPv6 interfaces	249
Configuring IPv6 routing	250
Configuring IPv6 security policies	251
Configuring IPv6 DNS	255
Configuring IPv6 DHCP	255
Configuring IPv6 over IPv4 tunneling	255
Configuring IPv6 IPsec VPNs	256

IPv6 troubleshooting	258
ping6	259
diagnose sniffer packet	262
diagnose debug flow	263
IPv6 specific diag commands	263
Additional IPv6 resources	263
Advanced FortiGate firewall concepts	265
Central NAT table.	265
Central NAT Table configuration settings	266
Stateful inspection of SCTP traffic	266
Configuring FortiGate SCTP filtering.	267
Adding an SCTP custom service.	268
Adding an SCTP policy route	268
Changing the session time to live for SCTP traffic.	269
Port pairing	269
Blocking port 25 to email server traffic.	270
Dedicated traffic	271
Restricting traffic on port 25	272
Blocking HTTP access by IP	273
ICMP packet processing	274
Adding NAT security policies in Transparent mode	274
Adding a static NAT virtual IP for a single IP address and port	277
Double NAT: combining IP pool with virtual IP.	279
Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping.	281
Traffic shaping and per-IP traffic shaping	283
Endpoint Security.	284
Logging traffic	284
Quality of Service (QoS)	285
Identity-based security policies	285
Identity-based policy positioning	285
Identity-based sub-policies	286
Chapter 3 System Administration	287
Using the web-based manager	289
Web-based manager overview.	289
Web-based manager menus and pages	289
Using information tables	290
Using column settings	291

Using online help	291
Online help search tips.	292
Using the keyboard to navigate in the online help	293
Entering text strings	293
Entering text strings (names).	293
Entering numeric values	294
Selecting options from a list	294
Enabling or disabling options	294
Dashboard	294
Adding dashboards	295
Adding widgets to a dashboard	295
System Information widget.	296
License Information widget	303
FortiGate unit Operation widget	305
System Resources widget	305
Alert Message Console widget.	305
Log and Archive Statistics widget	306
CLI Console widget	308
Session History widget.	308
Top Sessions widget	308
Traffic History widget.	308
RAID monitor widget	308
Top Application Usage widget	311
Storage widget	311
P2P Usage widget	311
Per-IP Bandwidth Usage widget	312
VoIP Usage widget	312
IM Usage widget	312
Network Protocol Usage	312
Basic configurations	312
Changing your administrator password (best practices).	312
Changing the web-based manager language	312
Changing administrative access	313
Changing the web-based manager idle timeout.	313
Switching VDOMs	313
Connecting to the CLI from the web-based manager	313
Logging out	313
Using the CLI	315
Connecting to the CLI	315
Connecting to the CLI using a local console.	316
Enabling access to the CLI through the network (SSH or Telnet)	317
Connecting to the CLI using SSH	318
Connecting to the CLI using Telnet	319

Command syntax	319
Terminology	319
Indentation	321
Notation	321
Sub-commands	323
Example of table commands.	325
Permissions.	326
Tips	327
Help	327
Shortcuts and key commands	327
Command abbreviation	328
Environment variables	328
Special characters	328
Using grep to filter get and show command output	329
Language support and regular expressions	329
Screen paging	332
Baud rate.	332
Editing the configuration file on an external host	332
Using Perl regular expressions.	333
Basic setup	335
Connecting to the FortiGate unit	335
Connecting to the web-based manager	335
Connecting to the CLI	336
Setup Wizard	336
FortiExplorer	336
Installation	337
Configuration options	337
Updating FortiExplorer and firmware	337
Configuring NAT mode	338
Configure the interfaces	338
Configure a DNS	340
Add a default route and gateway	341
Add security policies	341
Configuring transparent mode	343
Switching to transparent mode	343
Configure a DNS	343
Add security policies	344
Verifying the configuration	345
Additional configuration	346
Setting the time and date	346
Configuring FortiGuard.	347

Passwords	348
Password considerations	348
Password policy	348
Forgotten password?	349
Administrators	349
Administrator configuration	349
Regular (password) authentication for administrators	349
Management access	350
Tightening Security.	350
Disable interfaces	353
RADIUS authentication for administrators	353
Configuring LDAP authentication for administrators.	354
TACACS+ authentication for administrators	354
PKI certificate authentication for administrators	355
Administrator profiles	355
Adding administrators	356
LDAP Admin Access and Authorization	357
Monitoring administrators	358
Trusted hosts.	359
General Settings	359
Administrative port settings	360
Password policies	360
Display options.	360
Backing up the configuration.	360
Backup and restore a configuration file using SCP	361
Restoring a configuration	363
Configuration revisions.	364
Firmware	364
Downloading firmware	365
Upgrading the firmware - web-based manager	365
Reverting to a previous firmware version	365
Configuration Revision	366
Upgrading the firmware - CLI	367
Installing firmware from a system reboot using the CLI	370
Testing new firmware before installing.	372
Controlled upgrade.	374
Interfaces	375
Physical.	375
Interface settings	377

Interface configuration and settings	379
Switch Mode	382
Loopback interfaces	383
Redundant interfaces	383
DHCP on an interface	384
PPPoE on an interface	385
Administrative access	386
Wireless.	387
Interface MTU packet size	387
Secondary IP addresses to an interface	388
Software switch interfaces	388
Virtual domains	389
Example	389
Virtual LANs	390
Example	391
Zones	392
IPv6	393
Example	393
Central management	395
Adding a FortiGate to FortiManager	395
FortiGate configuration.	396
FortiManager configuration	397
Configuration through FortiManager	397
Global objects	397
Locking the FortiGate web-based manager	398
Firmware updates	398
FortiGuard	399
Backup and restore configurations.	399
Administrative domains.	399
Best practices	401
Hardware	401
Environmental specifications.	401
Grounding	402
Rack mount instructions	402
Shutting down	403
Performance	403
Firewall	404
Intrusion protection.	404
Antivirus	404

Web filtering	405
Email Filtering (Antispam).	405
Security	405
FortiGuard	407
FortiGuard Services	407
Support Contract and FortiGuard Subscription Services	408
FortiGuard Analysis Service Options.	408
Antivirus and IPS	409
Antivirus and IPS Options	409
Manual updates	409
Automatic updates	410
Web filtering	412
Web Filtering and Email Filtering Options	413
URL verification	413
Email filtering	414
Security tools	414
URL lookup.	414
IP and signature lookup	415
Online virus scanner	415
Malware removal tools	415
Troubleshooting	415
Web-based manager verification	415
CLI verification	417
Port assignment	418
Monitoring	421
Dashboard	421
Widgets.	421
FortiClient connections.	422
sFlow	422
Configuration.	423
Monitor menus	423
Logging	424
FortiGate memory	424
FortiGate hard disk.	424
Syslog server.	424
FortiGuard Analysis and Management service.	425
FortiAnalyzer	426
Sending logs using a secure connection.	426
Alert email	427

SNMP	428
SNMP configuration settings	429
Gigabit interfaces	432
SNMP agent	432
SNMP community	432
Enabling on the interface	434
Fortinet MIBs	434
SNMP get command syntax	436
Fortinet and FortiGate traps	436
Fortinet and FortiGate MIB fields	441
Fortinet MIB	441
FortiGate MIB	443
Multicast forwarding	471
Sparse mode	471
Dense mode	472
Multicast IP addresses	473
PIM Support	473
Multicast forwarding and FortiGate units	474
Multicast forwarding and RIPv2	474
Configuring FortiGate multicast forwarding	475
Adding multicast security policies	475
Enabling multicast forwarding	476
Multicast routing examples	478
Example FortiGate PIM-SM configuration using a static RP	478
FortiGate PIM-SM debugging examples	484
Example multicast destination NAT (DNAT) configuration	489
Example PIM configuration that uses BSR to find the RP	491
Virtual LANs	503
VLAN ID rules	504
VLAN switching and routing	504
VLAN layer-2 switching	504
VLAN layer-3 routing	507
VLANs in NAT mode	511
Adding VLAN subinterfaces	512
Configuring security policies and routing	514
Example VLAN configuration in NAT mode	515
General configuration steps	516
Configure the FortiGate unit	516
Configure the VLAN switch	521
Test the configuration	522

VLANs in transparent mode	522
VLANs and transparent mode	523
Example of VLANs in transparent mode	525
General configuration steps	526
Configure the FortiGate unit	526
Configure the Cisco switch and router	530
Test the configuration	531
Troubleshooting VLAN issues	531
Asymmetric routing	531
Layer-2 and Arp traffic	532
Forward-domain solution	534
NetBIOS	534
STP forwarding	535
Too many VLAN interfaces	535
PPTP and L2TP	537
How PPTP VPNs work	537
FortiGate unit as a PPTP server	539
Configuring user authentication for PPTP clients	539
Enabling PPTP and specifying the PPTP IP address range	540
Adding the security policy	541
Configuring the FortiGate unit for PPTP VPN	542
Configuring the FortiGate unit for PPTP pass through	542
Configuring a virtual IP address	543
Configuring a port-forwarding security policy	543
Testing PPTP VPN connections	544
Logging VPN events	544
Configuring L2TP VPNs	545
Network topology	546
L2TP infrastructure requirements	547
L2TP configuration overview	547
Authenticating L2TP clients	548
Enabling L2TP and specifying an address range	548
Defining firewall source and destination addresses	548
Adding the security policy	549
Configuring a Linux client	549
Monitoring L2TP sessions	550
Testing L2TP VPN connections	550
Logging L2TP VPN events	550
Session helpers	551
Viewing the session helper configuration	551

Changing the session helper configuration	552
Changing the protocol or port that a session helper listens on	552
Disabling a session helper	554
DCE-RPC session helper (dcerpc)	555
DNS session helpers (dns-tcp and dns-udp).	555
File transfer protocol (FTP) session helper (ftp)	556
H.245 session helpers (h245I and h245O)	556
H.323 and RAS session helpers (h323 and ras)	556
Alternate H.323 gatekeepers.	556
Media Gateway Controller Protocol (MGCP) session helper (mgcp).	557
ONC-RPC portmapper session helper (pmap).	557
PPTP session helper for PPTP traffic (pptp)	557
Remote shell session helper (rsh)	559
Real-Time Streaming Protocol (RTSP) session helper (rtsp)	559
Session Initiation Protocol (SIP) session helper (sip).	560
Trivial File Transfer Protocol (TFTP) session helper (tftp).	560
Oracle TNS listener session helper (tns)	560
Advanced concepts	561
Dual internet connections	561
Redundant interfaces	561
Load sharing	564
Link redundancy and load sharing	564
Single firewall vs. multiple virtual domains	564
Single firewall vs. vdoms	565
Modem	567
USB modem port.	568
Modes	568
Additional modem configuration	570
Modem interface routing	570
DHCP servers and relays.	570
DHCP Server configuration	570
Service	572
Reserving IP addresses for specific clients	573
DHCP options	573
DHCP Monitor	574
Assigning IP address by MAC address	574

DNS services	574
DNS queries	574
Additional DNS CLI configuration	575
DNS server	575
Recursive DNS	576
Dynamic DNS.	577
Aggregate Interfaces	577
IP addresses for self-originated traffic	578
Administration for schools	579
Security policies	579
DNS	580
Encrypted traffic (HTTPS)	580
FTP	580
Example security policies	580
UTM Profiles	580
Logging.	582
Tag management.	582
Adding and removing tags	583
Reviewing tags	583
Tagging guidelines	584
Software switch	585
Soft switch example	586

Replacement messages list	587
Replacement message images	587
Adding images to replacement messages	587
Modifying replacement messages	588
Replacement message tags	588
Mail replacement messages	590
HTTP replacement messages	591
Web Proxy replacement messages	594
FTP Proxy replacement message	595
FTP replacement messages	595
NNTP replacement messages	596
Alert Mail replacement messages	597
Spam replacement messages	598
Administration replacement message	599
Authentication replacement messages	599
Captive Portal Default replacement messages	602
FortiGuard Web Filtering replacement messages	603
IM and P2P replacement messages	604
Endpoint NAC replacement messages	605
NAC quarantine replacement messages	606
Traffic quota control replacement messages	607
SSL VPN replacement message	608
MM1 replacement messages	608
MM3 replacement messages	612
MM4 replacement messages	615
MM7 replacement messages	617
MMS replacement messages	618
Replacement message groups	618
Disk	619
Formatting the disk	619
Setting space quotas	619
CLI Scripts	619
Uploading script files	620
Rejecting PING requests	620
Opening TCP 113	621
Obfuscate HTTP headers	621

Chapter 4 Logging and Reporting 623

Logging overview 625

What is logging?	625
How the FortiGate unit records log messages	625

Log messages	626
Explanation of a log message	627
Explanation of a debug log message	629
Viewing log messages	629
Log files.	630
Best Practices: Log management	631
The SQLite log database	633
SQL overview.	633
SQLite database tables.	634
SQLite statement examples	634
Distribution of Applications by Type in the last 24 hours	634
Top 10 Application Bandwidth Usage Per Hour Summary	635
Example of how to create a dataset containing attack name instead of attack ID	636
Troubleshooting SQL issues	636
SQL statement syntax errors.	636
Connection problems	636
SQL database error	637
Log devices	639
Choosing a log device	639
Example: Setting up a log device and backup solution	640
Configuring the FortiGate unit to store logs on a log device.	641
Logging to the FortiGate unit's system memory.	642
Logging to the FortiGate unit's hard disk	642
Logging to a FortiAnalyzer unit.	643
Logging to a FortiGuard Analysis server.	644
Logging to a Syslog server.	645
Logging to a WebTrends server	646
Logging to multiple FortiAnalyzer units or Syslog servers	647
Troubleshooting issues.	650
Testing FortiAnalyzer and FortiGuard Analysis server connections	650
Testing the FortiAnalyzer configuration	651
Testing the FortiGuard Analysis server configuration	651
Using diag sys logdisk usage	651
Connecting to a FortiAnalyzer unit using Automatic Discovery	652
Uploading logs to a FortiAnalyzer or a FortiGuard Analysis server	652

Logging FortiGate activity	655
Logs	655
Traffic.	656
Event	656
Data Leak Prevention	657
Application control	657
Antivirus	658
Web Filter	658
IPS (attack).	658
Email filter	659
Archives (DLP)	659
Network scan.	659
Configuring logging of FortiGate activity on your FortiGate unit	659
Enabling logging within a firewall policy	660
Enabling logging of events	660
Enabling SQL logging	662
Configuring IPS packet logging	662
Configuring NAC quarantine logging.	663
Viewing log messages and archives	664
Viewing log messages from the web-based manager and CLI	665
Quarantine	666
Downloading log messages and viewing them from your computer	667
Viewing log messages using the log table	667
Monitoring the recording activity of logs on the FortiGate unit	668
Customizing the display of log messages	668
Filtering and customizing application control log messages example.	668
Alert email messages.	669
Configuring an alert email message	669
Configuring an alert email for notification of FortiGuard license expiry	671
Log message usage	673
Using log messages to help when issues arise	673
HA log messages indicate lost neighbor information	673
Alert email test configuration issues example	674
How to use log messages to help verify settings and for testing purposes	674
Verifying to see if a network scan was performed example	674
Testing for the FortiGuard license expiry log message example	675
Using diag log test to verify logs are sent to a log device	675
Reports	677
FortiOS reports	677

Configuring a FortiOS report	678
Modifying the default FortiOS report.	678
Configuring charts, datasets, themes and styles for a report	681
Configuring a report layout.	684
Charts	685
Importing images for the report	688
Viewing reports	688
Report example	689
Report for analyzing web activity on the FortiGate unit	689

Chapter 5 Troubleshooting 697

Life of a Packet 699	699
Stateful inspection	699
Connections over connectionless	700
What is a session?	700
Differences between connections and sessions.	700
Flow inspection.	701
Proxy inspection	702
Comparison of inspection layers	702
FortiOS functions and security layers	703
Packet flow	703
Packet inspection (Ingress)	704
Interface	704
DoS sensor.	705
IP integrity header checking	705
IPsec	705
Destination NAT (DNAT)	705
Routing	705
Policy lookup.	705
Session tracking	706
User authentication	706
Management traffic.	706
SSL VPN traffic.	706
Session helpers	706
Flow-based inspection engine	706
Proxy-based inspection engine	707
IPsec	707
Source NAT (SNAT)	707
Routing	707
Egress	707
Example 1: client/server connection	707
Example 2: Routing table update.	709

Example 3: Dialup IPsec VPN with application control.	710
Troubleshooting process	713
Establish a baseline	713
Define the problem	714
Gathering Facts	715
Search for a solution	715
Technical Documentation	715
Release Notes	715
Knowledge Base	715
Fortinet Technical Discussion Forums	715
Fortinet Training Services Online Campus	716
Create a troubleshooting plan	716
Providing Supporting Elements	716
Obtain any required additional equipment	716
Ensure you have administrator level access to required equipment	717
Contact Fortinet customer support for assistance	717
Troubleshooting tools	719
FortiOS diagnostics	719
Check date and time	720
Resource usage	720
Proxy operation	724
Hardware NIC	725
Conserve mode	727
Traffic trace.	728
Session table.	728
Firewall session setup rate.	732
Finding object dependencies	733
Flow trace	734
Packet sniffing and packet capture	737
FA2 and NP2 based interfaces.	740
Debug command.	741
Other commands.	743
FortiGate ports	745
Diagnostic commands	746
FortiAnalyzer/FortiManager ports	746
FortiGuard troubleshooting.	746
Troubleshooting process for FortiGuard updates	746
FortiGuard server settings	747
FortiGuard URL rating	749
Technical Support Organization Overview	751
Fortinet Global Customer Services Organization.	751

Creating an account	752
Registering a device	753
Reporting problems	755
Logging online tickets	755
Following up on online tickets	757
Telephoning a technical support center	757
Assisting technical support.	758
Support priority levels	758
Priority 1	758
Priority 2	758
Priority 3	759
Priority 4	759
Return material authorization process	759
Troubleshooting common issues	763
How to troubleshoot cabling	763
How to troubleshoot no Internet connection	764
How to troubleshoot intermittent connection problems	766
How to troubleshoot a connection in Transparent mode	767
Common issues and questions.	768
Check hardware connections	770
Check FortiOS network settings	770
Check CPU and memory resources	772
Check modem status	773
Run ping and traceroute	773
Check the logs	777
Verify the contents of the routing table (in NAT mode)	778
Check the bridging information in Transparent mode	778
Perform a sniffer trace	780
Debug the packet flow	782
Check number of sessions used by UTM proxy	783
Examine the firewall session list	787
Checking wireless information	788
Other diagnose commands	788
Troubleshooting advanced	789
Traffic shaping issues	789
Use traffic shapers to limit traffic in testing and network simulations	790
Monitoring traffic	790
Displaying configured traffic shaping	790
Troubleshooting protocols and users using traffic shaping	791
Displaying current bandwidth and dropped packets for a traffic shaper	792

User and administrator logon issues	793
User logon issues	793
Administrator logon issues	795
IPsec VPN issues	797
VPN negotiations appear to be slow	798
VPN tunnel proposal will not connect	798
VPN Tunnel up but no traffic going over it	801
Other useful VPN IKE related commands	801
Logging	802
Other diagnose commands.	803
Troubleshooting 'get' commands	805
exec tac report	806
get firewall iprope appctrl	808
get firewall iprope list	809
get firewall proute	811
get firewall shaper	812
get hardware cpu	813
get hardware nic	815
get hardware memory	817
get hardware npu list	819
get hardware npu performance	821
get hardware npu status	823
get hardware status	825
get ips session	826
get router info kernel	827
get router info routing-table all	828
get system arp	829
get system auto-update status.	830
get system auto-update versions	832
get system ha status	834
get system performance firewall	836
get system performance status	838
get system performance top	839
get system session-helper	841
get system session-info full-stat	842
get system session-info list.	847
get system session-info ttl	850

get system startup-error-log	851
get system status.	852
get test	853
Syntax	853
Parameters	853
get test urlfilter	857
get vpn ipsec stats crypto	859
get vpn ipsec stats tunnel	860
get vpn ipsec tunnel details	861
get vpn ipsec tunnel summary	863
get vpn status ssl hw-acceleration-status	864
get vpn status ssl list	865
get webfilter ftgd-statistics	866
get webfilter status	868
Troubleshooting bootup and FSSO	871
FortiGate unit bootup issues	871
Basic bootup troubleshooting	871
Advanced bootup troubleshooting.	871
A. You have text on the screen, but you have problems	872
B. You do not see the boot options menu	872
C. You have problems with the console text.	872
D. You have visible power problems.	873
E. You have a suspected defective FortiGate unit.	873
FSSO issues	873
A. Initial information gathering	874
B. The CA is not running and not connected	874
C. The CA is running but not connected.	874
D. The CA is connected	874
E. There are at least some users logged on	875
F. Test user does not appear on the FSSO list	875

Chapter 6	UTM Guide	877
	UTM overview	879
	UTM components	879
	AntiVirus	879
	Intrusion Protection System (IPS)	879
	Anomaly protection (DoS policies)	880
	One-armed IDS (sniffer policies)	880
	Web filtering	880
	Email filtering	880
	Data Leak Prevention (DLP)	880
	Application Control (for example, IM and P2P)	880
	UTM profiles/lists/sensors	881
	Network defense	883
	Monitoring	883
	Blocking external probes	883
	Address sweeps	884
	Port scans	884
	Probes using IP traffic options	884
	Evasion techniques	886
	Defending against DoS attacks	888
	The “three-way handshake”	888
	SYN flood	888
	SYN spoofing	889
	DDoS SYN flood	890
	Configuring the SYN threshold to prevent SYN floods	890
	SYN proxy	890
	Other flood types	891
	Traffic inspection	891
	IPS signatures	891
	Suspicious traffic attributes	892
	DoS policies	892
	Application control	892
	Content inspection and filtering	893
	AntiVirus	893
	FortiGuard Web Filtering	893
	Email filter	894
	DLP	894

AntiVirus	895
Antivirus concepts	895
How antivirus scanning works	895
Antivirus scanning order	896
Antivirus databases	898
Antivirus techniques	899
FortiGuard Antivirus	899
Enable antivirus scanning	900
Viewing antivirus database information	900
Changing the default antivirus database.	900
Overriding the default antivirus database	901
Adding the antivirus profile to a security policy	902
Configuring the scan buffer size	902
Configuring archive scan depth	902
Configuring a maximum allowed file size	903
Configuring client comforting	904
Enable the file quarantine.	905
General configuration steps	905
Configuring the file quarantine	905
Viewing quarantined files.	906
Downloading quarantined files.	906
Enable grayware scanning	906
Testing your antivirus configuration	906
Antivirus examples	907
Configuring simple antivirus protection	907
Protecting your network against malicious email attachments	908
AntiVirus interface reference	909
Profile.	910
Virus Database	911
Email filter	913
Email filter concepts	913
Email filter techniques	913
Order of spam filtering	915
Enable email filter.	916
Configure email traffic types to inspect	916
Configure the spam action	916
Configure the tag location	917
Configure the tag format	917

Configure FortiGuard email filters	918
Enabling FortiGuard IP address checking	918
Enabling FortiGuard URL checking	918
Enabling FortiGuard phishing URL detection	918
Enabling FortiGuard email checksum checking	919
Enabling FortiGuard spam submission	919
Configure local email filters.	920
Enabling IP address black/white list checking	920
Enabling HELO DNS lookup	921
Enabling email address black/white list checking	922
Enabling return email DNS checking.	923
Enabling banned word checking.	923
How content is evaluated	924
Email filter examples	926
Configuring simple antispam protection	926
Blocking email from a user.	927
Email Filter interface reference	928
Profile.	930
Banned Word.	933
IP Address	936
E-mail Address.	939
Intrusion protection	943
IPS concepts	943
Anomaly-based defense	943
Signature-based defense	943
Enable IPS scanning	945
General configuration steps	945
Creating an IPS sensor.	945
Creating an IPS filter	946
Updating predefined IPS signatures	948
Viewing and searching predefined IPS signatures.	948
Creating a signature entry	948
Creating a custom IPS signature.	949
Custom signature syntax and keywords.	949
IPS processing in an HA cluster	964
Active-passive	964
Active-active	965

Configure IPS options	965
Configuring the IPS engine algorithm	965
Configuring the IPS engine-count	965
Configuring fail-open.	965
Configuring the session count accuracy	966
Configuring the IPS buffer size.	966
Configuring protocol decoders.	966
Configuring security processing modules	966
Enable IPS packet logging	967
IPS examples	968
Configuring basic IPS protection.	968
Using IPS to protect your web server	969
Create and test a packet logging IPS sensor	971
Creating a custom signature to block access to example.com	972
Creating a custom signature to block the SMTP “vrfy” command	974
Configuring a Fortinet Security Processing module	975
Intrusion Protection interface reference	979
IPS Sensor	980
DoS sensor.	987
Predefined	991
Custom.	995
Protocol Decoder	996
Web filter	997
Web filter concepts.	997
Different ways of controlling access	999
Order of web filtering.	999
Web content filter.	999
General configuration steps	1000
Creating a web filter content list	1000
How content is evaluated	1001
Enabling the web content filter and setting the content threshold.	1002
URL filter	1003
URL filter actions.	1003
General configuration steps	1005
Creating a URL filter list	1006
Configuring a URL filter list.	1006
SafeSearch	1006

Advanced web filter configuration	1007
ActiveX filter	1007
Cookie filter.	1007
Java applet filter	1007
Web resume download block	1007
Block Invalid URLs	1008
HTTP POST action	1008
Web filtering example	1008
School district	1009
Web Filter interface reference	1012
Profile.	1012
Browser cookie-based FortiGuard Web Filtering overrides	1016
URL Filter.	1017
Local Ratings.	1021
FortiGuard Web Filter	1023
Before you begin	1023
FortiGuard Web Filter and your FortiGate unit	1024
Order of web filtering.	1024
Enable FortiGuard Web Filter.	1026
General configuration steps	1026
Configuring FortiGuard Web Filter settings	1026
Configuring FortiGuard Web Filter usage quotas	1027
Checking quota usage	1029
Advanced FortiGuard Web Filter configuration	1029
Provide Details for Blocked HTTP 4xx and 5xx Errors.	1029
Rate Images by URL (blocked images will be replaced with blanks)	1029
Allow Websites When a Rating Error Occurs	1029
Strict Blocking	1029
Rate URLs by Domain and IP Address	1030
Block HTTP Redirects by Rating.	1030
Daily log of remaining quota	1030
Add or change FortiGuard Web Filter ratings	1030
Create FortiGuard Web Filter overrides	1031
Understanding administrative and user overrides	1031
Customize categories and ratings	1031
Creating local categories.	1031
Customizing site ratings	1032
FortiGuard Web Filter examples	1032
Configuring simple FortiGuard Web Filter protection	1032
School district	1033

Data leak prevention	1035
Data leak prevention concepts	1035
DLP sensor	1035
DLP filter	1035
Fingerprint	1036
File filter	1036
File size	1036
Regular expression	1036
Advanced rule	1036
Compound rule	1036
Enable data leak prevention	1036
General configuration steps	1036
Creating a DLP sensor	1037
Adding filters to a DLP sensor	1037
DLP document fingerprinting	1040
Fingerprinted Documents	1041
File filter	1042
General configuration steps	1042
Creating a file filter list	1043
Creating a file pattern	1043
Creating a file type	1043
Advanced rules	1044
Understanding the default advanced rules	1044
Creating advanced rules	1045
Compound rules	1045
Understanding the default compound rules	1045
Creating compound rules	1046
DLP archiving	1046
DLP examples	1047
Blocking sensitive email messages	1047
Data Leak Prevention interface reference	1049
Sensor	1049
Document Fingerprinting	1054
File Filter	1056
DLP archiving	1060
Application control	1061
Application control concepts	1061

Enable application control	1062
General configuration steps	1062
Creating an application sensor.	1062
Adding applications to an application sensor	1062
Understanding the default application sensor	1066
Viewing and searching the application list	1066
Application traffic shaping	1067
Enabling application traffic shaping	1067
Reverse direction traffic shaping.	1067
Shaper re-use	1068
Application control monitor.	1068
Enabling application control monitor.	1069
Application control packet logging	1070
Application considerations	1070
IM applications	1071
Skype.	1071
Application control examples.	1071
Blocking all instant messaging.	1071
Allowing only software updates	1072
Application Control interface reference	1073
Application Sensor	1074
Application List.	1078
DoS policy	1081
DoS policy concepts	1081
Enable DoS.	1081
Creating and configuring a DoS sensor	1082
Creating a DoS policy	1083
Apply an IPS sensor to a DoS policy.	1084
DoS example	1084
DoS Policy interface reference	1085
Endpoint Control and monitoring	1089
Endpoint Control overview	1089
User experience	1089
Configuration overview.	1091
Configuring FortiClient required version and download location.	1091
About application detection and control	1093
FortiClient application rules	1093
Other application rules	1093
The All application rule.	1093
About predefined profiles.	1094

Creating an endpoint control profile	1094
Setting endpoint FortiClient requirements	1094
Setting the default action for applications	1096
Adding application detection entries.	1096
Viewing the application database	1097
Enabling Endpoint Control in firewall policies	1098
Monitoring endpoints.	1099
Endpoint status	1099
Endpoint Application Usage	1100
Endpoint Traffic	1100
Modifying Endpoint Security replacement pages	1100
Example	1101
Configuring FortiClient download source and required version	1101
Creating an endpoint control profile	1102
Configuring FortiClient application detection entries	1102
Configuring application detection entries for other applications	1102
Configuring the firewall policy	1104
Endpoint Control interface reference.	1104
Profile.	1105
Application Database	1107
Client Installers.	1110
Vulnerability Scan	1113
Overview	1113
Selecting assets to scan	1113
Discovering assets	1113
Adding assets manually	1114
Requirements for authenticated scanning	1116
Configuring scans	1117
Viewing scan results	1121
Viewing scan logs	1121
Viewing Executive Summary graphs.	1122
Creating reports	1122
Viewing reports.	1123
Vulnerability Scan interface reference	1123
Asset Definition.	1124
Scan Schedule	1125
Vulnerability Result.	1126
Sniffer policy	1127
Sniffer policy concepts	1127
The sniffer policy list	1127
Before you begin	1128

Enable one-arm sniffing	1129
General configuration steps	1129
Designating a sniffer interface	1130
Creating a sniffer policy	1130
Sniffer example	1131
An IDS sniffer configuration	1131
Sniffer Policy interface reference	1134
Other UTM considerations	1137
UTM and Virtual domains (VDOMs)	1137
Conserve mode.	1137
The AV proxy	1137
Entering and exiting conserve mode.	1138
Conserve mode effects	1138
Configuring the av-failopen command.	1139
SSL content scanning and inspection	1139
Setting up certificates to avoid client warnings	1140
SSL content scanning and inspection settings	1141
Viewing and saving logged packets	1144
Configuring packet logging options	1144
Using wildcards and Perl regular expressions	1145
Protocol Options interface reference.	1148
ICAP interface reference	1150
ICAP profile.	1151
ICAP server.	1153
Profile Group interface reference.	1154
Profile Group configuration settings	1154
Monitor interface reference.	1155
AV Monitor	1156
Intrusion Monitor	1156
Web Monitor	1157
Email Monitor.	1157
Archive & Data Leak Monitor.	1158
Application Monitor	1158
FortiGuard Quota.	1159
Endpoint Monitor.	1159

Chapter 7 User Authentication **1161**

Introduction to authentication	1163
What is authentication?	1163

Methods of authentication	1163
Local password authentication	1164
Server-based password authentication	1164
Certificate-based authentication	1165
Two-factor authentication	1166
Types of authentication.	1166
Firewall authentication (identity-based policies)	1166
VPN authentication.	1168
User's view of authentication.	1169
Web-based user authentication	1169
VPN client-based authentication.	1170
FortiGate administrator's view of authentication.	1170
Authentication and User in the web-based manager	1171
User	1171
Local user accounts	1171
IM users	1174
Authentication settings.	1176
User groups	1176
User Group.	1177
Firewall user groups	1179
Fortinet Single Sign-On (FSSO) user groups.	1180
SSL VPN user groups	1180
Dynamically assigning VPN client IP addresses from a user group.	1181
Remote	1181
Administrators	1182
RADIUS.	1182
LDAP	1184
TACACS+	1187
FortiToken	1189
FortiToken configuration settings	1189
Fortinet Single Sign On Agent (FSSO)	1191
Fortinet Single Sign-on Agent configuration settings	1192
PKI	1194
Peer users and peer groups	1195
Monitor	1196
Firewall monitor list.	1196
IM user monitor list.	1197
The Banned User list	1198
Authentication servers	1201
FortiAuthenticator servers	1201

RADIUS servers	1201
Configuring the FortiGate unit to use a RADIUS server	1204
LDAP servers	1207
Components and topology.	1207
LDAP directory organization	1208
Configuring the FortiGate unit to use an LDAP server	1209
Example — wildcard admin accounts - CLI	1211
Example of LDAP to allow Dial-in through member-attribute - CLI	1213
Troubleshooting LDAP	1214
TACACS+ servers	1215
Configuring a TACACS+ server on the FortiGate unit	1216
FSSO servers	1216
RSA ACE (SecurID) servers.	1217
Components	1217
Configuring the SecurID system	1217
Users and user groups	1223
Users	1223
Local users	1224
PKI or peer users.	1227
Two-factor authentication	1228
FortiToken	1230
Monitoring users	1234
User groups	1235
Firewall user groups	1236
FSSO user groups	1240
Configuring Peer user groups	1240
Viewing, editing and deleting user groups	1240
Configuring authenticated access	1243
Authentication timeout	1243
Security authentication timeout	1243
SSL VPN authentication timeout.	1243
Password policy	1244
Authentication protocols	1246

Authentication in security policies	1246
Enabling authentication protocols	1247
Authentication replacement messages	1248
Access to the Internet	1250
Configuring authentication security policies	1250
Identity-based policy	1253
FSSO authentication	1254
NTLM authentication	1255
Certificate authentication	1256
Dynamic profile	1257
Restricting number of concurrent user logons	1257
VPN authentication	1257
Configuring authentication of SSL VPN users	1258
Configuring authentication of remote IPsec VPN users	1258
Configuring authentication of PPTP VPN users and user groups	1260
Configuring authentication of L2TP VPN users/user groups	1261
Certificate-based authentication	1263
What is a security certificate?	1263
Certificates overview	1264
Certificates and protocols	1264
IPsec VPNs and certificates	1265
Certificate types on the FortiGate unit	1265
Certificate signing	1266
Managing X.509 certificates	1266
Generating a certificate signing request	1267
Generating certificates with CA software	1268
Obtaining a signed server certificate from an external CA	1269
Installing a CA root certificate and CRL to authenticate remote clients	1270
Troubleshooting certificates	1271
Online updates to certificates and CRLs	1272
Backing up and restoring local certificates	1274
Configuring certificate-based authentication	1275
Authenticating administrators with security certificates	1275
Authenticating SSL VPN users with security certificates	1275
Authenticating IPsec VPN users with security certificates	1276
Example — Generate a CSR on the FortiGate unit	1277
Example — Generate and Import CA certificate with private key pair on OpenSSL	1278
Assumptions	1278
Generating and importing the CA certificate and private key	1278

Example — Generate an SSL certificate in OpenSSL	1279
Assumptions	1279
Generating a CA signed SSL certificate	1279
Generating a self-signed SSL certificate.	1280
Import the SSL certificate into FortiOS	1280
FSSO integration with Windows AD or Novell	1283
Introduction to FSSO	1283
Using FSSO in a Windows AD environment	1284
Using FSSO in a Novell eDirectory environment.	1290
FSSO for Windows AD	1290
FSSO components for Windows AD.	1290
Standard versus Advanced mode	1291
Installing FSSO for Windows AD.	1291
Updating FSSO with Windows AD.	1294
Configuring Fortinet Single Sign On with Windows AD	1295
Configuring Windows AD server user groups	1295
Configuring Collector agent settings.	1296
Configuring Directory Access settings.	1299
Configuring the Ignore User List	1300
Configuring FortiGate group filters.	1301
Configuring FSSO ports	1302
Configuring alternate user IP address tracking	1303
Viewing FSSO component status	1303
Selecting Domain Controllers and working mode for monitoring	1304
FSSO for Novell eDirectory.	1304
FSSO components for Novell eDirectory	1305
Installing FSSO for Novell	1305
Configuring Fortinet Single Sign On with Novell networks.	1305
Configuring FSSO on FortiGate units	1308
Configuring LDAP server access.	1308
Specifying your Collector agents or Novell eDirectory agents.	1310
Selecting Windows user groups (LDAP only)	1311
Viewing information imported from the Windows AD server.	1311
Creating Fortinet Single Sign-On (FSSO) user groups.	1312
Creating security policies	1312
Enabling guest access through FSSO security policies	1315
FortiOS FSSO log messages	1315
Enabling authentication event logging	1315
Viewing FSSO log messages.	1316
Testing FSSO.	1317

Troubleshooting FSSO	1318
General troubleshooting tips for FSSO	1318
User status “Not Verified” on the Collector agent	1319
After initial configuration, there is no connection to the Collector agent	1319
Collector Agent service freezing and shutting down.	1320
FortiGate performance is slow on a large network with many users	1320
Users from the Windows AD network are not able to access the network	1321
Users on a particular computer (IP address) can not access the network	1321
Guest users do not have access to network.	1321
Can’t find the DCagent service	1322
User logon events not received by FSSO Collector agent.	1322
User list from Windows AD is empty.	1322
Mac OS X users can’t access external resources after waking from sleep mode.	1323
Dynamic profiles and end points	1325
Overview	1325
When to use FSSO or dynamic profiles	1326
End points	1327
Dynamic profiles and security policies.	1327
Accounting system RADIUS configuration.	1330
User context list	1330
Accepting sessions only from dynamic profile users	1331
Configuring dynamic profile	1332
Make dynamic profiles visible	1333
RADIUS server configuration for dynamic profiles.	1333
Configuring dynamic profile-based security policies.	1339
Configuration concepts	1340
Configuring end points	1341
Configuring end points - CLI	1341
Controlling MMS service access based on a user’s end point - FortiCarrier Only	1342
Blocking access to the network based on end points - FortiOS Carrier only	1344
Extracting carrier end points for notifications - FortiOS Carrier only	1347
Timeout options	1349
Log settings	1350
Carrier end point filters and blocking.	1351
Controlling access to MMS services based on a user’s carrier end point.	1352
Blocking network access for IP addresses based on carrier end points	1354
Troubleshooting dynamic profiles	1356
General dynamic profile troubleshooting	1356
Dynamic profile related diag commands.	1357

Monitoring authenticated users	1359
Monitoring firewall users	1359
Monitoring SSL VPN users	1359
Monitoring IPsec VPN users	1360
Examples and Troubleshooting	1361
Firewall authentication example	1361
Overview	1361
Creating a locally-authenticated user account.	1362
Creating a RADIUS-authenticated user account.	1362
Creating user groups.	1363
Defining policy addresses	1365
Creating security policies	1365
LDAP Dial-in using member-attribute	1367
Dynamic Profile example	1369
Assumptions	1369
Topology	1369
General configuration	1370
Configuring RADIUS	1370
Configuring FortiGate interfaces	1370
Configuring dynamic profile RADIUS server on FortiGate	1372
Configuring FortiGate regular and dynamic profile security policies	1373
Testing	1376
Troubleshooting	1377
 Chapter 8 IPsec VPNs	 1379
IPsec VPN concepts	1381
VPN tunnels	1381
VPN gateways	1382
Clients, servers, and peers	1384
Encryption	1385
Authentication	1385
Phase 1 and Phase 2 settings	1386
Phase 1.	1386
Phase 2.	1386
Security Association	1387

IPsec VPN Overview	1389
Types of VPNs	1389
Route-based VPNs.	1389
Policy-based VPNs.	1390
Comparing policy-based or route-based VPNs	1390
Planning your VPN	1390
Network topologies	1391
General preparation steps	1392
How to use this guide to configure an IPsec VPN	1392
IPsec VPN in the web-based manager	1393
Auto Key (IKE)	1393
Phase 1 configuration	1394
Phase 2 configuration	1399
FortiClient VPN.	1402
Manual Key.	1403
Concentrator	1404
IPsec Monitor.	1405
Auto Key phase 1 parameters	1407
Overview	1407
Defining the tunnel ends	1408
Choosing main mode or aggressive mode.	1408
Choosing the IKE version.	1409
Authenticating the FortiGate unit.	1409
Authenticating the FortiGate unit with digital certificates	1409
Authenticating the FortiGate unit with a pre-shared key.	1410
Authenticating remote peers and clients	1412
Enabling VPN access for specific certificate holders	1412
Enabling VPN access by peer identifier	1414
Enabling VPN access with user accounts and pre-shared keys.	1415
Defining IKE negotiation parameters	1417
Generating keys to authenticate an exchange	1417
Defining IKE negotiation parameters.	1418
Using XAuth authentication.	1422
Using the FortiGate unit as an XAuth server	1422
Using the FortiGate unit as an XAuth client	1423
Phase 2 parameters	1425
Basic phase 2 settings	1425

Advanced phase 2 settings	1425
P2 Proposals	1426
Replay detection	1426
Perfect forward secrecy (PFS)(.	1426
Keylife	1426
Auto-negotiate	1426
Autokey Keep Alive.	1427
DHCP-IPsec	1427
Quick mode selectors	1427
Configure the phase 2 parameters	1428
Specifying the phase 2 parameters	1428
Defining VPN security policies	1431
Defining policy addresses	1431
Defining VPN security policies	1432
Defining an IPsec security policy for a policy-based VPN	1433
Defining security policies for a route-based VPN	1436
Gateway-to-gateway configurations	1437
Configuration overview	1437
General configuration steps	1439
Configuring the two VPN peers	1439
Configuring Phase 1 and Phase 2 for both peers	1439
Creating security policies	1440
How to work with overlapping subnets	1444
Solution for route-based VPN	1445
Solution for policy-based VPN	1447
Testing	1449
Hub-and-spoke configurations	1453
Configuration overview	1453
Hub-and-spoke infrastructure requirements	1454
Spoke gateway addressing	1454
Protected networks addressing	1454
Authentication	1455
Configure the hub	1455
Define the hub-spoke VPNs	1455
Define the hub-spoke security policies	1456
Configuring communication between spokes (policy-based VPN)	1458
Configuring communication between spokes (route-based VPN)	1458
Configure the spokes	1460
Configuring security policies for hub-to-spoke communication	1460
Configuring security policies for spoke-to-spoke communication.	1462

Dynamic spokes configuration example	1463
Configure the hub (FortiGate_1)	1463
Configure the spokes	1466
Dynamic DNS configuration	1469
Dynamic DNS over VPN concepts	1469
Dynamic DNS (DDNS)	1469
Dynamic DNS over VPN	1470
Dynamic DNS topology.	1471
Assumptions	1472
General configuration steps	1472
Configure the dynamically-addressed VPN peer.	1473
Configuring branch_2 VPN tunnel settings	1473
Configuring branch_2 security policies	1474
Configure the fixed-address VPN peer	1478
Configuring branch_1 VPN tunnel settings	1478
Configuring branch_1 security policies	1479
Testing	1482
FortiClient dialup-client configurations	1483
Configuration overview	1483
Peer identification	1484
Automatic configuration of FortiClient dialup clients	1484
One button FortiGate - to - FortiClient Phase1 VPN.	1485
Using virtual IP addresses	1486
FortiClient dialup-client infrastructure requirements	1487
FortiClient-to-FortiGate VPN configuration steps	1488
Configure the FortiGate unit	1488
Configuring FortiGate unit VPN settings	1488
Configuring the FortiGate unit as a VPN policy server.	1491
Configuring DHCP service on the FortiGate unit.	1491
Configure the FortiClient Endpoint Security application	1493
Configuring FortiClient to work with VPN policy distribution	1493
Configuring FortiClient manually	1493
Adding XAuth authentication	1494
FortiClient dialup-client configuration example	1495
Configuring FortiGate_1	1495
Configuring the FortiClient Endpoint Security application	1499
FortiGate dialup-client configurations	1501
Configuration overview	1501
FortiGate dialup-client infrastructure requirements	1503
FortiGate dialup-client configuration steps	1504

Configure the server to accept FortiGate dialup-client connections	1504
Configure the FortiGate dialup client	1506
Supporting IKE Mode config clients	1509
Automatic configuration overview	1509
IKE Mode Config overview	1509
Configuring IKE Mode Config	1510
Configuring an IKE Mode Config client	1510
Configuring an IKE Mode Config server	1510
Example: FortiGate unit as IKE Mode Config server	1512
Example: FortiGate unit as IKE Mode Config client	1513
Internet-browsing configuration	1515
Configuration overview	1515
Creating an Internet browsing security policy	1516
Routing all remote traffic through the VPN tunnel	1517
Configuring a FortiGate remote peer to support Internet browsing	1517
Configuring a FortiClient application to support Internet browsing	1518
Redundant VPN configurations	1519
Configuration overview	1519
General configuration steps	1520
Configure the VPN peers - route-based VPN	1521
Redundant route-based VPN configuration example	1523
Configuring FortiGate_1	1523
Configuring FortiGate_2	1529
Partially-redundant route-based VPN example	1534
Configuring FortiGate_1	1536
Configuring FortiGate_2	1538
Creating a backup IPsec interface	1541
Transparent mode VPNs	1543
Configuration overview	1543
Transparent VPN infrastructure requirements	1546
Configure the VPN peers	1547
Manual-key configurations	1551
Configuration overview	1551
Specify the manual keys for creating a tunnel	1552
IPv6 IPsec VPNs	1555
Overview of IPv6 IPsec support	1555
Certificates	1555

Configuring IPv6 IPsec VPNs.	1556
Phase 1 configuration	1556
Phase 2 configuration	1556
Security policies	1556
Routing	1556
Site-to-site IPv6 over IPv6 VPN example	1557
Configure FortiGate A interfaces.	1557
Configure FortiGate A IPsec settings	1558
Configure FortiGate A security policies	1558
Configure FortiGate A routing	1559
Configure FortiGate B	1559
Site-to-site IPv4 over IPv6 VPN example	1560
Configure FortiGate A interfaces.	1561
Configure FortiGate A IPsec settings	1561
Configure FortiGate A security policies	1561
Configure FortiGate A routing	1562
Configure FortiGate B	1562
Site-to-site IPv6 over IPv4 VPN example	1563
Configure FortiGate A interfaces.	1563
Configure FortiGate A IPsec settings	1564
Configure FortiGate A security policies	1564
Configure FortiGate A routing	1565
Configure FortiGate B	1565
L2TP and IPsec (Microsoft VPN)	1567
Overview	1567
Layer 2 Tunneling Protocol (L2TP)	1567
Assumptions	1568
Configuring the FortiGate unit	1568
Configuring L2TP users and firewall user group	1568
Configuring L2TP.	1569
Configuring IPsec	1570
Configuring security policies	1572
Configuring the Windows PC.	1573
Troubleshooting	1574
Quick checks.	1575
Mac OS X and L2TP	1575
Setting up logging	1575
Understanding the log messages	1576
Using the FortiGate unit debug commands	1577
GRE over IPsec (Cisco VPN) configurations	1579
Overview	1579

Configuring the FortiGate unit	1580
Enabling overlapping subnets	1580
Configuring the IPsec VPN	1580
Configuring the GRE tunnel	1582
Configuring security policies	1582
Configuring routing	1585
Configuring the Cisco router	1585
Troubleshooting	1585
Quick checks	1586
Setting up logging	1586
Understanding the log messages	1587
Using diagnostic commands	1587
Protecting OSPF with IPsec	1589
Overview	1589
OSPF over IPsec configuration	1590
Configuring the IPsec VPN	1590
Configuring static routing	1591
Configuring OSPF	1591
Creating a redundant configuration	1595
Adding the second IPsec tunnel	1595
Adding the OSPF interface	1595
Hardware offloading and acceleration	1597
Overview	1597
IPsec session offloading requirements	1597
Packet offloading requirements	1598
IPsec encryption offloading	1598
HMAC check offloading	1598
IPsec offloading configuration examples	1599
Accelerated route-based VPN configuration	1599
Accelerated policy-based VPN configuration	1601
Monitoring and troubleshooting	1603
Monitoring VPN connections	1603
Monitoring connections to remote peers	1603
Monitoring dialup IPsec connections	1603
Testing VPN connections	1604
Testing VPN connection	1604
Troubleshooting VPN connections	1605
Logging VPN events	1607
VPN troubleshooting tips	1609
General troubleshooting tips	1609
A word about NAT devices	1610

Chapter 9	SSL VPN	1611
	Introduction to SSL VPN	1613
	SSL VPN modes of operation	1613
	Web-only mode	1614
	Tunnel mode	1614
	Port forwarding mode	1615
	SSL VPN and IPv6	1616
	Traveling and security	1616
	Host check	1616
	Cache cleaning.	1616
	Basic Configuration	1617
	User accounts and groups	1617
	Authentication	1618
	IP addresses for users	1618
	Authentication of remote users	1619
	Configuring SSL VPN web portals	1620
	SSL connection configuration	1621
	Portal configuration	1621
	Tunnel mode settings	1625
	The Session Information widget	1628
	The Bookmarks widget.	1628
	The Connection Tool widget	1630
	Configuring security policies	1631
	Firewall addresses	1631
	Create an SSL VPN security policy	1632
	Create a tunnel mode security policy	1634
	Split tunnel Internet browsing policy	1635
	Enabling a connection to an IPsec VPN	1636

Additional configuration options	1637
Routing in tunnel mode	1638
Changing the port number for web portal connections	1638
SSL offloading	1638
Customizing the web portal login page	1639
Host Check.	1640
Windows OS check	1643
Configuring cache cleaning	1643
Configuring virtual desktop	1644
Configuring client OS Check.	1646
Adding WINS and DNS services for clients	1647
Setting the idle timeout setting	1647
SSL VPN logs	1647
Monitoring active SSL VPN sessions	1648
Troubleshooting	1649
The SSL VPN client	1651
FortiClient.	1651
Downloading the SSL VPN tunnel mode client.	1652
Tunnel mode client configuration.	1653
Uninstalling the tunnel mode client.	1653
Setup examples	1655
Secure internet browsing.	1655
Creating an SSL VPN IP pool and SSL VPN web portal.	1655
Creating the SSL VPN user and user group	1656
Creating a static route for the remote SSL VPN user	1656
Creating security policies	1657
Results	1657
Split Tunnel.	1658
Creating a firewall address for the head office server	1658
Results	1660
Multiple user groups with different access permissions example	1661
General configuration steps	1661
Creating the firewall addresses	1662
Creating the web portals.	1663
Creating the user accounts and user groups	1664
Creating the security policies	1665
Create the static route to tunnel mode clients.	1667
Enabling SSL VPN operation.	1668

Chapter 10 Advanced Routing 1669

Advanced Static routing 1671

Routing concepts	1671
Routing in VDOMs	1671
Default route	1672
Routing table	1672
Building the routing table	1678
Static routing security	1679
Multipath routing and determining the best route	1681
Route priority	1682
Troubleshooting static routing	1683
ECMP route failover and load balancing	1685
Route priority	1685
Equal-Cost Multi-Path (ECMP).	1686
Configuring interface status detection for gateway load balancing	1687
Configuring spillover or usage-based ECMP	1689
Configuring weighted static route load balancing	1691
Static routing tips.	1692
Policy Routing	1693
Adding a policy route.	1694
Moving a policy route	1696
Transparent mode static routing	1696

Dynamic Routing Overview 1699

What is dynamic routing?	1699
Comparing static and dynamic routing	1699
Dynamic routing protocols.	1700
Minimum configuration for dynamic routing	1702
Comparison of dynamic routing protocols	1702
Features of dynamic routing protocols.	1702
When to adopt dynamic routing	1705
Choosing a routing protocol	1707
Dynamic routing terminology.	1708
IPv6 in dynamic routing	1713

Routing Information Protocol (RIP) 1715

RIP background and concepts.	1715
Background	1715
Parts and terminology of RIP.	1716
How RIP works.	1721

Troubleshooting RIP	1726
Routing Loops	1726
Split horizon and Poison reverse updates	1729
Debugging IPv6 on RIPng	1729
RIP routing examples.	1730
Simple RIP example	1730
Network layout and assumptions	1730
General configuration steps	1732
Configuring the FortiGate units system information	1732
Configuring FortiGate unit RIP router information	1740
Configuring other networking devices	1744
Testing network configuration	1744
RIPng — RIP and IPv6	1744
Network layout and assumptions	1745
General configuration steps	1746
Configuring the FortiGate units system information	1746
Configuring RIPng on FortiGate units	1749
Configuring other network devices.	1750
Testing the configuration.	1750
Border Gateway Protocol (BGP)	1751
BGP background and concepts	1751
Background.	1751
Parts and terminology of BGP	1751
BGP and IPv6	1752
Roles of routers in BGP networks	1752
Confederations.	1756
Network Layer Reachability Information (NLRI)	1757
BGP attributes	1757
How BGP works	1760
IBGP versus EBGP.	1760
BGP path determination — which route to use	1760
Troubleshooting BGP.	1763
Clearing routing table entries	1763
Route flap	1763
BGP routing examples	1767
Dual-homed BGP example.	1767
Network layout and assumptions	1769
General configuration steps	1770
Configuring the FortiGate unit	1771
Configuring other networking devices	1778
Testing this configuration	1778

Redistributing and blocking routes in BGP	1780
Network layout and assumptions	1780
Open Shortest Path First (OSPF)	1787
OSPF Background and concepts	1787
Background	1787
The parts and terminology of OSPF	1787
How OSPF works	1794
Troubleshooting OSPF	1799
Clearing OSPF routes from the routing table	1799
Checking the state of OSPF neighbors	1800
Passive interface problems	1800
Timer problems.	1800
Bi-directional Forwarding Detection (BFD).	1801
Authentication issues	1801
DR and BDR election issues	1801
OSPF routing examples	1801
Basic OSPF example	1802
Network layout and assumptions	1802
General configuration steps	1803
Configuring the FortiGate units	1804
Configuring OSPF on the FortiGate units	1806
Configuring other networking devices	1813
Testing network configuration	1813
Advanced inter-area OSPF example	1813
Network layout and assumptions	1813
General configuration steps	1815
Configuring the FortiGate units	1816
Configuring OSPF on the FortiGate units	1819
Configuring other networking devices	1823
Testing network configuration	1823
Controlling redundant links by cost	1823
Adjusting the route costs.	1824
Verifying route redundancy.	1826
Intermediate System To Intermediate System Protocol (IS-IS)	1827
IS-IS background and concepts	1827
Background	1827
How IS-IS works	1827
Troubleshooting IS-IS	1829
Routing Loops	1829
Split horizon and Poison reverse updates	1832

Simple IS-IS example.	1832
General configuration steps	1834
Configuring FortiGate hostnames, interfaces, and default routes	1834
Configuring FortiGate unit IS-IS router information	1838
Configuring other networking devices	1840
Testing network configuration	1840

Router Reference 1841

Static	1841
Static Route	1841
Default route and default gateway	1844
Policy Route	1845
Settings.	1847
Dynamic	1849
RIP	1849
OSPF	1854
BGP	1860
Multicast	1863
Bi-directional Forwarding Detection (BFD).	1867
Monitor	1869
Viewing routing information	1869
Searching the routing monitor table	1871

Chapter 11 Virtual Domains 1873

Virtual Domains 1875

Benefits of Virtual Domains.	1875
Enabling and accessing Virtual Domains.	1877
Enabling Virtual Domains.	1877
Viewing the VDOM list	1879
Global and per-VDOM settings	1881
Resource settings	1888
Virtual Domain Licensing.	1892
Logging in to VDOMs	1894
Configuring Virtual Domains	1895
Creating a Virtual Domain	1895
Disabling a Virtual Domain	1896
Deleting a VDOM.	1897
Removing references to a VDOM	1897
Administrators in Virtual Domains	1898

Virtual Domains in NAT/Route mode	1903
Virtual domains in NAT/Route mode	1903
Changing the management virtual domain.	1903
Configuring interfaces in a NAT/Route VDOM	1904
Configuring VDOM routing	1907
Configuring security policies for NAT/Route VDOMs	1909
Configuring UTM profiles for NAT/Route VDOMs	1910
Configuring VPNs for a VDOM	1910
Example NAT/Route VDOM configuration	1911
Network topology and assumptions	1911
General configuration steps	1912
Creating the VDOMs	1912
Configuring the FortiGate interfaces	1913
Configuring the vdomA VDOM	1915
Configuring the vdomB VDOM	1917
Testing the configuration.	1920
 Virtual Domains in Transparent mode	 1921
Before you begin	1921
Transparent operation mode	1922
Broadcast domains	1922
Forwarding domains	1922
Spanning Tree Protocol	1923
Differences between NAT/Route and Transparent mode	1923
Operation mode differences in VDOMs	1924
Configuring VDOMs in Transparent mode	1924
Switching to Transparent mode	1925
Adding VLAN subinterfaces	1926
Creating security policies	1926
Example of VDOMs in Transparent mode	1926
Network topology and assumptions	1926
General configuration steps	1927
Configuring common items	1927
Creating virtual domains	1928
Configuring the Company_A VDOM	1929
Configuring the Company_B VDOM	1933
Configuring the VLAN switch and router.	1937
Testing the configuration.	1938
 Inter-VDOM routing	 1941
Benefits of inter-VDOM routing.	1941

Getting started with VDOM links	1943
Viewing VDOM links	1943
Creating VDOM links	1944
Deleting VDOM links	1946
Inter-VDOM configurations	1946
Standalone VDOM configuration	1947
Independent VDOMs configuration	1948
Management VDOM configuration	1949
Meshed VDOM configuration	1950
Dynamic routing over inter-VDOM links	1950
HA virtual clusters and VDOM links	1951
What is virtual clustering?	1951
Example of inter-VDOM routing	1953
Network topology and assumptions	1953
General configuration steps	1954
Creating the VDOMs	1954
Configuring the physical interfaces	1955
Configuring the VDOM links	1957
Configuring the firewall and UTM settings	1958
Testing the configuration	1975
Troubleshooting Virtual Domains	1977
VDOM admin having problems gaining access	1977
FortiGate unit running very slowly	1977
General VDOM tips and troubleshooting	1978
Perform a sniffer trace	1978
Debug the packet flow	1980

Chapter 12 High Availability **1983**

Solving the High Availability problem	1987
FortiGate Cluster Protocol (FGCP)	1987
TCP session synchronization	1988
VRRP	1988
An introduction to the FortiGate Clustering Protocol (FGCP)	1991
About the FGCP	1992
FGCP failover protection	1993
Session Failover	1993
Load Balancing	1993
Virtual Clustering	1993
Full Mesh HA	1994
Cluster Management	1994

Configuring a FortiGate unit for FGCP HA operation	1995
Connecting a FortiGate HA cluster.	1996
Active-passive and active-active HA	1997
Active-passive HA (failover protection)	1998
Active-active HA (load balancing and failover protection)	1998
Identifying the cluster and cluster units	1999
Group name	1999
Password.	1999
Group ID	2000
Device failover, link failover, and session failover	2000
Primary unit selection.	2000
Primary unit selection and monitored interfaces.	2002
Primary unit selection and age	2003
Primary unit selection and device priority	2006
Primary unit selection and FortiGate unit serial number	2007
Points to remember about primary unit selection	2008
HA override	2008
Override and primary unit selection	2009
Controlling primary unit selection using device priority and override	2010
Points to remember about primary unit selection when override is enabled	2010
Configuration changes can be lost if override is enabled	2011
Override and disconnecting a unit from a cluster	2012
FortiGate HA compatibility with PPPoE and DHCP	2012
Hard disk configuration and HA	2013
HA Best practices	2013
Heartbeat interfaces	2014
Interface monitoring (port monitoring)	2014
Troubleshooting	2014
FGCP HA terminology	2015
HA web-based manager options.	2018
Configuring and connecting HA clusters	2021
About the procedures in this chapter	2021
Example: NAT/Route mode active-passive HA configuration	2021
Example NAT/Route mode HA network topology	2022
General configuration steps	2022
Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - web-based manager	2023
Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI	2027

Example: Transparent mode active-active HA configuration	2033
Example Transparent mode HA network topology	2033
General configuration steps	2034
Configuring a Transparent mode active-active cluster of two FortiGate-620B units - web-based manager.	2034
Configuring a Transparent mode active-active cluster of two FortiGate-620B units - CLI.	2039
Example: advanced Transparent mode active-active HA configuration	2046
Example Transparent mode HA network topology	2046
Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - web-based manager.	2047
Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - CLI.	2050
Example: converting a standalone FortiGate unit to a cluster	2054
Example: adding a new unit to an operating cluster	2055
Example: replacing a failed cluster unit	2056
Example: HA and 802.3ad aggregated interfaces	2057
HA interface monitoring, link failover, and 802.3ad aggregation	2058
HA MAC addresses and 802.3ad aggregation.	2058
Link aggregation, HA failover performance, and HA mode	2058
General configuration steps	2059
Configuring active-passive HA cluster that includes aggregated interfaces - web-based manager	2059
Configuring active-passive HA cluster that includes aggregate interfaces - CLI.	2065
Example: HA and redundant interfaces	2071
HA interface monitoring, link failover, and redundant interfaces	2071
HA MAC addresses and redundant interfaces.	2071
Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode	2072
Connecting multiple redundant interfaces to one switch while operating in active-active HA mode	2072
General configuration steps	2072
Configuring active-passive HA cluster that includes redundant interfaces - web-based manager	2072
Configuring active-passive HA cluster that includes redundant interfaces - CLI.	2078
Troubleshooting HA clusters	2083
Before you set up a cluster.	2084
Troubleshooting the initial cluster configuration	2084
More troubleshooting information	2086

Configuring and connecting virtual clusters	2089
Virtual clustering overview	2089
Virtual clustering and failover protection	2089
Virtual clustering and heartbeat interfaces	2090
Virtual clustering and HA override	2090
Virtual clustering and load balancing or VDOM partitioning	2091
Configuring HA for virtual clustering	2091
Example: virtual clustering with two VDOMs and VDOM partitioning	2093
Example virtual clustering network topology	2093
General configuration steps	2094
Configuring virtual clustering with two VDOMs and VDOM partitioning - web-based manager	2095
Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI . .	2100
Example: inter-VDOM links in a virtual clustering configuration	2107
Configuring inter-VDOM links in a virtual clustering configuration	2108
Troubleshooting virtual clustering	2109
Configuring and operating FortiGate full mesh HA	2111
Full mesh HA overview	2111
Full mesh HA and redundant heartbeat interfaces	2112
Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces	2113
Example: full mesh HA configuration	2113
FortiGate-620B full mesh HA configuration	2114
Full mesh switch configuration	2114
Full mesh network connections	2114
How packets travel from the internal network through the full mesh cluster and to the Internet	2114
Configuring FortiGate-620B units for HA operation - web-based manager . . .	2115
Configuring FortiGate-620B units for HA operation - CLI	2120
Troubleshooting full mesh HA	2125
Operating a cluster	2127
Operating a cluster	2127
Operating a virtual cluster	2128
Managing individual cluster units using a reserved management interface	2129
Configuring the reserved management interface and SNMP remote management of individual cluster units	2130
The primary unit acts as a router for subordinate unit management traffic	2134
Cluster communication with RADIUS and LDAP servers	2135
Clusters and FortiGuard services	2135
FortiGuard and active-passive clusters	2135
FortiGuard and active-active clusters	2135
FortiGuard and virtual clustering	2136

Clusters and logging	2136
Viewing and managing log messages for individual cluster units	2136
HA log messages.	2138
Example log messages.	2138
Fortigate HA message "HA master heartbeat interface <intf_name> lost neighbor information"	2141
Clusters and SNMP.	2143
SNMP get command syntax for the primary unit	2143
SNMP get command syntax for any cluster unit.	2145
Getting serial numbers of cluster units.	2146
SNMP get command syntax - reserved management interface enabled	2146
Clusters and file quarantine	2146
Cluster members list	2147
Virtual cluster members list.	2149
Viewing HA statistics	2150
Changing the HA configuration of an operating cluster	2151
Changing the HA configuration of an operating virtual cluster.	2151
Changing the subordinate unit host name and device priority.	2152
Upgrading cluster firmware.	2152
Changing how the cluster processes firmware upgrades	2153
Synchronizing the firmware build running on a new cluster unit.	2153
Downgrading cluster firmware	2154
Backing up and restoring the cluster configuration	2155
Monitoring cluster units for failover.	2155
Viewing cluster status from the CLI	2156
Examples.	2158
About the HA cluster index and the execute ha manage command.	2161
Managing individual cluster units	2163
Disconnecting a cluster unit from a cluster	2164
Adding a disconnected FortiGate unit back to its cluster	2165
HA and failover protection	2167
About active-passive failover.	2167
Device failure.	2168
Link failure	2168
Session failover	2168
Primary unit recovery.	2169
About active-active failover.	2169
Device failover	2169

HA heartbeat and communication between cluster units	2170
Heartbeat interfaces	2171
Connecting HA heartbeat interfaces	2172
Heartbeat interfaces and FortiGate switch interfaces	2172
Heartbeat packets and heartbeat interface selection	2172
Interface index and display order	2173
HA heartbeat interface IP addresses.	2173
Heartbeat packet Ethertypes.	2174
Modifying heartbeat timing.	2175
Enabling or disabling HA heartbeat encryption and authentication	2177
Cluster virtual MAC addresses	2177
Changing how the primary unit sends gratuitous ARP packets after a failover	2178
How the virtual MAC address is determined.	2179
Displaying the virtual MAC address	2181
Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain	2182
Synchronizing the configuration	2184
Disabling automatic configuration synchronization	2184
Incremental synchronization	2185
Periodic synchronization	2185
Console messages when configuration synchronization succeeds	2186
Console messages when configuration synchronization fails	2186
Comparing checksums of cluster units	2188
How to diagnose HA out of sync messages	2189
Synchronizing routing table updates	2191
Configuring graceful restart for dynamic routing failover	2191
Controlling how the FGCP synchronizes routing updates	2192
Synchronizing IPsec VPN SAs	2193
Link failover.	2194
If a monitored interface on the primary unit fails.	2196
If a monitored interface on a subordinate unit fails	2196
How link failover maintains traffic flow.	2197
Recovery after a link failover	2198
Testing link failover.	2198
Updating MAC forwarding tables when a link failover occurs	2198
Multiple link failures	2199
Example link failover scenarios	2199
Subsecond failover	2200

Remote link failover.	2200
Adding HA remote IP monitoring to multiple interfaces	2202
Changing the ping server failover threshold	2203
Monitoring multiple IP addresses from one interface	2204
Flip timeout.	2204
Detecting HA remote IP monitoring failovers	2204
Session failover (session pick-up)	2205
Improving session synchronization performance	2205
Session failover not supported for all sessions	2206
SIP session failover	2207
Session failover and explicit web proxy, WCCP, and WAN optimization sessions	2208
Session failover and SSL offloading and HTTP multiplexing	2208
IPsec VPN and SSL VPN sessions.	2208
PPTP and L2TP VPN sessions.	2208
Session failover and UDP, ICMP, multicast and broadcast packets	2208
FortiOS Carrier GTP session failover.	2209
Active-active HA subordinate units sessions can resume after a failover	2209
WAN optimization and HA	2209
Failover and attached network equipment	2210
Monitoring cluster units for failover.	2210
NAT/Route mode active-passive cluster packet flow	2210
Packet flow from client to web server	2211
Packet flow from web server to client	2211
When a failover occurs.	2212
Transparent mode active-passive cluster packet flow	2212
Packet flow from client to mail server	2213
Packet flow from mail server to client	2213
When a failover occurs.	2214
Failover performance.	2214
Device failover performance	2214
Link failover performance	2215
Reducing failover times	2215
HA and load balancing	2217
Load balancing overview	2217
Load balancing schedules	2218
Selecting which packets are load balanced	2219
More about active-active failover	2219
HTTPS sessions, active-active load balancing, and proxy servers	2220
Using FortiGate network processor interfaces to accelerate active-active HA performance	2220

Configuring load balancing settings	2221
Selecting a load balancing schedule.	2221
Load balancing UTM sessions and TCP sessions	2221
Configuring weighted-round-robin weights	2222
Dynamically optimizing weighted load balancing according to how busy cluster units are	2223
NAT/Route mode active-active cluster packet flow	2227
Packet flow from client to web server	2227
Packet flow from web server to client	2228
When a failover occurs.	2229
Transparent mode active-active cluster packet flow.	2229
Packet flow from client to mail server	2230
Packet flow from mail server to client	2231
When a failover occurs.	2231
HA with third-party products	2233
Troubleshooting layer-2 switches	2233
Forwarding delay on layer 2 switches	2234
Failover issues with layer-3 switches.	2234
Changing spanning tree protocol settings for some switches	2234
Spanning Tree protocol (STP)	2235
Bridge Protocol Data Unit (BPDU)	2235
Failover and attached network equipment	2235
Ethertype conflicts with third-party switches.	2235
LACP, 802.3ad aggregation and third-party switches	2236
VRRP	2237
Adding a VRRP virtual router to a FortiGate interface	2238
VRRP virtual MAC address.	2238
Configuring VRRP	2239
Example VRRP configuration: two FortiGate units in a VRRP group	2239
Example VRRP configuration: VRRP load balancing two FortiGate units and two VRRP groups	2241
Optional VRRP configuration settings	2242
TCP session synchronization	2243
Notes and limitations	2244
Configuring session synchronization	2244
Configuring the session synchronization link.	2245
Basic example configuration	2246

Chapter 13 Traffic Shaping	2249
The purpose of traffic shaping	2251
Quality of Service	2251
Traffic policing	2252
Bandwidth guarantee, limit, and priority interactions	2253
FortiGate traffic	2253
Through traffic	2254
Important considerations	2258
Traffic shaping methods	2261
Traffic shaping options	2261
Shared policy shaping	2262
Per policy	2262
All policies	2262
Maximum and guaranteed bandwidth	2262
Traffic priority	2262
VLAN, VDOM and virtual interfaces	2263
Shared traffic shaper configuration settings	2263
Per-IP shaping	2265
Per-IP traffic shaping configuration settings	2265
Application control shaping	2266
Example	2266
Enabling in the security policy	2267
Reverse direction traffic shaping	2267
Setting the reverse direction only	2267
Application control shaper	2268
Type of Service priority	2268
TOS in FortiOS	2269
Differentiated Services	2269
DSCP examples	2271
Tos and DSCP mapping	2275
Traffic Shaper Monitor	2276
Examples	2277
QoS using priority from security policies	2277
Sample configuration	2278
QoS using priority from ToS or differentiated services	2279
Sample configuration	2280

Example setup for VoIP	2280
Creating the traffic shapers	2281
Creating security policies	2282

Troubleshooting **2285**

Interface diagnosis	2285
Shaper diagnose commands	2285
TOS command	2285
Shared shaper	2285
Per-IP shaper	2286
Packet loss with statistics on shapers	2286
Packet lost with the debug flow	2287
Session list details with dual traffic shaper	2287
Additional Information	2288

Chapter 14 FortiOS Carrier **2289**

Overview of FortiOS Carrier features **2291**

Overview	2291
MMS	2291
GTP.	2291
MMS background	2292
MMS content interfaces	2292
How MMS content interfaces are applied	2293
How FortiOS Carrier processes MMS messages	2295
FortiOS Carrier and MMS content scanning	2296
FortiOS Carrier and MMS duplicate messages and message floods	2301
MMS protection profiles	2304
Bypassing MMS protection profile filtering based on user's carrier end points	2305
Applying MMS protection profiles to MMS traffic	2305
GTP basic concepts	2305
PDP Context	2306
GPRS security	2307
Parts of a GTPv1 network	2308
Radio access	2309
Transport	2309
Billing and records	2310
GPRS network common interfaces	2311
Packet flow through the GPRS network	2312

Carrier web-based manager settings	2315
MMS profiles	2315
MMS Content Checksum	2328
Notification List.	2329
Message Flood.	2331
Duplicate Message.	2333
Carrier Endpoint Filter Lists	2334
GTP Profile	2336
MMS UTM features	2353
Why scan MMS messages for viruses and malware?	2353
Example: COMMWARRIOR	2353
MMS virus scanning	2354
MMS virus monitoring	2354
MMS virus scanning blocks messages (not just attachments)	2355
Scanning MM1 retrieval messages	2355
Configuring MMS file filtering	2355
Removing or replacing blocked messages	2355
Carrier Endpoint Block	2356
MMS Content Checksum	2358
Passing or blocking fragmented messages	2359
Client comforting	2359
Server comforting	2360
Handling oversized MMS messages	2360
MM1 sample messages	2360
Configuring MMS virus scanning	2362
MMS file filtering	2363
Built-in patterns and supported file types	2364
MMS file filtering blocks messages (not just attachments)	2366
Configuring MMS file filtering	2366
Configuring sender notifications	2367
MMS content-based Antispam protection	2368
Overview	2369
Scores and thresholds	2370
Configuring content-based antispam protection	2370
Configuring sender notifications	2370
MMS DLP archiving	2371
Configuring MMS DLP archiving	2371
Viewing DLP archives	2372
Message flood protection	2373
Setting message flood thresholds	2374
Example	2374
Flood actions.	2375
Notifying administrators of floods	2375

Example — three flood threshold levels with different actions for each threshold	2375
Notifying message flood senders and receivers	2378
Responses to MM1 senders and receivers	2378
Forward responses for MM4 message floods	2379
Viewing DLP archived messages.	2379
Order of operations: flood checking before duplicate checking	2379
Bypassing message flood protection based on user's carrier end points	2380
Configuring message flood detection	2380
Sending administrator alert notifications	2381
Configuring how and when to send alert notifications.	2381
Configuring who to send alert notifications to	2383
Duplicate message protection	2385
Using message fingerprints to identify duplicate messages	2386
Messages from any sender to any recipient	2386
Setting duplicate message thresholds	2386
Duplicate message actions.	2387
Notifying duplicate message senders and receivers	2387
Responses to MM1 senders and receivers	2388
Forward responses for duplicate MM4 messages.	2388
Viewing DLP archived messages.	2389
Order of operations: flood checking before duplicate checking	2389
Bypassing duplicate message detection based on user's carrier end points	2389
Configuring duplicate message detection	2389
Sending administrator alert notifications	2390
Configuring how and when to send alert notifications.	2390
Configuring who to send alert notifications to	2391
MMS Replacement messages	2393
Changing replacement messages	2393
Multimedia content for MMS replacement messages	2395
MMS replacement message types	2396
Replacement message tags	2396
Replacement message groups.	2398

Configuring GTP on FortiOS Carrier	2401
GTP support on the FortiOS Carrier unit	2401
Packet sanity checking.	2402
GTP stateful inspection	2402
Protocol anomaly detection and prevention	2402
HA	2402
Virtual domain support.	2403
Configuring General Settings on the FortiOS Carrier unit	2403
Configuring Encapsulated Filtering in FortiOS Carrier	2403
Configuring Encapsulated IP Traffic Filtering	2403
Configuring Encapsulated Non-IP End User Address Filtering	2404
Configuring the Protocol Anomaly feature in FortiOS Carrier	2405
Configuring Anti-overbilling in FortiOS Carrier	2405
Overbilling in GPRS networks	2405
Anti-overbilling with FortiOS Carrier	2406
Logging events on the FortiOS Carrier unit	2406
GTP message type filtering	2409
Common message types on carrier networks	2409
GTP-C messages	2409
GTP-U messages	2410
Unknown Action messages	2411
Configuring message type filtering in FortiOS Carrier	2411
Message Type Fields.	2412
GTP identity filtering	2417
IMSI on carrier networks	2417
Other identity and location based information elements.	2418
When to use APN, IMSI, or advanced filtering	2419
Configuring APN filtering in FortiOS Carrier	2420
Configuring IMSI filtering in FortiOS Carrier	2421
Configuring advanced filtering in FortiOS Carrier	2422
Troubleshooting	2425
FortiOS Carrier diagnose commands	2425
GTP related diagnose commands	2425
Applying Intrusion and Prevention System (IPS) signatures to IP packets within GTP-U tunnels.	2426

GTP packets are not moving along your network	2427
Attempt to identify the section of your network with the problem	2427
Ensure you have an APN configured.	2427
Check the logs and adjust their settings if required	2428
Check the routing table	2428
Perform a sniffer trace	2429
Generate specific packets to test the network.	2431

Chapter 15 Deploying Wireless Networks 2433

Introduction to wireless networking 2435

Wireless concepts	2435
Bands and channels	2435
Power	2436
Antennas	2436
Security	2436
Whether to broadcast SSID	2436
Encryption	2436
Separate access for employees and guests	2437
Captive portal	2437
Power	2437
Monitoring for rogue APs	2437
Authentication	2438
Wireless networking equipment	2438
FortiWiFi units	2438
FortiAP units	2440
Third-party WAPs	2440
Deployment considerations	2440
Types of wireless deployment	2440
Deployment methodology	2440
Single access point networks	2442
Multiple access point networks	2442
Automatic Radio Resource Provisioning	2443

Configuring a WiFi LAN 2445

Overview of WiFi controller configuration	2445
About SSIDs on FortiWiFi units	2446
About automatic AP profile settings	2446
Process to create a wireless network	2447
Setting your geographic location.	2447
Creating a custom AP Profile.	2447

Defining a wireless network interface (SSID)	2448
Configuring DHCP for WiFi clients	2450
Configuring security	2450
Adding a MAC filter	2453
Configuring user authentication	2454
WPA-Enterprise authentication	2454
Creating a wireless user group.	2454
Configuring firewall policies for the SSID.	2455
Customizing captive portal pages	2456
Modifying the login page.	2456
Modifying the login failed page	2457
Configuring the built-in access point on a FortiWiFi unit.	2458
Access point deployment	2459
Overview	2459
Network topology for managed APs	2459
Discovering and authorizing APs.	2460
Configuring a managed AP	2462
Updating FortiAP unit firmware	2463
Advanced WiFi controller discovery	2464
Controller discovery methods	2464
Connecting to the FortiAP CLI	2466
Configuring a FortiWiFi unit as a WiFi AP	2466
Wireless network monitoring	2469
Monitoring wireless clients	2469
Monitoring rogue APs	2470
On-wire rogue AP detection technique	2470
Rogue AP scanning as a background activity	2471
Configuring rogue scanning	2471
Using the Rogue AP Monitor.	2472
Suppressing rogue APs	2473
Configuring wireless network clients	2475
Windows XP client	2476
Windows 7 client	2480
Mac OS client.	2481
Linux client	2483
Troubleshooting	2485
Checking that the client has received IP address and DNS server information	2485

Wireless network examples	2487
Basic wireless network	2487
Configuring authentication for wireless users	2487
Configuring the SSID.	2488
Configuring firewall policies	2489
Connecting the FortiAP units.	2490
A more complex example	2491
Scenario	2491
Configuration.	2492
Configuring authentication for employee wireless users.	2492
Configuring authentication for guest wireless users	2492
Configuring the SSIDs	2494
Configuring the custom AP profile	2496
Configuring firewall policies	2497
Connecting the FortiAP units.	2499
Using a FortiWiFi unit as a client	2501
Use of client mode	2501
Configuring client mode	2502
WiFi Reference	2503
Wireless radio channels	2503
WiFi Controller Reference	2505
WiFi Controller overview	2505
WiFi Network	2506
SSID list	2506
SSID configuration settings	2507
Rogue AP Settings	2509
Managed access points	2509
Local WiFi Radio configuration settings	2510
Managed FortiAP list	2510
Managed FortiAP configuration settings.	2511
Custom AP Profiles	2512
Custom AP Profile Settings	2513
Monitor	2514
Client Monitor	2514
Rogue AP Monitor	2515
Chapter 16 VoIP Solutions: SIP & FortiGate Voice	2517
FortiGate VoIP solutions: SIP	2519
SIP overview	2519

Common SIP VoIP configurations	2520
Peer to peer configuration	2520
SIP proxy server configuration	2521
SIP redirect server configuration	2522
SIP registrar configuration	2522
SIP with a FortiGate unit	2523
SIP messages and media protocols	2525
SIP request messages	2528
SIP response messages	2529
SIP message start line	2530
SIP headers	2531
The SIP message body and SDP session profiles	2533
Example SIP messages	2534
The SIP session helper	2536
SIP session helper configuration overview	2536
Configuration example: SIP session helper in Transparent Mode	2538
SIP session helper diagnose commands	2541
The SIP ALG	2541
SIP ALG configuration overview	2543
Conflicts between the SIP ALG and the session helper	2545
Stateful SIP tracking, call termination, and session inactivity timeout.	2547
SIP and RTP/RTCP	2549
How the SIP ALG creates RTP pinholes	2549
Configuration example: SIP in Transparent Mode	2550
RTP enable/disable (RTP bypass)	2553
Opening and closing SIP register and non-register pinholes	2554
Accepting SIP register responses	2554
How the SIP ALG performs NAT	2555
Source address translation.	2556
Destination address translation	2556
Call Re-invite messages	2556
How the SIP ALG translates IP addresses in SIP headers.	2557
How the SIP ALG translates IP addresses in the SIP body	2559
SIP NAT scenario: source address translation (source NAT)	2560
SIP NAT scenario: destination address translation (destination NAT)	2562
SIP NAT configuration example: source address translation (source NAT)	2564
SIP NAT configuration example: destination address translation (destination NAT)	2567
Additional SIP NAT scenarios	2570
NAT with IP address conservation	2573
Controlling how the SIP ALG NATs SIP contact header line addresses.	2574
Controlling NAT for addresses in SDP lines	2575
Translating SIP session destination ports	2575
Translating SIP sessions to multiple destination ports.	2577

Enhancing SIP pinhole security	2578
Hosted NAT traversal.	2580
Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B	2581
Hosted NAT traversal for calls between SIP Phone A and SIP Phone C	2584
Restricting the RTP source IP	2585
SIP over IPv6	2585
Deep SIP message inspection	2586
Actions taken when a malformed message line is found	2587
Logging and statistics	2587
Deep SIP message inspection best practices	2588
Configuring deep SIP message inspection	2588
Blocking SIP request messages	2590
SIP rate limiting.	2592
Limiting the number of SIP dialogs accepted by a security policy	2593
SIP logging and DLP archiving	2594
SIP and HA: session failover and geographic redundancy.	2594
SIP geographic redundancy	2595
Support for RFC 2543-compliant branch parameters	2596
SIP and IPS.	2596
SIP debugging	2596
SIP debug log format.	2596
SIP-proxy filter per VDOM	2597
SIP-proxy filter command	2598
SIP debug log filtering	2598
SIP debug setting	2598
SIP test commands	2599
Display SIP rate-limit data	2599
VoIP Profile options.	2600
Example FortiGate Voice branch office configuration	2603
General configuration steps	2604
Connecting the FortiGate Voice unit	2605
Configuring basic FortiGate Voice network and UTM settings.	2605
Configuring network settings for the devices on the Internal network	2608
Configuring the FortiGate Voice PSTN and PBX settings	2608
Configuring the FortiFones on the internal network	2615
Adding extensions and configuring FortiFones for users behind a NAT device	2616
FortiGate Voice IVR configuration	2618
Providing access to the company directory	2618

Adding a shortcut for checking voicemail	2619
Checking voicemail	2619
FortiGate Voice web-based manager configuration reference	2621
Unit operation dashboard widget	2621
Configuring interface settings to support VoIP PBX features	2621
Configuring an interface to accept SIP traffic	2622
Enabling access to the PBX user web portal	2622
SIP phone auto-provisioning	2623
Default FortiGate Voice auto-provisioning configuration	2624
Configuring SIP phones for auto-provisioning	2624
PBX configuration	2625
Configuring extensions.	2625
Configuring extension groups (ring groups)	2628
Configuring service providers (the FortiGuard Voice service)	2629
Configuring PSTN interfaces.	2629
Configuring the FortiGuard Voice service	2631
Adding SIP trunks	2633
Branch Office.	2635
Configuring dial plans	2636
Configuring voice menu options	2640
Configuring direct inward dialing.	2640
Configuring PBX global settings	2641
Importing a new voice prompt file	2643
Parking calls	2643
FortiFAX service	2644
Monitoring calls	2644
Monitoring recorded conference calls	2644
Monitoring voice mail storage	2644
Monitoring active phones	2644
Logging of PBX activities.	2645
Viewing log messages	2645
VoIP interface reference	2645
Profile.	2646
Using the PBX user web portal	2649
Logging into and out of the FortiGate Voice PBX user web portal.	2649
Configuring PBX extension settings	2649
Voicemail	2650
Configuring call forwarding.	2650
Sending a Fax using FortiFAX	2651
Conference calls	2651
Managing conference calls.	2652

FortiGate Voice VoIP, PBX, and PSTN CLI Reference	2653
config pbx dialplan	2653
config pbx did	2655
config pbx extension	2655
config pbx global	2657
config pbx ringgrp	2659
config pbx voice-menu	2660
config pbx sip-trunk	2661
config system pstn	2663
config system interface.	2665
execute pbx	2665
get pbx branch-office.	2667
get pbx dialplan	2667
get pbx did	2667
get pbx extension	2667
get pbx ftgd-voice-pkg	2668
get pbx global	2668
get pbx ringgrp	2669
get pbx sip-trunk	2669
get pbx voice-menu	2669
diagnose pbx restart	2670

Chapter 17 WAN Optimization, Web Cache, Explicit Proxy, and WCCP **2671**

WAN optimization, web cache, explicit proxy, and WCCP concepts	2673
WAN optimization topologies	2674
Basic WAN optimization topologies	2674
Out-of-path topology.	2675
Topology for multiple networks	2677
WAN optimization with web caching.	2678
WAN optimization and web caching with FortiClient peers	2679
Explicit Web proxy topologies	2680
Explicit FTP proxy topologies	2681
Web caching topologies	2682
WCCP topologies.	2684

WAN optimization client/server architecture	2685
WAN optimization peers	2686
Peer-to-peer and active-passive WAN optimization	2686
WAN optimization and the FortiClient application	2687
Operating modes and VDOMs	2687
WAN optimization tunnels	2687
Tunnel sharing	2688
Protocol optimization	2689
Byte caching	2689
WAN optimization and HA	2690
WAN optimization, web caching and memory usage	2690
Monitoring WAN optimization performance	2690
Traffic Summary	2691
Bandwidth Optimization	2691
Configuring WAN optimization traffic usage logs	2691
WAN optimization best practices.	2692
WAN optimization and Web cache storage	2693
Formatting the hard disk	2693
Configuring WAN optimization and Web cache storage	2694
Changing the amount of space allocated for WAN optimization and Web cache storage.	2694
Adjusting the relative amount of disk space available for byte caching and web caching	2694
WAN optimization peers and authentication groups	2697
Basic WAN optimization peer requirements	2697
Accepting any peers	2697
How FortiGate units process tunnel requests for peer authentication	2698
Configuring peers.	2699
Configuring authentication groups	2700
Secure tunneling	2702
Monitoring WAN optimization peer performance	2703
Configuring WAN optimization rules	2705
WAN optimization rules, security policies, and UTM protection	2705
WAN optimization transparent mode.	2706
WAN optimization rule list	2707
How list order affects rule matching	2708
Moving a rule to a different position in the rule list.	2709
WAN optimization address formats	2709

Configuring WAN optimization rules	2710
Processing non-HTTP sessions accepted by an HTTP rule	2714
Processing unknown HTTP sessions	2714
WAN optimization configuration examples	2715
Example: Basic peer-to-peer WAN optimization configuration	2715
Network topology and assumptions	2715
General configuration steps	2716
Configuring basic peer-to-peer WAN optimization - web-based manager	2716
Configuring basic peer-to-peer WAN optimization - CLI	2718
Testing and troubleshooting the configuration.	2719
Example: Active-passive WAN optimization	2721
Network topology and assumptions	2721
General configuration steps	2722
Configuring basic active-passive WAN optimization - web-based manager	2722
Configuring basic active-passive WAN optimization - CLI.	2725
Testing and troubleshooting the configuration.	2727
Example: Adding secure tunneling to an active-passive WAN optimization configuration	2728
Network topology and assumptions	2729
General configuration steps	2729
Configuring WAN optimization with secure tunneling - web-based manager	2729
Configuring WAN optimization with secure tunneling - CLI	2732
Web caching	2735
Web caching in security policies	2736
Example: Web caching of Internet content for users on an internal network.	2736
Web Caching only WAN optimization	2739
Example: Web Cache Only WAN optimization.	2739
Web caching for active-passive WAN optimization	2744
Example: Active-passive Web Caching	2744
Web caching for peer-to-peer WAN optimization	2748
Example: Peer-to-peer web caching.	2749
Exempting web sites from web caching	2752
Changing web cache settings	2753
Monitoring Web caching performance	2755

Advanced configuration example	2757
Out-of-path WAN optimization with inter-VDOM routing	2757
Network topology and assumptions	2757
Configuration steps	2758
Client-side configuration steps - web-based manager	2759
Server-side configuration steps - web-based manager	2766
Client-side configuration steps - CLI	2769
Server-side configuration steps - CLI	2776
SSL offloading for WAN optimization and web caching	2781
About SSL server full and half mode	2782
WAN optimization full mode SSL server configuration	2782
WAN optimization half mode SSL server configuration	2783
Reverse proxy web cache full mode SSL server configuration	2784
Reverse proxy web cache half mode SSL server configuration	2785
Example: SSL offloading for a WAN optimization tunnel.	2786
Network topology and assumptions	2786
General configuration steps	2787
Client-side configuration steps.	2787
Server-side configuration steps	2788
Example: SSL offloading and reverse proxy web caching for an Internet	
web server using static one-to-one virtual IPs	2789
Network topology and assumptions	2789
General configuration steps	2791
Configuration steps - web-based manager	2791
Configuration steps - CLI	2794
Example: SSL offloading and reverse proxy web caching for an Internet	
web server using a port forwarding virtual IP for HTTPS traffic	2795
Network topology and assumptions	2796
General configuration steps	2797
Configuration steps - web-based manager	2798
Configuration steps - CLI	2801
FortiClient WAN optimization	2805
Configuring FortiClient WAN optimization	2805
FortiClient configuration steps	2806
FortiGate unit configuration steps	2806
The FortiGate explicit web proxy	2807
Explicit web proxy configuration overview	2809
Proxy auto-config (PAC) configuration.	2812
Unknown HTTP version	2812
Authentication realm	2812
Other explicit web proxy options.	2813

Proxy chaining	2813
Adding a web proxy forwarding server	2814
Web proxy forwarding server monitoring and health checking	2814
Adding proxy chaining to an explicit web proxy security policy	2814
Explicit web proxy authentication	2815
IP-Based authentication	2816
Per session authentication	2816
UTM features and the explicit web proxy	2817
Explicit web proxy sessions and flow-based scanning	2818
Explicit web proxy sessions and protocol options.	2818
Explicit web proxy sessions web filtering and FortiGuard web filtering	2818
Explicit web proxy sessions and HTTPS deep scanning	2818
Explicit web proxy sessions and antivirus	2819
Web Proxy Services	2819
Web Proxy Service Groups	2820
Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning	2820
General configuration steps	2821
Configuring the explicit web proxy - web-based manager	2821
Configuring the explicit web proxy - CLI.	2823
Testing and troubleshooting the configuration.	2824
Explicit proxy sessions and user limits	2825
Explicit web proxy configuration options.	2827
Explicit Web Proxy Options	2827
Web Proxy Forwarding Servers Options.	2829
Adding Web Proxy Forwarding Servers	2829
Restricting the IP address of the explicit web proxy.	2830
Restricting the outgoing source IP address of the explicit web proxy.	2830
The FortiGate explicit FTP proxy	2831
How to use the explicit FTP proxy to connect to an FTP server.	2832
Explicit FTP proxy configuration overview	2834
Restricting the IP address of the explicit FTP proxy.	2837
Restricting the outgoing source IP address of the explicit FTP proxy.	2837
UTM features and the explicit FTP proxy	2837
Explicit FTP proxy sessions and protocol options.	2837
Explicit FTP proxy sessions and antivirus	2838

Example: users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning	2838
General configuration steps	2839
Configuring the explicit FTP proxy - web-based manager	2839
Configuring the explicit FTP proxy - CLI	2840
Testing and troubleshooting the configuration.	2842
Explicit FTP proxy sessions and user limits	2844
Explicit FTP proxy options	2844
FortiGate WCCP	2845
WCCP service groups, service numbers, service IDs and well known services	2846
Example WCCP server and client configuration for caching HTTP sessions (service ID = 0).	2846
Example WCCP server and client configuration for caching HTTPS sessions	2847
Example WCCP server and client configuration for caching HTTP and HTTPS sessions	2848
Other WCCP service group options	2848
WCCP configuration overview	2849
Example: caching HTTP sessions on port 80 using WCCP	2850
Configuring the WCCP server (WCCP_srv)	2851
Configuring the WCCP client (WCCP_client)	2852
Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP	2852
Configuring the WCCP server (WCCP_srv)	2853
Configuring the WCCP client (WCCP_client)	2854
WCCP packet flow	2854
Configuring the forward and return methods and adding authentication	2855
WCCP Messages.	2856
Troubleshooting WCCP	2856
Real time debugging	2856
Application debugging	2856
WAN optimization, web cache, explicit proxy and WCCP get and diagnose commands	2859
get test {wa_cs wa_dbd wad wad_diskd wccpd} <test_level>	2859
Examples.	2859
diagnose wad.	2862
Examples.	2862
diagnose wacs	2864
diagnose wadbd	2864

diagnose debug application {wa_cs | wa_dbd | wad | wad_diskd | wccpd} [<debug_level>]
2864

Chapter 18 Load Balancing 2867

Configuring load balancing 2869

Load balancing overview	2869
Load balancing, UTM, authentication, and other FortiOS features	2870
Configuring load balancing virtual servers	2870
Load balancing methods	2873
Session persistence	2874
Real servers	2874
Health check monitoring	2876
Monitoring load balancing	2879
Load balancing get command	2880
Load balancing diagnose commands	2880
Logging Diagnostics	2880
Real server diagnostics	2881
Basic load balancing configuration example	2882
HTTP and HTTPS load balancing, multiplexing, and persistence	2885
HTTP and HTTPS multiplexing	2885
HTTP and HTTPS persistence	2886
HTTP host-based load balancing	2888
SSL/TLS load balancing	2889
SSL offloading	2890
IP, TCP, and UDP load balancing	2897

Load balancing configuration examples 2899

Example: HTTP load balancing to three real web servers	2899
Web-based manager configuration	2900
CLI configuration	2903
Example: Basic IP load balancing configuration	2905
Example: Adding a server load balance port forwarding virtual IP	2905
Example: Weighted load balancing configuration	2907
Web-based manager configuration	2907
CLI configuration	2910
Example: HTTP and HTTPS persistence configuration	2911
CLI configuration: adding persistence for a specific domain	2914

Chapter 19 Hardware 2917

FortiGate installation 2919

Mounting the FortiGate unit	2919
Desk or table mounting	2919
Rack mounting	2919
Plugging in the FortiGate unit	2927
Connecting to the network.	2927
Turning off the FortiGate unit.	2927
Further configuration	2928

AMC module configuration 2929

Configuring AMC modules	2929
Auto-bypass and recovery for AMC bridge module	2930
Enabling or disabling bypass mode for AMC bridge modules	2931

FortiGate hardware accelerated processing 2933

How hardware acceleration alters packet flow.	2933
Network processors overview	2935
Network processor models	2935
Determining the network processors installed on your FortiGate unit.	2936
Content processors overview	2936
Determining the content processor in your FortiGate unit.	2938
Security processing modules overview	2938
Security processor module models	2939
Displaying information about security processing modules	2939
Setting switch-mode mapping on the ADM-XD4	2939
Configuring overall security priorities.	2940
Configuring traffic offloading	2940
Session fast path requirements	2940
Packet fast path requirements	2941
Fast path connections for specific FortiGate models	2941
Session offloading in HA active-active configuration	2946
Configuring traffic shaping offloading	2947
Checking that traffic is offloaded	2948
Disabling offloading	2948
Multicast offloading / acceleration	2949

Configuring IPsec VPN offloading	2949
IPsec offloading requirements	2949
Configuring HMAC check offloading.	2950
Configuring VPN encryption/decryption offloading	2950
Examples of ASM-FB4 accelerated VPNs	2951
Configuring IPS offloading	2955
Configuring pre-IPS anomaly detection	2955
Configuring policy-based IPS on SP modules	2956
Configuring interface-based IPS on SP modules	2956
Examples	2957
Accelerated tunnel mode IPsec	2958
Accelerated interface mode IPsec	2959
Configuring RAID	2961
RAID levels	2961
Configuring a RAID array	2962
Checking the status of a RAID array	2963
Rebuilding a RAID array	2964
Why rebuild a RAID array?	2964
How to rebuild the RAID array	2964
FortiBridge installation and operation	2967
Example FortiBridge application	2967
Connecting the FortiBridge unit	2968
Normal mode operation	2970
How the FortiBridge unit monitors the FortiGate unit	2970
Probes and FortiGate firewall policies	2971
Enabling probes to detect FortiGate hardware failure	2973
Enabling probes to detect FortiGate software failure	2973
Probe interval and probe threshold	2973
Bypass mode operation	2973
FortiBridge power failure	2974
Example FortiGate HA cluster FortiBridge application	2975
Connecting the FortiBridge-2002 (copper gigabit ethernet)	2976
Connecting the FortiBridge-2002F (fiber gigabit ethernet).	2976
Example configuration with other FortiGate interfaces.	2976

Completing the basic FortiBridge configuration	2979
Adding an administrator password	2979
Changing the management IP address	2979
Changing DNS server IP addresses	2980
Changing the default gateway and adding static routes.	2980
Allowing management access to the EXT1 interface	2981
Changing the system time and date	2981
Adding administrator accounts	2981
Resetting to the factory default configuration	2982
Installing FortiBridge unit firmware	2982
Changing firmware versions	2982
Installing firmware from a system reboot	2983
Example network configuration	2987
Configuring FortiBridge probes	2988
Probe settings	2989
Enabling probes	2990
Verifying that probes are functioning.	2992
Tuning the failure threshold and probe interval	2992
Configuring FortiBridge alerts	2992
Recovering from a FortiGate failure	2995
Manually switching between FortiBridge operating modes	2996
Backing up and restoring the FortiBridge configuration	2996

Chapter 20 Certifications and Compliances 2999

FIPS-CC operation of FortiGate units 3001	3001
Introduction to FIPS-CC	3001
Security level summary.	3001
Documentation.	3002
Overview of Common Criteria compliant operation	3002
Use of non-FIPS-CC compliant features.	3002
Effects of FIPS-CC compliant mode.	3002
Initial configuration of the FortiGate unit	3005
Installing the unit	3005
Configuration of units with AMC/FMC modules	3005
Downloading and installing FIPS-CC compliant firmware	3005
Verifying the firmware version of the unit	3006
A note about non FIPS-CC functionality.	3006
Enabling FIPS-CC mode	3007
Configuring interfaces	3008
FIPS-CC mode status indicators.	3008
Self-test settings	3008
Running self-tests manually	3008

Administration	3009
User guidance	3009
Remote access requirements	3009
Disclaimer access banner	3010
Administrator account lockout settings	3010
Scheduled administrator access	3011
Using custom administrator access keys (certificates)	3011
Configuration backup	3011
Firewall	3012
Security policies	3012
Firewall authentication	3012
Logging	3013
Logging to external devices	3013
Required logging settings	3013
Excluding specific logs (selective audit)	3015
Viewing log messages from the web-based manager	3016
Viewing log messages from the CLI	3016
Backing up log messages	3017
Viewing log file information.	3018
Deleting filtered log messages	3018
Deleting rolled log files	3018
Alarms	3018
Configuring alarms	3019
Alarm notifications	3021
Acknowledging alarms	3021
Alarm polling	3021
Error modes	3021
FIPS Error mode	3022
CC Error mode	3022
Disabling FIPS-CC mode.	3022
Configuring FortiGate units for PCI DSS compliance	3023
Introduction to PCI DSS	3023
What is PCI DSS?	3023
What is the Customer Data Environment	3023
PCI DSS objectives and requirements	3024
Network topology.	3027
Internet	3027
The CDE wired LAN	3028
The CDE wireless LAN	3028
Other internal networks	3028

Security policies for the CDE network	3028
Controlling the source and destination of traffic	3028
Controlling the types of traffic in the CDE	3029
The default deny policy	3029
Wireless network security	3029
Scanning for rogue access points	3029
Securing a CDE network WAP	3030
Protecting stored cardholder data	3031
Protecting communicated cardholder data	3031
Configuring IPsec VPN security	3031
Configuring SSL VPN security	3032
Protecting the CDE network from viruses	3032
Enabling FortiGate antivirus protection	3032
Configuring antivirus updates	3033
Enforcing firewall use on endpoint PCs	3033
Monitoring the network for vulnerabilities	3033
Using the FortiOS Network Vulnerability Scan feature.	3033
Monitoring with other Fortinet products	3033
Restricting access to cardholder data	3034
Controlling access to the CDE network	3034
Password complexity and change requirements	3034
Password non-reuse requirement	3035
Administrator lockout requirement.	3035
Administrator timeout requirement.	3036
Administrator access security	3036
Remote access security	3036
Appendix	3037
Index	3043



Introduction

This FortiOS™ Handbook v3 is the definitive guide to configuring and operating FortiOS 4.0 MR3. It contains concept and feature descriptions, as well as configuration examples worked out in detail for the web-based manager and the CLI. This document also contains operating and troubleshooting information.

This is the third version of the handbook. This version is still a work in progress but includes many improvements, corrections and additions. Among the additions are new chapters that describe troubleshooting and FortiGate hardware (including hardware installation, hardware acceleration, RAID, and FortiBridge). The new compliances and certifications chapter describes FortiOS PCI support and FIPS/CC support. New sections include more information about VRRP, SNMP, sFlow, advanced static routing, and more information about deploying wireless networks.

This introduction describes the following topics:


- [How this Handbook is organized](#)
- [Document conventions](#)
- [Registering your Fortinet product](#)
- [Fortinet products End User License Agreement](#)
- [Training Services](#)
- [Technical Documentation](#)
- [Customer service and support](#)

How this Handbook is organized

This handbook contains the following chapters:

- [Chapter 1, What's New](#) describes the new features in FortiOS 4.0 MR3 and includes some general upgrading information.
- [Chapter 2, Firewall](#) describes FortiOS firewall functionality on all FortiGate units. It includes the purpose of the firewall, how traffic moves through the FortiGate unit, the components involved in the firewall and its policies. This chapter also describes how to configure the basics and some more involved examples.
- [Chapter 3, System Administration](#) describes a number of administrative tasks to configure and setup the FortiGate unit for the first time. It also describes the best practices and sample configuration tips to secure your network and the FortiGate unit itself.
- [Chapter 4, Logging and Reporting](#) describes how to begin choosing a log device for your logging requirements, the types of log files, how to configure your chosen log device, including detailed explanations of each log type of log message.

- [Chapter 5, Troubleshooting](#) describes concepts of troubleshooting and solving issues that may occur with FortiGate units.
- [Chapter 6, UTM Guide](#) describes the Unified Threat Management (UTM) features available on your FortiGate unit, including antivirus, intrusion prevention system (IPS), anomaly protection (DoS), one-armed IPS (sniffer policies), web filtering, email filtering, data leak prevention (DLP), and application control. Also included is how to use the Endpoint features of FortiOS: endpoint Network Access Control (NAC), endpoint application detection, endpoint monitoring, and network vulnerability scanning. The chapter includes step-by-step instructions showing how to configure each feature. Example scenarios are included, with suggested configurations.
- [Chapter 7, User Authentication](#) defines authentication and describes the FortiOS options for configuring authentication for FortiOS.
- [Chapter 8, IPsec VPNs](#) provides a general introduction to IPsec VPN technology, explains the features available with IPsec VPN and gives guidelines to decide what features you need to use, and how the FortiGate unit is configured to implement the features.
- [Chapter 9, SSL VPN](#) provides a general introduction to SSL VPN technology, explains the features available with SSL VPN and gives guidelines to decide what features you need to use, and how the FortiGate unit is configured to implement the features.
- [Chapter 10, Advanced Routing](#) provides detailed information about FortiGate dynamic routing including common dynamic routing features, troubleshooting, and each of the protocols including RIP, BGP, and OSPF.
- [Chapter 11, Virtual Domains](#) describes FortiGate Virtual Domains (VDOMs) and is intended for administrators who need guidance on solutions to suit different network needs and information on basic and advanced configuration of VDOMs. Virtual Domains (VDOMs) multiply the capabilities of your FortiGate unit by using virtualization to partition your resources.
- [Chapter 12, High Availability](#) describes FortiGate HA, the FortiGate Clustering Protocol (FGCP), FortiGate support of VRRP, and FortiGate standalone TCP session synchronization.
- [Chapter 13, Traffic Shaping](#) describes how to configure FortiOS traffic shaping.
- [Chapter 14, FortiOS Carrier](#) describes FortiOS Carrier dynamic profiles and groups, Multimedia messaging service (MMS) protection, and GPRS Tunneling Protocol (GTP) protection.
- [Chapter 15, Deploying Wireless Networks](#) describes how to configure wireless networks with FortiWiFi, FortiGate, and FortiAP units.
- [Chapter 16, VoIP Solutions: SIP & FortiGate Voice](#) describes FortiOS SIP support.
- [Chapter 17, WAN Optimization, Web Cache, Explicit Proxy, and WCCP](#) describes how FortiGate WAN optimization, web caching, and web proxy work and also describes how to configure these features.
- [Chapter 18, Load Balancing](#) describes firewall HTTP, HTTPS, SSL or generic TCP/UDP or IP server load balancing.
- [Chapter 19, Hardware](#) describes how to mount, connect power, and turn on your FortiGate unit. Other topics include descriptions and configuration instructions for FortiGate accelerated hardware processing, RAID, and FortiBridge protection.
- [Chapter 20, Certifications and Compliances](#) explains how Fortinet products can help you comply with the Payment Card Industry Data Security standard and also describe FortiOS FIPS/CC support.



Chapter 1 What's New

This FortiOS Handbook chapter contains the following sections:

[Upgrading to FortiOS 4.0 MR3](#) provides information about upgrading to the new release.

[FortiOS 4.0 MR3 New Feature Highlights](#) describes the key new features available in FortiOS 4.0 MR3.

[Logging and reporting enhancements](#) describes new logging and reporting features.

[FortiOS 4.0 MR3 Usability improvements](#) describes FortiOS 4.0 MR3 usability enhancements and changes.

[More New Features](#) describes other general new FortiOS 4.0 MR3 features and lists what's new in FortiOS 4.0 MR3 patches 1 to 5.



Upgrading to FortiOS 4.0 MR3

This section explains how to properly upgrade to FortiOS 4.0 MR3. The following topics are included in this section:

- [General firmware upgrade steps](#)
- [Backing up and restoring your FortiGate configuration file](#)
- [Temporarily installing FortiOS 4.0 MR3](#)

General firmware upgrade steps

Regardless of whether you are installing the 4.0 MR3, patch release or GA firmware, you should use the following general procedure as a guideline for installing the firmware image. Upgrade the firmware during a low-traffic time period to avoid disrupting your network.

For more information about upgrading to FortiOS 4.0 MR3 see the FortiOS 4.0 MR3 Release Notes.

General procedure for upgrading current firmware - web-based manager

- 1 Verify what firmware image you need to upgrade to from the current firmware image that is running on the unit.
- 2 Download the new firmware image.
- 3 Back up your current configuration file.
- 4 Log into the web-based manager and go to *System > Dashboard > Status* and in the System Information Widget select *Update* beside the *Firmware Version*.
- 5 Select the firmware image file to install.
- 6 Clear your browser's cache after the installation process is finished.
After a few minutes you can log back into the web-based manager.
- 7 Manually update antivirus and intrusion protection definitions and engines to the current version.

Go to *System > Config > FortiGuard > Antivirus and IPS Options* and select *Update Now*.

The signatures included with a firmware image upgrade may be older than ones currently available from FortiGuard.

Backing up and restoring your FortiGate configuration file



Always back up your FortiGate configuration before upgrading or downgrading firmware, or resetting configuration to factory defaults. Then if required you can restore the configuration by uploading the backed up configuration file to your FortiGate unit.

Before installing any firmware image, you should back up the current FortiGate configuration file. This ensures that you have a current configuration file if the upgrade is not successful. You are also ensuring that all configuration settings are available if there are some that are not carried forward.

To back up your configuration file - web-based manager

- 1 Log into the web-based manager and go to *System > Dashboard > Status* and in the System Information Widget select *Backup* beside *System Configuration*.
- 2 Select where to save configuration file.
- 3 If you want to encrypt your configuration file to save VPN certificates, select *Encrypt configuration file* and enter and confirm a password. An encrypted configuration file can only be opened by uploading it to the same FortiGate unit and entering this password.
- 4 Select *Backup* and save the configuration file.

To restore your configuration file - web-based manager

You may need to restore your configuration file if you have experienced problems during a firmware upgrade.

- 1 Log into the web-based manager and go to *System > Dashboard > Status* and in the System Information Widget select *Restore* beside *System Configuration*.
- 2 Select to configuration file to restore.
- 3 If the configuration file is encrypted, enter the password.
- 4 Select *Restore* to restore the configuration to the one of the saved the configuration file.

The FortiGate unit uploads and installs the configuration.

- 5 Clear your browser's cache after the installation process is finished.
After a few minutes you can log back into the web-based manager.

Temporarily installing FortiOS 4.0 MR3

The following procedure describes how install temporarily install a firmware image to the system memory. When you reboot the FortiGate unit it will restart running the current firmware.

The procedure describes how to reboot the FortiGate unit and download firmware from a TFTP server and select the `Run image without saving` option to temporarily store the firmware image in memory without upgrading the firmware image stored on the FortiGate bootup device.

This procedure provides a way to become familiar with new FortiOS 4.0 MR3 new features and changes before committing to a full upgrade to the new version.

To temporarily install a new firmware image

- 1 Copy the new firmware image file to the root directory of a TFTP server.


- 2 Set up a console connection to the FortiGate unit CLI.
- 3 Make sure the FortiGate unit can connect to the TFTP server using the `execute ping` command.
- 4 Restart the FortiGate unit. For example, enter the following command:
`execute reboot`
- 5 As the FortiGate unit reboots, press any key to interrupt the system startup when the following message appears:
`Press any key to display configuration menu ...`



You have only three seconds to press any key. If you do not press a key soon enough, the FortiGate unit reboots and you must log in and repeat the `execute reboot` command.

- 6 If you successfully interrupt the startup process, a message similar to the following appears:
`[G]: Get firmware image from TFTP server`
`[F]: Format boot device`
`[B]: Boot with backup firmware and set as default`
`[C]: Configuration and information`
`[Q]: Quit menu and continue to boot with default firmware`
`[H]: Display this list of options.`

Enter G, F, Q, or H:
- 7 Type G to get the new firmware image from the TFTP server.
The following message appears:
Enter TFTP server address [192.168.1.168]:
- 8 Type the address of the TFTP server and press Enter.
The following message appears:
Enter Local Address [192.168.1.168]:
- 9 Type an IP address of the FortiGate unit to connect to the TFTP server.
The IP address must be on the same network as the TFTP server, but make sure you do not use the IP address of another device on the network.
The following message appears:
Enter File Name [image.out]:
- 10 Enter the firmware image file name and press Enter.
The TFTP server uploads the firmware image file to the FortiGate unit and the following appears:
`Save as Default firmware/Backup firmware/Run image without saving: [D/B/R]`
- 11 Type R.
The firmware image is installed to system memory and the FortiGate unit starts running the new firmware image, but with its current configuration.
- 12 When you are done, reboot the FortiGate unit and it will resume using the previous firmware image.



FortiOS 4.0 MR3 New Feature Highlights

This section describes the key new features available in FortiOS 4.0 MR3. In addition to the highlights described in this section see [“Logging and reporting enhancements” on page 127](#) for information about the new logging features in FortiOS 4.0 MR3 and [“FortiOS 4.0 MR3 Usability improvements” on page 139](#) for complete information about all of the usability improvements in FortiOS 4.0 MR3.

- [Flow-based UTM Extensions](#)
- [UTM Configuration and Inspection Enhancements](#)
- [Modem interface Improvements](#)
- [WiFi Extensions](#)
- [Strong Authentication Enhancements](#)
- [New PCI Compliance Features](#)
- [Feature Improvements to extend IPv6 support](#)
- [Explicit proxy and web caching improvements](#)

Flow-based UTM Extensions

Flow-based inspection can result in major performance improvements to UTM inspection. First introduced to improve antivirus performance in FortiOS 4.0 MR2, in MR3 flow-based inspection has been extended to web filtering and data leak prevention (DLP) and also includes the ability to virus scan compressed files.

This flow-based scanning performance improvements come from reduced memory requirements, high concurrent session count, high session start rates and low latency. In addition flow-based scanning is not affected by a maximum file size.

The trade-off for these advantages is that flow-based scanning may not be as accurate or comprehensive as proxy-based scanning although Fortinet is continuing to improve the accuracy and depth of coverage provided by flow-based UTM features.

Flow-based web filtering

Flow-based web filtering is a non-proxy solution which provides high concurrent session, high session rate, and low-latency web-filtering service. You can enable flow-based web filtering within a web filter profile.

You can enable flow-based web filtering in any Web Filter Profile by setting the *Inspection Mode* to *Flow-based*. Flow-based web filtering can be enabled in some web filtering profiles and not others, allowing you to apply flow-based web filtering to some traffic and proxy-based web filtering to other traffic.

Flow-based Data Leak Prevention (DLP)

Flow-based DLP is a non-proxy solution which provides high concurrent session, high session rate, and low-latency DLP services. You can enable flow-based DLP within a DLP sensor.

You can enable flow-based DLP in any DLP Sensor by setting the *Inspection Method* to *Flow-based Detection*. Flow-based DLP can be enabled in some DLP sensors and not others, allowing you to apply flow-based DLP to some traffic and proxy-based DLP to other traffic.

UTM Configuration and Inspection Enhancements

FortiOS 4.0 MR3 includes the following improvements to UTM functionality.

UTM profile and sensor configuration improvements

All UTM features including Antivirus, intrusion protection, web filtering, email filtering, data leak prevention, application control, VoIP and ICAP include one or more default profiles or sensors. In many cases you can add the default profile or sensor to a security policy to apply basic functionality for that UTM feature. You can also modify the default profiles and sensors to meet your requirements and create new profiles and sensors.

Within the Configuration Settings page for all UTM features you can do the following:

- view and edit the default profile or sensor
- view the current settings of a profile or sensor
- create or remove a new profile or sensor
- view a list of profiles or sensors that you created

In the upper-right corner of the Configuration Settings page there is a drop-down list, *Create New* icon and *View List* icon. These are shown in [Figure 1](#). You can use them in the following ways:

- Create a new profile or sensor by selecting *Create New*.
- View a specific profile or sensor by selecting it from the drop-down list.
- View the profiles or sensors that you have created by selecting *View List*.
- Remove the current profile or sensor that you are viewing by selecting *Delete*.

Figure 1: Example antivirus profile page

Edit AntiVirus Profile default [View List icon] [Create New icon]

Name: default

Comments: scan and delete virus (maximum 63 characters)

	Web		Email		File Transfer	
	HTTP	SMTP	POP3	IMAP	FTP	IM
Virus Scan and Removal	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
File Filtering with: builtin-patterns	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Quarantine	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Apply

Drop-down list for viewing a specific profile or sensor

The actual configuration operations for all UTM profiles and sensors has also been modified to make configuration of sensors and profiles faster and more effective.

Archive inspection for antivirus profiles

Within antivirus profiles, you have more control about how the FortiGate unit handles file archives (for example .zip files). These options have been added because some archives cannot be virus scanned (for example, encrypted archives).

From the CLI you can select options to block all encrypted archives, block corrupted archives, block multipart archives, and write log messages whenever an archive file is received that cannot be virus scanned.

The following is an example.

```
config antivirus profile
edit av_1
config http
set options block-encrypted-archive block-corrupted-
archive block-multipart-archives log-unhandled-archive
end
```

Improved IPS default block rate

The IPS default block rate was improved so that the critical level, high level and medium levels are now higher. The critical level now has an 80 percent default block rate or higher; high level has 70 percent or higher; and the medium level has 50 percent rate or higher.

IPS signature rate count threshold

The IPS signature threshold has been enhanced to allow you to configure a signature that will not be triggered until a rate count threshold is met. This provides a better, more controlled recording of attack activity. For example, multiple login-failed events are detected in a short period of time, and an alert is raised.

This enhancement is enabled from the CLI. Once you enter a value for the rate count you can configure the rate limit mode optionally the packet fields to track. The command syntax is:

```
config ips sensor
edit <sensor_name>
config override
edit 0
set rate-count <integer>
set rate-duration <integer_seconds>
set rate-mode {continuous | periodical}
set rate-track {dest-ip | dhcp-client-mac | dns-domain |
none | src-ip}
end
```

IPS Predefined signature viewer

When you are viewing predefined signatures in *UTM Profiles > Intrusion Protection > Predefined*, you can more easily view information about each signature using the IPS Signatures Viewer.

Web Filter profiles

In *UTM Profiles > Web Filter > Profiles*, the Web Filter Profile Configuration Settings page contains a complete redesign of what was previously there. Previously, there was FortiGuard Web Filtering and FortiGuard Web Filtering Overrides; these are now in the CLI. Filters can be added to the profile. A filter is a category or banned word.

The Web Filtering feature contains the following new features and changes to existing features. The Web Filter menu no longer contains the submenu Web Content Filter; however, these settings are still available in the CLI.

Figure 2: The default configuration settings on the Edit Web Filter Profile page

Previous web filter profiles are carried forward and those settings are merged into the new redesign. If you want to view the previous settings in FortiOS 4.0 MR3, use the `show webfilter profile <name>` command to view the entire previous settings.

Depending on what settings you need to configure within the web filter profile, you may need to have access to both the CLI and web-based manager. The FortiGuard Web Filtering, FortiGuard Web Filtering Overrides and web content filtering are now configured in the CLI.

In a web filter profile, you can also include keywords that may appear in a search engine that a user enters in a search engine. These keywords are logged in the web filter log. The keywords are entered, separated with a comma, in the *Search Engine Keyword Filter* field.

Web Filtering Overrides

Web filtering overrides are now simplified and are profile-based. Profile-based overrides are web filter profiles that contain only overrides, and these overrides allow a rule to be created that changes the web-filter profile that applies to a user. The override feature is extended to apply to all features within the web-filter profile, and an override link appears in all related blocked pages where the user can override the block and continue on.

When you want to create a web filtering override, create a new web filter profile that is specifically for overrides. These configuration settings are available only in the CLI. These settings in the new override profile are configured by the administrator, and the administrator has complete control over the changes for the user.

The command syntax for configuring a new profile-based web filter override is as follows:

```
config webfilter {override | override-user}
edit <id_number>
set expires <yyyy/mm/dd hh:mm:ss>
set initiator admin
```



```
set status {enable | disable}
set old-profile <old_profile_name>
set new-profile <new_profile_name>
set scope {ip | ip6 | user |user-group}
set user <user_name>
set user-group <user_group_name>
end
```



For the option `initiator`, its value is always `admin`.

In the above command syntax, you see the commands `old-profile` and `new-profile` and its these commands that control how the override rule is applied. For example, if a user is browsing using a session that contains an old web filter profile applied, then the new-profile is used instead. If a user browses using a session that has a profile other an old profile, their profile will not be changed to a new profile. A user may be able to create a new override rule if the configuration permits it; however, only one override rule is allowed per user/profile pair.

With a web filter profile-based override, you can modify the URL filter list or add local ratings to deal with the extraction of offsite URLs. However, the function of this new override may change a user's FortiGuard categories, URL-filter list and so on.

When upgrading, existing web-filter profiles are carried forward but will not work until the administrator modifies them to the new settings. Any existing rules, both user and administrative, are not carried forward because there is no way to change the old override type to profiles for each rule without running out of profiles. This concerns only administrative overrides.

Application Control Sensors and filters

Application Sensors are similar to IPS or DoS sensors and replace the application control lists available in FortiOS 4.0 MR2.

An Application sensor contains application filters which you can configure to select individual applications to control or you can use categories, vendor names, behavior types, technology types, protocols, and tags to select groups of related applications. Application filters allow you to monitor, block, and reset sessions for single applications or groups of applications.

Application filters also allow you to apply shared traffic shaping to applications in the filter. You can apply forward and reverse traffic shaping and if the traffic shaper includes DSCP (or DiffServ) settings, these are also applied to applications specified in the filter. You can also set the session TTL for different applications and enable packet logging for applications.

Fortinet is constantly adding more applications to application control. Recent additions include the ability to individually monitor and block many Facebook applications.

Geography-based filtering for firewall addresses

Geography-based filtering for firewall addresses allows you to create a firewall address consisting of the name of a country. You can then add this address to a security policy to match traffic from any IP address assigned to that country. The list of countries and IP addresses that the FortiGate unit uses to identify the country of origin of an address is based on historical data compiled from the FortiGuard network.

For example, to configure a security policy to allow connections to multiple branch offices in Brazil (headquarters are in United States); the source address in this particular policy is *Any*, destination address is Brazil (geographic firewall address) and the *Action* is *Allow*.

To add a geography-based firewall address from the web-based manager, go to *Firewall Objects > Address*, select *Create New*, set *Type* to *Geography* and select a country name.

Use the following command to add a geography-based firewall address for Brazil:

```
config firewall address
  edit <addr_name>
    set type geography
    set country BR
  end
```

In the command you set the `country` to the two-letter abbreviation for the country name. In the example, BR is the abbreviation for Brazil.

You can use the following command to view information about geography-based addressing. The command does not display information about the entire address database, but displays country and address information for the countries that have been added to firewall addresses.

```
diagnose firewall ipgeo {country-list | ip-list | ip2country}
```

Where:

`country-list` lists all of the countries that have been added to a firewall address.

`ip-list` lists the IP addresses of a specified country or all of the countries added to firewall addresses.

`ip2country` displays the country of origin for a specified IP address. The address must be assigned to one of the countries that has been added to a firewall address.

For example, use the following command to view the countries that have been added to a firewall address. The example command output shows that a firewall address has been added for Brazil.

```
diagnose firewall ipgeo country-list
Total countries loaded:1

BR
```

DLP document fingerprinting

DLP document fingerprinting is a new feature that allows you to better protect your network from the loss of specific documents. Document fingerprinting, in this sense, is a method of identifying a document. This method breaks up files into chunks, taking a checksum of those chunks and using that checksum as the fingerprint. The fingerprint is then applied to a DLP filter rule within a DLP sensor which is then used during the scanning process of DLP activity.

DLP document fingerprinting is configured in *UTM Profiles > Data Leak Prevention > Document Fingerprinting* and then a DLP filter rule is applied within a DLP Sensor in *UTM Profiles > Data Leak Prevention > Sensor* to instruct the FortiGate unit to look for document fingerprints when scanning DLP activity on a security policy. A percentage parameter set in the DLP sensor is used when the unit is trying to match the file chunks.

For example, you transfer a file that is on the server (or uploaded), it will match 100 percent; a truncated file on the server will be matched 100 percent except for possibly the first or last chunks that may have a different checksum because the boundaries are different. The same is true for a file that is partially copied into another file; if that part is large enough, it will match but at a low percentage.

All documents in the source, as well as the ones you uploaded individually, are pre-scanned. This means that the task of breaking the files into checksums occurs soon after creating them and are all put into the database on the FortiGate unit.

There is an option to upload archived files and have those archived files fingerprinted as well, however, this is for only individual files that are configured in *Manual Document Fingerprints* on the DLP Fingerprint page.



DLP document fingerprinting is available only on FortiGate models with internal hard drives or flash drive storage.

Internet Content Adaptation Protocol (ICAP)

The Internet Content Adaptation Protocol (ICAP) is supported in this release. ICAP is a light-weight response/request protocol that allows the FortiGate unit to offload HTTP traffic to external servers for different kinds of processing. ICAP is often used for offloading virus scanning and web filtering but has many other applications.

If you enable ICAP in a security policy, HTTP traffic intercepted by the policy is transferred to an ICAP server in the ICAP profile added to the policy. Responses from the ICAP server are returned to the FortiGate unit which forwards them to an HTTP client or server.

You can offload HTTP responses or HTTP requests (or both) to the same or different ICAP servers.

Example ICAP sequence for an ICAP server performing web URL filtering on HTTP requests

- 1 A user opens a web browser and sends an HTTP request to connect to a web server.
- 2 The FortiGate unit intercepts the HTTP request and forwards it to an ICAP server.
- 3 The ICAP server receives the request and determines if the request is for URL that should be blocked or allowed.
 - If the URL should be blocked the ICAP server sends a response to the FortiGate unit. The FortiGate unit returns this response to the user's web browser. This response could be a message informing the user that their request was blocked.
 - If the URL should be allowed the ICAP server sends a request to the FortiGate unit. The FortiGate unit forwards the request to the web server that the user originally attempted to connect to.

When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

Example of adding ICAP to a security policy

The following is an example of configuring the ICAP feature on the FortiGate unit and applying an ICAP profile to an existing security policy.

1 Log in to the CLI.

2 Enter the following to configure the ICAP server:

```
config icap server
edit icap_server
set ip-address 172.16.122.151
set ip-version 4
set max-connections 25
set port 453
end
```

3 Enter the following to configure the ICAP profile to then apply to a security policy:

```
config icap profile
edit icap_profile_1
set request enable
set request-failure error
set request-path 1220
set request-server icap_server
set response enable
set response-failure error
set response-path 1225
set response-server 172.16.122.151
set streaming-content-bypass enable
end
```

4 In the config firewall policy command, apply the ICAP profile to policy 1:

```
config firewall policy
edit 1
set icap-profile icap_profile_1
end
```

Troubleshooting ICAP

You can use the following diagnose commands when troubleshooting ICAP.

```
diag system icap server list <name>
```

Displays a list of all servers or specified servers.

```
diag system icap profile list <name>
```

Displays information concerning total sent and responses, last connection attempts and host-bypass count.

Profile Group

The Profile Group feature is now included in the web-based manager of some models. This feature was previously only found in the CLI. Profile groups are groups of UTM profiles and sensors, which includes protocol options, and are often applied to security policies. By configuring a group of profiles and sensors, you can easily apply them to a security policy at once, instead of enabling them one at a time.

In *UTM Profiles > Profile Group > Profile Group*, you can create profile groups and view the selected profiles from within the group, either when creating a new group or editing an existing one. You can view the profile or sensor that you are including in the group by selecting the *View* icon. You can also create a new profile or sensor from within the page by selecting the *Create New* icon (the plus sign icon).



In this release, profile groups cannot be applied to security policies within the web-based manager, only in the CLI.

Modem interface Improvements

The Modem interface feature has been updated to include settings for configuring 3G or 4G wireless modems, as well as other modems. A list of supported modems is available from FortiGuard and can be updated to include recently supported modems.

When you first go to *System > Network > Modem*, the configuration settings within the page have changed, and now there is a *General Settings* section and *External Modem* section. The *General Settings* section is for configuring the primary modem, for example an external modem. You can still configure the modem to be *Standalone* or *Redundant* from this page, as well the type of *Dial Mode*.

The *External Modem* section of the page allows you to configure the external modem or USB modem. When you select *Configure Modem*, you are automatically redirected to the Modem Configuration Settings page. This page displays the supported modems under the *Supported* section, and the *Custom* section displays the external modems that you have configured. By selecting *Update Now*, you can easily update the list of supported modems from FortiGuard.

3G/4G modem list available from FortiGuard

You can now access a list from FortiGuard that contains all support 3G and 4G modems. The list is available without a subscription. By default, a list of supported modems is available on the FortiGate unit; however, you can update this list at any time from FortiGuard.

The list is available within the Modem page, on the Modem Selection page. This page appears when you select *Configure Custom Modem*. You can either update this page or choose a modem from the list.

WiFi Extensions

FortiOS 4.0 MR3 includes many improvements and changes to the WiFi controller feature (formerly the wireless controller feature). Among the highlights are automatic AP provisioning (channel, power, etc.) using distributed ARRP, Rogue AP “on wire” detection, Rogue AP Suppression, Unified FortiWifi and FortiAP management, support for new FortiAP models, and the addition of the WiFi controller feature to FortiWiFi units.

WiFi controller redesign

On the web-based manager the former wireless controller has been renamed *WiFi Controller* and menus and submenus have been redesigned as follows:

- WiFi Network
 - SSID
 - Rogue AP Settings
- Managed Access Points
 - Managed FortiAP
 - Custom AP Profile
- Monitor
 - Client Monitor
 - Rogue AP Monitor

Within the WiFi Network menu, instead of configuring a virtual access point, you configure an SSID. The SSID can easily be configured to be up or down, using the *Administrative Status* option within the SSID.

When configuring an SSID from the web-based manager you can choose from any of the following consolidated security options:

SSID security option	Descriptions
WPA/WPA2-Personal	Supports both WPA and WPA2 for personal use. Select this option and add a 8 to 63 character preshared key.
WPA/WPA2-Enterprise	Supports both WPA and WPA2 for enterprises. Select this option and select a RADIUS server or authentication group to use to authenticate connections to the SSID.
Captive Portal	Supports captive portal authentication. Select this option and select user groups that can authenticate with the captive portal. You can also configure the appearance of the captive portal page.

Configure an SSID from the CLI using the `config wireless-controller vap` command. From the CLI the following additional security options are available:

- open
- wep128
- wep64
- wpa-only-enterprise
- wpa-only-personal
- wpa2-only-enterprise
- wpa2-only-personal

Captive portal enhancements

The captive portal security option was previously available; however, in this release it has been enhanced and given additional replacement messages so that you can have specific pages for specific actions, such as when a failed login attempt or a declined disclaimer. The captive portal security is now more streamlined and is applied within an SSID, instead of previously being applied within a security policy.

Rogue AP detection and reporting

The PCI DSS regulatory compliance requires quarterly site surveys for unauthorized wireless access points on their networks to prevent data leakage, as well as assurance that trusted access points are running the latest WPA or WPA2 enterprise encryption. In this release, you can now easily gather this information and view it on the FortiGate unit.

From the WiFi Controller menu, you can configure the FortiGate unit to gather information on rogue APs, monitor them, and then take the information gathered from logs and generate a report.

The Rogue AP Settings submenu enables detecting rogue wireless access points. The option, *Enable On-Wire Rogue AP Detection Technique*, is a special detection technique that allows the FortiGate unit to help identify rogue APs that may be performing a bridging function, routing or NAT.

Rogue AP Suppression

Rogue AP suppression is now supported on the FortiWiFi and FortiAP units. This feature is available only when there is at least one radio signal dedicated to Rogue AP detection. On a FortiWiFi unit, this feature is available only when in dedicated detection mode. Rogue AP suppression is also not available for background Rogue AP scans.

On-wire Rogue AP detection

The on-wire scan feature allows you to detect if a rogue AP is connected to a wired network. A rogue AP poses a higher security risk if that rogue AP is an unmanaged AP and connected to an organization or company's wireless network. A rogue AP is an AP that is not managed by the controller.

To enable on-wire scan rogue AP detection go to *WiFi Controller > Wireless Network > Rogue AP Settings* and select *Enable On-Wire Rogue AP Detection Technique*.

Custom AP profiles

Previously, there were access point (AP) profiles that you could configure from within the WiFi Controller menu. These profiles are still available, however, to view them in the web-based manager you must enable the feature using the following command.

```
config system global
    set gui-ap-profile enable
end
```

Under Managed Access Points, the previous AP Profile menu has been replaced with the Custom AP Profile menu that contains a selection of default AP profiles. These default profiles have been designed to be a good starting point for many wireless network applications using FortiWiFi or FortiAP units. You can customize the default profiles for your needs and create new profiles.

Distributed ARRP (Automatic Radio Resource Provisioning)

For FortiAP units, each unit needs to autonomously and periodically determine the best channel that is best suited for communication. The distributed ARRP feature allows FortiAP units to select their channel so that they do not interfere with each other in a larger square footage network scenario.

The distributed ARRP behaves in the following way:

- Each FortiAP unit independently scans the available channels, measuring interface and channel utilization, and then selecting the channel with the least interface and then lowest utilization for communication.

- The FortiAP unit periodically performs this scan in the background to determine if any conditions have changed. This periodical scan is every ten minutes (by default).
- If any conditions have changed, the FortiAP unit signals all clients to move to a newly selected channel.

Log messages are recorded to reflect when the channel was changed by the FortiAP unit, and debug logs are also recorded to reflect the decision of the distributed ARRP algorithm for all runs. The ARRP algorithm is automatically on by default if multiple channels are selected in the web-based manager. If a single channel is selected, the ARRP algorithm becomes benign.

WiFi monitor

The WiFi Monitor menu is a new WiFi Controller feature in this release. It merges the previous monitoring menus into the new Monitor menu. There are two submenus, Client Monitor and Rogue AP Monitor.

The Client Monitor submenu allows you to view information about wireless clients of your managed access points. On the Client Monitor page, several columns have changed, as well as a new column added, called Auth. The following columns are no longer available on the page:

- Bandwidth Rx
- Bandwidth Tx
- Idle Time

Previously, the Client Monitor was available in *WiFi Controller > Wireless Client > Wireless Client*.

The Rogue AP Monitor allows you to view information about access points that may be rogue APs. Several columns have been removed. There are two new columns, Manufacturer and Security Mode. From the monitor you can also mark and suppress APs.

New WiFi commands

The following are new commands regarding WiFi server certificates and user group authentication for WiFi.

Certificate commands:

```
config system global
    set wifi-certificate <cert_name>
    set wifi-ca-certificate <ca_cert_name>
end
```

Authentication user group commands:

```
config system interface
    edit <wlan>
        set wifi-auth usergroup <user_name>
    end
```

The `wifi-auth usergroup` command is available only if WPA-Enterprise or WPA2-Enterprise option is selected as the security mode.

The authentication user group commands also apply to WiFi Controller.

Strong Authentication Enhancements

FortiOS 4.0 MR3 new strong authentication features include support for the FortiToken two-factor (2-factor) authentication solution as well as two-factor authentication using Email or SMS. Additional new strong authentication features include improved multiple user group support, dynamic profiles, and PKI authentication enhancements.

FortiToken support

The FortiToken-200 device provides time-based, one-time passwords (OTP) that are based on the Open Authentication (OATH) standard. Using FortiToken allows organizations to deploy a two-factor authentication solution that reduces the risk of compromise created by alternative single-factor authentication systems relying on, for example, static passwords. The FortiToken enables administrators with the need for two-factor authentication to offer enhanced security for both remote and on-premise users. The FortiToken-200 is a part of Fortinet's broad multi-factor authentication product strategy; it ensures that only authorized individuals access your organization's sensitive information; enabling business, protecting your data, lowering IT costs, and boosting user productivity.

The FortiToken-200 provides a secure one-time password (OTP) that is entered along with regular login credentials whenever authentication is required.

Each FortiToken device contains a serial number (located on the back of the device), a six-digit LCD display, and a small button. The serial number is used to activate the hardware token generator. When you press the small button, the LCD displays a six-digit token password code that is used in two-factor authentication. Two-factor authentication is authentication that requires an additional password or code that a user must enter to successfully authenticate in addition to their own user name and password.

The FortiToken device must be activated and synchronized with the FortiGate unit before it can be used for authentication purposes.

The FortiToken behaves as follows:

- FortiToken's serial number is added in the list in *User > FortiToken > FortiToken*. This serial number is a number containing 16 case-sensitive characters which is located on the back of the device. The serial number is used only in this way.
- The FortiToken is activated by selecting *Activate* on the FortiToken page. During the activation process the serial number is encrypted and sent to FortiGuard where it is verified as a valid FortiToken, and then activates the FortiToken on the FortiGate unit.
If you have a file containing the seed used to generate a token password code, you can import that file to the FortiGate unit from *User > FortiToken > FortiToken*.
- Synchronize the FortiToken by selecting *Synchronize* on the FortiToken page. This synchronizes the FortiToken's system time with the unit's system time so that both contain the same time period. The correct time period is necessary to verify that the token password code that is being used by a user is valid.
- If you have more than one FortiToken device, you must enter each one in *User > FortiToken > FortiToken*; then select each one in the list and for each one activate and synchronize. FortiOS does not support the activation and synchronization of multiple FortiTokens at one time. For example, you cannot select four FortiTokens and then select *Activate* and immediately after select *Synchronization*.

The FortiToken works with two-factor authentication in the following way:

- FortiToken is assigned to a user (for example a local user in *User > User > User*)

- The user logs in with the token password code they received in an email or text message on their mobile phone

The token password code that is generated is from a seed that is unique within each FortiToken. When a user uses the correct token code for the current time period, that token code provides proof that the user is in possession of the physical FortiToken. The token code changes every 60 seconds on the FortiToken device, to prevent replay attacks. For example, a person steals some of a user's token code to reuse at a later time.

Two-factor authentication

Two-factor authentication provides a way to minimize security breaches due to stolen user credentials. Two-factor authentication requires the authenticating client to provide additional credentials beside a user name and password. You can add two-factor authentication to PKI users. SMS and Email token authentication is also supported. SSL VPN two-factor authentication supports FortiTokens.

Two-factor authentication is available for FortiGate administrators, local users and PKI users.

In *User > FortiToken > FortiToken*, the token password code from FortiToken is entered into the list by selecting *Create New*. You must activate and synchronize FortiToken so that you can use the generated token password code.

After synchronizing the token password code, you can then apply it where two-factor authentication is available within FortiOS. For example, FortiGate administrators can have the two-factor authentication enabled for their account in *System > Admin > Administrators*.

When FortiToken is used in a third-party IPsec client configuration, each user that has two-factor authentication enabled and configured must use the token password code when only a password is supported to gain access. This authentication using only a password is not supported when the password and token password code are sent in CHAP or MS-CHAP form, and the local user is authenticated using a remote server. This is because FortiOS is unable to extract back both the password and the token password code.

Example for configuring users with two-factor authentication

This example explains how to configure multiple users with the new feature two-factor authentication.

Your company requires remote access to the network for the following employees:

- two sales employees
- two employees that often work from home
- one remote FortiGate administrator for remote management

The company has purchased a FortiToken-200 for each employee that will be required to authenticate using two-factor authentication. The employees that will be logging in using two-factor authentication will be using SSL VPN. Each employee already has their own user account configured from a previous setup. You are only enabling two-factor authentication and notifying them of this new, additional log in credential.

To activate each FortiToken

- 1 Log in to the web-based manager of the FortiGate unit.
- 2 Go to *User > FortiToken > FortiToken*.

- 3 On the FortiToken page, enter the serial number of the first FortiToken, and repeat until all FortiTokens are entered.
- 4 Select *OK*.
- 5 For the first FortiToken in the list on the FortiToken page, select it and then select *Activate*.
- 6 After completing the activation, select the first FortiToken in the list and then select *Synchronization*.
- 7 Repeat steps 5 and 6 until all FortiTokens are synchronized.

The FortiTokens are now synchronized. Users can now be configured with the two-factor authentication. In the following procedure, the token password codes will be sent to users to their email accounts.

To configure employees with two-factor authentication

- 1 In the web-based manager, go to the location of the employee's account.
For example, *User > User > User*.
- 2 In the first employee's account, select *Enable Two-factor Authentication*.
- 3 Under *Deliver Token Code by*, select *FortiToken* and then select the FortiToken serial number of the FortiToken that the person will be using.
- 4 Select *Email to* and then enter the sales person's email address.
For example, *sales_1@example.com*
- 5 Select *OK*.
- 6 Repeat steps 3 and 4 to complete the rest of the employee's two-factor authentication settings.

The following procedure sends the token password code to the FortiGate administrator's mobile phone.

Enabling two-factor authentication for administrators

The new two-factor authentication is available for FortiGate administrator accounts in *System > Admin > Administrators*. Two-factor authentication is a way for you to add an additional log-in credential for users, which is a token password code. The token password code is provided by a device called the FortiToken.

FortiGate administrators with two-factor authentication must enter the token password code when logging in to the web-based manager. The token password code can be sent to the FortiGate administrator by either email or mobile phone in a text message. When an administrator with two-factor authentication first tries to log in to the web-based manager, a message similar to the following appears below *Password*.

An email message containing a Token Code will be sent to
<xxxxx@xxxxx.com> in a moment.

If *SMS* is enabled for sending the token password code to a mobile phone, the above message will reflect that. The administrator enters their token code in the *Token Code* field and then selects *Login*.

Figure 3: Example of an administrator who is logging in to the web-based manager for the first time who has two-factor authentication

To configure the FortiGate administrator with two-factor authentication

- 1 In the web-based manager, go to *System > Admin > Administrators*.
- 2 Edit a FortiGate administrator account.
- 3 In the Edit Administrator page, select *Enable Two-factor Authentication*.
- 4 Under *Deliver Token Code by*, select *FortiToken* and then select the FortiToken serial number that the administrator will be using.
- 5 Select *SMS*.
- 6 Select the mobile provider from the drop-down list beside *(Mobile Provider)*.
- 7 Enter the FortiGate administrator's phone number in the field beside *(Phone Number)*.
- 8 Select *OK*.

A text message containing the token password code is sent to their phone.

Multiple authentication group enforcement

Previously, when a user belonged to multiple user groups, this user could only access the group services that were within one group. With the multiple group enforcement feature, a user can now access the services within the groups that the user is part of. For example, userA belongs to user_group1, user_group2, user_group3, and user_group4; userA can access services within user_group1, user_group2, user_group3, and user_group4.

This feature is available only in the CLI and is enabled by default. The new command for this feature is `auth-multi-group` and checks all groups a user belongs to for firewall authentication. This new command is found in `config user settings`.

Dynamic Profiles

The Dynamic Profile feature, previously only found in FortiOS Carrier, is now available for all FortiGate models. Using the dynamic profile feature a FortiGate unit can dynamically assign a UTM profile group to a user authenticated with a RADIUS server. Dynamically assigning a UTM profile group means you can dynamically assign different levels of UTM inspection, web access and other UTM features. For information about dynamic profiles, see the [Authentication](#) chapter of the [FortiOS Handbook](#).

Hard-timeout enhancement

The new authentication hard-timeout feature ensures that users will always need to authenticate whenever their time expires. The timeout behavior is configured in the following ways:

- When the timeout behavior is set to be a hard timeout, this option forces all of the user's sessions to immediately end when the authentication timeout expires. This causes the user to re-authenticate.
- When the timeout behavior is set to be a hard timeout new sessions, this option keeps all existing sessions but forces new sessions on the same user, which that user then has to re-authenticate.

The following are the new commands you can use to configure authentication hard-timeout.

```
config user setting
  set auth-timeout <number_minutes>
  set auth-timeout-type {idle-timeout | hard-timeout | new-session}
end
```

PKI certificate authentication enhancement

PKI certificate authentication now supports the extraction of the user name from within the UPN field. This extraction allows users to log in without having to enter their user name.

This enhancement is available only in the CLI. The command syntax is as follows:

```
config user peer
  edit <peer_name>
    set ldap-mode {password | principal-name}
  end
```

The `principal-name` value extracts the user name from within the UPN field.

An option for “user group matching” is available in the `config user group` command as well. This option allows you to configure authentication to match PKI user groups. The command syntax to configure this feature is as follows:

```
config user group
  edit <group_name>
    config match
      edit <group>
        set server-name <server_name>
        set group-name <group_name>
      end
    end
```

The following is an example of how to configure this user group matching feature.

```
config user group
  edit sslvpn
    set sslvpn-portal full-access
    set member vmlg test
    config match
      edit 1
        set server-name vmlg
        set group-name
          cn=Internet,ou=test,dc=ay,dc=fortinet,dc=com
      end
    end
  next
  edit 2
    set server-name test
  end
```

```
set group-name
    CN=qa, OU=T1359, DC=AY, DC=FORTINET, DC=COM
end
end
end
```

NTLM authentication enhancements

There are two enhancements for NTLM, one for guest profile access and one for inspection of initial HTTP-User-Agent values. These two enhancements are configured in the CLI. The new commands are `ntlm-guest {enable | disable}` and `ntlm-enabled-browsers <browser_name>`, which are available under the `config firewall policy` command.

The `ntlm-guest` command provides guest access to users who fail NTLM authentication. The `ntlm-enabled-browsers` command allows users to access non-supported browsers without a prompt beforehand.

NTLM authentication is essentially enabled when you configure FSSO and enabled NTLM in the identity-based security policy. Any users and user groups associated with the security policy will use NTLM to authenticate without further configuration.

New PCI Compliance Features

FortiOS 4.0 MR3 improves PCI compliance support by enhancing WiFi rogue AP detection, adding Rogue AP suppression and enhancing Endpoint security features. For information about Rogue AP features, see [“Rogue AP detection and reporting” on page 113](#) and [“Rogue AP Suppression” on page 113](#).

Endpoint Security enhancements

In FortiOS 4.0 MR3 Endpoint NAC has been renamed Endpoint Control and is available on the web-based manager from *UTM Profiles > Endpoint Control* you can configure endpoint security profiles, view the Endpoint Security application database, and work with FortiClient installers. From *UTM Profiles > Monitor > Endpoint Monitor* you can perform endpoint monitoring.

Endpoint profiles include the warn option, which displays a “block” page but allows a user to choose to continue or not, as well as sends the information back to the client. Previously, when the FortiGate unit blocks a client, the unit quarantines the user but no information was sent back to the client. With this new option, within FortiClient, you can view a list of applications that the FortiGate unit requested, any applications that caused the client to be blocked by the unit, and any applications that cause a user to continue on even though a “block” page was triggered.

The *Client Installers* submenu provides information regarding FortiClient installation, version enforcement, and FortiGuard availability for updating to a recent FortiClient Endpoint versions. The *Profile* submenu provides configuration settings for endpoint profiles which are then applied to firewall policies. The settings within application sensor were merged into the settings that are available in the Profile submenu.

Network Vulnerability Scan

The Network Vulnerability Scan feature provides more granularity and options for network scanning. Network vulnerability scanning now includes Asset Definition, Scan Schedule and Vulnerability Result. You can access the network vulnerability scanner from *UTM Profiles > Vulnerability Scan*.

Asset Definition

The Asset Definition menu allows adding ranges, discovering assets or start scans. Ranges can easily be added by selecting *Create New* on the Asset Definition page. Assets are still configured the same as they were previously, but you can now add an IP address range to be scanned.

When a scan is being performed, the activity icon in the Scan Activity column displays the progress. The scan's results appear in the Discovered Hosts for <network> window. You can scroll down the list to view all discovered assets.

Scan Schedule

The Scan Schedule menu allows you to see the status of a scan (or to start a scan), the schedule settings, the type of vulnerability scan mode, and advanced settings as well. The scheduled scan applies to all assets or asset groups that are currently enabled.

The types of scans that you can schedule are quick, standard and full. Quick scanning examines a set of the most commonly used ports for vulnerabilities. Standard scanning examines a larger number of application ports, covering many known applications. This scanning covers TCP, Service Discovery and OS Discovery but UDP is disabled.

The full scan scans the full port range 1-65535 and looks for applications running on non-standard ports, examining them for vulnerabilities.

The advanced options that are available are as follows:

- *Enable TCP Port Scan*
- *Enable Service Detection*
- *Enable OS Detection*
- *Enable UDP Port Scan*

When a scan is processing, the following occur:

- All Host assets are discovered as specified in the asset definition
 - All discovered hosts are scanned for the configured sensors, port scans and so on.
- all IP Range assets are discovered as specified in the asset definition
 - authentication is not available, reducing scan capabilities
 - these IP ranges should be converted to Host assets as needed to perform a full scan
- scanning will run for each unique IP in the list, and up to the maximum number of IP addresses supported, per-platform.
- logs are recorded about the scanning activity

Vulnerability Result

The Vulnerability Result menu allows you to view, in both graphical and tabular form, the results of the network scan. Platforms that do not have SQL logging enabled, or no SQL logging available, will only have the graphical representation.

When viewing the table containing vulnerabilities, a table similar to the log viewer table, appears at the left side of the page.

Netscan asset authentication options

In the `config netscan assets` command, the following values are hidden when the value `addr-type` is set to `range`:

- `auth-unix`

- auth-windows
- unix-username
- unix-password
- win-username
- win-password

Feature Improvements to extend IPv6 support

Each new release of FortiOS brings more IPv6 feature support. FortiOS 4.0 MR3 is no exception. This release adds IPv6 firewall acceleration using XG2, XE2, CE4, and FE8 security processors, IPv6 support for the SSL VPN web portal, IPv6 support for firewall authentication, IPv6 support for SNMP, IPv6 over DHCP, Addition IPv6 features for OSPF NSSA (not so stubby area), and more information is displayed about IPv6 sessions in the dashboard session widget.

In FortiOS 4.0 MR3, IPv6 traffic can now be redirected for user authentication using local database, RADIUS, TACACS+, or LDAP

In this release, the IPv6 Policy page contains the option of including a section title within the IPv6 security policy list. IPv6 security policies support antivirus, web filter, email filter, DLP sensor, VoIP and ICAP UTM features. Local in security policies also support IPv6.

Top Session dashboard widget IPv6 support

The Top Session dashboard widget can now display IPv4 and IPv6 addresses. IPv6 addresses are displayed only when the *IPv6 Support on GUI* is enabled in *System > Admin > Settings*.

OSPFv3 NSSA extension

OSPFv3 NSSA now includes the `default-information-originate` and `external route summary` commands for IPv6. This helps you to configure the originating default information and the external route information for IPv6 addresses.

A `get` command was also introduced to show the OSPFv3 NSSA external LSAs in the database.

The following commands have been added to `config router ospf6`, in config area:

```
set type {regular | stub | nssa}
set nssa-translator-role {candidate | never | always}
set nssa-default-information-originate {disable | enable}
set nssa-default-information-originate-metric <integer>
set nssa-default-information-originate-metric-type {2 | 1}
set nssa- redistribute {enable | disable}
```

The following were added to the `config router ospf6` command:

```
set default-information-originate {disable | enable | always}
set default-information-metric <integer>
set default-information-metric-type {2 | 1}
set default-information-route-map {route-map-name}
```

The following were added to `config summary-address`:

```
set prefix6
set advertise {enable | disable}
set tag <tag_number>
```


The following `get` command was added: `get router info6 ospf database nssa-external`.

DHCP for IPv6

DHCP for IPv6 addresses is now supported in the CLI. DHCP IPv6 is similar to DHCP IPv4.

This release also introduces the rapid-commit option. Rapid-commit is the process whereby the DHCP client and the DHCP server use a rapid DHCP IPv6 two-message exchange. This provides a short cut and the messages that are exchanged are called the DHCP IPv6 “SOLICIT” and “REPLY” messages. The `rapid-commit` command is enabled or disabled in the CLI.

For more information about DHCP IPv6, see [RFC 3315](#).

Explicit proxy and web caching improvements

The explicit proxy feature provides additional options in this release, as well as new features, such as forwarding servers and a completely new explicit FTP proxy.

The web proxy feature also now provides two `diagnose` commands to list and clear web proxy users. The `diag wad user list` command lists existing users and the `diag wad user clear` clears all users or a specific user.

Explicit FTP proxy

An explicit FTP proxy can be configured from the web-based manager and the CLI. The explicit FTP proxy is in *System > Network > Explicit Proxy* and in the CLI it is `config ftp-proxy explicit` command syntax.

FTP users connect to the explicit proxy and then connect through the proxy to remote FTP servers.

To enable the explicit FTP proxy go to *System > Network > Explicit Proxy* and select *Enable Explicit FTP Proxy* and select *Apply*. Then go to *System > Network > Interface*, select the Interface on which to enable the explicit FTP proxy and select *Enable Explicit FTP Proxy*.

Enter the following command to enable the explicit FTP proxy from the CLI:

```
config ftp-proxy explicit
  set status enable
end
```

Enter the following command to enable the explicit FTP proxy on the internal interface:

```
config system interface
  edit internal
    set explicit-ftp-proxy enable
  end
```

Once the explicit FTP proxy is enabled on an interface you must create security policies with the *ftp-proxy* as the source interface to allow explicit FTP proxy traffic. For example, to allow FTP connections from the internal network to an FTP server on the Internet, enable the explicit FTP proxy on the FortiGate internal interface and add `ftp-proxy` to `wan1` security policies.

To connect to an FTP server through the explicit FTP proxy

The following steps are required when a user starts an FTP client to connect to an FTP server through the explicit FTP proxy. Any RFC-compliant FTP client can be used.

- 1 The user connects to the explicit FTP proxy by starting an FTP session with the explicit proxy. In this example, the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100.

For example:

```
ftp 10.31.101.100
```

- 2 The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:

```
Connected to 10.31.101.100.  
220 Welcome to Fortigate FTP proxy  
Name (10.31.101.100:user):
```

You can change the message by editing the FTP Proxy replacement message.

- 3 At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:

```
<proxy-user>:<proxy-password>:<server-user>@<server-address>
```

For example, if the proxy username and password are `p-name` and `p-pass` and a valid username for the FTP server is `s-name` and the server's IP address is `ftp.example.com` the syntax would be:

```
p-name:p-pass:s-name@ftp.example.com
```



If the FTP proxy accepts anonymous logins `p-name` and `p-pass` can be any characters.

- 4 The FTP proxy forwards the connection request, including the user name, to the FTP server.
- 5 If the user name is valid for the FTP server it responds with a password request prompt.
- 6 The FTP proxy relays the password request to the FTP client.
- 7 The user enters the FTP server password and the client sends the password to the FTP proxy.
- 8 The FTP proxy relays the password to the FTP server.
- 9 The FTP server sends a login successful message to the FTP proxy.
- 10 The FTP proxy relays the login successful message to the FTP client.
- 11 The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

Explicit Web Proxy Forwarding Servers (proxy chaining)

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with an already existing web proxy solution.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

If the explicit proxy feature is enabled and configured, you can apply a web proxy forwarding server to a web proxy security policy. Applying a forwarding server to the policy is the same as applying a UTM profile or sensor; select the *Web Proxy Forwarding Server* check box and then select a forwarding server from the drop-down list.

Authentication cookie for session-based authentication of explicit web proxy sessions

When configuring a web proxy security policy, you can now include a web-proxy cookie option which reduces the amount of authentication requests to authentication servers when session-based authentication is applied using explicit web proxy. The cookie will remember the user's session, which will then be used to map to an existing user, reducing the chance to require an authentication. This feature provides better load balancing, as well as latency.

The web authentication cookie is available only in the CLI. The `web-auth-cookie` command is used to configure this feature and is within the `config firewall policy` command.



The `web-auth-cookie` command is available only when session-based authentication is enabled.

Form-based user authentication for explicit web proxy

Previously, the explicit web proxy supported authentication only through the HTTP protocol using HTTP headers. A form-based user authentication for explicit web proxy is now available, which is similar to form-based authentication for regular security policies. A form-based authentication is used when a web page is returned to a web client which the user then authenticates with his or her user name and password. These credentials are then sent through HTTP Post request.

The form-based authentication for explicit web proxy authenticates the user and then redirects the user back to their own original URL, if the user authorizes access to the URL. This authentication is available only for IP-based authentication.

Web caching in security policies

Web caching can now be enabled in a security policy. When enabled, the FortiGate unit will apply web caching to HTTP traffic accepted by the security policy. This option is available only on FortiGate units that support WAN Optimization and web caching. Enabling web caching in a security policy is similar to enabling web caching in a WAN Optimization rule. However, enabling web caching in a security policy means you can also apply UTM options to web cached traffic in a single VDOM.

You can use this option to apply web caching for explicit web proxy traffic if the Source Interface/Zone is set to the web-proxy interface. Previously, web caching was enabled as part of the explicit proxy configuration. In this release, web caching does not need to be applied to all explicit proxy traffic.

Enabling web caching in a security policy can not apply web caching to HTTPS traffic. To apply web caching to HTTPS traffic you need to create a WAN optimization rule.

Web Caching in a security policy takes place before web caching in a WAN Optimization rule. So traffic accepted by a security policy that includes web caching will not be cached by the WAN optimization rule.

Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites on the FortiGate unit hard disk. Some HTTP content accepted by the security policy may not be cached. See [RFC 2616](#) for information about web caching for HTTP 1.1.



Logging and reporting enhancements

This section describes new FortiOS 4.0 MR3 Log and Reporting features including:

- [The FortiGate UTM Weekly Activity Report](#)
- [Log Access Improvements](#)
- [SQL logging enabled by default](#)
- [Sending DLP archives to multiple FortiAnalyzer units](#)
- [Remote logging configuration enhancements](#)
- [Log and Report Monitoring](#)
- [Log Message Enhancements](#)
- [SSL connection encryption level option over OFTP](#)
- [Uploading logs to a FTP server in text format](#)
- [Deleting all local logs, archives and user-configured report templates](#)
- [FortiGuard Analysis and Management Service \(FAMS\)](#)

The FortiGate UTM Weekly Activity Report

The FortiGate UTM Weekly Activity Report is available on FortiGate units with hard disks if logging to disk is enabled by going to *Log&Report > Log Config > Log Setting* and selecting Disk Logging and Archiving. When you enable Disk logging you can go to *Log&Report > Report Access* to view the FortiGate UTM Weekly Activity Report. You can browse through the sections of this report to view current bandwidth and application usage, web usage, email usage, threats, and VPN usage.

The data for the report is generated by saved SQL logging messages. By default logging to disk and SQL logging are enabled and the report is produced. If logging to disk is disabled the report does not appear. If logging to disk is enabled and SQL logging is not enabled the report appears but will not contain any data. If some report data appears and some does not the cause could be that only some types of SQL logging are enabled.

SQL logging is only enabled and configured from the CLI. Use the following command to enable SQL logging:

```
config log disk setting
  set status enable
config sql-logging
  set app-ctrl enable
  set attack enable
  set dlp enable
  set event enable
  set netscan enable
  set spam enable
  set traffic enable
  set virus enable
```

```
        set webfilter enable
    end
end
```

By default the UTM Weekly Activity Report is generated and saved weekly and from any report access page you can select Historical Reports to view previously generated reports. From the Historical reports list you can also download the reports in PDF format and delete them.

You can modify the default FortiGate UTM Weekly Activity Report to meet your requirements from any report page by selecting *Edit*. When you select *Edit*, the page refreshes in Editing mode. In editing mode can change the content and the appearance of the report pages, change the data displayed on individual report pages and add or delete report pages.

A modified report must be saved using the Save icon. Any modifications that are not saved, are lost.

A report consists of text, charts and images. Text elements are used to add titles and descriptive text to the report. Images are used to add graphics to the report. Charts are used to add text and graphical data to the report. You can add a bar chart, line chart, pie chart and table chart. When you add a chart you can customize its appearance as well as the data that the chart displays. To customize the data displayed you can choose from hundreds of predefined reports. Each report includes formatting settings and settings to extract data from the FortiGate log database.

FortiOS 4.0 MR3 includes the following new reports:

- traffic.bandwidth.apps.app_cat
- traffic.bandwidth.app_cats.user
- traffic.bandwidth.users
- traffic.sessions.apps.app_cat
- traffic.sessions.app_cats.user
- traffic.sessions.users
- traffic.bandwidth.apps.user
- traffic.bandwidth.users.app
- traffic.bandwidth.app_cats
- traffic.sessions.apps.user
- traffic.sessions.users.app
- traffic.sessions.app_cats
- traffic.bandwidth.wanopt
- web.allowed-request.sites.user
- web.allowed-request.users.web_cat
- web.allowed-request.web_cats
- web.blocked-request.sites.user
- web.blocked-request.users.web_cat
- web.blocked-request.web_cats
- web.requests.phrases.user
- web.requests.users.phrase
- web.requests.phrases

- web.allowed-request.users.site
- web.allowed-request.sites
- web.blocked-request.users.site
- web.blocked-request.sites
- web.bandwidth.sites.user
- web.bandwidth.users.site
- web.bandwidth.sites
- web.bandwidth.stream-sites.user
- web.bandwidth.users.stream-site
- web.bandwidth.stream-sites
- email.request.timeperiods.sender
- email.request.senders
- email.bandwidth.timeperiods.sender
- email.bandwidth.senders
- email.request.timeperiods.receiver
- email.request.receivers
- virus.count.viruses.user
- virus.count.users.virus
- virus.count.viruses
- virus.count.users
- virus.count.viruses.protocol
- virus.count.protocols
- attack.count.critical-attacks.user
- attack.count.users.critical-attack
- attack.count.critical-attacks
- attack.count.attacks.user
- attack.count.users.attack
- attack.count.attacks
- vpn.bandwidth.static-tunnels.user
- vpn.bandwidth.users.static-tunnel
- vpn.bandwidth.static-tunnels
- vpn.bandwidth.ssl-sources.user
- vpn.bandwidth.users.ssl-source
- vpn.bandwidth.ssl-sources
- vpn.bandwidth.dynamic-tunnels.user
- vpn.bandwidth.users.dynamic-tunnel
- vpn.bandwidth.dynamic-tunnels

Viewing the current and historical reports

Going to *Log&Report > Report Access* you can view current data in the FortiGate UTM Weekly Activity Report.

You can also select *Historical Reports* to view previously generated FortiGate UTM Weekly Activity Reports. When you select *Historical Reports* on the Viewing default layout page, you are automatically redirected to the Historical Reports page where you can view, download, and delete generated reports.

Historical Reports page	
Lists all generated reports. You can remove generated reports from this page or return to the previous page, the Viewing default layout page.	
Return to Layout	When selected, you are automatically redirected to the Viewing default layout page.
Delete	<p>Removes a report from within the list.</p> <p>To remove multiple reports from within the list, on the page, in each of the rows of the reports you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all reports from the list, on the page, select the check box in the check box column, and then select <i>Delete</i>.</p>
Report File	The report name that the FortiGate unit gave the report. This name is in the format <scheduledtype>-<report_title>-<yyyy-mm-dd>-<start_time>. For example, Once-examplerreport_1-2010-09-12-103044, which indicates that the report titled examplereport_1 was scheduled to generate only once and did on September 12 at 10:30 am. The hour format is hh:mm:ss.
Started	The time when the report began generating. The format is yyyy-mm-dd hh:mm:ss.
Finished	The time when the report finished generating. The format is yyyy-mm-dd hh:mm:ss.
Size (bytes)	The size of the report after it was generated. The size is in bytes.
Other Formats	The other type of format you chose the report to be in, for example PDF. When you select PDF in this column, the PDF opens up within the page. You can download the PDF to your local PC from this page as well.

Creating custom reports from the CLI

You can add additional reports from the CLI by adding datasets, charts, layout, style, summary, and themes for reports; however, these options are available only from the CLI. When you add a report from the CLI the report layout does not appear on web-based manager. You can review historical reports for CLI-configured reports in the same way as FortiGate UTM Weekly Activity Reports.

Log Access Improvements

The Log & Archive Access menu contains the following changes to existing features, as well as support for downloading log messages directly from the FortiGate unit to your PC.

- [Viewing log messages](#)
- [Filtering log messages](#)
- [Downloading log messages](#)

Viewing log messages

When you are viewing log messages within the Log & Archive Access menu, you will find detailed information about the log messages at the bottom of the page. For example, you are viewing event log messages in *Log&Report > Log & Archive Access > Event*, and you see the first log message, in detail, in a table below the *Log location*: `<log_storage_device>` and page controls.

By selecting the down arrow beside *Detailed Information*, you can view this detailed information about log messages either at the bottom of the page, or on the right side of the page. You can also select *Hidden*, which hides the table.

Figure 4: Viewing event log messages with the default Bottom viewing option selected

#	Date	Time	Level	Sub Type	ID	User Interface	Action	Message
1	2011-02...	11:01:...	informati...	admin	320...	http(172.20.120.23)	login	Administrator admin logged in successfully from http(172.20.120...
2	2011-02...	10:59:...	informati...	admin	320...	http(172.20.120.23)	logout	Administrator admin logged out from http(172.20.120.23)
3	2011-02...	10:59:...	alert	admin	324...	jsconsole		Configuration is changed in the admin session
4	2011-02...	10:59:...	informati...	admin	320...	jsconsole	logout	Administrator admin logged out from jsconsole
5	2011-02...	10:23:...	informati...	admin	321...			interface wan1 gets a DHCP lease, ip:172.20.120.231, mask:255...
6	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
7	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
8	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
9	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
10	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
11	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
12	2011-02...	09:55:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
13	2011-02...	09:49:...	informati...	dhcp	260...			Client requests IP address/configuration parameters
14	2011-02...	09:39:...	notice	auth	430...	UNKNOWN(2.2.2.2)		forticlient msg
15	2011-02...	09:39:...	notice	auth	430...	UNKNOWN(2.2.2.2)		forticlient msg

Log Location: Disk

1 / 21

Date	2011-02-04	Time	11:01:13
Level	information	Sub Type	admin
ID	32001	User	admin
User Interface	http(172.20.120.23)	Action	login
Status	success	Reason	none
Profile Name	super_admin	Message	Administrator admin logged in successfully from http(172.20.120.23)

At the bottom of the list of logs on the page, before page controls, the *Log location*: `<log_storage_device>` indicates where the logs are being stored, such as the local hard disk or memory.

Filtering log messages

Previously, when filtering log messages, you had to select the Filter icon within each column and then indicate the information that you wanted filtered. You can now use *Filter Settings*, providing an easier way to filter the information on the page without using the Filter icons. The Filter icons still indicate if filtering is enabled for that column.

Downloading log messages

Log messages can now be downloaded in Raw format directly from within the Log Access menu. For example, in *Log&Report > Log & Archive Access > Event*; within the Event page, select *Download Raw Log*.

All log messages, including archived log messages, can be downloaded from the FortiGate unit to the management computer at any time. The downloaded file is a text file, which can be viewed on a text editor, such as Notepad. The log file name is in the format <log name><number>.log. For example, elog0101.log. The last number changes to reflect the log type number, such as DLP, which is nine (for example, dlog0109.log).

New Unified UTM Log Access

The new UTM Log submenu in *Log&Report > Log & Archive Access > UTM Log* provides a central location for all UTM-related log messages. These include virus, attack, DLP, application control, email filter, and web filter log messages. On the UTM Log access page includes a new column called *UTM Type* which indicates if the UTM feature that generated the log message.

Figure 5: The UTM page in *Log&Report > Log Access > UTM*

#	Date	Time	Level	Sub Type	ID	UTM Type	Message	Src	Dst	User	Src Port	Dst Port	Packet Log
1	2011-01-27	09:36:15	warning	infected	8192	AntiVirus	File is infect...	1.1.1.1	2.2.2.2		2560	5120	
2	2011-01-27	09:36:15	warning	filename	0440	AntiVirus	File is blocke...	1.1.1.1	2.2.2.2			5120	
3	2011-01-27	09:36:15	warning	ftgd_bik	13056	Web Filter		1.1.1.1	2.2.2.2	user	2560		
4	2011-01-27	09:36:15	alert	signature	16384	Attack	N/A			user	2560	5120	
5	2011-01-27	09:36:15	notice	smtp	20480	Email Filter		1.1.1.1	2.2.2.2	user	2560		
6	2011-01-27	09:36:15	warning	dip	24576	DLP		1.1.1.1	2.2.2.2	user			
7	2011-01-27	09:36:15	informat...	app ctrl all	28672	Application Co...						5120	

Log location: Disk			
1 / 1			
Date	2011-01-27	Time	09:36:15
Level	warning	Sub Type	infected
ID	0192	Message	File is infected.
Status	passthrough	Service	mm1
Src	1.1.1.1	Dst	2.2.2.2
Src Port	2560	Dst Port	5120
Src Interface	lo	Dst Interface	eth0
Relaxed IDS	10048	Control threshold	245

When viewing logs in Raw format, the downloaded UTM log file is called ulog and contains all the UTM-related log types, such as virus and attack. There is no log field called UTM Type in the log message when viewing them in the Raw format. The type and subtype fields indicate which log file the log message is associated with. For example, type=virus and subtype=filename.

SQL logging enabled by default

SQL is not enabled by default on models with an internal hard disk, such as a FortiGate-60C, as well as models with a removable hard drive when the disk is inserted into the FortiGate unit.

After upgrading to this release, a window appears when logging into the web-based manager. In the window, you can enable SQL logging when you select Go. This option does not immediately send logs to the FortiGate unit's local hard disk or removable hard disk; traffic must be flowing through the unit as well as UTM profiles and/or sensors applied to security policies.

The window appears when both the following are present:

- the model contains an internal or removable hard disk
- no SQL logging options are enabled

If you decide you would rather enable SQL logging later, select *Remind me Later*, which will prompt you when you log into the web-based manager again.

When SQL is enabled from the window, the FortiGate unit converts an previous logs to SQL format, and all log categories are that were previously enabled for disk logging are written in SQL format.

Sending DLP archives to multiple FortiAnalyzer units

When configuring multiple FortiAnalyzer units, you can now include sending DLP archives to both the second and third FortiAnalyzer units. This enhancement allows you to ensure DLP archives are not lost when logging to multiple FortiAnalyzer units.

Remote logging configuration enhancements

Remote logging to a log device is now configured mostly in the CLI, except you can configure uploading logs to a FortiAnalyzer unit or FAMS in either the CLI or web-based manager. However, you must configure when to upload the logs from the CLI, since the time period is not supported in the web-based manager until after the time period is configured in the CLI.

SQL logging is enabled by default for those models that have SQL databases. If you want to disable or enable certain SQL logs, including archiving, you must use the CLI.

Figure 6: Log settings in FortiOS 4.0 MR3

When configuring logging to a FortiAnalyzer unit, you can control the buffer rate to the FortiAnalyzer unit. This is available only in the CLI. The buffer size is between 20 to 20 000.



Previously, you could upload logs to an AMC disk; however, this feature has been removed because of the new feature of remotely storing and uploading logs to a FortiAnalyzer unit and FAMS server.

Log and Report Monitoring

The new Monitor submenus allow you to view monitored network activity on the FortiGate unit. In *Log&Report > Monitor > Logging Monitor*, you can view the log activity being recorded by the FortiGate unit on a weekly basis.

Logging Monitor

The Logging Monitor allows you to view the log activity that is being recorded by the FortiGate unit. The information displays as a bar chart and contains information regarding the total number of logs recorded by the unit on each day of the week. For example, on Wednesday of this week there were a total of 30 log entries recorded by the event log.

When you select a bar in the bar chart, you are automatically redirected to the Log Activity for <day of week> page. On this page you can view the logs for that day and the number of entries for that log file that occurred. For example, you select Wednesday's bar on the Logging Monitor page; you are redirected to the Log Activity for Wednesday page, where the logs for that day display. When you want to return to the Logging Monitor page, select *Return*.

Log Message Enhancements

There are several enhancements, as well as changes, that occurred for logs in this release. For example, event logs contain a new subtype called DNS.

This topic includes the following:

- [Event logs](#)
- [Other-traffic logs](#)
- [Chat message log support for MSNP21](#)



In antivirus logs, the URL address now states the type of protocol used instead of always using "http://". For example, in an ftp-over-http traffic log, the URL starts with ftp://.

Event logs

There are two new subtypes that have been added to the event log file, config and dns. The following explains each one.

A new subtype was added to event logs, called DNS. This new log message provides information about any DNS look-up that occurred. The option is enabled within the Event Log page (*DNS lookup event*), or within the CLI.

There is only one log message that occurs within the event log.

The following is an example of an event-dns log message.

```
2010-08-13 20:05:43 log_id=0108050000 type=event subtype=dns
vd=root pri=information policyid=1 src=172.16.120.166
dst=10.10.1.10 src-intf="internal" dst-intf="wan1" user="user1"
group="group123" dns_name="xx.example.com" dns_ip="172.55.154.199"
```

The event-config log messages provides detailed information about what setting was changed by a user. For example, a user disabled the explicit web proxy event on the Event Log page.

You can enable this subtype within the Event Log page, by selecting *Configuration change event*, or in the CLI. By default, this option is disabled.

The following is an example of an event-config log message:

```
2010-09-15 10:15:55 log_id=010 type=event subtype=config vd=root
pri=information vd=root user="admin" ui="GUI(10.10.10.1)"
action="edit" cfg_tid=1179790 cfg_path="log.eventfilter"
cfg_attr="wan-opt[enable->disable]" msg="Edit log.eventfilter"
```

Within the event log file, a log message containing information about an explicit web proxy event is recorded when enabled in *Log&Report > Log Config > Event Log*. The check box beside *Explicit web proxy event* must be selected so that this log can be recorded by the unit. This option is also available in the CLI.

Event-system

There are now two specific log messages that indicate when the system starts up and when it shuts down. These log messages are included in the event-system logs, and log message 20202 indicates when the system started up, and log message 20203 indicates when the system shut down.

The following are examples of these two event-system log messages:

```
2010-09-12 10:24:02 log_id=0100020203 type=event subtype=system
vd=root pri=information action=daemon-shutdown daemon=getty pid=68
msg= "Daemon getty shut down"

2010-09-12 10:24:02 log_id=0100020203 type=event subtype=system
vd=root pri=information action=daemon-startup daemon=cauploadd
pid=94 msg "Daemon cauploadd started."
```

Traffic logs

There are two new enhancements for traffic logs. Additional information has been added to other-traffic logs and a new subtype introduced. The new webproxy-traffic subtype for traffic logs indicates activity regarding web proxy traffic that was detected using a web-proxy security policy.

Other-traffic logs

When viewing other-traffic logs in the web-based manager, you will see additional information such as IM and P2P application information, as well as two icons in the status field that indicate the status of the traffic logs of 6 and 5. The status icons that appear in the web-based manager are a green check mark or a circle with a line through it. When you move your mouse over the icon, it indicates what the icon is, either *accept* (which is the green check mark), or *deny* (the circle with a line through it).

Chat message log support for MSNP21

In Windows Live Messenger 2011, a new protocol was introduced called Microsoft Notification Protocol 21 (MSNP21) which handles chat messages. The FortiGate unit now supports logging of chat messages that use this new protocol.

The FortiGate unit detects the protocol by following the same path as previously for IM logging. These logs are found in the DLP archive logs.

SSL connection encryption level option over OFTP

The SSL connection encryption level option for SSL connections that occur over OFTP, such as FortiGate to FortiAnalyzer, is now available. This type of connection provides a way to customize the level of SSL encryption over OFTP for these connections.

The commands are as follows:

```
config log fortianalyzer setting
    set enc-algorithm {default | high | low | disable}
end
config log fortianalyzer2 setting
    set enc-algorithm {default | high | low | disable}
end
config log fortianalyzer3 setting
    set enc-algorithm {default | high | low | disable}
end
config log fortianalyzer override-setting
    set enc-algorithm {default | high | low | disable}
end
config log fortiguard setting
    set enc-algorithm {default | high | low | disable}
end
config system central-management
    set enc-algorithm {default | high | low | disable}
end
config log disk setting
    set upload-ssl-conn {default | high | low | disable}
end
```

When you select the `default` option, you are choosing to have the SSL communication encryption with high and medium encryption algorithms.

Uploading logs to a FTP server in text format

Logs can now be uploaded in text format to a FTP server. This provides more flexibility for saving logs in a specific format for viewing later on. This is available only for FortiGate units with hard disks and only for uploading to a FTP server.

Logs that are saved in text format can be viewed in a text editor, and these logs are in Raw format. Raw format is a type of format that displays log messages as they would appear in the log file.

Example for uploading logs to a FTP server in text format

In this example, an administrator is configuring logging to the FortiGate unit's disk, as well as specifying uploading logs to an FTP server in text format.

```
config log disk setting
    set status enable
    config sql-logging
        set app-crt1 enable
        set attack enable
        set dlp enable
        set event enable
        set netscan enable
        set traffic enable
        set spam enable
        set traffic enable
        set virus enable
        set webfilter enable
    end
    set ips-archive enable
```

```
set storage Internal
set diskfull overwrite
set log-quota 50
set report-quota 50
set upload enable
set upload-destination ftp-sesrver
set uploadip 172.16.120.154
set uploadport 443
set uploaduser user_1
set uploadpass 123456789
set uploadaddir c:\logs_fgt50B
set uploadtype appctrl attack dlp event spamfilter traffic
    virus webfilter
set uploadzip enable
set upload-format text
set uploadsched enable
set uploadtime 7
set drive-standby-time 19800
set upload-delete-files disable
set sql-max-size 65536
set sql-max-size-action overwrite
set sql-oldest-entry 1024
end
```

Deleting all local logs, archives and user-configured report templates

The new `execute` command, `execute log-report reset`, deletes all local logs, log archives and user-configured report templates on the FortiGate unit. However, this command also restores the default FortiOS UTM Activity report to its original default settings, if the default report has been modified. The user-configured templates are the themes that you have configured from scratch for reports.

FortiGuard Analysis and Management Service (FAMS)

Enhancements, such as support for FortiAnalyzer units, was introduced. The following explains this support and additional enhancements:

FortiAnalyzer with FAMS support

The FortiAnalyzer unit now provides support for FAMS. The FAMS subscription service allows you to backup and store logs on a FortiAnalyzer unit. This provides additional archival storage as well as a back up solution in the event the FortiAnalyzer unit becomes unavailable.

Logs on the FortiAnalyzer unit are sent on a regular basis, based on a scheduled time period, to the FAMS server.

FAMS enhancements

The FortiGate unit can now be configured to upload recently recorded logs to the FAMS servers on a regular basis, similar to how the unit uploads logs to the FortiAnalyzer unit on a regular basis.

When you are configuring to upload logs to FAMS, you can also test the connection between FAMS and the FortiGate unit by selecting *Test Connectivity*, in *Log&Report > Log Config > Log Settings*.

The FAMS server stores the logs for archival usage. The FortiGate unit stores the logs locally either in system memory or disk, and then uploads the logs to the FAMS server.



FortiOS 4.0 MR3 Usability improvements

A major effort has been made to improve the usability of the FortiOS 4.0 MR3 web-based manager experience. Changes have been made throughout to improve the visibility of information and make it easier and more efficient to view and change configurations and monitor network activity and FortiGate activities and processes.

This section contains the following topics:

- [High-level web-based manager menu changes](#)
- [New FortiGate Setup Wizard](#)
- [FortiExplorer enhancements](#)
- [Dashboard Widgets](#)
- [Chart display improvements](#)
- [Monitoring Improvements](#)
- [Filtering web-based manager lists](#)
- [Reference count column \(object usage visibility\)](#)
- [Configuration object tagging and coloring](#)
- [Security configuration object icons](#)
- [Access to online help](#)
- [Backing up and restoring configuration files per-VDOM](#)

High-level web-based manager menu changes

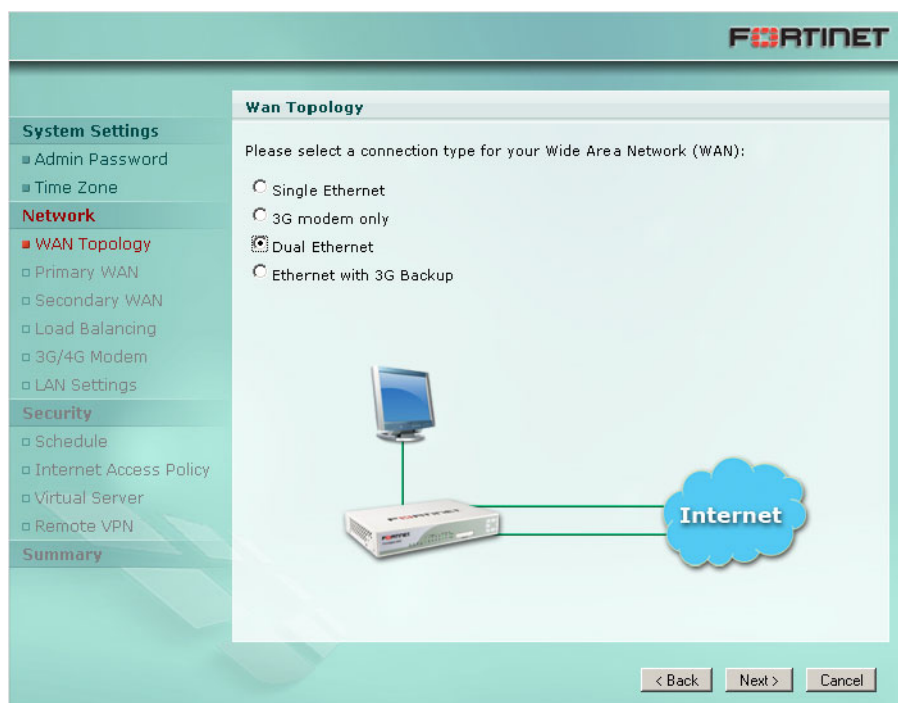
FortiOS 4.0 MR3 patch 1 introduces the following menu changes to the web-based manager. The CLI commands for these configuration items have not changed.

- The new *Policy* menu is used to configure IPv4 and IPv6 security policies, view the central NAT table, configure DoS policies, Sniffer policies, and protocol options. You can also monitor sessions and policy usage. Security policies are also called firewall policies. These options were available from *Firewall > Policy*.
- The *Firewall* menu has been renamed *Firewall Objects* and contains menus for configuring firewall addresses, services, schedules, traffic shapers, virtual IPs, firewall load balancing and monitoring load balancing and traffic shaping.
- The *UTM* menu has been renamed *UTM Profiles*.
- *Endpoint Control* and *Vulnerability Scan* have been moved under *UTM Profiles*. This functionality is now documented in the UTM Guide chapter of the FortiOS Handbook.
- The *System > Network > DNS* contains DNS fields formerly present in the *System > Network > Options* page. This page also includes DDNS settings.

New FortiGate Setup Wizard

Available on selected models, the new FortiGate setup wizard allows for quick and easy set-up of your FortiGate configuration. Within the wizard, you can configure the administrator password, FortiGate unit time zone, network settings (single or dual WAN interfaces, modem settings, DHCP and LAN settings), apply security features such as access schedules, UTM features, NAT, virtual servers, and remote SSL or IPsec VPN access.

Figure 7: Configuration wizard WAN topology setting



FortiExplorer enhancements

The most recent version of FortiExplorer is compatible with recent FortiGate models running FortiOS 4.0 MR3. You can use FortiExplorer to easily and quickly configure your FortiGate unit with basic settings. FortiExplorer also allows access to the web-based manager and CLI through a USB connection. FortiExplorer runs on all Windows platforms and on Mac OS X.

FortiExplorer contains improved setup wizard support, FortiGuard support, additional system improvements and a new security policy wizard which is similar to the FortiGate setup wizard but for security policies. FortiExplorer also contains a 3G/4G modem configuration page, for those units that have 3G/4G modem capabilities.

Dashboard Widgets

There are several enhancements to dashboard widgets in this release, as well as a new widget called Network Protocol Usage. You can view dashboard widgets from *System > Dashboard > Status*.

Traffic History

The Traffic History widget has been enhanced, allowing you to customize which line charts contain a specified time period range. For example, the second line chart displays the last seven days of traffic information.

You can choose to have the time period display in days, minutes or hours. You must enter a time period that is either in minutes or days, such as 10 minutes or 30 days. If you choose to enter zero, that specific time period is disabled.

Dashboard - Traffic History Settings	
Provides settings for modifying the default settings of the Traffic History widget.	
Custom Widget Name	Enter a new name for the widget. This is optional.
Select Network Interface	Select an interface (FortiGate unit's interfaces) from the drop-down list. The interface you choose displays the traffic occurring on it.
Enable Refresh	Select to enable the information to refresh.
Time Period 1	The time period for the first line chart. Enter a number in the first field, then select <i>Hour(s)</i> , <i>Minute(s)</i> or <i>Day(s)</i> from the drop-down list beside the field.
Time Period 2	The time period for the second line chart. Enter a number in the first field, then select <i>Hour(s)</i> , <i>Minute(s)</i> or <i>Day(s)</i> from the drop-down list beside the field.
Time Period 3	The time period for the third line chart. Enter a number in the first field, then select <i>Hour(s)</i> , <i>Minute(s)</i> or <i>Day(s)</i> from the drop-down list beside the field.

System Resources

The System Resource widget now only displays information concerning the CPU and memory usage amounts. You can view this information either in real-time or current information, or historical. If you want to view the information in historical view, you can also change the type of fill-line color.

Dashboard - Custom System Resource Display	
Provides settings to modify the default or current configuration of the System Resource widget.	
Custom Widget Name	Enter a new name for the widget. This is optional.

View Type	<p>Select which type you want to view the system resource information in.</p> <ul style="list-style-type: none"> Real-time – displays the current information as a dial gauge, along with a percent, located at the bottom. For example, Memory Usage 58%. Historical – displays the information in a fill-line chart for each CPU and memory. When you select <i>Historical</i>, the <i>Chart Color</i> option appears. When viewing CPU and memory usage in the web-based manager, only the information for core processes displays. CPU for management processes, is excluded. For example, HTTPS connections to the web-based manager.
Chart Color	Select <i>Change</i> to change add a new fill color to the chart. Select <i>Reset</i> to reset the color back to its default color.
Time Period	Select from the drop-down list, the period of time that the information will be displayed.

Network Protocol Usage

The Network Protocol widget allows you to view many different protocols over a period of time. This widget reflects what was previously found in the basic traffic report, located in *Log&Report > Report Access > Memory* in FortiOS 4.0 MR1 and lower.

The Network Protocol Usage widget allows you to view many different protocols over a period of time. You can view this information with either a line chart or bar chart style. Network protocol usage information can be viewed for up to the last 30 days, or as recent as the last 24 hours.


Custom Network Protocol Usage Display Provides settings for modifying the default settings of the Network Protocol Usage widget.		
Custom Widget Name	Enter a new name for the widget. This is optional.	
Chart Style	Select either <i>Line</i> or <i>Bar</i> style for the chart. The line chart is a fill-line chart style type.	
Time Period	Select a time period from the drop-down list. For example, if you choose <i>Last 24 hours</i> , only the information gathered in the last 24 hours displays.	
Protocols	You can choose from any of the following protocols:	
	• Browsing	• DNS
	• FTP	• Gaming
	• Newsgroups	• P2P
	• TFTP	• VoIP
	• Generic UDP	• Generic ICMP
	• E-mail	• Instant Messaging
	• Streaming	• Generic TCP

	<ul style="list-style-type: none"> Generic IP 	
	All protocols are enabled by default. If you do not want to include certain protocols, select the check box beside each protocol that should not be included.	

Chart display improvements

Charts within the web-based manager have a larger font size and chart style applied to them. These changes make it easier to read the information that displays. These chart improvements include the charts within widgets as well as within FortiOS reports.

Figure 8: Example of the display improvements to a chart



Monitoring Improvements

In each menu in the web-based manager, there is now a Monitor submenu containing one or multiple submenus that allow you to view the activity of a specific feature that is currently being monitored by the FortiGate unit. The information displayed is usually in a table or graphical format, providing a more user-friendly display of the monitored information.

The information is displayed in a similar manner as to how widgets display their information in charts or lists on a dashboard in *System > Dashboard*.

You must enable logging for certain features since the information that is compiled for certain Monitor submenus only comes from logs. These features are the UTM Monitor submenus, security policy (Policy Monitor), and the Logging Monitor submenu.

DHCP Monitor

The DHCP Monitor is available from *System > Monitor > DHCP Monitor*. Using this monitor you can view the DHCP servers and relays that are being monitored by the FortiGate unit.

On the DHCP Monitor page, you can also add IP addresses from the page to the IP reservation list. The IP reservation list is a list of reserved IP addresses on a DHCP network for a user who wants to always assign that same IP address to one of the DHCP network's hosts.

On the DHCP Monitor page, you can also refresh the information to ensure current information displays on the page.

Modem Monitor

The Modem monitor is available from *System > Monitor > Modem Monitor*. Using this monitor, you can view the unit's modem status and activity. The information on the page is displayed in a bar chart as well as in a table, located below the bar chart.

On the Modem Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Session Monitor

The Session Monitor is available from *Policy > Monitor > Session Monitor*. Using this monitor you can view all of the sessions that are currently being monitored by the FortiGate unit. On the Session Monitor page, you can filter the information, delete a session, or refresh the list.

This monitoring submenu is similar to the widget, Top Sessions, which is still available in *System > Dashboard*. The session information that displays in the widget can also be seen within the Session Monitor submenu.

On the Session Monitor page, you can refresh the information to ensure current information displays on the page. You can also filter the information using Filter Settings. If you want to delete a session, select the Delete icon in the row of the session you want removed.

Policy Monitor

The Policy Monitor submenu is available from *Policy > Monitor > Policy Monitor*. Using this monitor you to view the top security policy usage by the FortiGate unit. The information displays in a bar chart and details such as action and packets, are displayed in a table below the bar chart.

On the Policy Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Load Balance Monitor

The Load Balance Monitor is available from *Firewall Objects > Monitor > Load Balance Monitor*. this monitor display the status of each virtual server and real server, as well as the start or stop status of the real servers, is displayed on the Load Balance Monitor page.

Traffic Shaper Monitor

The Traffic Shaper Monitor is available from *Firewall Objects > Monitor > Traffic Shaper Monitor*. Using this monitor you can view traffic shaping activity that is being monitored by the FortiGate unit. This information displays in a bar chart. You can view the traffic shaper usage information by current bandwidth or by dropped packets. Use the *Report By* drop-down list on the page to view traffic shaper usage by selecting either *Current Bandwidth* or *Dropped packets*.

On the Traffic Shaper Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

AV Monitor

The AV Monitor is available from *UTM Profiles > Monitor > AV Monitor*. Using this monitor you can view activity concerning viruses detected by the FortiGate unit. This information displays on the AV Monitor page in a bar chart as well as in a table located below the bar chart.

On the AV Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Intrusion Monitor

The Intrusion Monitor is available from *UTM Profiles > Monitor > Intrusion Monitor*. Using this monitor you can view the attack activity detected by the FortiGate unit. This information displays on the Intrusion Monitor page, in a bar chart as well as in a table located below the bar chart.

On the Intrusion Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Web Monitor

The Web Monitor is available from *UTM Profiles > Monitor > Web Monitor*. Using this monitor you can view the web activity detected by the FortiGate unit. This information displays on the Web Monitor page, in a pie chart and bar chart. The total HTTP requests information displays in the pie chart and the blocked HTTP requests display in a bar chart. The total number of web requests display at the bottom of the charts, in *Total Web Requests (HTTP): <number>*.

On the Web Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Email Monitor

The Email Monitor is available from *UTM Profiles > Monitor > Email Monitor*. Using this monitor you can view the email activity detected by the FortiGate unit. This information displays on the Email Monitor page, similar to how the Web Monitor page displays its monitoring information, the total number of emails in a pie chart and the blocked emails in a bar chart. The total number of emails is located at the bottom of the charts, in *Total Emails: <number>*.

On the Email Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Archive & Data Leak Monitor

The Archive & Data Leak Monitor is available from *UTM Profiles > Monitor > Archive & Data Leak Monitor*. Using this monitor you can view the DLP usage performed that is being detected by the FortiGate unit. This information displays in a bar chart. You can view this information by security policy, DLP sensor, or by protocol using the *Report By* drop-down list. The total number of dropped DLP archives is located at the bottom of the chart, in *Total Dropped Archives: <number>*.

On the Archive & Data Leak Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

Application Monitor

The Application Monitor is available from *UTM Profiles > Monitor > Application Monitor*. Using this monitor you can view the application usage detected by the FortiGate unit. This information displays in a bar chart.

On the Application Monitor page, you can refresh the information to ensure current information displays on the page. You can also reset the information, which removes the information from the page and starts the monitoring process again.

IPsec Monitor

The IPsec Monitor is available from *VPN > Monitor > IPsec Monitor*. Using this monitor you can view the activity on IPsec VPN tunnels. The page also shows the start and stop of tunnel activity.

The list includes both dial-up IPsec users as well as static IP or Dynamic DNS VPNs. The list provides status and IP addressing information about VPN tunnels, which VPN tunnels are active or non-active, connecting to remote peers that have static IP addresses or domain names. If you want, you can also start and stop individual tunnels from the list as well.

SSL-VPN Monitor

The SSL Monitor is available from *VPN > Monitor > SSL-VPN Monitor*. Using this monitor you can view the activity of SSL VPN sessions. The list displays the user name of the remote user, the IP address of the remote client, and the time the connection was made. You can also see which services are being provided, and delete an active web or tunnel session from the unit.

Web Cache Monitor

The Web Cache Monitor is available from *WAN Opt. & Cache > Monitor > Cache Monitor*. Using this monitor you can view the activity of SSL VPN sessions. The web cache monitor includes two widgets that display information about web cache requests and web cache traffic. The Web Cache Requests widget displays the number of session that were cached and the number that were not in a pie chart. The Web Cache Traffic widget consists of a line graph that compares the amount of HTTP traffic in kbytes on the WAN with the amount of HTTP traffic in kbytes on the LAN. The difference between the LAN and WAN traffic shows how much traffic was cached.

WAN optimization Peer Monitor

The WAN optimization Peer Monitor is available from *WAN Opt. & Cache > Monitor > Peer Monitor*. Using this monitor you can view a list of WAN optimization peers that the FortiGate unit can communicate with. For each peer you can view the peer's name and IP address, the type of peer, and the amount of traffic reduction as a result of WAN optimization or web caching with that peer.

WAN optimization web cache monitor

To view the web cache monitor, go to *WAN Opt. & Cache > Monitor > Cache Monitor*.

The web cache monitor shows the percentage of web cache requests that retrieved content from the cache (hits) and the percentage that did not receive content from the cache (misses). A higher the number of hits usually indicates that the web cache is being more effective at reducing WAN traffic. To improve cache performance you can

The web cache monitor also shows a graph of web traffic on the WAN and LAN. A lower WAN line on the graph indicates the web cache is reducing traffic on the WAN. The web cache monitor also displays the total number of web requests processed by the web cache.

Filtering web-based manager lists

In previous releases, when you wanted to filter information within a web-based manager list you used the filter icons. Filter icons are still available for filtering, however, *Filter Settings* have been introduced, providing a central location to configure multiple filters at once. Previously, you had to configure filters one at a time.

When you select *Filter Settings*, a new *Filters:* pane appears at the top of the list. You can use this pane to add and remove multiple filters and configure settings for each one. Add a filter by selecting Add new filter or by selecting the filter icon for a column in the list. When you select a filter icon in a list column the Filters pane opens with that column added to it.

Figure 9: Filter Settings

Filters:

- ✖ Action: accept [\[Change\]](#)
- ✖ Source: 10.10.10.1 [\[Change\]](#)
- ✖ Destination: 5.5.5.5 [\[Change\]](#)

Schedule:

✖ Value: ☐ NOT

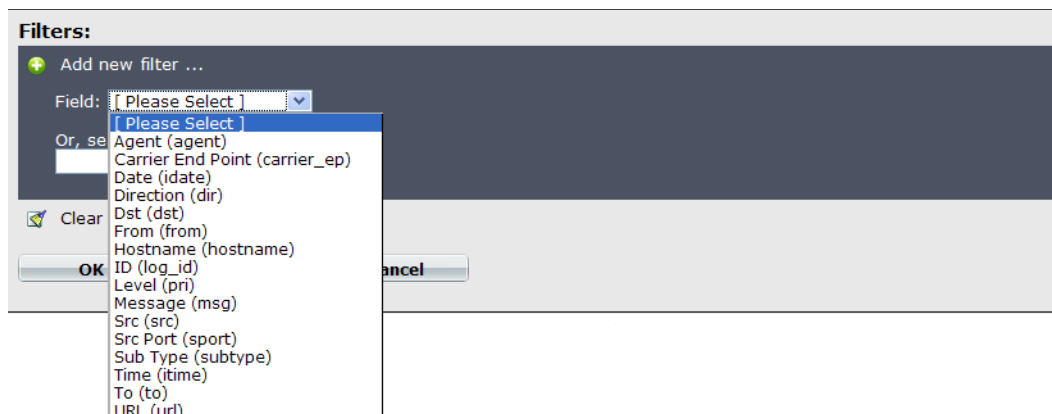
Use commas (,) to separate multiple values.
To filter entries that contain a specific prefix, use an * (asterisk).

[+ Add new filter ...](#)

[🗑 Clear all filters](#)

You must select **OK** when you are ready to apply the filters, otherwise the filter settings will not be applied to the information on the page. You can modify or remove a filter at any time.

Figure 10: Example of adding a log field from the Field drop-down list when filtering log messages



Reference count column (object usage visibility)

Within most web-based manager lists, a new column displays called the Reference count column, or *Ref.* This new column shows that a configuration object (for example an interface) is referenced to another object (for example a security policy) and how many times that object is referenced within FortiOS. For example, in [Figure 11](#) the default antivirus profile is referenced once. Finding a referenced object in previous releases was available only in the CLI.

Figure 11: The Ref. count column in the firewall address list

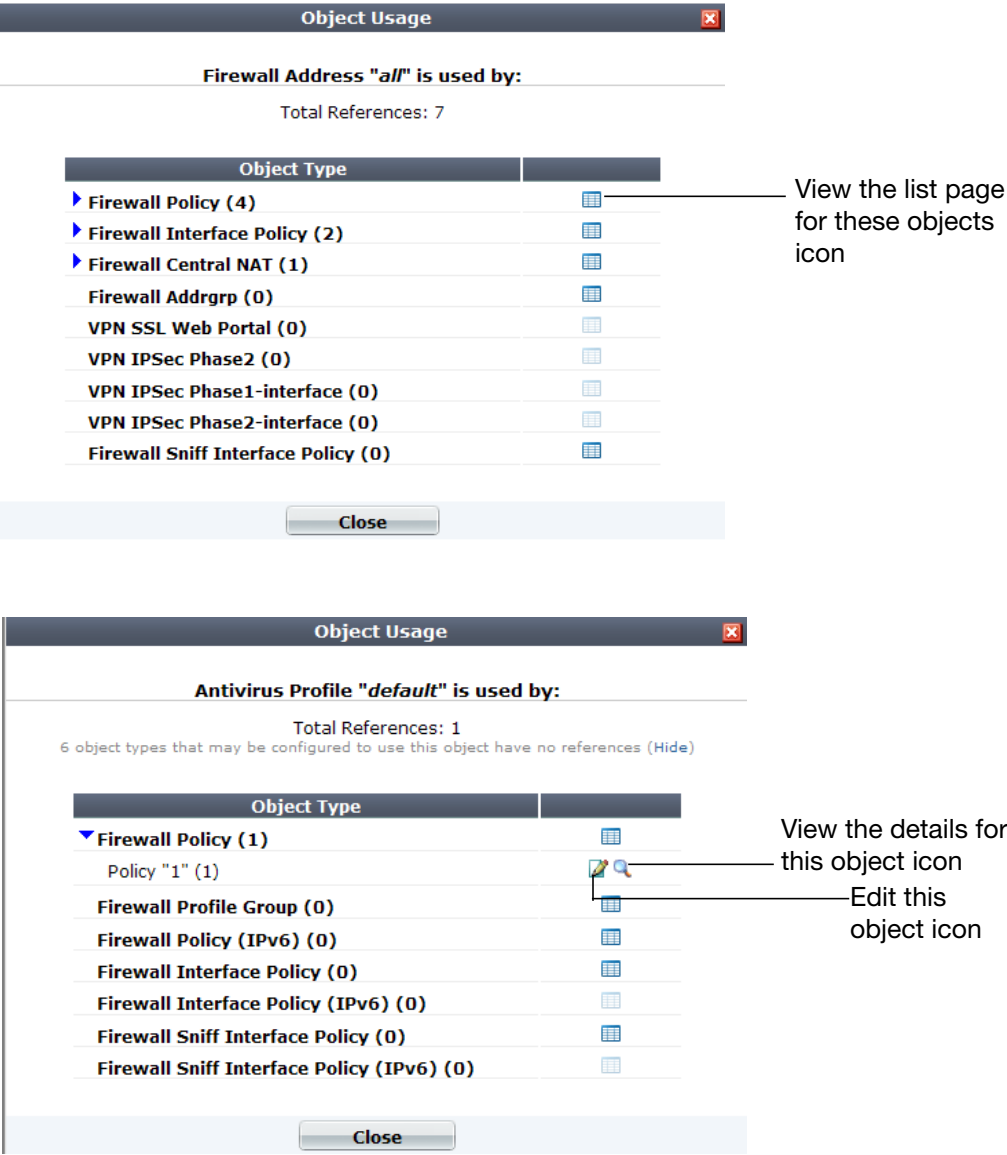
Create New Edit Delete					
	Name	Address/FQDN	Interface	Type	Ref.
<input type="checkbox"/>	all	0.0.0.0/0.0.0.0	Any	Subnet	10
<input type="checkbox"/>	SSLVPN_TUNNEL_ADDR1	10.0.0.1-10	Any	IP Range	2
<input type="checkbox"/>	all	::/0		IPv6	0

The *Ref.* column helps you to determine the object that is being referenced, and where it is referenced in. The *Ref.* column also helps you when you need to remove an object but are unable to because it is being referenced.

When you select the number within the reference count column, the Object Usage window appears, showing you exactly where the object is referenced within an object type. An object type in this usage is the location of where an object is referenced in. For example, in [Figure 12](#) the Object Usage window shows that the firewall address “all” is referenced within seven object types; in the default antivirus profile, it is referenced once in a security policy.

By selecting on the *View the list page for these objects* icon, you are automatically redirected to the page where the entry is referenced in. When you see the Expand Arrow beside some of the object types, it means that you can either view the location of the object in that particular object type, or modify the object type.

Figure 12: Two views of the Object Usage window, one without any expanded object types, and one with an object type expanded showing the available icons



The Object Usage window also provides a way to view the settings for an object, as seen in Figure 1.

Table 1: The Object Usage window displaying the object type table

Object Usage									
Firewall Address	<table> <tr> <td>policyid</td><td>1</td></tr> <tr> <td>status</td><td>enable</td></tr> <tr> <td>orig-port</td><td>5</td></tr> <tr> <td>nat-port</td><td>12-20</td></tr> </table>	policyid	1	status	enable	orig-port	5	nat-port	12-20
policyid	1								
status	enable								
orig-port	5								
nat-port	12-20								
Total Refer									
Object Type									
▶ Firewall Policy (4)									
▶ Firewall Interface Policy (2)									
▼ Firewall Central NAT (1)									
Central NAT Table "1" (1)									
Firewall Addrgrp (0)									
VPN SSL Web Portal (0)									
VPN IPSec Phase2 (0)									
VPN IPSec Phase1-interface (0)									
VPN IPSec Phase2-interface (0)									
Firewall Sniff Interface Policy (0)									
Close									



If you have selected *View the list page for these objects*, and are on the page where the entry is referenced, you can go back to the previous location by selecting the back option on your browser.

Configuration object tagging and coloring

The Tag Management menu provides a central location to view, search and manage tags that you created. Tags are keywords or a term that is assigned to a specific configuration that can be used for searching or filtering purposes.


From this central location in *System > Config > Tag Management*, you can do any of the following:

- a search to find a specific tag
- view where a tag is referenced, for example, a single tag could be referenced in a security policy, predefined signature and application
- go to where the tag is located, for example, a security policy
- view how many tags are currently unused
- remove tags.

The Tag Management page also provides a way to easily locate a specific object, such as a security policy, because of how tags work. For example, an SSL VPN security policy is tagged with the keywords `ssl vpn`, `SSL VPN`, `remote`, and `ssl branch office`; from the Tag Management page, enter `ssl` and the tags for that security policy appear; select one of the tags and within the Object Usage window, select to go to the SSL VPN security policy.

You can view detailed information about what object is using a tag by selecting one of the tags in the rectangular area that contains a gray background. The Object Usage window appears, which displays similar information as when you select a number in the *Ref.* column.

Figure 13: Tag Management page with the option to remove four unused tags



Adding tags to configuration objects

Tags can be created for security policies and firewall addresses. Tags are keywords or a term that is assigned to a specific piece of information, for example a firewall address, which can then be used for filtering or searching purposes.

Tags created within security policies and firewall addresses are used only for filtering and searching purposes. This provides a more concise output. For example, you have multiple VDOMs that contain multiple security policies; tags applied to these security policies allow you to find specific security policies within specific VDOMs.

Tags can also be added to predefined signatures and applications and are used within IPS and application sensors so that only those signatures are used.

The following example explains how to add tags to multiple security policies and then use Tag Management to find a security policy using the tags that were applied to the security policy. Tags are used in the same way for firewall addresses so the example can also be used as basis when configuring and applying tags for firewall addresses.



In the Add Tags window, you can select to add existing tags to the security policy or address list; however, these tags belong to predefined signatures and applications as well as to other security policies and address lists so the tags may not be applicable. You should make sure that the tag is valid for its use when applied to a security policy or other object otherwise it becomes redundant.

Example of how to find a security policy using Tag Management

Your FortiGate unit contains many security policies and the unit is currently in VDOM mode. You want to apply tags to only the SSL VPN security policies so that you can easily get to those policies. There are two SSL VPN security policies.

To add tags to multiple security policies

- 1 In vdom_1, go to *Policy > Policy > Policy*.
- 2 For the first security policy, select it in the row to highlight it.
- 3 Select the down arrow beside *Edit*, and then select *Add Tags*.
- 4 In the Add Tags window, enter `remote ssl, ssl vpn, remote, intranet, non-public, internal` and then select the plus sign.
By selecting the plus sign, the tag is automatically added. If you do not select the plus sign, the tag is not added and you have to enter the tag again.
- 5 Select *OK*.
- 6 In the second security policy, select it in the row to highlight it.
- 7 Select the down arrow beside *Edit*, and then select *Add Tags*.
- 8 In the Add Tags window, enter `internet, ssl vpn, remote, public, external` and then select the plus sign.
- 9 Select *OK*.

To search for a security policy from Tag Management

- 1 Go to *System > Config > Tag Management*.
- 2 On the Tag Management page, enter `remote` in the *Type to find tags: search* field.
The tag appears in the rectangular box with the gray background.
- 3 Select `remote` to view where the tag is currently being used.
- 4 In the Object Usage window, select the *View the list page for these objects* icon in the row of the *Object Type*.
You are redirected to the *Policy > Policy > Policy*, where you can select the security policy and then make changes to that policy.

Adding tags to predefined signatures and applications

Tags can be created for predefined signatures and applications which are then used in a sensor to provide a means to specify the use of only those tagged objects. Tags are keywords or a term that describes a piece of information and are assigned to that specific piece of information.

Tags that are created within a signature in *UTM Profiles > Intrusion Protection > Predefined* can be used within an IPS sensor by applying that same tag in an IPS filter entry. Tags that are created are not displayed within the IPS filter list; you must view them from within the IPS filter itself.

Tags that are created within an application in *UTM Profiles > Application Control > Application List* are applied to an application entry within an application sensor. Tags are used in the exact same as they are for IPS sensors. Tags that are used for predefined signatures cannot be used for applications within the application control list and vice versa.

Tags can also be added to security policies and address lists. When you want to view all tags that are configured for predefined signatures, application control list, security policies and addresses, go to *System > Config > Tag Management*.

















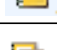

Security configuration object icons

Within the Firewall Objects menu, there are now firewall configuration object icons that you can change the color for. For example, within *Firewall Objects > Address > Address*, the configuration object icon for IP/Netmask was changed to pink.

These icons also include representing an action within the Action column in security policies. For example, for a deny security policy the Action column on the IPv6 Policy page shows a red circle with a line through it.

The following table explains the security policy configuration object icons that you can customize the color for.

Table 2: Security policy configuration object icons

Icon	Definition	Icon	Definition
	Allow		Recurring schedule
	Deny		One-time schedule
	IPsec		Schedule group
	SSL VPN		Pre-defined service
	IP/Netmask		Custom Service
	IP Range		Service Group
	IPv6 Address		Virtual IP
	FQDN Address		Virtual Server
	Address group		Virtual IP Group

Access to online help

Online help is stored and accessed from our Tech Docs web site; previously it was within the firmware image itself. Online help works in the exact same way as before, providing the same search capabilities as well.

Backing up and restoring configuration files per-VDOM

From the Global VDOM, you can now back up or restore a configuration file for a specific VDOM within the web-based manager. This provides a quick and easy way to back up or restore your configuration file within a specific VDOM. There is an option to back up or restore the full configuration, if needed.

You can back up or restore a specific VDOM configuration file from the System Information widget. When you are on the Backup or Restore page, the *VDOM Config* option is available and you can then choose the specific VDOM you want to back up or restore the configuration from by selecting a VDOM from the drop-down list. These options are only available when you are in the Global VDOM.

This feature is available only when VDOMs are enabled on the FortiGate unit.



More New Features

This section describes additional new features available in FortiOS 4.0 MR3 and contains the following sections:

- New features for FortiOS 4.0 MR3 Patch 5
- New features for FortiOS 4.0 MR3 Patch 4
- New features for FortiOS 4.0 MR3 Patch 3
- New features for FortiOS 4.0 MR3 Patch 2
- New features for FortiOS 4.0 MR3 Patch 1
- Login grace timer for SSH connections
- FortiManager automatic authorization
- Dynamic DNS commands
- New diagnose commands
- New get commands
- MTU configuration support on non-IPsec tunnel interfaces
- Customizing maximum number of invalid firewall authentication attempts
- Controlling the connection between a FortiManager unit and a FortiGate unit
- Bringing up or down IPsec tunnels
- Configuring active CPUs
- Formatting multiple disk partitions
- Transparent mode port pairs
- DNS server changes
- DHCP Server changes
- Installing firmware on a partition without a reboot
- SNMP enhancements
- Replacement message changes
- VDOM and global privileges for access profiles
- HA dynamic weighted load balancing
- VRRP virtual MAC address support
- FGCP HA subsecond failover
- Static Route enhancements
- Monitoring ISIS from the Routing Monitor page
- Security Policy and Firewall Object Enhancements
- Virtual IP source address filter support
- Virtual IP port forwarding enhancements
- Load balancing HTTP host connections
- Web Proxy Service and Web Proxy Service Group

- [SSL renegotiation for SSL offloading provides allow/deny client renegotiation](#)
- [SSL VPN Port forwarding support](#)
- [IKE negotiation](#)
- [SHA-384 and SHA-512 support for IKE](#)
- [FortiOS Carrier URL extraction feature](#)

New features for FortiOS 4.0 MR3 Patch 5

- Intrusion Protection (IPS) is now supported for load balancing virtual servers (load balancing virtual IPs (VIPs)). You can now enable UTM and select an IPS sensor in a firewall policy that contains a load balancing virtual server. This includes the case where the load balancing virtual server supports persistence. However, IPS does not work with virtual server load balancing of SSL sessions.

New features for FortiOS 4.0 MR3 Patch 4

- Combine IPS and vulnerability management service into one section
- Move disk management to *System > Config > Advanced > Disk Management*
- Disable memory logging for low-end models with large flash drives
- Enhance memory logging for low-end models with no log disk
- WAN Opt & Cache no longer available on the web-based manager for low-end FortiGate/FortiWiFi models since these features can affect performance of the low-end models. The features are still available from the CLI.

New features for FortiOS 4.0 MR3 Patch 3

- FortiGuard Web Filter category update
- Multiple email fields in log messages
- Weekly Report in PDF and Web Format
- Up to 100 VDOMs supported for the FortiGate-1240B
- [“WAN optimization web cache monitor” on page 146.](#)

New features for FortiOS 4.0 MR3 Patch 2

The following is a list of changes made to FortiOS 4.0 MR3 Patch 2:

- Central Management has been moved from *System > Admin > Central Management* to *System > Admin > Settings*.
- The web-based manager page for adding and editing security policies has been enhanced to make policies easier to configure and understand.
- [“WAN optimization Peer Monitor” on page 146.](#)
- FortiClient Connect is now called FortiClient.
- From *System > Certificates > CA Certificates*, the Fortinet_Wifi_CA certificate is now called PositiveSSL_CA.
- FortiGate-VM now has a 15-day trial evaluation license and upgrade available.

- Firewall address table size has been increased to 2000 on FortiGate-110C, 200A, 200B, and 80C series models.
- Static route entries table size has been increased to 5000 for the FortiGate-310B and 300C models.
- Support for load balancing on the FMG-XG2 card is now available.
- The Session widget now contains offload information in the Offload information column.
- In *Firewall Objects > Service > Service Groups*, the following are default service groups you can use:
 - *Exchange Server*
 - *Exchange Service OWA*
 - *Outlook*
 - *Windows AD*
- In *Firewall Objects > Service > Web Proxy Service*, there is one default web proxy service available for you to use.
- You can now capture packets within the web-based manager by going to *System > Config > Advanced* and creating a packet capture filter. You can start and stop a filter at any time.
- Support is now available for 64bit FortiOS on the FortiGate-1240B.
- When configuring wireless settings in an SSID for wireless networking, you can now choose to have both TKIP and AES.
- On the Application Control Monitor page, it has changed to look similar to a Dashboard page. There are three widgets and each display data using specific commands in the CLI. The three widgets are:
 - Top Applications by Bandwidth (use `diag stats app-bandwidth`)
 - Top Applications by Session Count
 - Top IP/User for `<ip address_application>` (use `diag stats app-usage-ip <address/application>`)
- SQL logging may or may not be, by default, enabled on certain models. You should verify that SQL logging is enabled after upgrading to FortiOS 4.0 MR3 Patch-2.
- Report function may be affected by the change of logging back to text-based logs.
- For FortiGate models that support SSD, the default database size is 10GB in a single VDOM environment. FortiGate models that support a flash drive, their default database size is 1.5GB. You can change this default size to meet your own network logging requirements.
- When searching for information in logs in the web-based manager, you may not be able to continue searching if your search resulted in less than 50 records. Use the CLI instead if you want to continue your search.
- When configuring RADIUS servers for dynamic profile configurations, you can now choose to close all sessions associated with an IP address when a RADIUS STOP message is received. You can also enable logging of these RADIUS message events. A dynamic profile group must be configured first, before these RADIUS configuration settings become available within the New RADIUS Server page.

- Load balance mode allows you increased flexibility in how you use the interfaces on the FortiGate unit. When enabled on a FortiGate-3140B or a FortiGate-3950B/3951B with an installed FMC-XG2, traffic between any two interfaces (excluding management and console) is accelerated. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.

New features for FortiOS 4.0 MR3 Patch 1

The following is a list of changes made to FortiOS 4.0 MR3 Patch 1.

- The configuration of Web Filtering local ratings and local categories has been simplified.
- Support FSSO and sniffer policies: Log messages recording information gathered by a sniffer policy include a user name if the IP address in the log message corresponds to the IP address of a user who has been authenticated with FSSO.
- Web Filter Profile, IPS and application control pages of the web-based manager have been changed to enhance usability.
- [“High-level web-based manager menu changes” on page 139.](#)
- FortiGate unit Central Management Locking: FortiGate configuration changes cannot be made of the CLI or web-based manager if the unit is being remotely managed from FortiManager.
- FortiOS 4.0 MR3 patch 1 is compatible with FortiClient. FortiGate units support up to 10 FortiClient Connect connections.
- The new FortiGate UTM Weekly Activity Report now includes support for data based on geographic locations. For example, the default report includes a graph of Top Destination Countries by session
- BGP dynamic routing now supports AS override. With `as-override` enabled, while advertising an AS path to a peer, all leading occurrences of the peer's AS number are replaced with the AS number of the advertising router.

```
config router bgp
config neighbor
edit 192.168.1.112
...
set as-override disable|enable
set as-override6 disable|enable
...
```

- Forward and reverse traffic shaping can now be set independently in security policies and in application control sensors
- The WiFi controller feature in a FortiWiFi unit can manage local WiFi functions in the same manner as a remote FortiAP or FortiWiFi unit.
- SMTP virus scanning now supports scanning of STARTTLS messages.
- [“Web Cache Monitor” on page 146.](#)
- Control whether to bypass or block SSL sessions that cannot be decrypted by SSL content scanning and inspection. This behavior is controlled from the CLI in a protocol options profile. For example, for POP3S:

```
config firewall profile-protocol-options
```

```
edit new_profile
  config pop3s
    set unsupported-ssl {bypass | block}
  ...
```

- When adding LDAP, RADIUS or TACAS+ authentication servers you can select Test to verify that the configuration is correct. You can also use `diagnose test authserver` commands to test a number of aspects of authentication server configuration.
- LDAP group checking is now supported using the following command. You can set LDAP group checking to perform group object checking or user attribute checking.

```
config user ldap
  edit new_ldap
    set group-member-check {group-object | user-attr}
    set unsupported-ssl {bypass | block}
```

- FortiOS Carrier supports GTPv1 release 7.15.0 and GTPv1 release 8.12.0

Login grace timer for SSH connections

A grace timer has been introduced which allows control over the login time limits of SSH connections to the FortiGate unit. The grace timer can close open but unauthenticated SSH connections to the FortiGate unit. For example, if the timer is set to 60 seconds, any open, unauthenticated SSH session is closed after 60 seconds.

The default value of the allowed time is 120 seconds but can be configured for 10 to 3600 seconds.

This feature is available only in the CLI. The CLI command syntax used is:

```
config system global
  set admin-ssh-grace-time <seconds>
end
```

FortiManager automatic authorization

In previous releases, the `authorize-manager-only` command restricted access to authorized FortiManager units. This authorization is now automatically found during the communication exchanges between the FortiGate and FortiManager units.

This automatic authorization behaves as follows:

- On the FortiManager unit, an administrator enters the management IP or FQDN of the FortiManager unit.
- During the protocol exchange between the two units, on the FortiManager unit's side that management IP or FQDN is sent to the FortiGate unit.
- The FortiGate unit, after receiving the management IP or FQDN, determines that is valid, saves that management IP or FQDN as the FortiManager unit's

Dynamic DNS commands

The following DDNS commands were removed from the `config system interface` command. The DDNS commands are now found under the new `config system ddns` command.

```
set ddns {enable | disable}
```

```
set ddns-server <server>
set ddns-domain <server>
set ddns-username <username>
set ddns-password <password>
```

The DDNS commands above now in the `config system ddns` command:

```
set monitor-interface <interface_name>
set ddns-server <server>
set ddns-domain <domain_name>
set ddns-username <username>
set ddns-password <password>
```

New diagnose commands

Real-time session, traffic shaper bandwidth and CP6 statistics

The following CLI commands display real-time session set up rate statistics, accurate current traffic shaper bandwidth, and CP6 statistics information.

```
diag hardware ipsec
diag hardware deviceinfo cp6 {brief | cmdq | cmdqdis | rng |
task}
```

The CLI command `diag firewall shaper traffic-shaper list` now displays the accurate current traffic shaper bandwidth.

diag sys session filter proto-state

A new `diag` command has been introduced, which is an enhancement to the `get sys session-info {stat|full-stat}` command. The new command includes counts of the various TCP states, which the `get sys session-info stat|full-stat` command previously did not have.

The `diag sys session filter proto-state` command allows you to view the counts of various TCP states. This command can help in enterprise-type environments when tuning various protocol timers, for example, there are 60 percent of sessions in syn-sent state in comparison to the established sessions.

diag log-stats show

The new `diag log-stats show` command displays the number log messages that were discarded by the unit.

New get commands

IPsec get commands

There are several new `get` commands that help you to view IPsec VPN tunnel information as well as IKE gateway information and IPsec tunnel statistics.

The following `get` commands for IPsec VPN are as follows:

```
get vpn ike gateway
get vpn ipsec stats crypto
get vpn ipsec stats tunnel
get vpn ipsec tunnel summary
get vpn ipsec tunnel details
get vpn ipsec tunnel name
```

An example of the output of the `get vpn ipsec tunnel summary` is:

```
gateway
  name: 'phase1'
  type: policy-based
  local-gateway: 0.0.0.0:0 (dynamic)
  remote-gateway: 10.10.5.24:0 (static)
  mode: ike-v1
  interface: 'wan2' (4)
  rx packets: 0 bytes: 0 errors: 0
  tx packets: 0 bytes: 0 errors: 0
  dpd: enabled/unnegotiated
  selectors
    name: 'phase2'
    auto-negotiate: disable
    mode: tunnel
    src: 0:0.0.0.0/0.0.0.0:0
    dst: 0:0.0.0.0/0.0.0.0:0
```

The `get vpn ipsec stats tunnel` command gives statistics about the total number of IPSec tunnels and their types, including the number of selectors, how many are up, and any errors.

The following `get` commands for IPsec are now removed:

```
get vpn status concentrators
get vpn status ike config
get vpn status ike errors
get vpn status ike routers
get vpn status ike status detailed
get vpn status ike status summary
get vpn status ipsec
get vpn status tunnel stat
get vpn status tunnel dialup-list
get vpn status tunnel number
```

The following commands were changed:

- `get vpn status ike gateway` is replaced by `get vpn ike gateway`
- `get vpn status tunnel list` is replaced by `get vpn ipsec tunnel summary` and `get vpn ipsec tunnel details`
- `get vpn status tunnel name` is replaced by `get vpn ipsec tunnel name`
- `get vpn status ike crypto` is replaced by `get vpn ipsec stats crypto`

Traffic shaper and per-IP shaper

You can now view traffic shaper and per-IP shaper information from within the CLI. The commands display general information about shapers which includes their current bandwidth.

These commands are within the `get` command branch:

```
get firewall shaper traffic
get firewall shaper per-ip-shaper
```

Management checksum configuration information for FortiManager

There are now three new `get` commands that allow you to view the configuration checksum information to the FortiManager unit.

The management checksum commands are:

```
get system mgmt-csum global
get system mgmt-csum vdom <vdom_name>
get system mgmt-csum all
```

If you have no VDOMs enabled, entering `get system mgmt-csum` allows you to view the overall checksum information. The following is an example:

```
get system mgmt-csum
debugzone
global: 5c d6 08 fd e5 52 b3 18 e3 4d be 7f dc 40 86 66
root: 04 02 f0 e5 f2 21 36 63 72 05 f5 dc 31 94 c5 63
all: 24 90 19 d0 e4 67 7a c1 81 99 67 ae 77 fa bb 01

checksum
global: 5c d6 08 fd e5 52 b3 18 e3 4d be 7f dc 40 86 66
root: 04 02 f0 e5 f2 21 36 63 72 05 f5 dc 31 94 c5 63
all: 24 90 19 d0 e4 67 7a c1 81 99 67 ae 77 fa bb 01
```

MTU configuration support on non-IPsec tunnel interfaces

MTU configuration for non-IPsec tunnel interfaces is now supported. This allows you to customize the transmission amount for each interface on the FortiGate unit.

MTU is configured only in the CLI. The MTU setting is hidden until you enable the `mtu-override` setting.

Customizing maximum number of invalid firewall authentication attempts

A new option in the `config user setting` command allows you to customize the maximum number of invalid firewall authentication attempts before the FortiGate unit blocks them. This provides a way to tune CPU usage against invalid authentication connections.

The new option in the `config user setting` command is `auth-invalid-max`, and you can set the value between 1 and 100. For example, entering five allows five invalid authentication attempts before the unit blocks the user.

The following is an example of using this feature.

```
config user setting
    set auth-invalid-max 3
end
```

Controlling the connection between a FortiManager unit and a FortiGate unit

In the `config system interface` command, you can now configure whether an interface lets a FortiManager unit connect with a FortiGate on that unit's interface. For example, port 2 on the FortiGate unit does not allow the FortiManager unit to connect to it. When the FortiManager unit tries to connect to the FortiGate unit, the FortiGate unit refuses the connection.

If this feature is configured to not allow the FortiGate unit to connect to a FortiManager unit, the FortiGate unit will not allow an administrator to input the FortiManager unit's serial number into the central management configuration.

The command syntax for configuring this feature is as follows:

```
config system interface
  edit <interface>
    set allowaccess {ping | http | https | ssh | telnet | snmp |
      fgfm}
  end
```

Bringing up or down IPsec tunnels

Previously, you could activate or shut down IPsec tunnels using the `diag vpn tunnel {up | down}` commands. You can now use the following `execute` commands to help you bring up or down, and IPsec tunnel.

```
execute vpn ipsec tunnel down <phase2> <phase1> <serial>
execute vpn ipsec tunnel up <phase2> <phase1> <serial>
```

When using these `execute` commands, you can optionally use the phase 1 name, phase 2 or serial number to shut down or bring up the tunnel. However, if you are bringing down a tunnel, and that is a dial-up tunnel, phase 1 name is required. Bringing up a tunnel using the `execute vpn ipsec tunnel up` command cannot be used to activate a dial-up tunnel.

Configuring active CPUs

The new global command, `num-cpus`, allows you to configure a set number of active CPUs. This new feature is available only on platforms with eight or more CPUs.

The following is an example of how to configure five active CPUs.

```
config system global
  set num-cpus 10
end
```

Formatting multiple disk partitions

On a FortiGate unit with multiple disk partitions, you can now format multiple partitions at one time. This provides a quick and easy way to format multiple disk partitions.

Formatting multiple disk partitions uses the `execute disk format` command. The formatting process behaves in the following ways:

- If the formatting requires a reboot because one of the partitions is currently in use, all partitions are formatted before the reboot.
- If no reboot is required and an error occurs in the formatting process, the error is written to the event log.
- If an error occurs in formatting and a reboot is required, the error is logged to the event log.
- RAID (enable or disable) and RAID rebuilds take place before the reboot

- The `execute disk format` command requires that you enter each of the reference numbers of the partitions you want formatted. The reference numbers are found using the `execute disk list` command.

Table 3: Explanation of the encryption levels available

Encryption level	Explanation of encryption level	Algorithm associated with encryption level
High	Encryptions with key lengths larger than 128 bits, and some Cipher suites with 128-bit keys.	DHE-RSA-AES256-SHA AES256-SHA EDH-RSA-DES-CBC3-SHA DES-CBC3-SHA DES-CBC3-MD5 DHE-RSA-AES128-SHA AES128-SHA
Medium	Encryptions that are using 128 bit encryption	RC4-SHA RC4-MD5 RC4-MD
Low	Encryptions using 64 or 56 bit encryption but excluding export Cipher suites.	EDH-RSA-DES-CDBC-SHA DES-CBC-SHA DES-CBC-MD5

Transparent mode port pairs

Port pairing is an option in transparent mode to bind two ports together. In doing this, you can create security policies that regulate traffic only between two specific ports, VLANs or VDOMs. In its simplest form, this enables an administrator to create security policies that are only between these two ports. Traffic is captured between these ports. No other traffic can enter or leave a port pairing.

For example, a FortiGate unit has three ports, where port 1 and port 2 are paired together, because the two networks only need to communicate with each other. If packet arrives on port 1, the FortiGate unit needs to figure out whether the packet goes to port 2 or port 3. With port pairing configured, it is more simple. If packet arrives on port 1, then the FortiGate automatically directs the packet to port 2. The opposite is also true in the other direction. This can be ideal when to groups only need to transfer data between each other.

To configure port pairing - web-based manager

- 1 Go to System > Network > Interface.
- 2 Select the arrow next to the Create New button and select Port Pair.
- 3 Enter a Name for the port pair.
- 4 Select the physical or virtual ports from the Available Members list and select the right-facing arrow to add the ports to the Selected Members.
There can be only two ports added.
- 5 Select OK.

To configure port pairing - CLI

```
config system port-pair
  edit <pair_name>
    set member <port_names>
  end
```

When configuring security policies with the port pairs, selecting the Source Interface automatically populates the Destination Interface, and vice versa. All other aspects of the security policy configuration remain the same.

DNS server changes

Previously, when a DNS request was locally matched to a defined zone with no answer defined, it was not recursively forwarded. In this release, the DNS request is now forwarded when it cannot find a local answer in a non-authoritative zone, provided that the ingress interface has recursive DNS query enabled, using the authoritative option. This option is available within a DNS zone, in *System > Network > DNS Server*.

This new option, *Authoritative*, controls the DNS server's behavior so that it is more flexible. You can enable or disable this option in the web-based manager or CLI.



Fortinet recommends not using a FortiGate unit as an authoritative domain server.

DHCP Server changes

The DHCP Server information in the web-based manager is now located within the Network menu, *System > Network > DHCP Server*. The Network menu also contains the DHCP feature IP Reservation which is located in *System > Network > IP Reservation*. IP Reservation allows you to reserve an IP address that is on a DHCP network for a user who wants to always assign that same IP address to one of the DHCP network's hosts. The DHCP feature also includes support for IPv6.

When you create a new DHCP server, you can configure additional options under the Advanced section of the service page. There can be up to three options configured for a service. You can also add excluded ranges when configuring a DHCP server.

DHCP IP Reservation

Within the DHCP pool of addresses, you can ensure certain computers will always have the same address. This can be to ensure certain users always have an IP address when connecting to the network, or if you want a device that connects occasionally to have the same address for monitoring its activity or use.

In the example below, the IP address 172.20.19.69 will be matched to MAC address 00:1f:5c:b8:03:57.

- 1 Go to *System > Network > DHCP Server*.
- 2 Select the DHCP server from the list or add a new DHCP server.
- 3 Select IP Reservation and select Create New.
- 4 Enter an IP address of 172.20.19.69
- 5 Enter the MAC address of 00:1f:5c:b8:03:57.
- 6 Select OK.

You can also select *Add from DHCP Client List* and select the MAC and IP address pairs to add.

Installing firmware on a partition without a reboot

When you are upgrading the firmware on your FortiGate unit, you now have the option of installing the firmware on a partition without having to reboot the unit and run the image as the active firmware that is running on the unit. You can easily upgrade or downgrade to the firmware of your choice by using this new feature.

The following is an example of how to install the firmware image on a partition and not have the firmware running as the active firmware on the unit. The following also explains how to install the new firmware from the non-active partition and then make it the current active firmware running on the unit.

Example of installing a firmware on a partition without rebooting

You have decided to install FortiOS 4.0 MR3 release on the unit but still want to be able to easily switch back to a FortiOS 4.0 MR2 Patch release afterwards. You currently have two partitions on the unit's local hard disk and would like to be able to switch between the two firmware images at any given time.

The following procedures do not include backing up the configuration file since it is assumed that the back up has already been done.

To install a firmware image on a partition without a reboot

- 1 Go to *System > Dashboard > Status* and locate the System Information widget.
- 2 In the System Information widget, select *Update* in the *Firmware Version* row.
- 3 On the Firmware Upgrade/Downgrade page, select *Local Hard Disk* from the drop-down list beside *Upgrade From*.
- 4 Select *Browse* beside the *Upgrade File* field to locate the firmware image.
- 5 Clear the check box beside *Boot the New Firmware*.

This disables the reboot process that occurs when a firmware is being installed on the FortiGate unit.

- 6 Select *OK*.

A message similar to the following appears:

```
Software upload has completed. To use the new firmware, please
select it under System > Maintenance > Firmware, and use the
'Upgrade' option.
```

- 7 Go to *System > Maintenance > Firmware*.

In the table on the Firmware page, you can see that Partition 1 has the firmware image FortiOS 4.0 MR3 release.

The following procedure assumes that you are already in *System > Maintenance > Firmware*.

To install the new firmware from the partition

- 1 On the Firmware page, select the check box in the row of the firmware image FortiOS 4.0 MR3.

- 2 Select the Upgrade icon, located above the table.

The following message appears:

The page at `http://172.16.177.153` says:

System will reboot immediately and the current non-active partition will be set as the default boot partition. Continue?

- 3 Select OK.

The Boot alternate firmware page appears with the following message:

Please wait for system reboot to the new partition. Refresh your browser after a few minutes.

The unit reboots with the FortiOS 4.0 MR3 as the firmware image actively running on the unit.

SNMP enhancements

The SNMP feature contains several enhancements, as well as changes to the SNMP menu in the web-based manager.

Previously, for SNMP OIDs in FortiOS 4.0, the FortiOS OIDs were re-numbered so that each separate Fortinet product has its own root in the FortiOS OID space. SNMP 3.0 OIDs were still supported in FortiOS 4.0; however, both types of OIDs appear during an SNMP walk. In FortiOS 4.0 MR3 release, they are no longer supported so that only 4.0 SNMP OIDs appear.



IPv6 is now supported for SNMP.

WAN optimization, Web Cache and Explicit proxy MIBs

There are new MIBs for the new web proxy and caching features as well as explicit proxy. Some MIBs are supported by transparent proxy and these can be ported to the explicit proxy. There are also special OIDs for specific models.

SNMPv3

SNMPv3 is now supported. Within the SNMP configuration settings, you can configure SNMPv3 users, and include the events as well. The SNMP menu (*System > Config > SNMP*) provides all the configuration settings to create multiple SNMPv3 users. Each user can have multiple events enabled for them, as well as their specific security level. Multiple notification hosts can also be configured for each user.

SNMPv3 is usually included for additional security and remote configuration enhancements to SNMP. SNMPv3 provides confidentiality, integrity and authentication. For additional information about SNMPv3, see RFC 3411-3418.

Replacement message changes

There are several changes to replacement messages, as well as a new feature that allows you to upload images and use them in certain replacement messages. In this document, uploading images and using them in replacement messages is referred to as image embedding.

Replacement messages that contain authentication pages are now updated using the color scheme and image embedding feature. The types of replacement messages that are updated to the new color scheme and image embedding are:

- disclaimer pages
- login pages
- declined disclaimer pages
- login failed page
- login challenge page
- keepalive page

Endpoint NAC download portal and recommendation portal replacement messages are also updated to the new color scheme and image embedding feature. HTTP replacement messages are also updated.

Archive replacement messages and FTP proxy replacement message

The archive replacement messages and the FTP proxy replacement message are introduced because of the changes that occurred to the antivirus profile with regards to log archival options, and the new FTP proxy.

The following are the archive replacement messages and the FTP proxy replacement message.

- FTP Explicit-banner (under FTP Proxy)
- Archive block message (under HTTP)

Successful firewall authentication replacement message

The new Success message within the Authentication replacement messages provides a message indicating to the authenticating user that they have successfully authenticated their Telnet session. This replacement message is a text-only message.

Web filtering disclaimer replacement message

The web filtering disclaimer page allows users to bypass an override whenever they try to access a blocked page. The FortiGuard Web Filtering override form replacement message contains information so that the user can override the blocked page by authenticating with their user name and password. This replacement message is available in *System > Config > Replacement Message*, under FortiGuard Web Filtering.

Video chat block replacement message

The video chat block replacement message displays when a video chat has been blocked by the FortiGate unit. This message is available in *System > Config > Replacement Message*, under IM and P2P.

Replacement message images

The Replacement Message Image menu allows you to upload your organization or company's image to include in a replacement message. You can upload GIF, JPEG, TIFF, or PNG files, and give the file a unique name as well. The maximum image size that can be uploaded is 6000 bytes.

There are three default Fortinet images that you can choose from: the logo_fg_guard_wf, logo_fnet and logo_fw_auth. The following is a special tag to indicate that an image from the replacement message image list should be used in the replacement message.

```
<img src=%%IMAGE: <config_image_name>%% size=<btyes> >
```

When you include an image in a replacement message, it is referenced by the FortiGate unit. This reference number is displayed in the reference column of the Replacement Message Image page.

VDOM and global privileges for access profiles

Access profiles can now be configured with a VDOM or global privilege. These two privileges allow the FortiGate administrator access to either a specific VDOM or global access. Global access allows access to all VDOMs and global settings. When an administrator's account contains an access profile with a VDOM privilege, that administrator can access only the VDOM that is specified in their account. For example, admin_1 has the access profile admin_vdom; admin_vdom contains read and write privileges for logging and VDOM access; admin_1's account is associated with vdom_1. The admin_1 accessibility is limited to vdom_1 and the ability to configure only log settings.

Previously, when administrator accounts were configured, the VDOM was specified in the administrator account and access permissions were specified in an admin profile. By using this new access profile privilege, you can apply an access profile to an administrator that is specific for VDOM configuration.

These new access profile privileges are available only in the CLI. A new command, `scope`, provides the ability to have an access profile contain VDOM privileges or global privileges.

Example of incorporating the new access profile to existing administrator accounts

Company_A's branch office requires two administrators to access their FortiGate unit and they currently have VDOMs configured. An administrator with global access must be configured and an administrator with VDOM access that can configure reports are required.

There are currently two administrator accounts that contain global access and VDOM access to the FortiGate unit. However, management wants to apply the new privileges to the existing accounts.

This example explains how to incorporate the new privileges into two existing administrator accounts. The existing administrator accounts are admin_vdom and admin_global. You need to configure a global access profile because you cannot modify the super_admin access profile.

To modify the existing access profiles

- 1 Log in to the CLI and then in to the global level.
`config global`
- 2 Enter the following command within the global VDOM:
`config system accprofile`
- 3 Modify the vdom access profile first:
`edit vdom`
`set scope vdom`
`next`

The admin_vdom account will now be only able to access VDOMs within the configuration.

- 4 Enter the following commands to modify the global access profile:

```
edit global
  set scope global
  set admingrp read-write
  set authgrp read-write
  set endpoint-control-grp read-write
  set fwgrp read-write
  set loggrp read
  set mntgrp read-write
  set netgrp read-write
  set routegrp read-write
  set sysgrp read-write
  set updategrp read-write
  set utmgrp read-write
  set vpngrp read-write
  set wanoptgrp read-write
end
```

- 5 Enter the following command to apply the new global access profile to the existing admin_global administrator account:

```
config system admin
  edit admin_global
    set accprofile global
  end
```

You can verify that the access profiles have their global and VDOM privileges by going to *System > Admin > Administrators* and viewing the *Scope* column. In the *Scope* column, admin_vdom contains VDOM: vdom_1, and in the admin_global.

HA dynamic weighted load balancing

The following explains the weighted failover feature that is supported in this release. It is explained in two parts; the configuration of weighted-round-robin weights and weighted load balancing.

Configuring weighted-round-robin weights

You can configure weighted round-robin load balancing for a cluster and configure the static weights for each of the cluster units according to their priority in the cluster. When you set `schedule` to `weight-round-robin` you can use the `weight` option to set the static weight of each cluster unit. The static weight is set according to the priority of each unit in the cluster. A FortiGate HA cluster can contain up to 16 FortiGate units so you can set up to 16 static weights.

The priority of a cluster unit is determined by its device priority, the number of monitored interfaces that are functioning, its age in the cluster and its serial number. Priorities are used to select a primary unit and to set an order of all of the subordinate units. Thus the priority order of a cluster unit can change depending on configuration settings, link failures and so on. Since weights are also set using this priority order the weights are independent of specific cluster units but do depend on the role of the each unit in the cluster.

You can use the following command to display the priority order of units in a cluster. The following example displays the priority order for a cluster of 5 FortiGate-620B units:

```
get system ha status
Model: 620
Mode: a-p
```



```

Group: 0
Debug: 0
ses_pickup: disable
Master:150 head_office_cla FG600B3908600825 0
Slave :150 head_office_clb FG600B3908600705 1
Slave :150 head_office_clc FG600B3908600702 2
Slave :150 head_office_cld FG600B3908600605 3
Slave :150 head_office_cle FG600B3908600309 4
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
Slave :2 FG600B3908600702
Slave :3 FG600B3908600605
Slave :4 FG600B3908600309

```

The cluster units are listed in priority order starting at the 6th output line. The primary unit always has the highest priority and is listed first followed by the subordinate units in priority order. The last 5 output lines list the cluster units in vcluster 1 and are not always in priority order.

The default static weight for each cluster unit is 40. This means that sessions are distributed evenly among all cluster units. You can use the `set weight` command to change the static weights of cluster units to distribute sessions to cluster units depending on their priority in the cluster. The weight can be between 0 and 255. Increase the weight to increase the number of connections processed by the cluster unit with that priority.

You set the weight for each unit separately. For the example cluster of 5 FortiGate-620B units you can set the weight for each unit as follows:

```

config system ha
  set mode a-a
  set schedule weight-round-robin
  set weight 0 5
  set weight 1 10
  set weight 2 15
  set weight 3 20
  set weight 4 30
end

```

If you enter the `get` command to view the HA configuration the output for `weight` would be:

```
weight 5 10 15 20 30 40 40 40 40 40 40 40 40 40 40
```

This configuration has the following results if the output of the `get system ha status` command is that shown above:

- The first five connections are processed by the primary unit (host name head_office_cla, priority 0, weight 5). From the output of the
- The next 10 connections are processed by the first subordinate unit (host name head_office_clb, priority 1, weight 10)
- The next 15 connections are processed by the second subordinate unit (host name head_office_clc, priority 2, weight 15)
- The next 20 connections are processed by the third subordinate unit (host name head_office_cld, priority 3, weight 20)

- The next 30 connections are processed by the fourth subordinate unit (host name head_office_cle, priority 4, weight 30)

Dynamic weighted load balancing

You can configure active-active HA weighted round robin load balancing to load balance sessions according to individual cluster unit CPU usage, memory usage, and number of UTM proxy sessions. If any of these system loading indicators increases above configured high watermark thresholds, weighted load balancing sends fewer new sessions to the busy unit until it recovers.

For example, if you set a CPU usage high watermark, when a cluster unit's CPU usage reaches the high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to a low watermark threshold. When this happens the cluster resumes load balancing sessions to the cluster unit as normal.

You can set different high and low watermark thresholds for CPU usage and memory usage, and for the number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP UTM proxy sessions. For each loading indicator you set a high watermark threshold a low watermark threshold and a weight. When you first enable this feature the weighted load balancing configuration is synchronized to all cluster units. Subsequent changes to the weighted load balancing configuration are not synchronized so you can configure different weights on each cluster unit.

The CPU usage, memory usage, and UTM proxy weights determine how the cluster load balances sessions when a high watermark threshold is reached and also affect how the cluster load balances sessions when multiple cluster units reach different high watermark thresholds at the same time. For example, you might be less concerned about a cluster unit reaching the memory usage high watermark threshold than reaching the CPU usage high watermark threshold. If this is the case you can set the weight lower for memory usage. Then, if one cluster unit reaches the CPU usage high watermark threshold and a second cluster unit reaches the memory usage high watermark threshold the cluster will load balance more sessions to the unit with high memory usage and fewer sessions to the cluster unit with high CPU usage.

Use the following command to set thresholds and weights for CPU and memory usage and UTM proxy sessions:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold <weight> <low> <high>
  set memory-threshold <weight> <low> <high>
  set http-proxy-threshold <weight> <low> <high>
  set ftp-proxy-threshold <weight> <low> <high>
  set imap-proxy-threshold <weight> <low> <high>
  set nntp-proxy-threshold <weight> <low> <high>
  set pop3-proxy-threshold <weight> <low> <high>
  set smtp-proxy-threshold <weight> <low> <high>
end
```

For each option, the weight range is 0 to 255 and the default weight is 5. The low and high watermarks are a percent (0 to 100). The default low and high watermarks are 0 which means they are disabled. The high watermark must be greater than the low watermark.

For CPU and memory usage the low and high watermarks are compared with the percentage CPU and memory use of the cluster unit. For each of the UTM proxies the high and low watermarks are compared to a number that represents percent of the max number of proxy sessions being used by a proxy. This number is calculated using the formula:

$$\text{proxy usage} = (\text{current sessions} * 100) / \text{max sessions}$$

where:

`current sessions` is the number of active sessions for the proxy type.

`max sessions` is the session limit for the proxy type. The session limit depends on the FortiGate unit and its configuration.


You can use the following command to display the maximum and current number of sessions for a UTM proxy:

```
get test {ftpd | http | imap | nntp | pop3 | smtp} 4
```

Example weighted load balancing configuration

Consider a cluster to three FortiGate-620B units with host names 620_ha_1, 620_ha_2, and 620_ha_3 as shown in Figure 14. This example describes how to configure weighted load balancing settings for CPU and memory usage for the cluster and then to configure UTM proxy weights for each cluster unit.

Figure 14: Example HA weighted load balancing configuration



Use the following command to set the CPU usage threshold weight to 30, low watermark to 60, and high watermark to 80. This command also sets the memory usage threshold weight to 10, low watermark to 60, and high watermark to 90.

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold 30 60 80
  set memory-threshold 10 60 90
end
```

The static weights for the cluster units remain at the default values of 40. Since this command changes the mode to `a-a` and the schedule to `weight-round-robin` the weight settings are synchronized to all cluster units.



For FortiOS 4.0 MR3, the static weights assigned to cluster units using the `set weight` have changed. The default value is 40 and the range is now 0 to 255.

As a result of this configuration, if the CPU usage of `620_ha_1` reaches 80% the static weight for `620_ha_1` is reduced from 40 to 10 and correspondingly fewer sessions are load balanced to it. If the memory usage of this same cluster unit also reaches 90% the static weight further reduces to 0 and no new sessions are load balanced to it. If the memory usage of a `620_ha_2` reaches 90% the static weight of `620_ha_2` reduces to 30 and 30 and fewer new sessions are load balanced to it.

Now that you have set the basic weighted load balancing configuration for the cluster you can configure different settings on each cluster unit. For example, to set the HTTP usage threshold weight to 20, low watermark to 60, and high watermark to 90 for `620_ha_2` use the `execute ha manage` command to log into the `620_ha_2` CLI. Then enter the following command:

```
config system ha
    set http-proxy-threshold 20 60 90
end
```

To set the pop3 usage threshold weight to 20, low watermark to 60, and high watermark to 90 for `620_ha_3` use the `execute ha manage` command to log into the `620_ha_3` CLI. Then enter the following command:

```
config system ha
    set pop3-proxy-threshold <weight> <low> <high>
end
```

VRRP virtual MAC address support

Previously in FortiOS 4.0 MR2, the VRRP virtual MAC address (also known as the virtual router MAC address) feature, as described in [RFC 3768](#), **was supported**. The VRRP virtual MAC address is a shared MAC address adopted by the VRRP master. If the VRRP router group master fails the same virtual MAC master fails over to the new master of the group. As a result, all packets for VRRP routers can continue to use the same virtual MAC address. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

Each VRRP router is associated with its own virtual MAC address. The last part of the virtual MAC depends on the VRRP virtual router ID using the following format:

```
00-00-5E-00-01-<VRID_hex>
```

Where `<VRID_hex>` is the VRRP virtual router ID in hexadecimal format in internet standard bit-order. For more information about the format of the virtual MAC see RFC 3768.

Some examples:

- If the VRRP virtual router ID is 10 the virtual MAC would be 00-00-5E-00-01-05.
- If the VRRP virtual router ID is 200 the virtual MAC would be 00-00-5E-00-01-c8.

The VRRP virtual MAC address feature is disabled by default. When you enable the feature on a FortiGate interface, all of the VRRP routers added to that interface use their own VRRP virtual MAC address. Each virtual MAC address will be different because each virtual router has its own ID.

Use the following command to enable the VRRP virtual MAC address on the port2 interface and add a VRRP virtual router with ID 5, IP address 10.31.101.120 and priority 255.

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  config vrrp
    edit 5
      set vrip 10.31.101.120
      set priority 255
    end
  end
end
```

The port2 interface will now accept packets sent to the MAC address 00-00-5E-00-01-05.

FGCP HA subsecond failover

FGCP HA subsecond failover (that is a failover time of less than one second) can reduce the failover time after a device or link failover. In FortiOS 4.0 MR3 the CLI option for configuring subsecond failover has been removed and the feature is available for interfaces that include:

- Network processors: NP2, NP4
- Content processors: CP4, CP5, CP6
- Accelerated interfaces, for example the ASM-FB4, ADM-FB8, ADM-XB2, ADM-XD4, RTM-XD2
- Security processor modules: ASM-CE4, ASM-XE2

Subsecond failover can accelerate HA failover depending on the FortiGate unit HA and hardware configuration and the network configuration. Network devices that respond slowly to an HA failover can prevent this feature from reducing failover times to less than a second. Also, subsecond failover can normally only be achieved for a cluster of two units operating in Transparent mode with only two interfaces connected to the network.

For best subsecond failover results, the recommended heartbeat interval is 100ms and the recommended lost heartbeat threshold is 5.

```
config system ha
  set hb-lost-threshold 5
  set hb-interval 1
end
```

Static Route enhancements

Static routes now provides *Priority* and *Distance* settings in the Advanced section on the New Static Route page. The priority and distance settings can be displayed on the Static Route page using *Column Settings*. The priority and distance columns do not appear by default.

Figure 15: The Static Route page with the priority and distance columns displayed

Create New Edit Delete Column Settings						
	IP/Mask	Gateway	Device	Comment	Priority	Distance
<input type="checkbox"/>	0.0.0.0/0.0.0.0	0.0.0.0	wan2		0	10
<input type="checkbox"/>	172.20.120.0/255.255.255.0		ssl.root	for VPN SSL	400	155

When configuring a static route (or when modifying its settings), you can now include a comment within the static route. If you want to configure the priority and/or distance within a static route, you must select *Advanced...* to display priority and distance options.

Monitoring ISIS from the Routing Monitor page

You can now view ISIS routes from the Routing Monitor page. ISIS, introduced in FortiOS 4.0 MR2, is a routing protocol described in RFC 1142. ISIS is configured within the CLI.

Security Policy and Firewall Object Enhancements

There are several enhancements to firewall policies, including the Policy page (in *Policy > Policy > Policy*), which provides more flexibility and granularity. These enhancements also include page controls on the Address page in *Firewall Objects > Address > Address* to easily navigate through the list of addresses on the page.

The Firewall menu also provides more granularity when configuring a schedule. When configuring a schedule, you can now specify minutes in five minute intervals, for example, 5, 10, 15, 20, and all the way up to 55.

In FTP proxy security policies, FSSO guest user groups are now supported. FSSO authentication is IP-based authentication.



Traffic shaping bandwidth is now in kbits.

Source IP addresses for FortiGate-originating traffic

Previously, the source IP address feature was introduced in FortiOS 4.0MR2. In this release the source-ip address is extended, adding more options for configuring a source IP address to self-originating traffic. For example, NTP.

The source-ip address feature allows you to specify the source IP address of self-originating traffic.

This feature is configured only in the CLI. A source IP address can be configured for NTP FortiGuard, DNS, RADIUS, TACACS+, and FSSO.

You can use the `get system source-ip status` command to view the services that force their communication to use a specific source IP address.

Example of using the source IP address feature to track logs at a syslog server

Management wants to be able to track logs at a syslog server. There are five log devices; two FortiAnalyzer units that are being used for archival purposes, and three Syslog servers that store all other log files. All log devices have been configured and you must edit the existing Syslog server configuration for the Syslog server that management wants tracked, `syslog_2`. The source IP address is 172.20.120.155.

To include the source IP address to track logs in the existing configuration

- 1 Log in to the CLI.
- 2 Enter the following command:

```
config log syslog2 setting
```
- 3 In the `syslog_2` configuration, enter the following commands:

```
set source-ip 172.20.120.155
end
```
- 4 View the services for `syslog_2` using the following command:

```
get system source-ip status
```

Local-in security policies

Local-in security policies are policies that are designed for traffic that is FortiGate-oriented. For example, central management. There are already local-in policies, which are automatically set up by the FortiGate unit. These policies include central-management, update announcement, and Netbios forward.

When configuring security policies for local-in traffic, the destination address is limited to the FortiGate interface IP and secondary IP addresses. Local-in policies are used in a backward compatible way with allow-access. These security policies are configured only in the CLI. You can configure local-in security policies for both IPv4 and IPv6.

The following are the commands used to configure a local-in security policy:

```
config firewall policy
edit <integer>
set intf <source_interface_name>
set srcaddr <source_address_name>
set dstaddr <destination_address_name>
set action {accept | deny}
set service <service_name>
set schedule <schedule_type>
set auto-asic-offload {enable | disable}
set status {enable | disable}
end
```

Protocol Options

When accessing *Policy > Policy > Protocol Options*, you will notice that you are directed to the Edit Protocol Options page. This page is referred to as the Configuration Settings page, similar to how the UTM profiles and sensors are accessed. A default protocol options list is available as well as your configured protocol options lists.

You can create a new protocol option list from the Configuration Settings page by selecting the *Create New* icon. If you want to view a list of all protocol option lists, select the *View List* icon. You can access a protocol option list at any time on the Settings page by selecting one from the drop-down list beside the *Create New* icon.

FTPS support

FTPS is now supported within the Protocol Options page as well as within the UTM features. This support extends the SSL proxy so that decrypted FTPS data can be examined by the proxies.

Virtual IP source address filter support

In *Firewall Objects > Virtual IP > Virtual IP*, you can now add multiple source IP addresses for filtering purposes. This feature allows packets from different sources to be translated to different VIPs. By default, the filter is set to 0.0.0.0, which means that all source IP addresses provide a backward compatibility. The mapped IP address range is also set to 0.0.0.0 by default.

When you enter the mapped IP address for the mapped range for source address filter, the FortiGate unit automatically calculates the range.

Virtual IP port forwarding enhancements

The VIP port forwarding feature (*Firewall Objects > Virtual IP > Virtual IP*) has been enhanced so that you can easily enter the external service port range first. The FortiGate unit calculates the mapped port range after you enter the start of the port range.

You must select *Port Forwarding* to reveal the configuration settings for port forwarding as well as enable it.

Load balancing HTTP host connections

Load balancing for HTTP host connections can be used for load balancing across multiple real servers using the host's HTTP header to guide the connection to the correct real server, providing better load balancing for those connections. The HTTP host can be configured either in the CLI or web-based manager, in *Firewall Objects > Load Balance > Virtual Server*.

The load balancing method used is called http-host. When selected in the CLI, this allows a real server to specify a http-host attribute which is the domain name of the traffic for that real server. For example, a FortiGate unit is load balancing traffic to three real servers; traffic for www.example.com should go to 10.10.10.1, traffic for www.example.org should go to 10.10.10.5, and traffic for any other domain should go to 10.10.10.100.

Web Proxy Service and Web Proxy Service Group

There are two new menus in *Firewall Objects > Service*: Web Proxy Service and Web Proxy Group.

The Web Proxy Service menu provides configuration settings for web proxy services that can then be applied to a security policy. Web proxy services are similar to custom services, where you can configure the services to define one or more protocols and port numbers that are associated with each web proxy service. Web proxy services can also be grouped, in *Firewall Objects > Service > Web Proxy Service Group*.

The Web Proxy Service Group menu, similar to the Group menu, provides configuration settings for grouping the configured web proxy services. By grouping web proxy services, you can apply multiple services to a security policy.

SSL renegotiation for SSL offloading provides allow/deny client renegotiation

FortiOS now supports SSL offloading that either allows or denies client renegotiation. This feature helps to resolve the issue that affects all SSL and TLS servers that support renegotiation, which was identified by the Common Vulnerabilities and Exposures system, in [CVE-2009-3555](#). The IETF is working on a TLS protocol change that will permanently fix the issue and until they implement the change, the allow and deny client renegotiation feature in FortiOS provides a workaround. This workaround allows you to disable support for SSL/TLS renegotiation in a server, for the SSL offloading feature.

The configuration is in the CLI:

```
config firewall vip
    set ssl-client-renegotiation {allow | deny}
end
```

The allow option is enabled by default for backwards capability. If you choose deny, as soon as a “ClientHello” message (indicating a renegotiation) is received from the client, the server terminates the TCP connection.

You can test the renegotiation behavior using OpenSSL. The OpenSSL client application has a request feature that it can do renegotiation, by typing “R”. When you use this feature, the `diag debug appl vs -1` can be used to view the renegotiation where deny is used.

SSL VPN Port forwarding support

You can now configure port forwarding for Citrix, native RDP and general port forwarding for portals for web mode. The configuration settings are found in the Portal Settings page, in the Settings Window. These port forwarding settings are also available in the CLI.

IKE negotiation

The IKE negotiation process now provides options for how the negotiation is controlled when there is no traffic, as well as how long the FortiGate unit waits for the negotiation to occur. Within the CLI, two new commands help you to configure the `negotiation-timeout` (which is new) and `auto-negotiation` which now replaces `auto-keepalive` or `set keepalive {enable | disable}`.

The `auto-negotiation` command controls whether IKE is negotiation even when there is no traffic. This command would usually be used where there is multiple redundant or overlapping tunnels and there is a need to have the primary connection established. When enabled, the FortiGate unit keeps trying to negotiate IKE event if the link is down and traffic is flowing over a secondary tunnel.

For `auto-negotiation`, if the previous configuration has DPD enabled, the upgrade process automatically enables auto-negotiation so that the behavior is the same as previously configuration.

The `negotiation-timeout` command controls how long the FortiGate unit waits for IKE to negotiate, similar to the web-based manager’s timeout settings. The default time is 30 seconds. If DPD was enabled in a previous configuration, the `negotiate-timeout` settings will be that of the `dpd-retrycount` and `dpd-retryinterval` so that the FortiGate unit will time out connections at the same rate as they would have in the previous build.

SHA-384 and SHA-512 support for IKE

For IPsec, you can now choose either SHA-384 or SHA-512 when configuring IPsec. These authentication algorithms are available for IKE (including phase 1 and phase 2), and manual key configurations.

In the web-based manager, both *Authentication Algorithm* and *Encryption Algorithm* drop-down lists provide the SHA-384 and SHA-512 options for IPsec.

FortiOS Carrier URL extraction feature

The URL extraction feature extracts the embedded Uniform Resource Identifier (URI) within the path for only the host that is specified. The feature applies to HTTP requests for URLs. For example, the URI “http://example.proxy.com/http://www.example.com”; when the URI is broken down, you find the FQDN (example.proxy.com) and the path (http://www.example.com).

The URL extraction feature, however, does not extract the URL if its a regular HTTP request, such as http://example.proxy.com/examples/example.html. The feature also does not extract a URL if the request does not match the FQDN of the proxy server.

This feature is available within a web filter profile, under *URL Extraction*. You must select the *Enable URL Extraction* check box to enable it and access the other settings. The settings that you can choose from are:

- *URL Extraction proxy server FQDN* – the proxy server hostname, such as FQDN, for which the URL extraction will apply. The proxy server hostname must be entered in the field.
- *Blocked page redirect header name* – HTTP header name that is used for client redirect on blocked requests.
- *Blocked page redirect header value (URL)* – HTTP header value that is used for client redirect on blocked requests.

You can also use the CLI command `redirect-no-content` which behaves in the following way:

- `enabled` – if extracted URL is blocked by this feature, the HTTP response contains no content, for example message body is no present.
- `disabled` – the value from Blocked Page redirect header name configuration includes both the redirect header and the message body.



Chapter 2 Firewall

-

This FortiOS Handbook chapter contains the following sections:

[Understanding the FortiGate firewall](#) provides general information about what the FortiGate firewall does, what it is comprised of, and explains how a packet travels through the FortiGate unit.

[Working with NAT in FortiOS](#) provides information about how NAT works in FortiOS and the combinations of NAT that you can use in your configuration. This section explains how the different modes, such as Transparent mode, work and how the FortiGate unit behaves when in each of these modes.

[Firewall components](#) provides in-depth information about the firewall components that help in creating a FortiGate firewall configuration.

[Security policies](#) explains what security policies are, as well as how these rules work to help protect your network. This section also explains the importance of how security policies are ordered within the security policy list, and describes the different policies that can be created for different firewall configurations.

[Monitoring firewall traffic](#) explains how you can monitor traffic within the web-based manager using the Session and Policy Monitoring pages.

[Internet Protocol version 6 \(IPv6\)](#) explains how IPv6 can be implemented in FortiOS, as well as what features support IPv6, such as IPsec VPN and dynamic routing. This section also explains a high-level summary of IPv6.

[Advanced FortiGate firewall concepts](#) explains the advanced firewall features that you may want to configure for your network, as it expands. This section explains advanced firewall features that include stateful inspection of SCTP traffic, port pairing (Transparent mode only), and adding NAT security policies in Transparent mode.



Understanding the FortiGate firewall

The FortiGate firewall is one of the most important features on the FortiGate unit, allowing not only traffic to flow through, but also, with the help of security policies, scan the traffic for vulnerabilities and misuse and abuse. This type of firewall provides flexibility for expansion in a growing network environment.

This section helps to explain the FortiGate firewall and its role in protecting your network. This section also explains the life of a packet, which helps you to understand how the traffic flows through the FortiGate unit and the role the FortiGate firewall plays in the life of a packet.

The following topics are included in this section:

- [What is the FortiGate firewall?](#)
- [FortiGate firewall components](#)
- [Understanding how a packet travels through the FortiGate unit](#)

What is the FortiGate firewall?

A firewall is, in the simplest of terms, a device that permits or denies network traffic based on a set of rules. For the FortiGate firewall, it can do this and much more. The FortiGate firewall scans the network traffic, and based on the set of rules (in Fortinet, however, these rules are called security policies), determines what action needs to be taken. The action may be to quarantine a virus that the FortiGate unit finds, or to record the activity, or both. These security policies provide the information the FortiGate unit needs to determine what to do with the incoming and outgoing traffic.

At the heart of these networking security functions, is the security policies. Security policies control all traffic attempting to pass through the FortiGate unit, and between FortiGate interfaces, zones, and VLAN subinterfaces. They are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. It is through these policies that the FortiGate unit permits or denies the packets to pass through to the network, who gets priority (bandwidth) over other users, and when the packets can come through.

FortiGate firewall components

The FortiGate firewall is comprised of many different features that provides flexibility for the specific needs of your network, both now and as it grows. These features are:

- interfaces (including VLANs)
- zones
- unified threat management (UTM)
- firewall addresses (this includes IPv4 and IPv6, IP pools, . wildcard addresses and netmasks, and geography-based addresses)
- monitoring traffic
- traffic shaping and per-ip traffic shaping (advanced)
- firewall schedules
- services (such as AOL, DHCP and FTP)
- logging traffic (advanced)
- QoS (advanced)
- identity-based policies (advanced)
- endpoint security (advanced)

All of these components each provide an important role in configuring your FortiGate firewall. For example, the administrator applies the PING admin access to the wan1 interface so that he or she can ping this external interface and verify that Internet traffic is hitting the internal to wan1 security policy. If there was no PING admin access applied to the external interface, the administrator could not properly verify if traffic is hitting the policy.

For more in-depth explanations of these components, see the [“Firewall components” on page 195](#).

How the firewall components create a FortiGate firewall and help in protecting your network

The firewall components each help in protecting your network, as well as helping traffic to flow better through the network, for example traffic shaping helps to load balance traffic on your network.

The following explains how all of the firewall components get combined to create the FortiGate firewall.

1 In *System > Network > Interface*, create VLAN subinterfaces for each department: sales, marketing and engineering.

These VLAN subinterfaces will be grouped into a zone and the zone will then be applied to a security policy.

2 Create a zone for the VLAN subinterfaces.

3 In *Firewall Objects > Address > Address*, create the IP address ranges that are required: one for sales, one for marketing, and one for engineering.

Each of these ranges corresponds to the departments that have these IP address ranges. For example, sales has 172.16.120.100 - 172.16.120.200.

4 Create a firewall schedule that allows sales and marketing Internet access all day; create another firewall schedule that allows engineering access to the Internet only during their lunch break.

By creating two different firewall schedules, you can block access for one group for a specified time period, and allow another group all day access.

5 Group the firewall schedules together so that you can apply them both to a security policy.

6 Create a virtual IP address that will be used to allow Internet users access to a web server on your DMZ network.

7 In *Policy > Policy > Policy*, create the following:

- a security policy that allows Internet users access to the web server
- a security policy that applies the firewall schedule group for Internet access for the sales, marketing and engineering departments (this applies the zone)
- a deny policy that blocks FTP downloads

8 With all the policies now in the list, arrange them so that the most important policies are first, and least important are last. The list order is:

- deny policy
- security policy that allows Internet users access to the web server
- security policy for sales, engineering and marketing that allows Internet access

Now that all the policies are in the correct order, you need to test that all are working properly.

9 To verify that traffic is hitting the policies, verify that there is a packet count increase occurring in the Count column of each of the policies in the policy list. Troubleshoot any issues using the *diagnose sniffer* and *diagnose debug flow* commands in the CLI.

By testing that traffic is hitting the policies that you just created, you can see whether you need to solve any issues or not. When you use the `diagnose` commands, you can see detailed information about the traffic hitting the policy.

10 Back up the configuration after testing and troubleshooting.

By backing up the changes you made to the configuration, you ensure that a current configuration of this FortiGate firewall configuration is available at any time.

Understanding how a packet travels through the FortiGate unit

Directed by security policies, a FortiGate unit screens network traffic from the IP layer up through the application layer of the TCP/IP stack. The FortiGate firewall plays an important role in how the packet travels through the FortiGate unit out to its destination. The following explains how the packet travels through the FortiGate unit and how the FortiGate firewall plays a role in the life of a packet.

The FortiGate unit performs three types of security inspection:

- stateful inspection, that provides individual packet-based security within a basic session state
- flow-based inspection, that buffers packets and uses pattern matching to identify security threats
- proxy-based inspection, that reconstructs content passing through the FortiGate unit and inspects the content for security threats.

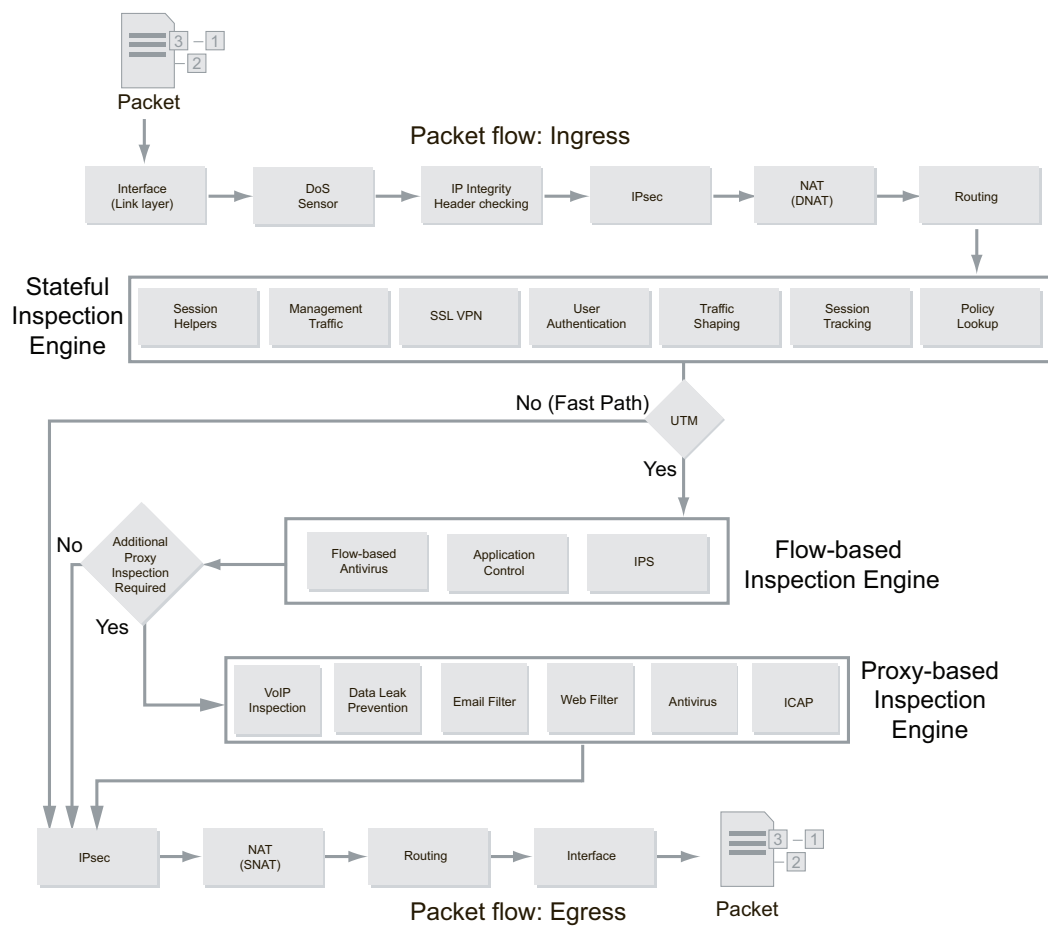
Each inspection component plays a role in the processing of a packet as it traverses the FortiGate unit en route to its destination. When you understand these inspections, you will understand the packet's journey through the FortiGate unit and how the FortiGate firewall helps the packet along to its destination.

For more information about how packets travel through the FortiGate unit, see the Troubleshooting chapter in the *FortiOS Handbook*. The following explains, in a high-level description, of how a packet travels through the FortiGate unit.

How packets flow in and out of the FortiGate unit

The following provides a high-level description of the steps a packet takes when it enters the FortiGate unit, travelling to its destination, the internal network. Similar steps occur for outbound traffic; they are just in reverse.

- 1 An incoming packet enters the external interface of the FortiGate unit to start its journey through to the internal network. This is called ingress. During ingress, the following processes occur:**
 - DoS Sensor
 - IP integrity header checking
 - IPsec
 - Destination NAT (DNAT)
 - Routing
- 2 After the Routing process finishes, the stateful inspection engine processes the packet, and does the following:**
 - Session Helpers
 - Management Traffic
 - SSL VPN
 - User Authentication
 - Traffic Shaping
 - Session Tracking
 - Policy lookup
- 3 If nothing comes from the stateful inspection engine, then the packet travels to the UTM scanning process. This process may have either a flow-based or proxy-based inspection engine that also processes the packet.**
- 4 If nothing matches the UTM rules, the packet then travels to other processing steps, which include:**
 - IPsec
 - NAT (Source NAT)
 - Routing
 - Internal Interface
- 5 After step 4 is finished, the packet travels out of the internal interface of the FortiGate unit, heading towards its final destination, the internal network. This is referred to as Egress.**

Figure 16: Packet flow



Working with NAT in FortiOS

This section explains NAT and the NAT/Route mode of the FortiGate unit, as well as Transparent mode and its role with NAT. This section also explains the types of NAT that FortiOS supports, including combinations of NAT that you can configure in FortiOS.

This section also includes information about Route mode and how it behaves in FortiOS.

The following topics are included in this section:

- [NAT in FortiOS](#)
- [Types of NAT in FortiOS](#)
- [Combining types of NAT](#)

NAT in FortiOS

Network address translation (NAT) translates one IP address (either a source IP address or destination IP address) for another IP address. NAT in FortiOS, however, can translate IP addresses in many different ways, providing the flexibility you need for your specific network requirements. For example, you can use the Central NAT table to help in translating multiple IP addresses.

When configuring NAT in FortiOS, you should also know how it works within the different modes that the FortiGate unit can be configured in.

This topic contains the following:


- [NAT/Route mode](#)
- [Route mode](#)
- [Transparent mode](#)

NAT/Route mode

In NAT/Route mode, the FortiGate unit is visible to the network that is connected to. All of its interfaces are on different subnets. Each interface it is connected to a network that must be configured with an IP address that is valid for that subnetwork.

NAT/Route mode is typically used when the FortiGate unit is deployed as a gateway between private and public networks. In its default NAT mode configuration, the FortiGate unit functions as a firewall. Security policies control communications through the FortiGate unit to both the Internet and between internal networks. In NAT/Route mode, the FortiGate unit performs network address translation before IP packets are sent to the destination network. For example, a company has a FortiGate unit as their interface to the Internet. The FortiGate unit also acts as a router to multiple sub-networks within the company.

In [Figure 17](#), the FortiGate unit is set to NAT/Route mode and is connected to a network. By using this mode, the FortiGate unit can have a designated port for the Internet, and the internal segments are behind the FortiGate unit, which are invisible to the public access. The FortiGate unit translates IP addresses passing through it to route the traffic to the correct subnet on the Internet.


Figure 17: An example of a FortiGate unit in NAT/Route mode on a network

Route mode

In Route mode, the FortiGate unit is only routing traffic, not translating the IP addresses. In this mode, the FortiGate unit acts similar to a switch, passing the packet along to the destination network. This mode is not to be confused with Transparent mode, which is invisible on the network; rather, in Route mode, the FortiGate unit is visible to the network, but does only routing.

The FortiGate unit is used in Route mode whenever no NAT translation needs to be done. For example, you want to connect two separate subnets without using NAT.

You must select NAT/Route mode when configuring the FortiGate unit for Route mode.

Figure 18: An example of a FortiGate unit in Route mode on a network


Transparent mode

In Transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. If you want to configure the FortiGate unit in Transparent mode, all you need to do is to configure a management IP address and a default route.

You would typically use Transparent mode on a private network behind an existing firewall or behind a router. In Transparent mode, the FortiGate unit functions as a firewall and can even perform NAT. Security policies control communications through the FortiGate unit to the Internet and internal network. Traffic cannot pass through until you add security policies when the FortiGate unit is in Transparent mode.

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools. If you want NAT to be performed in Transparent mode, you must configure two management IP addresses that are on different subnets.

Figure 19: A FortiGate unit in Transparent mode



Types of NAT in FortiOS

There are many types of NAT that are available, some you may already know such as port address translation (PAT). The following explains these types of NAT that are available in FortiOS.

This topic contains the following:

- [Static NAT \(SNAT\)](#)
- [Dynamic NAT \(DNAT\)](#)

Static NAT (SNAT)

Static NAT, or source address translation (SNAT), is when a static source IP address is translated by NAT to another source IP address. In FortiOS, when a packet with a specific source address is accepted by a security policy with NAT enabled, the source address is swapped with another IP address. For example, you want to allow a web server on a private network that is protected by a FortiGate unit to connect to the Internet; the web server has a static IP address of 10.10.30.10 and the external interface of the unit is 172.20.120.233; when the packet is received at the FortiGate's internal interface, it is translated from 10.10.30.10 to 172.20.120.133, and forwards the packet out to the Internet.

Static NAT is used when configuring basic security policies. For example, you want users on a private network to connect to the Internet.

When configuring static NAT security policies, there are several steps that must be configured prior to configuring the actual security policy. For example, for static DNAT, you must configure a virtual IP address that maps to a specific destination address.

Static Destination NAT (SDNAT)

As stated for static NAT, the same is true for static destination address translation, or SDNAT, whereby a packet with a specific destination address is accepted by a security policy with NAT enabled, the destination address is swapped with another destination address.

Static NAT port forwarding

There is also static port forwarding, which acts similarly to static DNAT, translating a destination address and port number to another destination address and port number. The difference is that port forwarding requires a virtual IP address so that the FortiGate unit can properly translate the port number.

When a packet with a destination address to be translated is accepted by a security policy (with DNAT enabled), and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

Dynamic NAT (DNAT)

As subnets grow larger, more work is required to set network address translation with each additional client. Rather than assigning static addresses, an administrator may want to set up IP pools. IP pools are ranges of addresses that clients on a subnet can use to send and receive packets, as well as which FortiGate units can use to translate the addresses of packets going through them. This type of translation is known as dynamic NAT, when address translation is done on a flexible or "many-to-one" basis using IP pools.

IP pools do not randomly assign addresses, rather, each IP pool is a prioritized list of IP addresses. When a client is assigned an IP address from the IP pool, it retains that address. Another client that requires an address is then assigned the next IP address from that IP pool list. When the range of virtual IPs are used instead of IP pools, these virtual IPs are prioritized in the same type of list.

Dynamic source address translation

Dynamic source address translation has economies of scale for larger subnets and more flexible subnets, enabling network infrastructure to change without the hassle of reconfiguring addresses after every change. Dynamic source address NAT or DNAT translates many source addresses as defined by an IP pool. Whenever a packet with the specific source address to be translated is accepted by a security policy with source NAT enabled, the FortiGate unit swaps the packet's source address with the other IP address selected from the IP pool.

For example, an organization may want packets leaving the FortiGate unit for the Internet to have source IPs in the range of 172.16.0.1-10. This means that packets accepted by a firewall policy must have their source addresses translated to an address in this range before being forwarded to the Internet. So if the server on the private network with the address 10.0.0.1 has its source IP translated to 172.16.0.1, then the next available source IP in the IP pool will be 172.16.0.2, which a server with address 10.0.0.2 can use.

Dynamic destination address

Dynamic destination address NAT (or DDNAT) translates one range of destination addresses to another range of destination addresses. Whenever a packet within the specified range of destination addresses to be translated is accepted by a security policy with DNAT is enabled, the FortiGate unit swaps the packet's destination address with one of the addresses from the other specified range.

For example, to allow customers from the Internet to connect to several web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10), enough for each protected web server, and a range of real addresses (for example, 10.0.0.1-10) for each web server. When a packet is received at the external interface of the FortiGate unit with a destination IP address within the Internet range of addresses, the FortiGate unit translates the destination address of the packet to the real address and forwards the packet to the web server on the network protected by the FortiGate unit.

Dynamic port forwarding

Dynamic port forwarding translates one range of destination addresses and ports to another range of destination addresses and ports. Whenever a packet with a specified destination address to be translated is accepted by a security policy with destination NAT enabled, and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

For example, to allow customers from the Internet to connect to web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10) and a range of port numbers (for example, 80-89), and a range of ports numbers to be mapped to (for example, 8080-8089). When a packet is received at the external interface of the FortiGate unit with the 172.16.0.3 destination IP address and port number 8082, the FortiGate unit translates that address to 10.0.0.3 and port number to 82, and then forwards the packet to the web server.

Combining types of NAT

In FortiOS, you can combine a number of NAT features to get the best firewall configuration possible for your network requirements. NAT combinations include Double NAT, which is combining IP pool with virtual IP, and using VIP range for SNAT and static one-to-one mapping.

These combinations can help you when creating your FortiGate firewall configuration. The combinations help when you have multiple addresses (IP pools) and when you need to use a virtual IP address with the IP pool. An example of this combination is called Double NAT.

You can also combine dynamic NAT types, such as dynamic source address translation, to help you with creating the FortiGate firewall using dynamic NAT. An example of this combination is using the Central NAT table.

When considering your FortiGate firewall configuration, you should also consider how to combine NAT types. By combining NAT types, you can easily use multiple addresses when configuring security policies, as well as when you want to provide specific NAT translations, such as using dynamic source NAT that will not change the source port; this combination allows for the handling of specific protocols or services that function only if they use a specific port and that port does not change.

The following are some combinations of NAT that you can use in your FortiGate firewall configuration:

- Double NAT
- Central NAT table (similar to IP pools)
- virtual IP range for SNAT
- static one-to-one mapping
- dynamic source NAT (also known as one-to-one source NAT)
- dynamic source NAT (this uses Dynamic IP pool and a virtual IP)



Firewall components

The FortiGate unit's primary purpose is to act as a firewall to protect your networks from unwanted attacks and to control the flow of network traffic. The firewall consists of many different and important components so that you can better protect your network as your network requirements grow. This section explains these components.

The following topics are included in this section:

- [Using Interfaces and zones in the FortiGate firewall](#)
- [Understanding the firewall address component](#)
- [UTM profiles](#)

Using Interfaces and zones in the FortiGate firewall

Interfaces and zones are used when configuring security policies to define incoming and outgoing traffic. For example, in an internal to wan 1 security policy, the internal interface is where traffic is coming in, and the wan 1 interface is where the traffic is going out to. When the FortiGate unit sees that traffic came in using the internal interface, and needs to leave using the wan 1 (or external interface), the security policy internal to wan1 is matched to the traffic and additional rules are applied to the traffic as well.

Interfaces, either virtual or physical, can be applied to security policies. VLAN subinterfaces are virtual interfaces that can be applied to security policies to control and direct traffic on those subinterfaces. VLAN subinterfaces are interfaces that are part of one of the main interfaces, for example, wan1. For more information about VLAN subinterfaces and how to configure them, see the System Admin chapter of the [FortiOS Handbook](#).

Zones provide the option of grouping multiple FortiGate interfaces, both virtual and physical, that you can then apply to security policies to control the incoming and outgoing traffic on those interfaces. By using zones, you can easily group multiple interfaces and VLAN subinterfaces together to help simplify creating security policies where a number of network segments can use the same policy and UTM settings.

How to apply VLANs and zones and to a security policy

The following explains how to create three VLAN subinterfaces, grouping these subinterfaces into a zone, and then applying the zone to a security policy. The security policy will control the traffic for these VLAN subinterfaces.

1 Create three VLANs in *System > Network > Interface* for engineering, sales and marketing on the internal interface.

These three VLANs will be grouped together to create a zone which will then be applied to the security policy. The zone will be applied to the policy instead of the individual VLANs.

2 Group the VLANs into a zone.

3 Create DHCP servers for each of the VLAN subinterfaces in *System > Network > DHCP*.

4 Create the security policy for the zone to control traffic in *Policy > Policy > Policy*.

In the *Source Interface/Zone* list, you would instead choose the zone. The *Destination Interface/Zone* is the external interface, wan1. By choosing the zone, you apply all the subinterfaces at once.

5 Select *Enable NAT* and *Use Destination Interface Address*; ensure that *Log Allowed Traffic* is also enabled so that you can use the logs to help determine if traffic is hitting the security policy.

Understanding the firewall address component

Firewall addresses in FortiOS provide flexibility when configuring access control over the network traffic. When this document talks about firewall addresses, this encompasses:

- IP addresses and netmasks
- IP pools (this can include the Central NAT table)
- virtual IP addresses
- geography-based addresses
- IPv4 addresses
- wildcard addresses and netmasks
- Fully Qualified Domain Name addresses (FQDN)
- IP address groups

Firewall addresses help define the network addresses that you use when configuring a security policy's source and destination address. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

A firewall address can contain one or more network addresses. Network addresses can be represented by an IP address with a netmask, an IP address range, or a fully qualified domain name (FQDN).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- a single computer, such as 192.45.46.45
- a subnetwork, such as 192.168.1.0 for a class C subnet
- 0.0.0.0, which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- netmask for a single computer: 255.255.255.255, or /32
- netmask for a class A subnet: 255.0.0.0, or /8
- netmask for a class B subnet: 255.255.0.0, or /16
- netmask for a class C subnet: 255.255.255.0, or /24
- netmask including all IP addresses: 0.0.0.0

Valid IP address and netmask formats include:

- x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

- x.x.x.x/x, such as 192.168.1.0/24



An IP address of 0.0.0.0 with a netmask 255.255.255.255 is not a valid firewall address.

When representing hosts by an IP address range, the range indicates hosts with continuous IP addresses in a subnet, such as 192.168.1.[2-10], or 192.168.1.* to indicate the complete range of hosts on that subnet. Valid IP Range formats include:

- x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120
- x.x.x.[x-x], such as 192.168.110.[100-120]
- x.x.x.*, such as 192.168.110.*

When representing hosts by an FQDN, the domain name can be a subdomain, such as mail.example.com. A single FQDN firewall address may be used to apply a security policy to multiple hosts, as in load balancing and high availability (HA) configurations. FortiGate units automatically resolve and maintain a record of all addresses to which the FQDN resolves. Valid FQDN formats include:

- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as mail.example.com
- <host_name>.<top_level_domain_name>



Be cautious when employing FQDN firewall addresses. By using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. If the DNS server should ever be compromised, security policies requiring domain name resolution may no longer function properly.

This topic contains the following:

- [IP addresses for self-originated traffic](#)
- [IP pools](#)
- [IP Pools for security policies that use fixed ports](#)
- [Source IP address and IP pool address matching](#)
- [Geography-based addressing](#)
- [Wildcard addresses](#)
- [Fully Qualified Domain Name addresses](#)
- [Address groups](#)
- [Virtual IP addresses](#)

IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog

- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Configuration of these services is performed in the CLI. In each instance, there is a command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
  set ntpsyn enable
  set syncinterval 5
  set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

IP pools

An IP pool defines a single IP address or a range of IP addresses. A single IP address in an IP pool becomes a range of one IP address. For example, if you enter an IP pool as 1.1.1.1, the IP pool is actually the address range, 1.1.1.1 to 1.1.1.1. Use IP pools to add NAT policies that translate source addresses to addresses randomly selected from the IP pool, rather than the IP address assigned to that FortiGate interface. You can use the Central NAT table as a way to configure IP pools. For more information, see [“Central NAT table” on page 265](#).

If a FortiGate interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

- (1.1.1.0-1.1.1.255) and (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20

The port2 interface overlap IP range with IP_pool_2 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20

The port2 interface overlap IP range with IP_pool_3 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select *Enable NAT* in a security policy and then select *Dynamic IP Pool*. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool.

IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

IP Pools for security policies that use fixed ports

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable `fixedport` when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
  edit policy_name
    ...
    set fixedport enable
    ...
end
```

However, enabling `fixedport` means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select *Dynamic IP pool* in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Source IP address and IP pool address matching

When the source addresses are translated to the IP pool addresses, one of the following three cases may occur:

Scenario 1: The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Original address	Change to
192.168.1.1	172.16.30.1
192.168.1.2	172.16.30.2
.....
192.168.1.254	172.16.30.254

Scenario 2: The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
.....
192.168.1.10	172.16.30.19
192.168.1.11	172.16.30.10
192.168.1.12	172.16.30.11
192.168.1.13	172.16.30.12
.....

Scenario 3: The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
192.168.1.3	172.16.30.12
No more source addresses	172.16.30.13 and other addresses are not used

Geography-based addressing

An option is available to add a geography-based address scheme. With this type of addressing, you indicate the geographic region, or country. The FortiGate unit includes an internal list of countries and IP addresses based on historical data from the FortiGuard network.



IPv6 does not support geography-based addressing. This feature is for IPv4 addresses only.

When used in security policies, traffic originating or going to a particular country can be logged, blocked or specific filtering applied.

In the following examples, an geographic-based address for China is added for the WAN1 port.

To add a geography-based address - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.
- 2 Enter the *Name* of China
- 3 For the *Type*, select *Geography*.
- 4 From the *Country* list, select *China*.
- 5 Select the *Interface* of WAN1.
- 6 Select *OK*.

To add a geography-based address - CLI

```
config firewall address
  edit China
    set type geography
    set country CN
    set interface wan1
  end
```

You can use a `diagnose` command to view more information about geography-based addressing. The command displays country and address information for the countries that have been added to firewall addresses.

```
diagnose firewall ipgeo {country-list | ip-list | ip2country}
```

Where:

- `country-list` shows all of the countries that have been added to a firewall address.
- `ip-list` shows the IP addresses of a specified country or all of the countries added to firewall addresses.
- `ip2country` shows the country of origin for a specified IP address. The address must be assigned to one of the countries that has been added to a firewall address.

Wildcard addresses

Wildcard addresses are addresses that identify ranges of IP addresses, reducing the amount of firewall addresses and security policies required to match some of the traffic on your network. Wildcard addresses are an advanced feature, usually required only for complex networks with complex firewall filtering requirements. By using these wildcard addresses in the firewall configuration, administrators can eliminate creating multiple, separate IP addresses and then grouping them to then apply to multiple security policies.

A wildcard address consists of an IP address and a wildcard netmask, for example, 192.168.0.56 255.255.0.255. In this example, the IP address is 192.168.0.56 and the wildcard netmask is 255.255.0.255. The IP address defines the networks to match and the wildcard netmask defines the specific addresses to match on these networks.

In a wildcard netmask, zero means ignore the value of the octet in the IP address, which means the wildcard firewall address matches any number in this address octet. This also means that the number included in this octet of IP address is ignored and can be any number. Usually, if the octet in the wildcard netmask is zero, the corresponding octet in the IP address is also zero.

In a wildcard netmask, a number means match addresses according to how the numbers translate into binary addresses. For example, the wildcard netmask is 255; the wildcard address will only match addresses with the value for this octet that is in the IP address part of the wildcard address. So, if the first octet of the IP address is 192 and the first octet of the wildcard netmask is 255, the wildcard address will only match addresses with 192 in the first octet.

In the above example, the wildcard address 192.168.0.56 255.255.0.255 would match the following IP addresses:

192.168.0.56, 192.168.1.56, 192.168.2.56, ..., 192.168.255.56

The wildcard addresses 192.168.0.56 255.255.0.255 and 192.168.1.56 255.255.0.255 define the same thing since the 0 in the wildcard mask means to match any address in the third octet.

If we use the wildcard address 172.0.20.10 255.0.255.255, it would match the following IP addresses:

172.1.20.10, 172.2.20.10, 172.3.20.10, ..., 172.255.20.10

In a wildcard netmask, a number other than 255 matches multiple addresses for this octet. You can perform a binary conversion to calculate the addresses that would be matched by a given value. For example, to create the IP address and wildcard netmask to match the following network addresses:

192.168.32.0/24
192.168.33.0/24
192.168.34.0/24
192.168.35.0/24
192.168.36.0/24
192.168.37.0/24
192.168.38.0/24
192.168.39.0/24

Table 4 shows how to write the third octet for these networks according to the octet bit position and address value for each bit.

Table 4: Octet bit position and address value for each bit

Decimal	128	64	32	16	8	4	2	1
32	0	0	1	0	0	0	0	0

Table 4: Octet bit position and address value for each bit

33	0	0	1	0	0	0	0	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
36	0	0	1	0	0	1	0	0
37	0	0	1	0	0	1	0	1
38	0	0	1	0	0	1	1	0
39	0	0	1	0	0	1	1	1
	M	M	M	M	M	D	D	D

Since the first five bits match, the networks can be summarized into one network (192.168.32.0/21 or 192.168.32.0 255.255.248.0). All eight possible combinations of the three low-order bits are relevant for the network ranges. The wildcard address that would match all of these subnet addresses can be written as 192.168.32.0 255.255.248.0.

Wildcard addresses are similar to routing access list wildcard masks. You add routing access lists containing wildcard masks using the `config router access-list` command. However, router access list wildcard masks use the inverse of the masking system used for firewall wildcard addresses. For the router access list wildcard masks, zero (0) means match all IP addresses and one (1) means ignore all IP addresses. So to match IP addresses 192.168.0.56, 192.268.1.56, 192.168.2.56, ... 192.168.255.56 you would use the following router access IP address prefix and wildcard mask: 192.168.0.56 0.0.255.0.

Wildcard firewall addresses are configured only in the CLI. The following is an example of how to configure a wildcard firewall address.

```
config firewall address
  edit example_wildcard_address
    set type wildcard
    set wildcard 192.168.0.56 255.255.0.255
  end
```

Using wildcard addresses in the firewall configuration

The following example shows how wildcard addresses can be applied to network traffic. This example consists of a security policy where both the source and destination addresses are firewall wildcard addresses.

Source Address: 10.129.5.0 255.127.7.0

Destination Address: 10.129.0.10 255.127.7.255

A security policy with these source and destination addresses would permit:

- A device with IP address 10.129.5.100 to connect through the FortiGate unit to IP address 10.129.0.10
- A device with IP address 10.129.13.100 to connect through the FortiGate unit to IP address 10.129.8.10
- A device with IP address 10.129.21.100 to connect through the FortiGate unit to IP address 10.129.0.10

In another example of wildcard addresses, the following shows how only odd numbered addresses get allowed through:

- 1 Create wildcard address 4.2.2.0/255.255.255.1.

This is configured in the CLI.

- 2 Create a deny security policy that uses the wildcard address, 4.2.2.0.

The results are that only the odd-numbered 4.2.2.0 addresses are allowed in; all other addresses are blocked.

Fully Qualified Domain Name addresses



Be cautious when employing FQDN firewall addresses. Using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Using Fully Qualified Domain Name (FQDN) addresses in security policies has the advantage of causing the FortiGate unit to keep track of DNS TTLs and adapt as records change. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache. The FortiGate unit will query the DNS for an amount of time specified, in seconds, and update the cache as required. This feature can reduce maintenance requirements for changing firewall addresses for dynamic IP addresses. This also means that you can create security policies for networks configured with dynamic addresses using DHCP.

You specify the TTL time in the CLI only. For example, to set the TTL for 30 minutes on an FQDN of `www.example.com` on port 1, enter the following commands:

```
config firewall address
  edit FQDN_example
    set type fdqn
    set associated-interface port 1
    set fqdn www.example.com
    set cache-ttl 1800
  end
```

Address groups

Similar to zones, if you have a number of addresses or address ranges that require the same security policies, you can put them into address groups, rather than creating multiple similar policies. Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to *Any* — addresses whose selected interface is *Any* are bound to a network interface during creation of a security policy, rather than during creation of the firewall address.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of *Any*, they can be grouped, even if the addresses involve different networks.

You cannot mix IPv4 firewall addresses and IPv6 firewall addresses in the same address group.

Virtual IP addresses

In FortiOS, virtual IP addresses (VIPs) can be used when configuring security policies to translate IP addresses and ports of packets received by a network interface. When the FortiGate unit receives inbound packets matching a security policy whose *Destination Address* field is a virtual IP, the FortiGate unit applies NAT, replacing packets's IP addresses with the virtual IP's mapped IP address.

VIPs can specify translation of packets' port numbers and/or IP addresses for both inbound and outbound connections. In Transparent mode, virtual IPs are available only in the CLI.

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

Grouping virtual IPs

You can organize multiple virtual IPs into a virtual IP group to simplify your security policy list. For example, instead of having five identical policies for five different but related virtual IPs located on the same network interface, you might combine the five virtual IPs into a single virtual IP group, which is used by a single security policy.

Security policies using VIP Groups are matched by comparing both the member VIP IP addresses) and port numbers).

Match-vip

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. Use the following command syntax to enable this feature. By default, it is disabled.

```
config firewall policy
  edit <vip_policy_name>
    set match-vip {disable | enable}
  end
```

How to use match-vip

In this example, a deny security policy has already been configured that blocks FTP sessions. A virtual IP address will be configured in this example and then applied to a security policy that allows Internet users access to a web server on the company's DMZ network.

1 Create the virtual IP address in *Firewall Objects > Virtual IP > Virtual IP*.

This address is called vip-dmz. You can configure the virtual IP address solely in the CLI. This would eliminate having to go back and forth.

2 Log in to the CLI and enter the following commands:

```
config firewall policy
  edit vip-dmz
    set match-vip enable
  end
```

3 Create the virtual IP security policy.

For this security policy, you need to turn on logging within the security policy.

4 Test the policy to view the activity that is occurring with the `match-vip` command enabled.

Services

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined in firewall services and protocols on the FortiGate unit. These predefined services and protocols are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

To view the predefined servers, go to *Firewall Objects > Service > Predefined*.

If there is a service that does not appear on the list, or you have a unique service or situation, you can create your own custom service. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create the custom service.

Predefined service list

Many well-known traffic types have been predefined in firewall services. These predefined services are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

Predefined services are located in *Firewall Objects > Service > Predefined*. [Table 5](#) lists the FortiGate firewall predefined services.

Table 5: Predefined services

Service name	Description	Protocol	Port
AFS3	Advanced File Security Encrypted File, version 3, of the AFS distributed file system protocol.	TCP	7000-7009
		UDP	7000-7009
AH	Authentication Header. AH provides source host authentication and data integrity, but not secrecy. This protocol is used for authentication by IPSec remote gateways set to aggressive mode.	IP	51
ANY	Matches connections using any protocol over IP.	all	all
AOL	America Online Instant Message protocol.	TCP	5190-5194
BGP	Border Gateway Protocol. BGP is an interior/exterior routing protocol.	TCP	179

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
CVSPSERVER	Concurrent Versions System Proxy Server. CVSPServer is very good for providing anonymous CVS access to a repository.	TCP	2401
		UDP	2401
DCE-RPC	Distributed Computing Environment / Remote Procedure Calls. Applications using DCE-RPC can call procedures from another application without having to know on which host the other application is running.	TCP	135
		UDP	135
DHCP	Dynamic Host Configuration Protocol. DHCP allocates network addresses and delivers configuration parameters from DHCP servers to hosts.	UDP	67 68
DHCP6	Dynamic Host Configuration Protocol for IPv6.	UDP	546, 547
DNS	Domain Name Service. DNS resolves domain names into IP addresses.	TCP	53
		UDP	53
ESP	Encapsulating Security Payload. ESP is used by manual key and AutoIKE IPsec VPN tunnels for communicating encrypted data. AutoIKE VPN tunnels use ESP after establishing the tunnel by IKE.	IP	50
FINGER	A network service providing information about users.	TCP	79
FTP	File Transfer Protocol.	TCP	21
FTP_GET	File Transfer Protocol. FTP GET sessions transfer remote files from an FTP server to an FTP client computer.	TCP	21
FTP_PUT	File Transfer Protocol. FTP PUT sessions transfer local files from an FTP client to an FTP server.	TCP	21
GOPHER	Gopher organizes and displays Internet server contents as a hierarchically structured list of files.	TCP	70
GRE	Generic Routing Encapsulation. GRE allows an arbitrary network protocol to be transmitted over any other arbitrary network protocol, by encapsulating the packets of the protocol within GRE packets.	IP	47
GTP (FortiOS Carrier only)	GPRS Tunneling protocol (GTP). GTP is used with GSM and UMTS networks to carry user data within GPRS core networks. FortiOS Carrier can accept and process IPv4 GTP packet.	UDP	2123,2152,3386

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
H323	H.323 multimedia protocol. H.323 is a standard approved by the International Telecommunication Union (ITU) defining how audiovisual conferencing data can be transmitted across networks. For more information, see the FortiGate Support for H.323 Technical Note .	TCP	1720, 1503
		UDP	1719
HTTP	Hypertext Transfer Protocol. HTTP is used to browse web pages on the World Wide Web.	TCP	80
HTTPS	HTTP with secure socket layer (SSL). HTTPS is used for secure communication with web servers.	TCP	443
ICMP_ANY	Internet Control Message Protocol. ICMP allows control messages and error reporting between a host and gateway (Internet).	ICMP	Any
IKE	Internet Key Exchange. IKE obtains authenticated keying material for use with the Internet Security Association and Key Management Protocol (ISAKMP) for IPSEC.	UDP	500, 4500
IMAP	Internet Message Access Protocol. IMAP is used by email clients to retrieve email messages from email servers.	TCP	143
IMAPS	IMAP with SSL. IMAPS is used for secure IMAP communication between email clients and servers. IMAPS is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	993
INFO_ADDRESS	ICMP information request messages.	ICMP	17
INFO_REQUEST	ICMP address mask request messages.	ICMP	15
IRC	Internet Relay Chat. IRC allows users to join chat channels.	TCP	6660-6669
Internet-Locator-Service	Internet Locator Service. ILS includes LDAP, User Locator Service, and LDAP over TLS/SSL.	TCP	389
L2TP	Layer 2 Tunneling Protocol. L2TP is a PPP-based tunnel protocol for remote access.	TCP	1701
		UDP	1701
LDAP	Lightweight Directory Access Protocol. LDAP is used to access information directories.	TCP	389
MGCP	Media Gateway Control Protocol. MGCP is used by call agents and media gateways in distributed Voice over IP (VoIP) systems.	UDP	2427, 2727

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
MMS (FortiOS Carrier only)	MMS tunneling protocol. MMS is used when sending and receiving multimedia content to a mobile phone.	TCP UDP	1755 1024-5000
MS-SQL	Microsoft SQL Server is a relational database management system (RDBMS) produced by Microsoft. Its primary query languages are MS-SQL and T-SQL.	TCP	1433, 1434
MYSQL	MySQL is a relational database management system (RDBMS) which runs as a server providing multi-user access to a number of databases.	TCP	3306
NFS	Network File System. NFS allows network users to mount shared files.	TCP	111, 2049
		UDP	111, 2049
NNTP	Network News Transport Protocol. NNTP is used to post, distribute, and retrieve Usenet messages.	TCP	119
NTP	Network Time Protocol. NTP synchronizes a host's time with a time server.	TCP	123
		UDP	123
NetMeeting	NetMeeting allows users to teleconference using the Internet as the transmission medium.	TCP	1720
ONC-RPC	Open Network Computing Remote Procedure Call. ONC-RPC is a widely deployed remote procedure call system.	TCP	111
		UDP	111
OSPF	Open Shortest Path First. OSPF is a common link state routing protocol.	IP	89
PC-Anywhere	PC-Anywhere is a remote control and file transfer protocol.	TCP	5631
		UDP	5632
PING	Ping sends ICMP echo request/replies to test connectivity to other hosts.	ICMP	8
PING6	Ping6 sends ICMPv6 echo request/replies to network hosts to test IPv6 connectivity to other hosts.	ICMP6	58
POP3	Post Office Protocol v3. POP retrieves email messages.	TCP	110
POP3S	Post Office Protocol v3 with secure socket layer (SSL). POP3S is used for secure retrieval of email messages. POP3S is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	995

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
PPTP	Point-to-Point Tunneling Protocol. PPTP is used to tunnel connections between private network hosts over the Internet. Note: Also requires IP protocol 47.		47
		TCP	1723
QUAKE	Quake multi-player computer game traffic.	UDP	26000, 27000, 27910, 27960
RADIUS	Remote Authentication Dial In User Service. RADIUS is a networking protocol that provides centralized access, authorization and accounting management for people or computers to connect and use a network service.	TCP	1812, 1813
RAUDIO	RealAudio multimedia traffic.	UDP	7070
RDP	Remote Desktop Protocol is a multi-channel protocol that allows a user to connect to a networked computer.	TCP	3389
REXEC	Rexec traffic allows specified commands to be executed on a remote host running the rexecd service (daemon).	TCP	512
RIP	Routing Information Protocol. RIP is a common distance vector routing protocol. This service matches RIP v1.	UDP	520
RLOGIN	Remote login traffic.	TCP	513
RSH	Remote Shell traffic allows specified commands to be executed on a remote host running the rshd service (daemon).	TCP	514
RTSP	Real Time Streaming Protocol is a protocol for use in streaming media systems which allows a client to remotely control a streaming media server, issuing VCR-like commands such as play and pause, and allowing time-based access to files on a server.	TCP	554, 7070, 8554
		UDP	554
SAMBA	Server Message Block. SMB allows clients to use file and print shares from enabled hosts. This is primarily used for Microsoft Windows hosts, but may be used with operating systems running the Samba daemon.	TCP	139
SCCP	Skinny Client Control Protocol. SCCP is a Cisco proprietary standard for terminal control for use with voice over IP (VoIP).	TCP	2000

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
SIP	Session Initiation Protocol. SIP allows audiovisual conferencing data to be transmitted across networks. For more information, see the Voice Solutions: SIP chapter of the FortiOS Handbook.	UDP	5060
SIP-MSNmessenger	Session Initiation Protocol used by Microsoft Messenger to initiate an interactive, possibly multimedia session.	TCP	1863
SMTP	Simple Mail Transfer Protocol. SMTP is used for sending email messages between email clients and email servers, and between email servers.	TCP	25
SMTPS	SMTP with SSL. Used for sending email messages between email clients and email servers, and between email servers securely. SMTPS is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook .	TCP	465
SNMP	Simple Network Management Protocol. SNMP can be used to monitor and manage complex networks.	TCP	161-162
		UDP	161-162
SOCKS	SOCKeT S. SOCKS is an Internet protocol that allows client-server applications to transparently use the services of a network firewall.	TCP	1080
		UDP	1080
SQUID	A proxy server and web cache daemon that has a wide variety of uses that includes speeding up a web server by caching repeated requests; caching web, DNS and other computer network lookups for a group of people sharing network resources; aiding security by filtering traffic.	TCP	3128
SSH	Secure Shell. SSH allows secure remote management and tunneling.	TCP	22
		UDP	22
SYSLOG	Syslog service for remote logging.	UDP	514
TALK	Talk allows conversations between two or more users.	UDP	517-518
TCP	Matches connections using any TCP port.	TCP	0-65535
TELNET	Allows plain text remote management.	TCP	23

Table 5: Predefined services (Continued)

Service name	Description	Protocol	Port
TFTP	Trivial File Transfer Protocol. TFTP is similar to FTP, but without security features such as authentication.	UDP	69
TIMESTAMP	ICMP timestamp request messages.	ICMP	13
TRACEROUTE	A computer network tool used to determine the route taken by packets across an IP network.	TCP	33434
		UDP	33434
UDP	Matches connections using any UDP port.	UDP	0-65535
UUCP	Unix to Unix Copy Protocol. UUCP provides simple file copying.	UDP	540
VDOLIVE	VDO Live streaming multimedia traffic.	TCP	7000-7010
VNC	Virtual Network Computing. VNC is a graphical desktop sharing system which uses the RFB protocol to remotely control another computer.	TCP	5900
WAIS	Wide Area Information Server. WAIS is an Internet search protocol which may be used in conjunction with Gopher.	TCP	210
WINFRAME	WinFrame provides communications between computers running Windows NT, or Citrix WinFrame/MetaFrame.	TCP	1494
WINS	Windows Internet Name Service is Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	TCP	1512
		UDP	1512
X-WINDOWS	X Window System (also known as X11) can forward the graphical shell from an X Window server to X Window client.	TCP	6000-6063

Service groups

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single address group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single service group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

Firewall schedules

When you add security policies on a FortiGate unit, those policies are always on, policing the traffic through the device. Firewall schedules control when policies are in effect, that is, when they are on. You can create one-time schedules which are schedules that are in effect only once for the period of time specified in the schedule. You can also create recurring schedules that are in effect repeatedly at specified times of specified days of the week.

You can create a recurring schedule that activates a policy during a specified period of time. For example, you might prevent game playing during office hours by creating a recurring schedule that covers office hours.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next. For example, to prevent game playing except at lunchtime, you might set the start time for a recurring schedule at 1:00 p.m. and the stop time at 12:00 noon. To create a recurring schedule that runs for 24 hours, set the start and stop times to 00.

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and on-time schedules. Schedule groups cannot contain other schedule groups.

Schedule expiry

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the `config firewall` command enter the command:

```
set schedule-timeout enable
```

By default, this is set to disable.

UTM profiles

Where security policies provide the instructions to the FortiGate unit as to what traffic is allowed through the device, the Unified Threat Management (UTM) profiles provide the screening that filters the content coming and going on the network. The UTM profiles enable you to instruct the FortiGate unit what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A UTM profile is a group of options and filters that you can apply to one or more firewall policies. UTM profiles can be used by more than one security policy. You can configure sets of UTM profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same UTM profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

UTM profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure UTM profiles in the *UTM* menu and applied when creating a security policy by selecting the UTM profile type.

For more information about configuring profiles that will be used in a security policy, see the UTM chapter of the [FortiOS Handbook](#).

How to use UTM profiles to monitor and protect your network

In this example, UTM profiles help you in monitoring and protecting your network from viruses, email filtering and web filtering. This example uses the default UTM profiles.

- 1 Locate the security policy that allows access to the Internet (internal -> wan 1) in *Policy > Policy > Policy*.**
- 2 On the Edit Policy page, select UTM and then select these options: *Enable Antivirus*, *Enable Web Filter* and *Enable Email Filter*.**

The FortiGate unit will apply the antivirus, web filter, and email filter settings to the packet if a match is found.

- 3 Select *OK*.**

When packets enter the FortiGate unit's internal interface, if a packet matches the internal -> wan 1 policy, the FortiGate unit now scans for viruses and applies any web filtering and email filtering rules if there are matches as well.

- 4 Go to the eicar.org web site and download the eicar test file.**

By downloading the eicar test file, you can determine that the antivirus profile is working properly, as well as to see this activity on the AV Monitor page. When attempting to download the file, a web page appears, stating that you are not permitted to download the file. This indicates that the antivirus profile is working properly.

5 Go to *UTM Profiles > Monitor > AV Monitor* to view the virus activity that just occurred.

On the page, you should see that the eicar test file was detected by the FortiGate unit; you can select the bar in the chart to see more details. This takes you directly to the FortiGuard Virus Encyclopedia.

6 Go to *UTM Profiles > Monitor > Web Monitor* to view the Internet activity that is occurring on your network.

On the page, you will see a pie chart that displays all HTTP requests and a bar chart that displays all blocked HTTP requests. If you want to view more detailed information about the blocked requests, hover your mouse over a bar; a tool-tip appears stating how many blocked requests occurred for that item. For example, for Virus, it is one blocked request because you tried to download the eicar test file.

7 Go to *UTM Profiles > Monitor > Email Monitor* to view the email activity that is occurring on your network.

On the page, you will see both a pie chart and a bar chart, similar to the Web Monitor page. The pie chart displays all the email activity and the bar chart displays all the blocked emails for SMTP, POP3, IMAP, and NNTP.



Security policies

Security policies control all traffic attempting to pass through the FortiGate unit, between FortiGate interfaces, zones, and VLAN subinterfaces.

This section explains what security policies are and how they affect all traffic to and from your network. This section also describes how to configure basic policies which are used as a building block to more complex policies, but they enable you to get the FortiGate unit running on the network quickly.

The following topics are included in this section:

- [Security policy overview](#)
- [Policy order](#)
- [Security policies](#)
- [Creating a basic security policy](#)

Security policy overview

Security policies are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. The FortiGate unit requires one security policy per traffic flow. For example, network traffic must flow from the internal network to the Internet; a security policy is created (internal interface -> external interface) that allows packets to flow freely from the Internet to the internal network, and from the internal network to the Internet.

Policy instructions may include network address translation (NAT), or port address translation (PAT), or by using virtual IPs or IP pools to translate source and destination IP addresses and port numbers.

Policy instructions may also include UTM profiles, which can specify application-layer inspection and other protocol-specific protection and logging, as well as IPS inspection at the transport layer.

You configure security policies to define which sessions will match the policy and what actions the FortiGate unit will perform with packets from matching sessions.

Sessions are matched to a security policy by considering these features of both the packet and policy:

- Source Interface/Zone
- Source Address
- Destination Interface/Zone
- Destination Address

- Schedule and time of the session's initiation
- Service and the packet's port numbers.

If the initial packet matches the security policy, the FortiGate unit performs the configured Action and any other configured options on all packets in the session.

Packet handling actions can be *ACCEPT*, *DENY*, *IPSEC* or *SSL-VPN*.

- **ACCEPT** policy actions permit communication sessions, and may optionally include other packet processing instructions, such as requiring authentication to use the policy, or specifying one or more UTM profiles to apply features such as virus scanning to packets in the session. An ACCEPT policy can also apply interface-mode IPsec VPN traffic if either the selected source or destination interface is an IPsec virtual interface.
- **DENY** policy actions block communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped, therefore it is not required to configure a DENY security policy in the last position to block the unauthorized traffic. A DENY security policy is needed when it is required to log the denied traffic, also called "violation traffic".
- **IPSEC** and **SSL-VPN** policy actions apply a tunnel mode IPsec VPN or SSL VPN tunnel, respectively, and may optionally apply NAT and allow traffic for one or both directions. If permitted by the firewall encryption policy, a tunnel may be initiated automatically whenever a packet matching the policy arrives on the specified network interface, destined for the local private network.

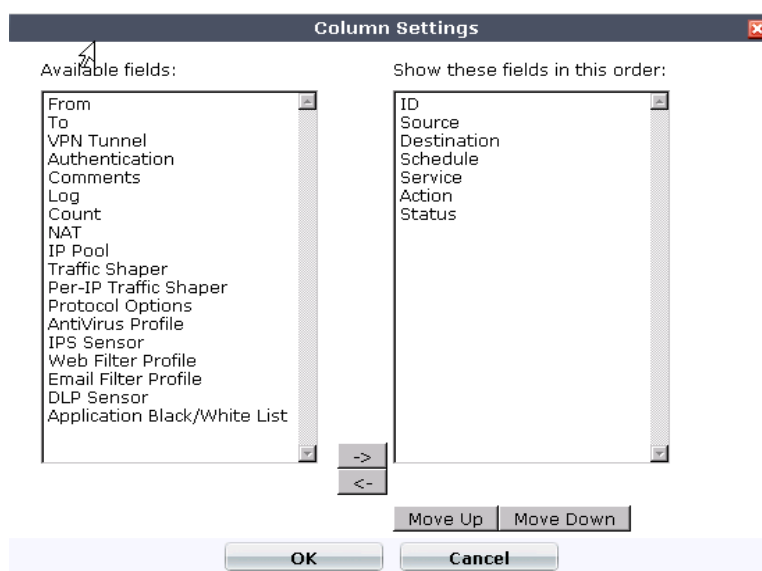
You need to create security policies based on how the network traffic is going to be flowing through the FortiGate unit. For example, a policy for POP3, where the email server is outside of the internal network, traffic should be from an internal interface to an external interface rather than the other way around. It is typically the user on the network requesting email content from the email server and thus the originator of the open connection is on the internal port, not the external one of the email server. This is also important to remember when view log messages as to where the source and destination of the packets can seem backwards.



If you make any changes to existing policies, those changes take effect immediately.

Security policy list details

The security policy table includes, by default, a number of columns to display information about the policy, for example, source, destination, service, and so on. You can add a number of additional columns to the table to view more information about the policies and what is in their configuration. By going to *Policy > Policy > Policy* and selecting the *Column Settings* link, you can add or remove a number of different columns of information to the policy list, and arrange their placement within the table.

Figure 20: Security policy column selection

Viewing security policies

When viewing security policies in the security policy list, you can view them in either *Section View* or *Global View*. In *Section View*, policies are grouped by how the traffic is directed by interface, for example, internal -> wan1. In *Global View*, policies are listed in one large list with no groupings, referred to as interface pairings.

The FortiGate unit will automatically change the view on the policy list page to *Global View* whenever a policy containing *any* in the *Source interface/zone* or *Destination interface/zone* is created. This occurs because the FortiGate unit understands that this particular policy allows or denies traffic on any FortiGate interface, which breaks the original policy sequence order.

Policies are ordered by fixed policies (ones that contain static interfaces) with each interface pairing (for example, port1 -> port2) and each pairing has their own specific policy order, which does not cause any conflict. However, this interface pairing creates a conflict when a policy containing an ANY interface is created, because the FortiGate unit is now unable to determine which policy set to use and which, in the pair's ordering, should traffic be blocked. The FortiGate unit uses Global View to represent its own understanding of the global policy that was created, using this to help determine the action to take.

Policy order

Each time a FortiGate unit receives a connection attempting to pass through one of its interfaces, the unit searches its security policy list for a matching security policy.

The search begins at the top of the policy list and progresses in order towards the bottom. The FortiGate unit evaluates each policy in the security policy list for a match until a match is found. When the FortiGate unit finds the first matching policy, it applies the matching policy's specified actions to the packet, and disregards subsequent security policies. Matching security policies are determined by comparing the security policy and the packet's:

- source and destination interfaces
- source and destination firewall addresses

- services
- time/schedule.

If no policy matches, the connection is dropped.

As a general rule, you should order the security policy list from most specific to most general because of the order in which policies are evaluated for a match, and because only the **first** matching security policy is applied to a connection. Subsequent possible matches are not considered or applied. Ordering policies from most specific to most general prevents policies that match a wide range of traffic from superseding and effectively masking policies that match exceptions.

















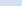
For example, you might have a general policy that allows all connections from the internal network to the Internet, but want to make an exception that blocks FTP. In this case, you would add a policy that denies FTP connections above the general policy.

Figure 21: Example: Blocking FTP — Correct policy order

<input type="checkbox"/>	Status	ID	Source	Destination	Schedule	Service	Action	
Internal -> wan1 (2)								
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	all	all	always	FTP		}Exception }General
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	all	all	always	ANY		
Implicit (1)								
<input type="checkbox"/>	Implicit		all	all	always	ANY		

FTP connections would immediately match the deny policy, blocking the connection. Other kinds of services do not match the FTP policy, and so policy evaluation would continue until reaching the matching general policy. This policy order has the intended effect. But if you reversed the order of the two policies, positioning the general policy before the policy to block FTP, all connections, including FTP, would immediately match the general policy, and the policy to block FTP would never be applied. This policy order would not have the intended effect.

Figure 22: Example: Blocking FTP — Incorrect policy order

	Status	ID	Source	Destination	Schedule	Service	Action	
internal -> wan1 (2)								
		1	 all	 all	 always	 ANY		}General }Exception
		2	 all	 all	 always	 FTP		
Implicit (1)								
	Implicit		all	all	always	ANY		

Similarly, if specific traffic requires authentication, IPsec VPN, or SSL VPN, you would position those policies above other potential matches in the policy list. Otherwise, the other matching policies would always take precedence, and the required authentication, IPsec VPN, or SSL VPN might never occur.



A default security policy may exist, which accepts all connections. You can move, disable or delete it. If you move the default policy to the bottom of the security policy list and no other policy matches the packet, the connection will be accepted. If you disable or delete the default policy and no other policy matches the packet, the connection will be dropped.

You can arrange the security policy list to influence the order in which policies are evaluated for matches with incoming traffic. When more than one policy has been defined for the same interface pair, the first matching security policy will be applied to the traffic session.

How to arrange policies

In this example, there are four policies that the FortiGate unit must use when packets enter the FortiGate unit's interface. These policies are IPsec VPN, DENY, Internet access, and an identity-based policy. You need to make sure that the policies are arranged so that the policies that are important do not get left out.

1 On the policy list, select *Global* view to view all policies in the list.

By viewing the list using Global view, you can easily see all policies in the list regardless of the sections that they are in. This helps you to see where in the list you need to move the policies, without having to expand each section to view the policies.

2 Move the IPsec VPN policy to the first line in the table.

You want the IPsec VPN policy to come first so that the process matches this policy first. If the IPsec VPN policy is not first, other policies would always take precedence and the authentication required for IPsec may never occur.

3 Move the DENY policy to the third line of the table.

This DENY policy contains information that denies all FTP traffic.

4 Move the identity-based policy to the fourth line in the table.

5 Move the Internet access policy after the identity-based policy.

Security policies

There are many different security policies that you can configure for the FortiGate firewall. These policies include SSL VPN, wireless, and identity-based policies. With different configurations come different security policies, and each contain different information for processing the packets coming into the FortiGate unit.

The following explain each type of security policy that can be configured and the reason for configuring such a security policy.

This topic contains the following:

- [Identity-based policies](#)
- [SSL VPN policies](#)
- [IPsec policies](#)
- [Accept policies](#)
- [Deny policies](#)
- [IPv6 policies](#)
- [Security policy 0](#)
- [Local-in policies](#)



If you make any changes to existing policies, those changes take effect immediately.

Identity-based policies

If you enable *Enable Identity Based Policy* in a security policy, network users must send traffic involving a supported firewall authentication protocol to trigger the firewall authentication challenge, and successfully authenticate, before the FortiGate unit will allow any other traffic matching the security policy.

User authentication can occur through any of the following supported protocols:

- HTTP
- HTTPS
- FTP
- Telnet

Authentication can also occur through automatic login using NTLM and FSSO receiverships, to bypass user intervention.

The authentication style depends on which of these supported protocols you have included in the selected firewall services group and which of those enabled protocols the network user applies to trigger the authentication challenge. The authentication style will be one of two types. For certificate-based (HTTPS or HTTP redirected to HTTPS only) authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users, which the FortiGate unit matches. For user name and password-based (HTTP, FTP, and Telnet) authentication, the FortiGate unit prompts network users to input their firewall user name and password.

For example, if you want to require HTTPS certificate-based authentication before allowing SMTP and POP3 traffic, you must select a firewall service (in the security policy) that includes SMTP, POP3 and HTTPS services. Prior to using either POP3 or SMTP, the network user would send traffic using the HTTPS service, which the FortiGate unit would use to verify the network user's certificate; upon successful certificate-based authentication, the network user would then be able to access his or her email.

In most cases, you should ensure that users can use DNS through the FortiGate unit without authentication. If DNS is not available, users will not be able to use a domain name when using a supported authentication protocol to trigger the FortiGate unit's authentication challenge.



If you do not install certificates on the network user's web browser, then network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate, which the network user's web browser may then deem as invalid.



When you use certificate authentication, if you do not specify any certificate when you create a security policy, the FortiGate unit will use the default certificate from the global settings. If you specify a certificate, the per-policy setting will override the global setting.

Authentication requires that *Action* is ACCEPT or SSL-VPN, and that you first create users, assign them to a firewall user group, and assign UTM profiles to that user group.

For additional information about identity-based-policy positioning and identity-based sub-policies, see ["Identity-based security policies" on page 285](#).

Identity-based policy example

With this basic identity-based policy example, the security policy will allow HTTPS traffic passing from the external interface (WAN1) to the internal interface (Internal) at all times, as soon as the network user enters their user name and password. For simplicity, the policy will request the firewall authentication. This authentication can be set up for users by going to *User > User > User* and their groupings by going to *User > User Group > User Group*. For this example, the group “accounting” is used. When a user attempts to browse to a secure site, they will be prompted for their log in credentials.

To create a identity-based policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following:
- 3 Select *Enable Identity Based Policy*.
- 4 *Firewall* authentication is enabled by default.
- 5 Select *Add*.
- 6 From the *Available User Groups* list, select the *Accounting* user group and select the right arrow to move it to the *Selected User Groups* area.
- 7 From the *Available Services* list, select the *HTTPS* and select the right arrow to move it to the *Selected Services* area.
- 8 For the *Schedule*, select *Always*.
- 9 Select *OK*.

To create a identity-based policy - CLI

```
config firewall policy
  edit 1
    set srcintf internal
    set srcaddr 10.13.20.22
    set dstintf wan1
    set dstaddr 172.20.120.141
    set action accept
    set schedule always
    set identity-based enable
    config identity-based-policy
      edit 1
        set group accounting
        set service HTTPS
        set schedule always
      end
    end
  end
```

SSL VPN policies

SSL VPN security policies are created for permitting SSL VPN clients, web-mode or tunnel-mode, access to the protected network behind the FortiGate unit. These security policies also contain authentication information that will authenticate the users and user group or groups.

IPsec policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.



For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

Accept policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

For information about how to configure accept policies, see [“Security policy list details” on page 218](#).

Deny policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.

For information about how to configure DENY policies, see [“Security policy list details” on page 218](#).

How to allow DNS queries to only one DNS server

In this example, a specific DNS server is used for all DNS queries. All other requests for DNS is not allowed. A deny security policy is used to restrict this access.

1 In *Firewall Objects > Address > Address*, create an IP address for the DNS server.

This address will be used for the policy that allows DNS requests from this DNS server.

2 Create a new security policy that blocks all DNS sessions to the Internet.

This policy would have the Action set to DENY and the Service set to DNS. In this policy, the FortiGate unit restricts all requests for any DNS queries.

3 Create a new policy that allows access to only the DNS server.

This policy is used by the FortiGate unit to allow DNS requests to the DNS server that is specified.

4 Move the policies so that they are in the correct order.

If the policies are not in the correct order, the FortiGate unit will not process the instructions properly and the policies will not work properly. The allowed policy needs to be first and the deny policy needs to come right after.

5 Test the policies.

You can test the policies by using diagnose debug command in the CLI or view the packet count in the Count columns of the policies. For more information about how to test and/or verify if traffic is hitting a policy, see [“How to create a basic security policy for Internet access” on page 227](#).

IPv6 policies

IPv6 security policies are created both for an IPv6 network, and a transitional network. A transitional network is a network that is transitioning over to IPv6, but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature in *System > Admin > Settings*.

For more information about IPv6 in FortiOS, see [“Internet Protocol version 6 \(IPv6\)” on page 235](#).

Security policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic
subtype=violation pri=warning vd=root SN=179089 duration=0
user=N/A group=N/A rule=0 policyid=0 proto=17 service=137/udp
app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A
dst_int="Internal" sent=0 rcvd=0 src_port=137 dst_port=137 vpn=N/A
tran_ip=0.0.0.0 tran_port=0
```

Local-in policies

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
edit <policy_number>
set intf <source_interface>
set srcaddr <source_address>
set dstaddr <destination_address>
set action {accept | deny}
set service <service name>
set schedule <schedule_name>
end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12 using SSH on port 3 (192.168.21.77) using the Weekend schedule which defines the time of access.

```
config firewall local-in-policy
edit <1>
set intf port3
set srcaddr 192.168.21.12
set dstaddr 192.168.21.77
set action accept
set service SSH
set schedule Weekend
end
```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```
config firewall local-in-policy
edit <policy_number>
set status disable
end
```

Use the same commands with a status of `enable` to use the policy again.

Local-in policies are also supported for IPv6 by entering the command `config firewall local-in-policy6`.

Creating a basic security policy

The following describes how to configure a basic security policy as well as how to test and verify that traffic hitting the policy.

This topic includes the following:

- [How to create a basic security policy for Internet access](#)
- [How to verify if traffic is hitting the basic security policy](#)
- [How to test the basic security policy](#)

How to create a basic security policy for Internet access

The following explains how a basic security policy is created, as well as how to test and verify that the policy is working properly. Testing a policy and verifying if traffic is hitting a policy are two ways to ensure that the policy that you created is working properly.

1 In the web-based manager, go to *Policy > Policy > Policy* and select *Create New*.

2 The source interface should be *internal* and the destination interface should *wan1*.

This indicates to the FortiGate unit that the incoming packets will be coming from the internal network and proceeding to the public network or Internet. The interfaces are also understood in reverse: packets that are coming from the outside or Internet and are destined for the internal network.

3 The source and destination addresses should *all*.

This is the default IP address range in *Firewall Objects > Addresses > Address*. This default IP address range indicates that any IP address is accepted within the range. This is written as 0.0.0.0/0.0.0.0.

4 For this policy, you must choose the default *always* schedule for *Schedule*, the *ANY* service for *Service*, and the *Action* to *ACCEPT*.

The default schedule always provides the time limitation, which is none, for the policy. A time limitation can limit the access users have to the Internet or can allow users to access resources at any time of the day or night.

5 Select *Log Allowed Traffic* to view the traffic activity using either *Policy > Monitor > Policy Monitor*, or traffic logs. Select *OK* to save the security policy.

You should test the policy after it has been created. To test a security policy, go to a web site; if you are able to get to the web site, the policy is working properly. You can also view the Count column on the Policy page. The Count column displays the number of packets that have recently passed through, which increases as the packets pass through the FortiGate unit.

How to test the basic security policy

After a security policy has been configured, you can test to see if the policy is working. This should be done after you create a security policy so that you can modify the policy's settings, if required, before backing up the configuration. You should always back up the configuration after making modifications to the FortiGate configuration; by doing so, you will have a current configuration whenever you need it.

1 On a computer that is on the internal network, open a web browser and access any web site.

You should be able to get to that web site.

2 If you are unable to get to a web site, use the following to help troubleshoot the problem:

- Is the policy order correct?
- Using the diag debug flow command, see if traffic is hitting the policy. If not, use the diag sniffer command to determine what is going on
- View the Count column; if no number appears, traffic is not hitting the policy.

3 After troubleshooting the problem, browse to a web site and if you can access it, and then save the current configuration.

How to verify if traffic is hitting the basic security policy

After configuring a security policy, you will want to verify that it is working properly. The following explains how to verify that traffic is hitting the basic security policy that you configured in “[How to verify if traffic is hitting the basic security policy](#)” on page 228.

1 In the web-based manager, go to *Policy > Policy > Policy* and locate the internal to wan1 policy.

2 In the Count column, verify that there is packets hitting the security policy.

The Count column displays the amount of packets that are hitting the security policy. In the beginning this count will be low, be increase as the packets come through the FortiGate unit.

3 Go to *Policy > Monitor > Policy Monitor* to view the security policy.

On the Policy Monitor page, you can see the active sessions, bytes or packets that are occurring from the bar chart and table. By selecting the bar within the chart, you can view more detailed information.

4 Go to the CLI, log in, and use `diag debug flow` commands to show traffic is hitting the security policy.

The `diag debug flow` commands show packet flow through the FortiGate unit. The following is an example of what the information gives when you use the `diag debug flow` commands to see if traffic is hitting a policy.

```
diagnose debug enable
diagnose debug flow show console enable
diagnose debug flow filter add 192.168.1.110
diagnose debug flow trace start 50

id=36871 trace_id=1 msg="vd-root received a packet(proto=6,
    192.168.1.110:3152->172.16.100.148:80) from internal."
id=36871 trace_id=1 msg="allocate a new session-0000724b"
id=36871 trace_id=1 msg="find a route: gw-172.20.120.2 via wan1"
id=36871 trace_id=1 msg="find SNAT: IP-172.20.120.11, port-
    40156"
id=36871 trace_id=1 msg="Allowed by Policy-3: SNAT"
id=36871 trace_id=1 msg="SNAT 192.168.1.110-
    >172.20.120.11:40156"
```



Monitoring firewall traffic

You can easily monitor the network traffic on the FortiGate firewall from either the Dashboard or the Monitor menus in the Policy and Firewall Objects menus. By using these monitors, you can understand how to improve your firewall, or resolve issues.

The following explains the various features that you can use to monitor firewall traffic.

The following topics are included in this section:

- [Session tables](#)
- [Monitoring security policy traffic activity](#)

Session tables

Firewall session tables include entries to record source and destination IP addresses and port numbers. For each packet received by a FortiGate unit, it references the session table for a match. Packets of an established session are checked against the session table continually throughout the communication. The performance of depends on the performance of processing session table.

Firewall sessions clear from the table based on the timeout, that is, Time-to-live (TTL) setting. Equally, a completely inactive session with no FIN or RESET will be flushed by the by the session TTL timer. Sessions are not closed based on FIN or a RESET. A FIN that is acknowledged with a FIN ACK would slush the session.

Viewing session tables in the web-based manager

Firewall sessions are viewable in the web-based manager using the *Top Sessions* widget. If this widget is not on the Dashboard, select the *Widget* link at the top of the web-based manager and select it from the pop-up dialog box.

While this view shows a graph of the connecting users and IP addresses, double-clicking on a bar in the graph will display the complete session information for that user.

You can clear a session from the table by scrolling to the right and selecting the delete icon for a given session.

Sessions Monitor

Session information display in *Policy > Monitor > Session Monitor*. You can delete sessions, refresh so that you are viewing current sessions, and you can also filter the session information on the page as well. Filtering allows you to view specific information. For example, you want to view only TCP sessions.

Session Monitor page

Displays the sessions that are currently being monitored by the unit.

Refresh

Select to refresh the information in the list.

Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify the settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filters settings. Select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column. <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p>
IPv4	Select to display only IPv4 addresses.
IPv6	Select to display only IPv6 addresses.
Both	Select to display both IPv4 and IPv6 addresses.
Total Concurrent Sessions: <number>/ New Sessions per Second: <number>	Indicates the total number of concurrent sessions, as well as new sessions that are occurring each second.
Page Controls	Use to navigate through the list.
Total: <number>	The total number of current sessions.
#	The number of the session within the list.
Protocol	The service protocol of the connection, for example, UDP.
Src Address	The source IP address of the connection.
Src Port	The source port of the connection.
Src NAT IP	The source NAT IP address.
Src NAT Port	The source NAT IP port.
Dst Address	The destination address of the connection.
Dst Port	The destination port of the connection.
Policy ID	The security policy identification number.
Expiry (sec)	The time, in seconds, before the connection expires.
Duration (sec)	The duration, in seconds, of the session.
Delete	Select to remove a session from within the list.

Viewing session tables in the CLI

Session tables and information is also viewable from the CLI. More information on sessions are available from the CLI where various diagnose commands reveal more granular data. To view the session information enter the following CLI command:

```
diagnose sys session list
```

Output will look something similar to:

```

session info: proto=17 proto_state=01 duration=121 expire=58
              timeout=0 flags=00000000 sockflag=00000000 sockport=0
              av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0 tunnel=/
state=may_dirty br
statistic(bytes/packets/allow_err): org=63/1/1 reply=133/1/1
              tuples=2
origin->sink: org pre->post, reply pre->post dev=6->2/2->6
              gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 172.20.120.85:51167-
              >8.8.8.8:53(0.0.0.0:0)
hook=post dir=reply act=noop 8.8.8.8:53-
              >172.20.120.85:51167(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0
              vd=0
serial=000171db tos=ff/ff app_list=0 app=0
dd_type=0 dd_rule_id=0
per_ip_bandwidth meter: addr=172.20.120.85, bps=1984
total session 189

```

To clear a session enter the following command:

```
diagnose sys session clear
```

State	Meaning
log	Session is being logged
local	Session is originating from, or destined for, a local stack.
ext	Session is created by a firewall session helper.
may_dirty	Session is created by a policy. For example, the session for FTP channel control will have this state but the FTP data channel will not.
ndr	Session will be checked by an IPS signature.
nds	Session will be checked by an IPS anomaly.
br	Session is being bridged, that is, in transparent mode.
npu	Session will possibly be offloaded to NPU.
wccp	Session is handled by WCCP.

Proto_state fields: TCP

The proto_state field value has two digits. This is because the FortiGate unit keeps track of the original direction and the reply direction.

State	Value	Expire Timer Default (seconds)
NONE	0	10
ESTABLISHED	1	3600
SYN_SENT	2	120

SYN & SYN/ACK	3	60
FIN_WAIT	4	120
TIME_WAIT	5	120
CLOSE	6	10
CLOSE_WAIT	7	120
LAST_ACK	8	30
LISTEN	9	120

Proto_state fields: SCTP

State	Value	Expire Timer Default (seconds)
SCTP_S_NONE	0	60
SCTP_S_ESTABLISHED	1	3600
SCTP_S_CLOSED	2	10
SCTP_S_COOKIE_WAIT	3	5
SCTP_S_COOKIE_ECHOED	4	10
SCTP_S_SHUTDOWN_SENT	5	30
SCTP_S_SHUTDOWN_REC'D	6	30
SCTP_S_ACK_SENT	7	3
SCTP_S_MAX	8	120

Proto_state fields: UDP

UDP is a sessionless protocol, however the FortiGate unit still monitors two different states:

- Reply Not Seen - 0
- Reply Seen - 1

In the example output below, a state of 00, the UDP packet has been seen and a session will be created, but no reply packet has been seen:

```
session info: proto=17 proto_state=00 expire=179 timeout=3600
use=3
```

In this example, the UDP packet has been seen and a session created. Reply packets have also been seen:

```
session info: proto=17 proto_state=01 expire=22 timeout=3600
use=3
```

Proto_state field for ICMP

There are no states for ICMP traffic; it will always appear as `proto_state=00`.

Monitoring security policy traffic activity

The Policy Monitor page provides information about the activity of security policies. This activity can be viewed at a high level, or in much more detail, by drilling down to get more specific information.

The Policy Monitor page allows you to view the information in either a graphical format, or in a table. The graphical format, or chart, provides an easy and user-friendly view of the traffic activity that is occurring. The chart also provides a way to drill-down to more information; you can view this information by selecting on a bar within the chart. The drill-down information can be displayed by source address or destination address or by destination port.

Below the chart, a table provides information as well about each policy include the type of action the policy

Policy Monitor page

Displays information about the security policy traffic occurring on the unit.

Tip: View additional and more detailed information by selecting a bar within the chart.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Policy Usage	Displays the top security policy usage in a bar chart.
Report By	Select to view information by the current active sessions, bytes or packets.

(Table explaining detailed information about the top policy usage)

Policy ID	The security policy identification number.
Source Interface/Zone	The source address or zone used within that security policy.
Destination Interface/Zone	The destination address of zone used within that security policy.
Action	The type of action that is specified in the security policy. For example <i>Action</i> is set to <i>DENY</i> . The action displays as an icon; for example, a green check mark is <i>ALLOW</i> .
Bytes	The number of bytes used by the security policy. This is reflected in the bar chart.
Packets	The number of packets.



Internet Protocol version 6 (IPv6)

This section explains IPv6 in FortiOS. This section does not explain IPv6 in its entirety, only a high-level summary of IPv6 and how IPv6 is supported in FortiOS. For any additional information about IPv6, see the ipv6.com web site.

The following topics are included in this section:

- [What is IPv6?](#)
- [IPv6 in FortiOS](#)
- [Dual stack routing configuration](#)
- [IPv4 tunneling configuration](#)
- [Remotely connecting to an IPv6 network over the Internet](#)
- [IPv6 overview](#)
- [Transition from IPv4 to IPv6](#)
- [Configuring FortiOS to connect to an IPv6 tunnel provider](#)
- [FortiGate IPv6 configuration](#)
- [IPv6 troubleshooting](#)
- [FortiGate IPv6 configuration](#)
- [IPv6 troubleshooting](#)
- [Additional IPv6 resources](#)

What is IPv6?

Internet Protocol version 6 (IPv6) is the next-generation version of IP addressing. This updated version of IP addressing provides many advances, such as more routing efficiency and reducing the need for NAT. IPv6 also provides better security and mobility support, as well as stateless auto-reconfiguration of hosts which allows IPv6 hosts to automatically configure when connected to a routed IPv6 network.

IPv6 uses 128-bit addressing, which is written in hexadecimal digits separated by a colon. For example, 2001:DB8::6334. This revised version of IP addressing has the potential to provide trillions and trillions of addresses, or an address for each device on the Internet.



For IPv6 address examples, documents use the IPv6 special address 2001:DB8::/32 to indicate that the address is an example. This is stated in RFC 3849. For more information about the specific addresses that are used in IPv6, see ipv6.com.

IPv6 in FortiOS

By default, the FortiGate unit is not enabled to use IPv6 options and settings; however, they are there. To enable IPv6, go to *System > Admin > Settings* and select *IPv6 Support on GUI*. When enabled, you can use IPv6 addressing on any of the address-dependant components of the FortiGate unit, including security policies, interface addressing and DNS servers. IPv6 addressing can be configured on the web-based manager and in the CLI.

There are many different features that FortiOS supports in IPv6. The following is what FortiOS supports in IPv6:

- Static routing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- DNS
- Network interface addressing
- Routing access lists and prefix lists
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Security policies
- Authentication
- IPv6 over SCTP
- UTM protection
- Packet and network sniffing
- IPsec VPN
- SSL VPN
- UTM protection
- NAT/Route and Transparent mode
- Logging and reporting
- SNMP
- Virtual IPs and groups
- IPv6-specific troubleshooting, such as ping6

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network.

IPv6 works almost the same as IPv4 in FortiOS. The only main difference is the IP addresses, since you are using IPv6 addressing instead of IPv4. There is also no NAT, unless you are configuring a dual stack routing or IPv4 tunnelling configuration.

Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

Dual stack routing configuration

A dual stack routing configuration implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. The FortiOS dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary.

Devices that are on this type of configured network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. The dual stack architecture of FortiOS provides all the features that you need for protecting your network, such as UTM security for the traffic, and routing.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunneling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

IPv4 tunneling configuration

In an IPv4 tunneling configuration, IPv6 packets are encapsulated within IPv4 headers, which carry these IPv6 packets over IPv4 tunnels. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

Remotely connecting to an IPv6 network over the Internet

Similar to the IPv4 tunneling configuration, FortiOS supports IPv6 tunneling over IPv4 across the Internet between two IPv6 networks that are protected by FortiGate units.

All traffic between the IPv6 networks are tunnelled over IPv4, which in this case is the Internet. Each FortiGate unit extracts the IPv6 traffic from the IPv4 tunnel and traffic on the internal networks uses IPv6.

In FortiOS, you configure this type of network configuration using IPsec VPN because IPv6 is supported for IPsec VPNs. The VPN provides higher security for the data transmitted between the IPv6 networks. This configuration includes an interface-based IPsec VPN between IPv6 interfaces on each FortiGate unit.

IPv6 overview

IP version 6 handles issues that weren't around decades ago when IPv4 was created such as running out of IP addresses, fair distributing of IP addresses, built-in quality of service (QoS) features, better multimedia support, and improved handling of fragmentation. A bigger address space, bigger default packet size, and more optional header extensions provide these features with flexibility to customize them to any needs.

IPv6 has 128-bit addresses compared to IPv4's 32-bit addresses, effectively eliminating address exhaustion. This new very large address space will likely reduce the need for network address translation (NAT) since IPv6 provides more than a billion IP addresses for each person on Earth. All hardware and software network components must support this new address size, an upgrade that may take a while to complete and will force IPv6 and IPv4 to work side-by-side during the transition period. During that time FortiOS supports IPv4 and IPv6 will ensure a smooth transition for networks.

Differences between IPv4 and IPv6

Table 6: IPv4 and IPv6 differences

Property	IPv4	IPv6
Address size	32 bits	128 bits
Network size	8 - 30 bits	64 bits
Packet header size	20 - 60 bytes	40 bytes
Header-level extension	Limited number of small IP options.	Unlimited number of IPv6 extension headers.

Table 6: IPv4 and IPv6 differences

Property	IPv4	IPv6
Fragmentation	Sender or any intermediate router allowed to fragment.	Only sender may fragment.
Control Protocols	Mixture of non-IP (ARP), ICMP and other protocols.	All control protocols based on ICMPv6.
Minimum MTU	567 bytes	1280 bytes
Address assignment	one address per host	multiple addresses per interface.
Address types	Use of unicast, multicast and broadcast address types.	Broadcast addressing no longer used, use of unicast, multicast and anycast address types
Address configuration	Devices configured manually or with host configuration protocols such as DHCP.	Devices configure themselves independently using stateless auto configuration or use DHCP.

IPv6 addresses are assigned to interfaces rather than nodes, thereby recognizing that a node can have more than one interface, and you can assign more than one IPv6 address to an interface. In addition, the larger address space in IPv6 addresses allows flexibility in allocating addresses and routing traffic, and simplifies some aspects of address assignment and renumbering when changing Internet Service Providers.

With IPv4, complex Classless Inter-Domain Routing (CIDR) techniques were developed to make the best use of the small address space. CIDR facilitates routing by allowing blocks of addresses to be grouped together into a single routing table entry. With IPv4, renumbering an existing network for a new connectivity provider with different routing prefixes is a major effort (see [RFC 2071, Network Renumbering Overview: Why would I want it and what is it anyway?](#) and [RFC 2072, Router Renumbering Guide](#)). With IPv6, however, it is possible to renumber an entire network ad hoc by changing the prefix in a few routers, as the host identifiers are decoupled from the subnet identifiers and the network provider's routing prefix.

The size of each subnet in IPv6 is 2^{64} addresses (64 bits), which is the square of the size of the entire IPv4 Internet. The actual address space utilized by IPv6 applications will most likely be small in IPv6, but both network management and routing will be more efficient.

IPv6 MTU

Maximum Transmission Unit (MTU) refers to the size (in bytes) of the largest packet or frame that a given layer of a communications protocol can pass onwards. A higher MTU brings higher bandwidth efficiency. IPv6 requires an MTU of at least 1280 bytes. With encapsulations (for example, tunneling), an MTU of 1500 or more is recommended.

IPv6 address format

The IPv6 address is 128 bits long and consists of eight, 16-bit fields. Each field is separated by a colon and must contain a hexadecimal number. In [Figure 23](#), an X represents each field.

The IPv6 address is made up of two logical parts:

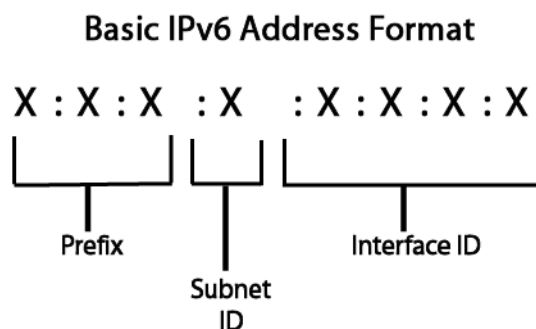
- 64-bit (sub)network prefix

- 64-bit host

The (sub)network prefix part contains the site prefix (first three fields, 48 bits) and the subnet ID (next two fields, 16-bits), for a total of 64-bits. The information contained in these fields is used for routing IPv6 packets. The (sub)network prefix defines the site topology to a router by specifying the specific link to which the subnet has been assigned. The site prefix details the public topology allocated (usually by an Internet Service Provider, ISP) to your site. The subnet ID details the private topology (or site topology) to a router that you assign to your site when you configure your IPv6 network.

The host part consists of the interface ID (or token) which is 64-bits in length and must be unique within the subnet. The length of the interface ID allows for the mapping of existing 48-bit MAC addresses currently used by many local area network (LAN) technologies such as Ethernet, and the mapping of 64-bit MAC addresses of IEEE 1394 (FireWire) and other future LAN technologies. The host is either configured automatically from the MAC address of the interface, or is manually configured.

Figure 23: IPv6 Address Format



IP address notation

IPv6 addresses are normally written as eight groups of four hexadecimal digits each, separated by a colon, for example:

2001:db8:3c4d:0d82:1725:6a2f:0370:6234

is a valid IPv6 address.

There are several ways to shorten the presentation of an IPv6 address. Most IPv6 addresses do not occupy all of the possible 128 bits. This results in fields that are “padded” with zeros or contain only zeros. If a 4-digit group is 0000, it may be replaced with two colons (::), for example:

2001:db8:3c4d:0000:1725:6a2f:0370:6234

is the same IPv6 address as:

2001:db8:3c4d::1725:6a2f:0370:6234

Leading zeroes in a group may be omitted, for example (in the address above):

2001:db8:3c4d::1725:6a2f:370:6234

The double colon (::) must only be used once in an IP address, as multiple occurrences lead to ambiguity in the address translation.

The following examples of shortened IP address presentations all resolve to the same address.

19a4:0478:0000:0000:0000:0000:1a57:ac9e
 19a4:0478:0000:0000:0000::1a57:ac9e
 19a4:478:0:0:0:0:1a57:ac9e

```
19a4:478:0:0::1a57:ac9e
19a4:478::0:0:1a57:ac9e
19a4:478::1a57:ac9e
```

All of these address presentations are valid and represent the same address.

For IPv4-compatible or IPv4-mapped IPv6 addresses (see [“Address types” on page 240](#)), you can enter the IPv4 portion using either hexadecimal or dotted decimal, but the FortiGate CLI always shows the IPv4 portion in dotted decimal format. For all other IPv6 addresses, the CLI accepts and displays only hexadecimal.

Netmasks

As with IP addresses, hexadecimal notation replaces the dotted decimal notation of IPv4. IPv4 Classless Inter-Domain Routing (CIDR) notation can also be used. This notation appends a slash (“/”) to the IP address, followed by the number of bits in the network portion of the address.

Table 7: IPv6 address notation

IP Address	3ffe:ffff:1011:f101:0210:a4ff:fee3:9566
Netmask	ffff:ffff:ffff:ffff:0000:0000:0000:0000
Network	3ffe:ffff:1011:f101:0000:0000:0000:0000
CIDR IP/Netmask	3ffe:ffff:1011:f101:0210:a4ff:fee3:9566/64

Address scopes

Address scopes define the region where an address may be defined as a unique identifier of an interface. The regions are: local link (link-local), site network (site-local), and global network. Each IPv6 address can only belong to one zone that corresponds to its scope.

Address types

IPv6 addresses are classified into three groups - [Unicast](#), [Multicast](#), and [Anycast](#).

Unicast

Identifies an interface of an individual node. Packets sent to a unicast address are sent to that specific interface. Unicast IPv6 addresses can have a scope reflected in more specific address names - global unicast address, link-local address, and unique local unicast address.

Multicast

Multicast addresses are assigned to a group of interfaces that typically belong to different nodes. A packet that is sent to a multicast address is delivered to all interfaces identified by the address.

IPv6 multicast addresses are distinguished from unicast addresses by the value of the high-order octet of the addresses. A value of 0xFF (binary 11111111) identifies an address as a multicast address. Any other value identifies an address as a unicast address.

The four least significant bits of the second address octet identify the address scope or the span over which the multicast address is propagated.

Anycast

Anycast addresses are assigned to a group of interfaces usually belonging to different nodes. A packet sent to an anycast address is delivered to just one of the member interfaces, typically the 'nearest' according to the router protocols' choice of distance. They cannot be identified easily as their structure is the same as a normal unicast address, differ only by being injected into the routing protocol at multiple points in the network. When a unicast address is assigned to more than one interface (making it an anycast address), the address assigned to the nodes must be configured in such a way as to indicate that it is an anycast address.

Interfaces configured for IPv6 must have at least one link-local unicast address and additional ones for site-local or global addressing. Link-local addresses are often used in network address autoconfiguration where no external source of network addressing information is available.

Special addresses

Special IPv6 addresses include unspecified and loopback addresses. For more information about IPv6 addresses, see [RFC 4921, IP Version 6 Addressing Architecture](#)

The IPv6 address space is split into scopes, or address scopes. The table below indicates which IPv6 address is used.

Table 8: IPv6 addresses with prefix information

Address Type	Binary Prefix	IPv6 Notation	Uses
Embedded IPv4 address	00...1111 1111 1111 1111 (96 bits)	::FFF/96	Prefix for embedding IPv4 address in an IPv6 address.
Loopback	00...1 (128 bits)	::1/128	Used as a node to send an IPv6 packet to itself. Seen as link-local unicast address of a virtual interface (loopback interface) to an imaginary link that goes nowhere. Must never be assigned to a physical interface, or as the source address of IPv6 packets that are sent outside of the single node. IPv6 destination address of loopback should not be sent outside a single node, and never forwarded by an IPv6 router. Equivalent to 127.0.0.1 in IPv4. RFC 246022
Global unicast	001	2000::3	Global unicast and anycast. RFC 429120
Global unicast	01 - 1111 1000 0	4000::/2 - FC00::/9	Global unicast and anycast (unallocated)

Table 8: IPv6 addresses with prefix information

Address Type	Binary Prefix	IPv6 Notation	Uses
Teredo	0010 0000 000 0001 0000 0000 000 0000	2001::/32	Teredo - RFC 438023
Nonroutable	0010 0000 0000 0001 1101 1000 1000 0000	2001:D88::/32	Nonroutable. Documentation purposes only - RFC 384924
6to4	0010 0000 0000 0010	2002::/16	Used for communication between two nodes running both IPv4 and IPv6 over the Internet. Formed by combining the IPv6 prefix with the 32-bits of the public IPv4 address of the node, creating a 48-bit address prefix. - RFC 3056
6Bone	0011 1111 1111 1110	3FFE::/16	Deprecated. 6Bone testing assignment 1996 to mid-2006 RFC 370125
Local-link unicast	1111 1110 10	FE80::/10	Used for addressing on a single link for automatic address configuration, neighbor discovery, or when no routers are present. Routers must not forward packets with link-local source or destination addresses.
Reserved	1111 1110 11	FEC0::/10	Used for addressing inside of a site without needing a global prefix. Routers must not forward packets with site-local source or destination addresses outside of the site. RFC 387926
Local IPv6 address	1111 110	FC00::/7	Unicast unique local address space, unicast and anycast - RFC 419327
Multicast	1111 1111	FF00::/8	Multicast address space - RFC 4291 For more information, see “Multicast” on page 240 .

Header Extension

The base header of an IPv6 address is fixed for efficient processing. Header extensions are indicated by the next header value in the next header field. Header extensions are optional and do not need to be present in all IPv6 packets. The sequence for the next header in order is represented by the diagram below.

IPv6 Header	Hop-by-Hop Options Header	Destination Options Header Router	Routing Header	Fragment Header	Authentication on Header	Encapsulation Security Payload	Destination Options Header Destination	Mobility Header (MIPv6)	TCP/UDP/SCTP	Payload
-------------	---------------------------	-----------------------------------	----------------	-----------------	--------------------------	--------------------------------	--	-------------------------	--------------	---------

The last header extension is the value of either 6 for TCP, 17 for UDP, 132 for SCTP or any other transport protocol defined by the IETF.

Header extensions appear in the following sequence:

- Hop-by-Hop Options Header
 - First Extension Header
 - Next Header value of 0 indicates the Hop-by-Hop Options Extension Header
 - All nodes along the route or path must process this extension header
- Routing Header
 - Second Extension Header
 - Next Header value of 43 indicates the Routing Extension Header
 - All nodes along the route or path must process this extension header
 - Note that the Routing Header Type 0 is due to security reason depreciated
- Fragmentation Header
 - Third Extension Header
 - Next Header value of 44 indicates the Fragmentation Extension Header
 - Used in case of transmitting payload longer than a IPv6 packet can carry
- Authentication Header (AH)
 - Fourth Extension Header
 - Next Header value of 50 indicates the Authentication Extension Header
 - Used to provide protection against replay, origin authentication and connectionless integrity
- Encapsulating Security Payload (ESP) Header
 - Fifth Extension Header
 - Next Header value of 51 indicates the ESP Extension Header
 - Used to provide protection against replay, origin authentication and connectionless integrity
- Destination Options Header
 - Sixth Extension Header
 - Next Header value of 60 indicates the Destination Options Header
 - Used to provide additional information for the end systems node
- Mobility Header
 - Seventh Extension Header
 - Next Header value of 135 indicates the Mobility Header
 - Used by mobile nodes to exchange information for Mobile IP nodes (MIPv6)

Table 9: Header and Protocol Types

Extension Header	Type
------------------	------

Table 9: Header and Protocol Types

Hop-by hop Options	0
Routing	43
Fragment	44
Destination Options	60
Authentication Header (AH)	50
Encapsulating Security Payload	51
Mobility	135
Protocol	Type
TCP	6
UDP	17
IPv6-in-IPv6	41
GRE	47
ICMPv6	58
No next header	59
OSPF	89
PIM	103
SCTP	132

IPv6 neighbor discovery

IPv6 Neighbor Discovery (ND) is a set of messages and processes that determine relationships between neighboring nodes. Neighboring nodes are on the same link. The IPv6 ND protocol replaces the IPv4 protocols Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMPv4), Router Discovery (RDISC), and ICMP Redirect, and provides additional functionality. The IPv6 ND protocol facilitates the autoconfiguration of IPv6 addresses. Autoconfiguration is the ability of an IPv6 host to automatically generate its own IPv6 address, making address administration easier and less time-consuming.

Hosts use ND to:

- discover addresses, address prefixes, and other configuration parameters
- discover neighboring routers.

Routers use ND to:

- advertise their presence, host configuration parameters, and on-link prefixes
- inform hosts of 'better' next-hop address to forward packets for a specified destination.

Nodes use ND to:

- resolve link-layer address of a neighboring node to which an IPv6 packet is being forwarded and determine whether the link-layer address of a neighboring node has altered
- determine whether IPv6 packets can be sent to and received from a neighbor
- automatically configure IPv6 addresses for its interfaces.

To facilitate neighbor discovery, routers periodically send messages advertising their availability. This communication includes lists of the address prefixes for destinations available on each router's interfaces.

ND defines five different Internet Control Message Protocol (ICMP) packet types: a pair of Neighbor Solicitation and Neighbor Advertisement messages, a pair of Router Solicitation and Router Advertisement messages, and a Redirect message.

A Neighbor Solicitation is sent by a node to determine the link-layer address of a neighbor, or to verify that a neighbor is still reachable via a cached link-layer address. Also used for Duplicate Address Detection (how a node determines that an address it wants to use is not already in use by another node). The Neighbor Advertisement message is a response to a Neighbor Solicitation message. A node may also announce a link-layer address change by sending unsolicited Neighbor Advertisements.

A host may send a Router Solicitation when an interface becomes enabled, requesting routers to generate a Router Advertisement immediately rather than at their next scheduled time.

Routers advertise their presence together with various link and Internet parameters according to a specific schedule or in response to a Router Solicitation message. A Router Advertisement contains prefixes used for on-link determination and/or address configuration, a suggested hop limit value, etc.

The Redirect message is used by routers to inform hosts of a better first-hop for a destination.

For more information, see RFC 2461, *Neighbor Discovery for IP Version 6 (IPv6)*.

Transition from IPv4 to IPv6

If the Internet is to take full advantage of the benefits of IPv6, there must be a period of transition to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach the IPv6 Internet over the IPv4 infrastructure.

RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers* and RFC 2185, *Routing Aspects of IPv6 Transition* define several mechanisms to ensure that IPv6 hosts and routers maintain interoperability with the existing IPv4 infrastructure, and facilitate a gradual transition that does not impact the functionality of the Internet. The mechanisms, known collectively as Simple Internet Transition (SIT), include:

- dual-stack IP implementations for hosts and routers that must interoperate between IPv4 and IPv6
- embedding of IPv4 addresses in IPv6 addresses. IPv6 hosts are assigned addresses that are interoperable with IPv4, and IPv4 host addresses are mapped to IPv6
- IPv6-over-IPv4 tunneling mechanisms to encapsulate IPv6 packets within IPv4 headers to carry them over IPv4 infrastructure
- IPv4/IPv6 header translation, used when implementation of IPv6 is well-advanced and few IPv4 systems remain.

FortiGate units are dual IP layer IPv6/IPv4 nodes and they support IPv6 over IPv4 tunneling. For more information, see [RFC 2893, Transition Mechanisms for IPv6 Hosts and Routers](#) and [RFC 2185, Routing Aspects of IPv6 Transition](#).


Configuring FortiOS to connect to an IPv6 tunnel provider

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses on the Internet. FortiOS supports IPv6 tunnelling over service provider IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their IPv6 destination. The internal network is running IPv6. The FortiGate unit creates an IPv6-over-IPv4 tunnel to the IPv6 tunnel broker. From the tunnel broker, your network can access IPv6 addresses on the Internet.

In this example the internal network is small and directly connected to the FortiGate unit. There is no need for routing on the internal network since everything is connected and on the same subnet. For this example, consider the following:

- Before configuring your FortiGate unit for IPv6-over-IPv4 tunneling, you need to choose an IPv6 tunnel broker and get their information.
- The addresses used in this example are for example use only.
- VDOMs are not enabled.
- The tunnel broker IPv4 address is 78.35.24.124.
- The tunnel broker IPv6 end of the tunnel is 2001:4dd0:ff00:15e::1/64.
- The FortiGate unit external IPv4 address is 172.20.120.17.
- The FortiGate unit IPv6 address of the tunnel is 2001:4dd0:ff00:15e::2/64.
- port1 of the FortiGate unit is connected to the internal network.
- port2 of the FortiGate unit is connected to the external network (Internet).

Figure 24: Connecting to an IPv6 tunnel broker



Steps to connect to an IPv6 tunnel broker

- 1 Create a SIT-Tunnel Interface.
- 2 Create a static IPv6 Route into the Tunnel-Interface.

- 3 Assign your IPv6 Network to your FortiGate.
- 4 Create a Firewall-Policy to allow Traffic from LAN to the Tunnel-Interface.

Create a SIT-tunnel interface

Creating the SIT-tunnel creates a virtual interface in the form of a tunnel, much like a VPN interface. The end points of the tunnel are the FortiGate unit and the tunnel broker's server addresses.

In the example, the external address of the FortiGate unit is DHCP-based and may change to any value on that subnet, so the source address allows for that.

```
config system sit-tunnel
  edit HE_ip6_broker
    set destination 78.35.24.124
    set interface port2
    set ip6 2001:4dd0:ff00:15e::2/64
    set source 172.20.120.0
  next
end
```

Now that the tunnel exists, some additional interface commands are required. Such as enabling ping6 for troubleshooting and allow HTTPS and SSH administration connections to the interface.

```
config system interface
  edit HE_ip6_broker
    config ipv6
      set ip6-allowaccess ping https ssh
    end
  next
end
```

Create a static IPv6 route into the tunnel-Interface

With the tunnel up and the security policies in place, all that remains is to add a default route for IPv6 traffic to go over the tunnel. As there will only be one static routing entry, there is no need for a priority. This may change in the future if other routes are added.

```
config router static6
  edit 1
    set device HE_ip6_broker
  next
end
```

Assign your IPv6 network to your FortiGate

This step assigns an IPv6 address to the internal interface on the FortiGate unit. That way all IPv6 traffic entering on this interface will be routed to the tunnel. Systems with addresses within this prefix are reachable on the subnet in question without help from a router, so the `onlink-flag` is enabled. Hosts can create an address for themselves by combining this prefix with an interface identifier, so the `autonomous-flag` is enabled.

```
config system interface
  edit port1
    config ipv6
      set ip6-address 2001:4dd0:ff42:72::1/64
      set ip6-allowaccess ping https ssh
      config ip6-prefix-list
        edit 2001:4dd0:ff42:72::/64
```

```

        set autonomous-flag enable
        set onlink-flag enable
        set preferred-life-time 3600
        set ip6-send-adv enable
    next
end
next
end

```

At this point any PCs on your internal network that are set to auto-configure, should have their addresses. To test this you can ping6 from the PC to the FortiGate unit. See [“IPv6 ping description” on page 260](#).

Create a security policy to allow traffic from port1 to the tunnel interface

With the tunnel configured, it will appear as an interface in the Network interface list. That means the next step is to add a security policies to allow traffic to and from the tunnel.

```

config firewall policy6
edit 2
    set srcintf port1
    set dstintf HE_ip6_broker
    set srcaddr "::/0"
    set dstaddr "::/0"
    set action accept
    set schedule "always"
    set service "ANY"
    set logtraffic enable
next
end

```

Test the connection

To test the tunnel, try to connect to an external IPv6 address such as <http://ipv6.google.com>.

If you want to see the path the IPv6 traffic takes, do a traceroute from a PC on the internal network to an external address. You will see the traffic enter the FortiGate unit, enter the tunnel, pass through the tunnel broker server, and on out over the Internet.

If you are entering an IPv6 address into your web browser, you have to type: [https://\[2001:4dd0:ff42:72::1\]](https://[2001:4dd0:ff42:72::1]). The square brackets are to discriminate between the address part and a port, like in [https://\[2001:4dd0:ff42:72::1\]:8080](https://[2001:4dd0:ff42:72::1]:8080)

FortiGate IPv6 configuration

FortiOS (4.0 MR2) supports the following FortiOS IPv6 features (all configurable from the web-based manager or CLI):

- Static routing and dynamic routing
- Network interface addressing
- DHCP Server (CLI only)
- Routing access lists and prefix lists
- IPv6 tunnel over IPv4, IPv4 tunnel over IPv6
- Security policies and identity-based security policies

- Local-in security policies
- IPv6 over SCTP
- Packet and network sniffing
- IPsec VPNs
- UTM protection including
- NAT/Route and Transparent mode
- Logging and reporting
- IPv6 specific troubleshooting such as ping6

Displaying IPv6 options on the web-based manager

Before configuring IPv6 using the web-based manager, you must first turn on IPv6 display by going to *System > Admin > Settings* and selecting the IPv6 support option. Once turned on, IPv6-related options and pages appear throughout the web-based manager. For example, you can add IPv6 addresses to any FortiGate interface, you can add IPv6 DNS server IP addresses, IPv6 security policies, IPv6 firewall addresses and so on.

UTM protection for IPv6 networks

FortiOS uses IPv6 security policies to provide UTM protection for IPv6 traffic. Antivirus, web filtering, FortiGuard Web Filtering, email filtering, FortiGuard Email Filtering, data leak prevention (DLP), and VoIP protection features can be enabled in IPv6 security policies using normal FortiOS UTM profiles for each UTM feature.

Configuring IPv6 interfaces

The dual stack architecture is most obvious when configuring IPv6 on interfaces on your FortiGate unit.

IPv6 interfaces - web-based manager

In the *Addressing mode* section of the *Create New* or *Edit* screen, there are two fields instead of one. Without IPv6 enabled, there is only the *IP/Netmask* field for IPv4 addresses. With IPv6 enabled, there is an additional field called *IPv6 Address*.

With both addresses configured for an interface, that interface will accept both IPv4 and IPv6 traffic. Each protocol will be handled differently, depending on the security policies and routing in place for it. This allows traffic from IPv6 to be sent to other IPv6 devices, and IPv4 traffic to be sent only to other IPv4 devices. This separation of the traffic is required because if IPv6 traffic is sent to devices that don't support it, that traffic will not reach its destination.

You should enable IPv6 Administrative Access to connect to the IPv6 address of an interface for administration.

IPv6 interfaces - CLI

In the CLI, there are a number of IPv6 specific interface settings. These are found as part of the `config system interface` command under `config ipv6`. In the CLI there are many more settings available, although many are optional. The settings that are required or recommended are highlighted.

```
config system interface
edit <interface_string>
config ipv6
```

```

set ip6-address <ipv6_addr>
set ip6-allowaccess <http https ping ssh telnet>
set ip6-link-mtu <bytes_int>
set ip6-send-adv <enable | disable>
set autoconf <enable | disable>
set ip6-default-life <seconds_int>
set ip6-hop-limit <count_int>
set ip6-manage-flag <enable | disable>
set ip6-max-interval <integer>
set ip6-min-interval <integer>
set ip6-other-flag <enable | disable>
set ip6-reachable-time <integer>
set ip6-retrans-time <integer>
  config ip6-extra-addr
    edit <ipv6_addr>
  end
  config ip6-prefix-list
    set autonomous-flag <enable | disable>
    set onlink-flag <enable | disable>
    set preferred-life-time <integer>
    set valid-life-time <integer>
  end
end
end
end

```

Configuring IPv6 routing

IPv6 routing is supported in both static and dynamic routing. The main difference from a configuration point of view is in the addresses.

Static routing

Static routing for IPv6 is essentially the same as with IPv4. From a configuration point of view, the only difference is the type of addresses used. When both IPv4 and IPv6 static routes are configured, they are displayed under two separate headings on the static routing page - *Route* and *IPv6 Route*. Use the arrows next to each heading to expand or minimize that list of routes.

To configure IPv6 static routes - web-based manager

- 1 Go to *Router > Static > Static Route*.
- 2 Select arrow to expand the *Create New* menu.
- 3 Select *IPv6 Route*.
- 4 Enter *Destination IP/Mask*, *Device*, *Gateway*, *Distance*, and *Priority* as with normal static routing using IPv6 addresses.
- 5 Select *OK*.

To configure IPv6 static routes - CLI

Use the following command to add an IPv6 static route:

```

config router static6
  edit 1
    set dst <ipv6_addr>
    set gateway <ipv6_addr>
    set device <interface>
  end
end

```



```

    set priority <integer>
end

```

Dynamic routing

As with static routing, the dynamic routing protocols all have IPv6 versions. Both IPv4 and IPv6 dynamic routing can be running at the same time due to the dual stack architecture of the FortiGate unit. IPv6 dynamic routing must be configured using CLI commands.

Table 10: Dynamic routing protocols, IPv6 versions, CLI command, and RFCs

Dynamic Routing	IPv6	CLI command	IPv6 RFC
RIP	RIP next generation (RIPng)	config router ripng	RFC 2080
BGP	BGP4+	config router bgp All parts of bgp that include IP addresses have IPv4 and IPv6 versions.	RFC 2545 and RFC 2858
OSPF	OSPFv3	config router ospf6	RFC 2740

Configuring IPv6 security policies

Configuring IPv6 security policies is similar to configuring IPv4 security policies. On the web-based manager go to *Policy > IPv6 Policy*. From the CLI use the command `config firewall policy6`. You must also add IPv6 firewall addresses (*Firewall Objects > Address* or `config firewall address6`) and address groups (*Firewall Objects > Address > Group* or `config firewall addgrp6`).

Under the security policies for IPv6, you can also define SSL-VPN actions and authentication policies.

IPv6 Policy configuration settings

The following are IPv6 security policy configuration settings in *Policy > Policy > IPv6 Policy*.

New Policy page	
Source Interface/Zone	Select the name of the FortiGate network interface, virtual domain (VDM) link, or zone on which IP packets are received. Interfaces and zones are configured on the System Network page. You can also create a web proxy firewall proxy by selecting <i>web-proxy</i> in Source Interface/Zone.
	If you select <i>any</i> as the source interface, the security policy matches all interfaces as source. When you select <i>any</i> as the source interface, that security policy list is displayed only in global view.
	If <i>Action</i> is set to <i>IPSEC</i> , the interface is associated with the local private network.
	If <i>Action</i> is set to <i>SSL-VPN</i> , the interface is associated with connections from remote SSL VPN clients.
Source Address	Select the name of a firewall address to associate with the <i>Source Interface/Zone</i> .

		<p>You can also create firewall addresses by selecting <i>Create New</i> from this list.</p> <p>If you want to associate multiple firewall addresses or address groups with <i>Source Interface/Zone</i>, from <i>Source Address</i>, select <i>Multiple</i>. In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i>.</p>
	Destination Interface/Zone	Select the name of the FortiGate network interface, virtual domain (VDM) link, or zone to which IP packets are forwarded. Interfaces and zones are configured on the System Network page.
		If you select <i>any</i> as the source interface, the security policy matches all interfaces as source. When you select <i>any</i> as the source interface, that security policy list is displayed only in <i>Global View</i> .
	Destination Address	Select the name of a firewall address to associate with <i>Destination Interface/Zone</i> . Only packets whose header contains an IP address matching the selected firewall address will be subject to this security policy.
		You can also create firewall addresses by selecting <i>Create New</i> from this list.
		If you want to associate multiple firewall addresses or address groups with the <i>Destination Interface/Zone</i> , from <i>Destination Address</i> , select <i>Multiple</i> . In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i> .
		If you select a virtual IP, the unit applies NAT or PAT. The applied translation varies by the settings specified in the virtual IP, and whether you select NAT (below).
	Schedule	<p>Select a one-time or recurring schedule or a schedule group that controls when the security policy is in effect.</p> <p>You can also create schedules by selecting <i>Create New</i> from this list.</p>
	Service	<p>Select a firewall service or create a new custom service.</p> <p>If you are creating a web proxy security policy, <i>Web Proxy Service</i> appears and you can choose either a web proxy service or web proxy group.</p>
	Action	Select how you want the firewall to respond when a packet matches the conditions of the security policy.
	Log Allowed Traffic	<p>Select to record security policy traffic activity whenever the security policy processes a connection. These log messages are located in the traffic log.</p> <p>You must also enable traffic log for a logging location and set the logging severity level to <i>Notification</i> or lower using the Log&Report menu.</p> <p>This option is not available for web-proxy security policies.</p>
	Log Violation traffic	<p>Select to record security policy traffic activity whenever the security policy processes a violation. These log messages are located in the traffic log.</p> <p>Appears only when <i>Action</i> is <i>DENY</i>.</p>

Enable web cache	<p>Select to enable web caching for HTTP traffic accepted by the security policy. This option is available only on FortiGate units that support WAN Optimization and web caching. Enabling web caching in a security policy is similar to enabling web caching in a WAN Optimization rule. However, enabling web caching in a security policy means you can also apply UTM options to web cached traffic in a single VDOM.</p> <p>You can use this option to apply web caching for explicit web proxy traffic if the Source Interface/Zone is set to the web-proxy interface.</p> <p>Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites on the FortiGate unit hard disk. Some HTTP content accepted by the security policy may not be cached. See RFC 2616 for information about web caching for HTTP 1.1.</p>
Enable NAT	<p>Available only if <i>Action</i> is set to <i>ACCEPT</i> or <i>SSL-VPN</i>. Enable or disable Network Address Translation (NAT) of the source address and port of packets accepted by the security policy. When <i>NAT</i> is enabled, you can also configure <i>Dynamic IP Pool</i> and <i>Fixed Port</i>.</p> <p>If you select a virtual IP as the <i>Destination Address</i>, but do not select the <i>NAT</i> option, the unit performs destination NAT (DNAT) rather than full NAT. Source NAT (SNAT) is not performed.</p>
Use Destination Interface Address	Select to use the destination interface address. If <i>Central NAT Table</i> is enabled, you can choose between this option and using the central NAT table.
Use Central NAT Table	Select to enabling logging using the Central NAT table that you configured in the Central NAT Table menu.
Use Dynamic IP Pool	<p>Available only when <i>Enable NAT</i> is selected.</p> <p>Select the check box, then select an IP pool to translate the source address to an IP address randomly selected from addresses in the IP Pool.</p> <p><i>IP Pool</i> cannot be selected if the destination interface, VLAN subinterface, or one of the interfaces or VLAN subinterfaces in the destination zone is configured using DHCP or PPPoE.</p>
Enable Identity Based Policy	Select to configure security policies that require authentication.
Resolve User Names Using FSSO Agent	Select to resolve user names when using the Fortinet Single Sign-On Agent feature.

Enable Dynamic Profile	<p>Select to configure a dynamic profile security policy. Dynamic profile is a method for users to use a RADIUS server for a single sign-on access to network resources.</p> <p>The <i>Enable Dynamic Profile</i> option does not display by default on the web-based manager; you must first enable it in <i>System > Admin > Settings</i>. If you have VDOMs enabled, you can configure one RADIUS server and security policy for dynamic profile per VDOM. With multiple VDOMs, you can have each one with their own profile group on their own RADIUS server with their own custom level of access.</p> <p>After selecting the check box beside <i>Enable Dynamic Profile</i>, the following options appear below:</p> <ul style="list-style-type: none"> • <i>Profile Group</i> – select a dynamic profile group from the drop-down list • <i>Dynamic Profile Users Only</i> – select to only accept sessions with source addresses that are in the user context list
UTM	Select an UTM option to apply to the security policy. You must enable UTM before you can select the available UTM options. When selecting an option, select a profile from the list, or select <i>Create New</i> from the list to build a profile.
Web Proxy Forwarding Server	Select a web proxy forwarding server from the drop-down list. This appears only when configuring a web proxy security policy.
GTP Profile (FortiOS Carrier only)	Select a GTP profile from the drop-down list. Select <i>Create New</i> to create a new GTP profile. Select <i>View</i> to view the GTP profile.
Traffic Shaping	Select a traffic shaper for the security policy. You can also create a new shared traffic shaper. Shared traffic shapers control the bandwidth available to and set the priority of the traffic as its processed by, the security policy.
Reverse Direction Traffic Shaping	Select to enable reverse traffic shaping and select a shared traffic shaper. For example, if the traffic direction that a security policy controls is from port1 to port2, select this option will also apply the security policy shaping configuration to traffic from port2 to port1.
Dynamic Profile Users Only	Select to configure the security policy to only accept sessions with source addresses that are in the dynamic profile user context list. Sessions with source addresses that are not in the user context list do not match the security policy. For sessions that do not match the security policy, the unit continues searching down the security policy list for a match.
Enable Endpoint Security	<p>Select to enable the Endpoint NAC feature and select the Endpoint NAC profile to apply.</p> <ul style="list-style-type: none"> • You cannot enable Endpoint in security policies if <i>Redirect HTTP Challenge to a Secure Channel (HTTPS)</i> is enabled in <i>User > Options > Authentication</i>. • If the security policy involves a load balancing virtual IP, the Endpoint check is not performed.
Enable Disclaimer	Select to include a disclaimer page. Select <i>Edit</i> to modify the disclaimer replacement message.

Tags	Applies tags to the security policy. Tags can be viewed on the Policy page in the <i>Tags</i> column.
Applied Tags	Displays the tags that you have added to the security policy.
Add Tags	Enter the tag in the field and select the plus (+) sign to add the tag to the security policy. This also adds the tag to the <i>Applied Tags</i> list.
Comments	Add information about the security policy. The maximum length is 63 characters.

Configuring IPv6 DNS

Configuring DNS servers with IPv6 addresses is located in the same location as IPv4, by going to *System > Network > DNS*. There is a separate area for adding IPv6 addresses for DNS. From the CLI, use the command `config system dns`, where additional commands `ip6-primary` and `ip6-secondary` are available.

Configuring IPv6 DHCP

Configuring DHCP servers with IPv6 is performed using the CLI only. While similar to IPv4, there are a few exceptions:

- There is no gateway to define. A host learns the gateway using router advertisement messages
- There is no WINS servers defined for dhcpv6, as it is obsolete.

To configure DHCP use the following command set:

```
config system dhcp6 server
  edit 1
    set domain example.com
    set interface port3
    config ip-range
      edit 1
        set end-ip 2800:68:15:3::10
        set start-ip 2800:68:15:3::1
      end
    set option1 50 'AABB'
    set subnet 2800:68:15:3::/64
    set dns-server1 2800:68:15:3::2
    set dns-server2 2800:68:15:3::29
    set dns-server3 2800:68:15:3::28
    set enable enable
  end
end
```

For more information on the commands, see the [CLI Reference](#).

Configuring IPv6 over IPv4 tunneling

IPv6 over IPv4 tunneling can only be configured in the CLI using the `config system sit-tunnel` command. When you configure an IPv6 over IPv4 tunnel, you are creating a virtual interface that can be used in configurations just like any other virtual interface such as VLANs.

The name of the command `sit-tunnel` comes from Simple Internet Transition (SIT) tunneling. For the period while IPv6 hosts and routers co-exist with IPv4, a number of transition mechanisms are needed to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach the IPv6 Internet over the IPv4 infrastructure.

These techniques, collectively called Simple Internet Transition, include:

- dual-stack IP implementations for interoperating hosts and routers
- embedding IPv4 addresses in IPv6 addresses
- IPv6-over-IPv4 tunneling mechanisms
- IPv4/IPv6 header translation

The syntax for the IPv6 over IPv4 tunneling CLI command is:

```
config system sit-tunnel
  edit <name_string>
    set destination <ipv4_addr>
    set interface <interface_string>
    set ip6 <ipv6_addr>
    set source <ipv4_addr>
  next
end
```

<name_string>	This will be the name of the tunnel, and appear in the network interface list. It should be descriptive such as <code>my_ip6_tunnel</code> . The maximum length allowed is 15 characters.
destination <ipv4_addr>	This is the tunnel broker's IPv4 server address. It is one of the two ends of the tunnel.
interface <interface_string>	This interface is the interface the tunnel piggy backs on. Generally this should be the external interface of the FortiGate unit. This setting is optional if you don't have a fixed IP address from your ISP.
ip6 <ipv6_addr>	The IPv6 address of the tunnel.
source <ipv4_addr>	This is the FortiGate unit end of the tunnel. It is just like any other FortiGate unit interface address. If this address is DHCP-based, it will change. In that case you should ensure the netmask covers the possible range of addresses. It is possible to use 0.0.0.0 to cover all possible addresses if you have a DDNS or PPPoE connection where the address changes.

For more configuration of tunnels see [“Configuring FortiOS to connect to an IPv6 tunnel provider” on page 246](#).

Configuring IPv6 IPsec VPNs

The FortiGate unit supports route-based IPv6 IPsec, but not policy-based.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

IPv4 over IPv6	<p>The VPN gateways have IPv6 addresses.</p> <p>The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.</p>
IPv6 over IPv4	<p>The VPN gateways have IPv4 addresses.</p> <p>The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.</p>

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported. This is because FortiOS 3.0 does not support IPv6 DNS.
- You cannot use RSA certificates in which the common name (cn) is a domain name that resolves to an IPv6 address. This is because FortiOS 3.0 does not support IPv6 DNS.
- DHCP over IPsec is not supported, because FortiOS 3.0 does not support IPv6 DHCP.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies and routing.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6 Support on GUI*.

Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the *Advanced* settings. Enable the *IPv6 Version* check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses.

For example:

```
config vpn ipsec phase1-interface
  edit tunnel6
    set ip-version 6
    set remote-gw6 0:123:4567::1234
    set interface port3
    set proposal 3des-md5
  end
```

Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default 0.0.0.0/0 address for Source address and Destination address to the IPv6 value ::/0. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, ::/0 for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like the following:

```
config vpn ipsec phase2-interface
edit tunnel6_p2
set phase1name tunnel6
set proposal 3des-md5
set src-addr-type subnet6
set dst-addr-type subnet6
end
```

Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network's port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

To create a static route in the web-based manager, go to *Router > Static > Static Route*. Select the drop-down arrow for *Create New* and select *IPv6 Route*. Enter the information and select *OK*. In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

```
config router static6
edit 1
set device port2
set dst 0::/0
next
edit 2
set device toB
set dst fec0:0000:0000:0004::/64
next
end
```

If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

IPv6 troubleshooting

There are a number of troubleshooting methods that can be used with IPv6 issues.

ping6

The main method of troubleshooting IPv6 traffic is using the IPv6 version of ping.

You can use the IPv6 ping command to:

- send an ICMP echo request packet to the IPv6 address that you specify.
- specify a source interface other than the one from which the probe originates by using the source interface keywords.
- specify a source IP address other than the one from which the probe originates by using the source address keywords

You can specify the following options:

packetCount	Number of packets to send to the destination IPv6 address. If you specify a zero, echo requests packets are sent indefinitely.
data-pattern	Sets the type of bits contained in the packet to all ones, all zeros, a random mixture of ones and zeros, or a specific hexadecimal data pattern that can range from 0x0 to 0xFFFFFFFF. The default is all zeros.
extended header attributes	Set the interface type and specifier of a destination address on the system that is configured for external loopback; the command succeeds only if the specified interface is configured for external loopback.
sweep interval	Specifies the change in the size of subsequent ping packets while sweeping across a range of sizes. For example, you can configure the sweep interval to sweep across the range of packets from 100 bytes to 1000 bytes in increments specified by the sweep interval. By default, the system increments packets by one byte; for example, it sends 100, 101, 102, 103, ... 1000. If the sweep interval is 5, the system sends 100, 105, 110, 115, ... 1000.
sweep sizes	Enables you to vary the sizes of the echo packets being sent. Used to determine the minimum sizes of the MTUs configured on the nodes along the path to the destination address. This reduces packet fragmentation, which contributes to performance problems. The default is to not sweep (all packets are the same size).
timeout	Sets the number of seconds to wait for an ICMP echo reply packet before the connection attempt times out.
hop limit	Sets the time-to-live hop count in the range 1-255; the default is 255.

The following characters may appear in the display after the ping command is issued:

- ! - reply received
- . - timed out while waiting for a reply
- ? - unknown packet type
- A - admin unreachable
- b - packet too big
- H - host unreachable
- N - network unreachable
- P - port unreachable
- p - parameter problem
- S - source beyond scope

t - hop limit expired (TTL expired)

IPv6 ping description

Ping uses the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. ECHO_REQUEST datagrams ("pings") have an IP and ICMP header, followed by a strict timeval and then an arbitrary number of "pad" bytes used to fill out the packet.

See also

IPv6 ping options

-a	Audible ping.
-A	Adaptive ping. Interpacket interval adapts to round-trip time, so effectively no more than one (or more, if preload is set) unanswered probe is present in the network. Minimal interval is 200msec for any user other than administrator. On networks with low rtt this mode is essentially equivalent to flood mode.
-b	Allow pingging of a broadcast address.
-B	Do not allow ping to change source address of probes. The address is bound to one selected when the ping starts.
-c count	Stop after sending count ECHO_REQUEST packets. With deadline option, ping waits for count ECHO_REPLY packets, until the timeout expires.
-d	Set the SO_DEBUG option on the socket being used. This socket option is not used by a Linux kernel.
-F flow label	Allocate and set 20 bit flow label on echo request packets (only ping6). If value is zero, kernel allocates random flow label.
-f	Flood ping. For every ECHO_REQUEST sent a period "." is displayed, while for ever ECHO_REPLY received a backspace is displayed. This provides a rapid display of how many packets are being dropped. If interval is not specified, it is set to zero and packets are output as fast as they come back or one hundred times per second, whichever is faster. Only the administrator may use this option with zero interval.
-i interval	Wait a specified interval of seconds between sending each packet. The default is 1 second between each packet, or no wait in flood mode. Only an administrator can set the interval to a value of less than 0.2 seconds.
-I interface address	Set source address to specified interface address. Argument may be numeric IP address or name of device. This option is required when you ping an IPv6 link-local address.
-l preload	If preload is specified, ping sends this number of packets that are not waiting for a reply. Only the administrator may select a preload of more than 3.
-L	Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.
-n	Numeric output only. No attempt will be made to look up symbolic names for host addresses.

-p pattern	You may specify up to 16 "pad" bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, -p ff will cause the sent packet to be filled with all ones.
-Q tos	Set Quality of Service -related bits in ICMP datagrams. tos can be either decimal or hex number. Traditionally (RFC1349), these have been interpreted as: 0 for reserved (currently being redefined as congestion control), 1-4 for Type of Service and 5-7 for Precedence. Possible settings for Type of Service are: minimal cost: 0x02, reliability: 0x04, throughput: 0x08, low delay: 0x10. Multiple TOS bits should not be set simultaneously. Possible settings for special Precedence range from priority (0x20) to net control (0xe0). You must be root (CAP_NET_ADMIN capability) to use Critical or higher precedence value. You cannot set bit 0x01 (reserved) unless ECN has been enabled in the kernel. In RFC 2474, these fields has been redefined as 8-bit Differentiated Services (DS), consisting of: bits 0-1 of separate data (ECN will be used, here), and bits 2-7 of Differentiated Services Codepoint (DSCP).
-q	Quiet output. Nothing is displayed except the summary lines at startup time and when finished
-R	Record route. (IPv4 only) Includes the RECORD_ROUTE option in the ECHO_REQUEST packet and displays the route buffer on returned packets. Note that the IP header is only large enough for nine such routes. Many hosts ignore or discard this option.
-r	Bypass the normal routing tables and send directly to a host on an attached interface. If the host is not on a directly-attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it provided the option -I is also used.
-s packetsize	Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.
-S sndbuf	Set socket sndbuf (send buffer). If not specified, it is selected to buffer not more than one packet.
-t ttl	Set the IP Time to Live.
-T timestamp option	Set special IP timestamp options. May be either tsonly (only timestamps), tsandaddr (timestamps and addresses) or tsprespec host1 [host2 [host3 [host4]]] (timestamp prespecified hops).
-M hint	Select Path MTU Discovery strategy. hint may be either do (prohibit fragmentation, even local one), want (do PMTU discovery, fragment locally when packet size is large), or don't (do not set DF flag).
-U	Print full user-to-user latency (the old behavior). Normally ping prints network round trip time, which can be different f.e. due to DNS failures.
-v	Verbose output.
-V	Show version and exit.

-w deadline	Specify a timeout, in seconds, before ping exits regardless of how many packets have been sent or received. In this case ping does not stop after count packet are sent, it waits either for deadline expire or until count probes are answered or for some error notification from network.
-W timeout	Time to wait for a response, in seconds. The option affects only timeout in absence of any responses, otherwise ping waits for two RTTs.

Examples

Ping a global V6 address with a 1400 byte packet from FortiGate CLI:

```
execute ping6 -s 1400 2001:480:332::10
```

Ping a multicast group using a ping6 command on FortiGate CLI (-I and port name must be specified for CLI ping6 command to ping v6 multicast group):

```
execute ping6 -I port1 ff02::1
```

Ping a localnet v6 address from FortiGate CLI:

```
execute ping6 FE80:0:0:0:213:e8ff:fe9e:ccf7
```

This address would normally be written as FE80::213:e8ff:fe9e:ccf7.

diagnose sniffer packet

The FortiOS built in packet sniffer also works with IPv6. The following are some examples using an IPv6-over-IPv4 tunnel called test6.

```
diagnose sniffer packet test6 'none' 4
interfaces=[test6]
filters=[]
pcap_lookupnet: test6: no IPv4 address assigned
34.258651 test6 -- 2001:4dd0:ff00:15d::2 -> 2001:4dd0:ff00:15d::1:
icmp6: echo request seq 1
34.324658 test6 -- 2001:4dd0:ff00:15d::1 -> 2001:4dd0:ff00:15d::2:
icmp6: echo reply seq 1
35.268581 test6 -- 2001:4dd0:ff00:15d::2 -> 2001:4dd0:ff00:15d::1:
icmp6: echo request seq 2
35.334230 test6 -- 2001:4dd0:ff00:15d::1 -> 2001:4dd0:ff00:15d::2:
icmp6: echo reply seq
```

```
diagnose sniffer packet any 'ip6 and tcp port 80' 4 10
interfaces=[any]
filters=[ip6 and tcp port 80]
1 LAN in 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: syn 2298823882
2 test6 out 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: syn 2298823882
3 test6 in 2a00:1450:8007::63.80 ->
2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037: syn 4218782319
ack
4 LAN out 2a00:1450:8007::63.80 ->
2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037: syn 4218782319
ack
5 LAN in 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: ack 4218782320
```

```
6 test6 out 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: ack 4218782320
```

diagnose debug flow

The `diagnose debug flow` CLI command is the same for IPv6 or IPv4. The output format is the same, however the command is only slightly different in that it uses `filter6` and an IPv6 address.

To enable diag debug flow for IPv6 - CLI

```
# diagnose debug enable
# diagnose debug flow show console enable
# diagnose debug flow show func enable
# diagnose debug flow filter6 addr 2001:4dd0:ff42:12::24
# diagnose debug flow trace start6
```

IPv6 specific diag commands

To list all the sit-tunnels that are configured:

```
diagnose ipv6 sit-tunnel list
total tunnel = 1:
devname=test6 devindex=4 ifindex=22 saddr=0.0.0.0
daddr=88.25.29.134 proto=41 vfid=0000 ref=2
```

To list all the IPv6 routes:

```
diagnose ipv6 route list
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst:::1/128 gwy::: prio=0
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst:2001:4dd0:ff00:75d::2/128 gwy::: prio=0
vf=0 type=01 protocol=kernel flag=00240021 oif=22(sixxs)
dst:2001:4dd0:ff00:75d::/64 gwy::: prio=100
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst:2001:4dd0:ff42:68::1/128 gwy::: prio=0
vf=0 type=01 protocol=kernel flag=01040001 oif=19(LAN)
dst:2001:4dd0:ff42:68:225:ff:feee:5314/128
gwy:2001:4dd0:ff42:68:225:ff:feee:5314 prio=0
.....
```

Some other IPv6 diagnose commands include:

<code>diagnose ipv6 neighbor-cache</code>	Add, delete, flush, or list the IPv6 ARP table or table entry.
<code>diagnose sys session6</code>	Clear, filter, full-stat, list, stat IPv6 sessions.
<code>tree diagnose ipv6</code>	View all the diagnose IPv6 commands.

Additional IPv6 resources

There are many RFCs available regarding IPv6. The following table lists the major IPv6 articles and their Internet Engineering Task Force (IETF) web locations.

RFC	Subject	Location
RFC 1933, <i>Transition Mechanisms for IPv6 Hosts and Routers</i>	Describes IPv4 compatibility mechanisms that can be implemented by IPv6 hosts and routers	http://www.ietf.org/rfc/rfc1933
RFC 2185, <i>Routing Aspects of IPv6 Transition</i>	Provides an overview of the routing aspects of the IPv6 transition	http://www.ietf.org/rfc/rfc2185
RFC 2373, <i>IP Version 6 Addressing Architecture</i>	Defines the addressing architecture of the IP Version 6 protocol [IPv6]	http://www.ietf.org/rfc/rfc2373
RFC 2402, <i>IP Authentication Header</i>	Describes functionality and implementation of IP Authentication Headers (AH)	http://www.ietf.org/rfc/rfc2402
RFC 2460, <i>Internet Protocol, Version 6 (IPv6) Specification</i>	Describes functionality, configuration of IP version 6 (IPv6) and differences from IPv4.	http://www.ietf.org/rfc/rfc2460
RFC 2461, <i>Neighbor Discovery for IP Version 6 (IPv6)</i>	Describes the features and functions of IPv6 Neighbor Discovery protocol	http://www.ietf.org/rfc/rfc2461
RFC 2462, <i>IPv6 Stateless Address Autoconfiguration</i>	Specifies the steps a host takes in deciding how to autoconfigure its interfaces in IPv6	http://www.ietf.org/rfc/rfc2462
RFC 2893, <i>Transition Mechanisms for IPv6 Hosts and Routers</i>	Specifies IPv4 compatibility mechanisms that can be implemented by IPv6 hosts and routers	http://www.ietf.org/rfc/rfc2893
RFC 3306, <i>Unicast-Prefix-Based IPv6 Multicast Addresses</i>	Describes the format and types of Ipv6 multicast addresses	http://www.ietf.org/rfc/rfc3306
RFC 3484, <i>Default Address Selection for Internet protocol version 6 (IPv6)</i>	Describes the algorithms used in IPv6 default address selection	http://www.ietf.org/rfc/rfc3484
RFC 3513, <i>Internet Protocol version 6 (IPv6) Addressing Architecture</i>	Contains details about the types of IPv6 addresses and includes examples	http://www.ietf.org/rfc/rfc3513
RFC 3587, <i>IPv6 Global Unicast Address Format</i>	Defines the standard format for IPv6 unicast addresses	http://www.ietf.org/rfc/rfc3587



Advanced FortiGate firewall concepts

The FortiGate firewall has advanced firewall component options, which allows for greater flexibility when these advanced options are needed to help with your growing network. These advanced firewall components include traffic shaping, QoS and identity-based policies.

The following topics are included in this section:

- [Central NAT table](#)
- [Stateful inspection of SCTP traffic](#)
- [Port pairing](#)
- [Blocking port 25 to email server traffic](#)
- [Blocking HTTP access by IP](#)
- [ICMP packet processing](#)
- [Adding NAT security policies in Transparent mode](#)
- [Adding a static NAT virtual IP for a single IP address and port](#)
- [Double NAT: combining IP pool with virtual IP](#)
- [Using VIP range for Source NAT \(SNAT\) and static 1-to-1 mapping](#)
- [Traffic shaping and per-IP traffic shaping](#)
- [Endpoint Security](#)
- [Logging traffic](#)
- [Quality of Service \(QoS\)](#)
- [Identity-based security policies](#)

Central NAT table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The NAT table also functions in the same way as the security policy table. That is, the FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well, the same way as security policies. NAT policies are applied to network traffic after a security policy.

To view the Central NAT configuration page, and use them in a security policy, you need to first enable it.

To enable Central NAT - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 In the Display Options on GUI section, select the check box beside *Central NAT table*.
- 3 Select *Apply*.

To enable Central NAT - CLI

```
config system global
    set gui-central-nat-table
end
```

NAT policies are created in the web-based manager by going to *Policy > Policy > Central NAT Table*. The NAT policies are enabled when you configure the security policy by selecting the *Use Central NAT Table* option.

NAT policies are created in the CLI by using the commands under `config firewall central-nat`. To enable the policies use the commands

```
config security policy
    edit <policy_number>
        set central-nat enable
    end
```

Central NAT Table configuration settings

To configure the Central NAT table, go to *Policy > Policy > Central NAT Table* and select *Create New*.

New NAT page

Source Address	Select the source IP address from the drop-down list. You can optionally create a group of source IP addresses when you select <i>Multiple</i> in the drop-down list. You can also create a new source IP address when you select <i>Create New</i> in the drop-down list.
Translated Address	Select the dynamic IP pool from the drop-down list.
Original Source Port	Enter the source port that the address is originating from.
Translated Port	Enter the translated port number. The number in the <i>From</i> field must be greater than the lower port number that is entered in the <i>To</i> field.

Stateful inspection of SCTP traffic

The Stream Control Transmission Protocol (SCTP) is a transport layer protocol similar to TCP and UDP. SCTP is designed to provide reliable, in-sequence transport of messages with congestion control. SCTP is defined in [RFC 4960](http://tools.ietf.org/html/rfc4960).

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a reliable transport protocol that runs on top of a connectionless packet network (IP). SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure-detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Configuring FortiGate SCTP filtering

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to ANY. FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism

- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Adding an SCTP custom service

This example creates a custom SCTP service that accepts SCTP traffic using destination port 2905. SCTP port number 2905 is used for SS7 Message Transfer Part 3 (MTP3) User Adaptation Layer (M3UA) over IP.

To add the SCTP custom service - web-based manager

- 1 Go to *Firewall Objects > Service > Custom* and select *Create New*.
- 2 Enter the following and select *OK*.

Name	M3UA_service
Protocol Type	TCP/UDP/SCTP
Protocol	SCTP
Source Port (Low)	1
Source Port (High)	65535
Destination Port (Low)	2905
Destination Port (High)	2905

To add the SCTP custom service - CLI

```
config firewall service custom
  edit M3UA_service
    set protocol TCP/UDP/SCTP
    set sctp-portrange 2905
  end
```

Adding an SCTP policy route

You can add policy routes that route SCTP traffic based on the SCTP source and destination port as well as other policy route criteria. The SCTP protocol number is 132.

The following example directs all SCTP traffic with SCTP destination port number 2905 to the next hop gateway at IP address 1.1.1.1.

To add the policy route - web-based manager

- 1 Go to *Router > Static > Policy Route*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*.

Protocol	132
Incoming interface	internal
Source address / mask	0.0.0.0 0.0.0.0
Destination address / mask	0.0.0.0 0.0.0.0
Destination Ports	From 2905 to 2905

Force traffic to:

Outgoing interface	external
Gateway Address	1.1.1.1

To add the policy route - CLI

```
config router policy
  edit 1
    set input-device internal
    set src 0.0.0.0 0.0.0.0
    set dst 0.0.0.0 0.0.0.0
    set output-device external
    set gateway 1.1.1.1
    set protocol 132
    set start-port 2905
    set end-port 2905
  end
```

Changing the session time to live for SCTP traffic


Use the following command to change the session timeout for SCTP protocol M3UA on port 2905 to 3600 seconds.

```
config system session-ttl
  config port
    edit 1
      set protocol 132
      set start-port 2905
      set end-port 2905
      set timeout 3600
    end
  end
```

Port pairing

Port pairing is an option in Transparent mode to bind two ports together. In doing this, you can create security policies that regulate traffic only between two specific ports, VLANs or VDOMs. In its simplest form, this enables an administrator to create security policies that are only between these two ports. Traffic is captured between these ports. No other traffic can enter DNS services or leave a port pairing.

For example, a FortiGate unit has three ports, where port 1 and port 2 are paired together, because the two networks only need to communicate with each other. If packet arrives on port 1, the FortiGate unit needs to figure out whether the packet goes to port 2 or port 3. With port pairing configured, it is more simple. If packet arrives on port 1, then the FortiGate automatically directs the packet to port 2. The opposite is also true in the other direction. This can be ideal when to groups only need to transfer data between each other.

Figure 25: Port pairing**To configure port pairing - web-based manager**

- 1 Go to *System > Network > Interface*.
- 2 Select the arrow beside *Create New*, and select *Port Pair*.
- 3 Enter a *Name* for the port pair.
- 4 Select the physical or virtual ports from the *Available Members* list and select the right-facing arrow to add the ports to the *Selected Members* list.
There can be only two ports added.
- 5 Select *OK*.

To configure port pairing - CLI

```
config system port-pair
  edit <pair_name>
    set member <port_names>
  end
```

When configuring security policies with the port pairs, selecting the *Source Interface* automatically populates the *Destination Interface*, and vice versa. All other aspects of the security policy configuration remains the same.

Blocking port 25 to email server traffic

Port 25 is the default port for SMTP traffic. Certain types of malware can install themselves on an unsuspecting user's computer and send spam using its own email server. By blocking port 25, this prevents a host system, and potentially your network or company, from being deemed a spam source.

This does, however limit your corporation from using a web server. You have a few options for this:

- if the email server is on a dedicated port, such as a DMZ port, security policies can ensure no traffic goes out from this port except the email server.
- Block all traffic on port 25 except the specific address of the email server.

Dedicated traffic

This example shows the steps to ensure only traffic exits from the DMZ where the email server is connected. The internal port is connected to the internal network and the WAN1 port connects to the Internet.

First, create a security policy that will not allow any traffic through port 25 from the internal interface, which connects to the internal network. Place this policy at the top of the security policy list.

To block traffic on port 25 - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Set the following options and select *OK*.

Source Interface	Internal
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY
Comments	Prevent Malware spam.

You may also want to enable *Log Violation Traffic* to see if there is any potential malware or other user sending email using the non-corporate email server.

To block traffic on port 25 - CLI

```
config security policy
edit <policy_number>
set srcintf Internal
set srcaddr all
set dstintf wan1
set dstaddr all
set schedule always
set service smtp
set action deny
set comment "Prevent Malware spam."
end
```

Next, create a security policy for the email server, IP address 10.10.11.29 that only allows SMTP traffic from the email server on port 25.

To allow traffic on port 25 for the email server - web-based manager

- 3 Go to *Policy > Policy > Policy* and select *Create New*.
- 4 Set the following options and select *OK*.

Source Interface	DMZ
Source Address	10.10.11.29
Destination Interface	WAN1
Destination Address	ALL

Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf dmz
    set srcaddr 10.10.11.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Restricting traffic on port 25

This example shows how to limit traffic on port 25 on the wan port to only traffic from the email server. The web server's address is 10.10.10.29.

To allow traffic on port 25 for the email server - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	10.10.10.29
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr 10.10.10.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Next, add a deny security policy that blocks all SMTP traffic from the Internal port to the WAN1 port. Ensure this policy is directly after the policy created above.

To block SMTP traffic on port 25 for the rest of the company - web-based manager

- 3 Go to *Policy > Policy > Policy* and select *Create New*.

- 4 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY

To block SMTP traffic on port 25 for the rest of the company - CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr all
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action deny
  end
```

Blocking HTTP access by IP

To block a web site using the IP, create a URL filter entry, using the additional information below. Note that this is only effective with HTTP or FortiGate units running Deep Inspection.

You need to create two URL filter entries. The first filter only allowing a text string containing two or more sets of text separated by a period. This is to match the various domain possibilities for web sites, for example:

- example.org
- www.example.com
- www.example.co.jp

The second filter blocks any IP address lookup.

To add the URL filter entries

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select *Create New* to add a filter group, give it a name and select *OK*.
- 3 Select *Create New* for a new filter.
- 4 Enter the *URL* of `^([a-z0-9-]+\.)\{1,\}[a-z]+\`
- 5 Set the *Type* to *Regex*.
- 6 Set the *Action* to *Allow*.
- 7 Select *OK*.
- 8 Select *Create New*.
- 9 Enter the *URL* of `[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}`

10 Set the *Type* to *Regex*.

11 Set the *Action* to *Block*.

12 Select *OK*.

Position these at the end of the URL filter list so that any exemptions or blocks before that are still effective.

Both of these filter entries are required. If you only enter the second one, the FortiGate unit will also catch a URL lookup as they both behave in a similar fashion after the URL is resolved to an IP. The first entry is needed to break out of the URL filter and allow the web site before it does the second check if they entered text.

ICMP packet processing

ICMP messages are used to relay feedback to the traffic source that the destination IP is not reachable. ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

Common error messages include:

- destination unreachable messages
- time exceeded messages
- redirect messages

For example, a security policy that allows TFTP traffic through the FortiGate unit. User1 (192.168.21.12) attempts to connect to the TFTP server (10.11.100.1), however, the UDP port 69 has not been opened on the server. The corresponding sniffer trace occurs:

```
diagnose sniffer packet any "host 10.11.100.1 or icmp 4"
3.677808 internal in 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
3.677960 wan1 out 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
3.678465 wan1 in 10.11.100.1.132 -> 192.168.21.12: icmp:
10.11.100.1 udp port 69 unreachable
3.678519 internal out 10.11.100.1 -> 192.168.21.12: icmp:
192.168.182.132 udp port 69 unreachable
```

Adding NAT security policies in Transparent mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.

- Add IP pools as required for source address translation

For NAT firewall policies to work in NAT mode you must have two interfaces on two different networks with two different subnet addresses. Then you can create firewall policies to translate source or destination addresses for packets as they are relayed by the FortiGate unit from one interface to the other.

A FortiGate unit operating in Transparent mode normally has only one IP address, the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

In the example shown in [Figure 26](#), all of the PCs on the internal network (subnet address 192.168.1.0/24) are configured with 192.168.1.99 as their default route. One of the management IPs of the FortiGate unit is set to 192.168.1.99. This configuration results in a typical NAT mode firewall. When a PC on the internal network attempts to connect to the Internet, the PC's default route sends packets destined for the Internet to the FortiGate unit internal interface. Similarly on the DMZ network (subnet address 10.1.1.0/24) all of the PCs have a default route of 10.1.1.99.

This example describes adding an internal to WAN1 security policy to relay these packets from the internal interface out the WAN1 interface to the Internet. Because the WAN1 interface does not have an IP address of its own, you must add an IP pool to the WAN1 interface that translates the source addresses of the outgoing packets to an IP address on the network connected to the wan1 interface.


The example describes adding an IP pool with a single IP address of 10.1.1.201. So all packets sent by a PC on the internal network that are accepted by the Internal to WAN1 policy leave the WAN1 interface with their source address translated to 10.1.1.201. These packets can now travel across the Internet to their destination. Reply packets return to the WAN1 interface because they have a destination address of 10.1.1.201. The Internal to WAN1 NAT policy translates the destination address of these return packets to the IP address of the originating PC and sends them out the internal interface to the originating PC.

Use the following steps to configure NAT in Transparent mode

- Add two management IPs
- Add an IP pool to the WAN1 interface
- Add an Internal to WAN1 security policy



You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

Figure 26: Example NAT in Transparent mode configuration**To add a source address translation NAT policy in Transparent mode**

- 1 Enter the following command to add two management IPs.

The second management IP is the default gateway for the internal network.

```
config system settings
  set manageip 10.1.1.99/24 192.168.1.99/24
end
```

- 2 Enter the following command to add an IP pool to the WAN1 interface:

```
config firewall ippool
  edit nat-out
    set interface "wan1"
    set startip 10.1.1.201
    set endip 10.1.1.201
  end
```


- 3 Enter the following command to add an Internal to WAN1 security policy with NAT enabled that also includes an IP pool:

```
config security policy
  edit 1
    set srcintf "internal"
    set dstintf "wan1"
    set scraddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set ippool enable
    set poolname nat-out
  end
```

Adding a static NAT virtual IP for a single IP address and port

In this example, the wan1 interface of the FortiGate unit is connected to the Internet and the Internal interface is connected to the DMZ network. The IP address 192.168.37.4 on port 80 on the Internet is mapped to 10.10.10.42 on port 8000 on the private network. Attempts to communicate with 192.168.37.4 from the Internet are translated and sent to 10.10.10.42 by the FortiGate unit. The computers on the Internet are unaware of this translation and see a single computer at 192.168.37.4 rather than a FortiGate unit with a private network behind it.

Figure 27: Static NAT virtual IP for a single IP address example



To add a static NAT virtual IP for a single IP address and port - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.
- 3 Complete the following and select *OK*.

Name	static_NAT
External Interface	wan1
Type	Static NAT
External IP Address/Range	192.168.37.4.
Mapped IP Address/Range	10.10.10.42
Port Forwarding	Selected
Protocol	TCP
External Service Port	80
Map to Port	8000

To add a static NAT virtual IP for a single IP address and port - CLI

```
config firewall vip
  edit static_NAT
    set extintf wan1
    set type static-nat
    set extip 192.168.37.4
    set mappedip 10.10.10.42
    set portforward enable
    set extport 80
    set mappedport 8000
  end
```

Add a external to dmz1 security policy that uses the virtual IP so that when users on the Internet attempt to connect to the web server IP address packets pass through the FortiGate unit from the external interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the external IP to the DMZ network IP address of the web server.

To add a static NAT virtual IP for a single IP address to a security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following:

Source Interface/Zone	wan1
Source Address	All
Destination Interface/Zone	Internal
Destination Address	static_nat
Schedule	always
Service	HTTP
Action	ACCEPT

- 3 Select *NAT*.
- 4 Select *OK*.

To add a static NAT virtual IP for a single IP address to a security policy - CLI

```
config security policy
  edit 1
    set srcintf wan1
    set dstintf internal
    set srcaddr all
    set dstaddr static_nat
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
```


Double NAT: combining IP pool with virtual IP

In this example, a combination of virtual IPs, IP pools and security policies will allow the local users to access the servers on the DMZ. The example uses a fixed port and IP pool to allow more than one user connection while using virtual IP to translate the destination port from 8080 to 80. The security policy uses both the IP pool and the virtual IP for double IP and/or port translation.

For this example:

- Users in the 10.1.1.0/24 subnet use port 8080 to access server 172.20.120.1.
- The server's listening port is 80.
- Fixed ports must be used.

Figure 28: Double NAT



To create an IP pool - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > IP Pool*.
- 2 Select *Create New*.
- 3 Enter the *Name* pool-1.
- 4 Enter the *IP Range/Subnet* 10.1.3.1-10.1.3.254.
- 5 Select *OK*.

To create an IP pool - CLI

```

config firewall ippool
  edit pool-1
    set startip 10.1.3.1
    set endip 10.1.3.254
  end

```

Next, create the virtual IP with port translation to translate the user internal IP used by the network users to the DMZ port and IP address of the server.

To create a Virtual IP with port translation - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.

- 3 Enter the following information and select OK.

Name	server-1
External Interface	Internal
Type	Static NAT
External IP Address/Range	172.20.120.1
	Note: This address is the same as the server address.
Mapped IP Address/Range	172.20.120.1
Port Forwarding	Enable
Protocol	TCP
External Service Port	8080
Map to Port	80

To create a Virtual IP with port translation - CLI

```
config firewall vip
  edit server-1
    set extintf internal
    set type static-nat
    set extip 172.20.120.1
    set mappedip 172.20.120.1
    set portforward enable
    set extport 80
    set mappedport 8080
  end
```

Add an internal to DMZ security policy that uses the virtual IP to translate the destination port number and the IP pool to translate the source addresses.

To create the security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following and select *OK*:

Source Interface/Zone	internal
Source Address	all
Destination Interface/Zone	dmz
Destination Address	server-1
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select
Dynamic IP Pool	Select, and select the <i>pool-1</i> IP pool.

To create the security policy - CLI

```
config security policy
```

```

edit 1
  set srcintf internal
  set dstintf dmz1
  set srcaddr all
  set dstaddr server-1
  set action accept
  set schedule always
  set service HTTP
  set nat enable
  set ippool enable
  set poolname pool-1
end


```

Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

This example shows how to use VIP ranges to perform source NAT (SNAT) with a static 1-to-1 mapping from internal to external IP addresses. This is similar to using an IP pool with the advantage of having predictable and static 1-to-1 address mapping.

Figure 29: Network diagram



This example will associate each internal IP address to one external IP address for the Source NAT (SNAT) translation.

Using the diagram above, the translations will look like the following:

Traffic from Source IP Translated to Source IP (SNAT)

10.10.10.42	192.168.37.4
10.10.10.43	192.168.37.5

```
...
10.10.10.46      192.168.37.8
```

First, configure the virtual IP.

To configure the virtual IP - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New*.
- 2 Enter the *Name* of `Static_NAT_1to1`.
- 3 Select the *External Interface* of *port 1* from the drop-down list.
- 4 Enter the *External IP Address* of `192.168.37.4`.
- 5 Enter the *Mapped IP Address* range of `10.10.10.42 to 10.10.10.46`.
- 6 Select OK.

To configure the virtual IP - CLI

```
config firewall vip
  edit "Static_NAT_1to1"
    set extip 192.168.37.4
    set extintf "port1"
    set mappedip 10.10.10.42-10.10.10.46
  next
end
```

Next, configure the firewall policies. Even if no connection needs to be initiated from external to internal, a second security policy number is required to activate the VIP range. Otherwise the IP address of the physical interface is used for NAT. In this example it is set as a “DENY” security policy for security purpose.

To configure the firewall policies - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following and select OK:

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
NAT	Select

- 3 Complete the following and select OK:

Source Interface/Zone	port 1
Source Address	all
Destination Interface/Zone	port 2
Destination Address	Static_NAT_1to1

Schedule	always
Service	ALL
Action	deny
Comments	Used to activate static Source NAT 1-to-1

To configure the firewall policies - CLI

```

config firewall policy
  edit 1
    set srcintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  next
  edit 2
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr Static_NAT_1to1
    set schedule always
    set service ANY
    set action deny
    set comments (Used to activate static Source NAT 1-to-1)
  next
end
end

```

Traffic shaping and per-IP traffic shaping

Traffic shaping helps to optimize traffic flow through the FortiGate unit, and per-IP traffic shaping does much the same, however, it applies traffic shaping per IP address instead of per policy or per shaper. Traffic shaping, when included in a security policy, controls the bandwidth available to the policy, and sets the priority of the traffic processed by the policy. Traffic shaping makes it possible to control which policies have the highest priority when large amounts of data are moving through the FortiGate unit. For example, the policy for the corporate web server might be given higher priority than the policies for most employee's computers. An employee who needs extra high speed Internet access could have a special outgoing policy set up with higher bandwidth.

Traffic shaping is available for security policies whose Action is ACCEPT, IPSEC, or SSL VPN. It is also available for all supported services, including H.323, TCP, UDP, ICMP, and ESP.

Traffic shaping is used to improve the quality of bandwidth-intensive and sensitive traffic; it also cannot increase the total amount of bandwidth available. The bandwidth available for traffic set in a traffic shaper is used to control data sessions for traffic in both directions.

For more information about traffic shaping, see the Traffic Shaping chapter in the [FortiOS Handbook](#).

Endpoint Security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the UTM chapter in the [FortiOS Handbook](#).

Logging traffic

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13 05:23:47 log_id=4 type=traffic subtype=other pri=notice
vd=root status="start" src="10.41.101.20" srcname="10.41.101.20"
src_port=58115 dst="172.20.120.100" dstname="172.20.120.100"
dst_country="N/A" dst_port=137 tran_ip="N/A" tran_port=0
tran_sip="10.31.101.41" tran_sport=58115 service="137/udp"
proto=17 app_type="N/A" duration=0 rule=1 policyid=1 sent=0 rcvd=0
shaper_drop_sent=0 shaper_drop_rcvd=0 perip_drop=0
src_int="internal" dst_int="wan1" SN=97404 app="N/A" app_cat="N/A"
carrier_ep="N/A"
```

If you want to know more about logging, see the Logging and Reporting chapter in the [FortiOS Handbook](#). If you want to know more about traffic log messages, see the [FortiGate Log Message Reference](#).

Quality of Service (QoS)

The Quality of Service (QoS) feature is an advanced firewall component that applies bandwidth limits and prioritization to traffic. QoS is the capability of the network to adjust some quality aspects for selected flows within your overall network traffic, and may include such techniques as priority-based queuing and traffic policing.

QoS can be implemented for services that include H.323, TCP, UDP, ICMP, and ESP. QoS uses the following techniques:

Traffic policing	Drops packets that do not conform to bandwidth limitations
Traffic shaping	This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.
Queuing	This transmits packets in order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

QoS can be helpful for organizations that are trying to manage their voice and streaming multi-media traffic, which can rapidly consume bandwidth. Both voice and streaming multi-media are sensitive to latency.

For additional information about QoS, see the Traffic Shaping chapter in the [FortiOS Handbook](#).

Identity-based security policies

Identity-based security policies, also known as authentication policies, match traffic that requires a supported authentication protocol to trigger the firewall authentication challenge and successfully authenticate network users. Network users authentication can occur using HTTP, HTTPS, FTP, and Telnet protocols as well as through automatic login using NTLM and FSSO, to bypass user intervention.

Identity-based security policies are usually configured for IPsec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

When configuring identity-based policies, you can use schedules to limit network users authentication sessions. For example, example.com has a schedule policy to use P2P applications between noon and 1:00 pm, and a user authentication timeout of 30 minutes. When a user logs in at 12:15 pm, their authentication time logs them off at 12:45 (30 minutes later). You can configure this type of authentication by using the `schedule-timeout` field in the `config firewall policy` command in the CLI.

Identity-based policy positioning

With identity-based security policies, positioning is extremely important. For a typical security policy, the FortiGate unit matches the source, destination and service of the policy. If matched, it acts on that policy. If not, the FortiGate unit moves to the next policy.

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-rules for the user groups and services. That is, it acts on the authentication and completes the remainder of that policy and goes no further in the policy list.

The way identity based policies work is that once src/dest are matched, it will process the identity based sub-rules (for lack of a better term) around the user groups and services. It will never process the rest of your rulebase. For this reason, unique security policies should be placed **before** an identity-based policy.

For example, consider the following policies:

Seq. No.	Source	Destination	Schedule	Service	Action	Status	Authentication
1	all	all	always	DNS	ACCEPT	<input checked="" type="checkbox"/>	
2	all	all	always	HTTP HTTPS	ACCEPT	<input checked="" type="checkbox"/>	FSAE_Guest_Users
3	all	all	always	ANY	DENY	Implicit	

DNS traffic goes through successfully as does any HTTP traffic after being authenticated. However, if there was FTP traffic, it would not get through. As the FortiGate unit processes FTP traffic, it skips rule one since it's matching the source, destination and service. When it moves to rule two it matches the source and destination, it determines there is a match and, sees there are also processes the group/service rules, which requires authentication and acts on those rules. Once satisfied, the FortiGate unit will never go to rule three.

In this situation, where you would want FTP traffic to traverse the FortiGate unit, create a security policy specific to the services you require and place it above the authentication policy.

Identity-based sub-policies

When adding authentication to a security policy, you can add multiple authentication rules, or sub-policies. Within these policies you can include additional UTM profiles, traffic shaping and so on, to take affect on the selected services.

Figure 30: Authentication sub-policies

☒ Enable Identity Based Policy

Rule ID	User Group	Service	Schedule	UTM	Traffic Shaping	Logging	
1	FSAE_Guest_Users	HTTP,HTTPS	always	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
2	FSAE_Guest_Users	FTP	always	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

☒ Firewall ☒ Directory Service(FSAE) ☐ NTLM Authentication

These sub-policies work on the same principle as normal security policies, that is, top down until the criteria has been met. As such, if there is no matching policy within the list, the packet can still be dropped even after authentication is successful.



Chapter 3 System Administration

This guide contains the following sections:

[Using the web-based manager](#) provides an overview of the web-based manager interface for FortiOS. If you are new to the FortiOS web-based manager, this chapter provides a high level overview of how to use this method of administration.

[Using the CLI](#) provides an overview of the command line interface (CLI) for FortiOS. If you are new to the FortiOS CLI, this chapter provides a high level overview of how to use this method of administration.

[Basic setup](#) describes the simple setup requirements an Administrator should do to get the FortiGate unit on the network and enabling the flow of traffic.

[Interfaces](#) describes FortiGate interface settings.

[Central management](#) describes how to configure the FortiGate unit to use FortiManager as a method of maintaining the device and other features that FortiManager has to facilitate the administration of multiple devices.

[Best practices](#) discusses methods to make the various components of FortiOS more efficient, and offer suggestions on ways to configure the FortiGate unit.

[FortiGuard](#) discusses the FortiGuard network services and configuration examples.

[Monitoring](#) describes various methods of collecting log data and tracking traffic flows and trends.

[Multicast forwarding](#) describes multicasting (also called IP multicasting) and how to configure it on the FortiGate unit.

[Virtual LANs](#) discusses their implementation in FortiOS and how to configure and use them.

[PPTP and L2TP](#) describes these VPN types and how to configure them.

[Session helpers](#) describes what they are and how to view and configure various session helpers.

[Advanced concepts](#) describes more involved administrative topics to enhance network security and traffic efficiency.



Using the web-based manager

This section describes the features of the web-based manager administrative interface (sometimes referred to as a graphical user interface, or GUI) of your unit. This section also explains common web-based manager tasks that an administrator does on a regular basis, as well as online help.

The following topics are included in this section:

- [Web-based manager overview](#)
- [Web-based manager menus and pages](#)
- [Using online help](#)
- [Entering text strings](#)
- [Basic configurations](#)

Web-based manager overview

The web-based manager is a user-friendly interface for configuring settings and managing the unit. Accessing the web-based manager is easy; by using HTTP or a secure HTTPS connection from any management computer using a web browser. The recommended minimum screen resolution for properly displaying the web-based manager is 1280 by 1024. Some web browsers do not correctly display the windows within the web-based manager interface. Verify that you have a supported web browser by reviewing the Knowledge Base articles, [Microsoft Windows web browsers supported by Fortinet products web-based manager \(GUI\) web browsers](#), and [Mac OS browsers for use with Fortinet hardware web-based manager \(GUI\)](#).

The web-based manager also provides the CLI Console widget, which enables you to connect to the command line interface (CLI) without exiting out of the web-based manager.

Web-based manager menus and pages

The web-based manager provides access to configuration options for most of the FortiOS features from the main menus. The web-based manager contains the following main menus:

System	Configure system settings, such as network interfaces, virtual domains, DHCP services, administrators, certificates, High Availability (HA), system time and set system options.
Router	Configure static, dynamic and multicast routing and view the router monitor.
Policy	Configure firewall policies, protocol options and Central NAT Table.
Firewall Objects	Configure supporting content for firewall policies including scheduling, services, traffic shapers, addresses, virtual IP and load balancing.

UTM Profiles	Configure antivirus and email filtering, web filtering, intrusion protection, data leak prevention, and application control. This menu also includes endpoint security features, such as FortiClient configuration and application detection patterns.
VPN	Configure IPsec and SSL virtual private networking.
User	Configure user accounts and user authentication including external authentication servers.
WAN Opt. & Cache	Configure WAN optimization and web caching to improve performance and security of traffic passing between locations on your wide area network (WAN) or from the Internet to your web servers.
WiFi Controller	Configure the unit to act as a wireless network controller, managing the wireless Access Point (AP) functionality of FortiWiFi and FortiAP units.
Log&Report	Configure logging and alert email as well as reports. View log messages and reports.
Current VDOM	Appears only when VDOMs are enabled on the unit to switch between VDOMs.

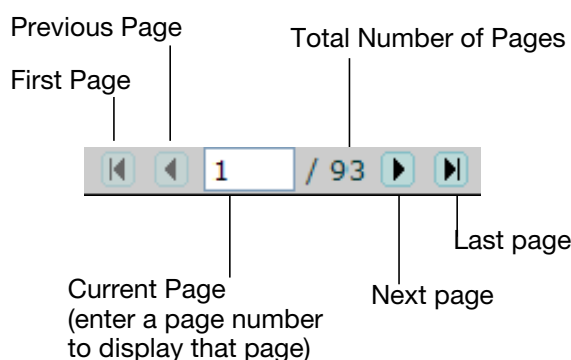
Using information tables

Many of the web-based manager pages contain tables of information which you can filter to display specific information. Administrators with read and write access can define the filters.

Using page navigation

The web-based manager pages that contain information and lists that span multiple pages. At the bottom of the page is the page navigation controls that enable you to move between pages.

Figure 31: Page controls



Adding filters to web-based manager lists

To locate a specific set of information or content within multiple pages, you use filters. These are especially useful in locating specific log entries. Depending on the type of information, the filtering options vary.

To create a filter, select *Filter Settings*, or a filter icon in a column heading. When a filter is applied to a column, the filter icon becomes green. Filter settings are stored in the unit's configuration and will be maintained the next time that you access any list for which you have added filters.

Filtering variable can include a numeric range such as 25-50 or an IP address or part of an address, or any text string combination, including special characters.

Note that the filtering ignores characters following a "<" unless the followed by a space. For example, the filtering ignores <string but not < string. Filtering also ignores matched opening and closing (< and >) characters and any characters between them. For example, filtering will ignore <string>.

For columns that can contain only specific content, such as log message severity, you can only select a single item from a list.

Using column settings

On pages where large amounts of information is available, not all content can be displayed, or some content may not be of use to you. Using column settings, you can display only that content which is important to your requirements.

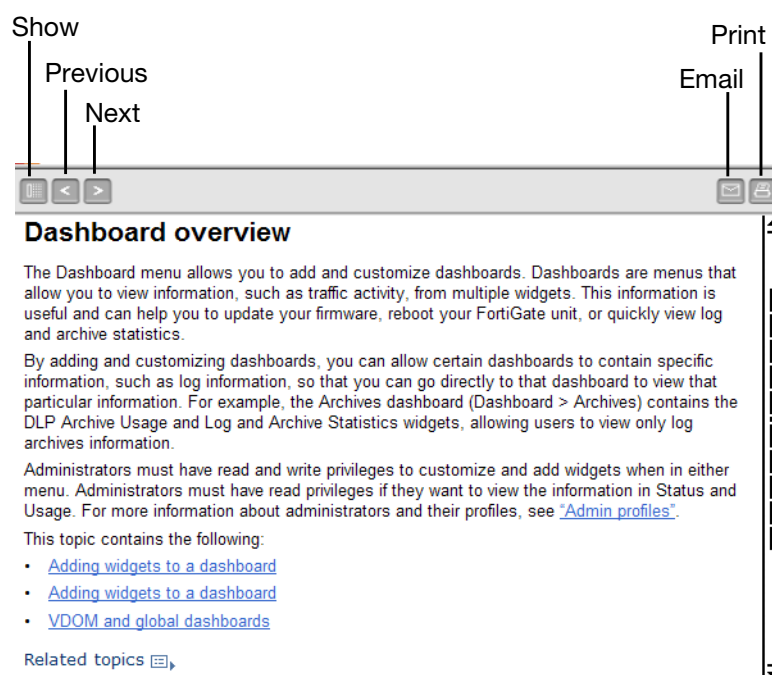
To configure column settings, select the *Column Settings* link at the top right of the page.

Any changes that you make to the column settings of a list are stored in the unit's configuration and will display the next time that you access the list.

Using online help

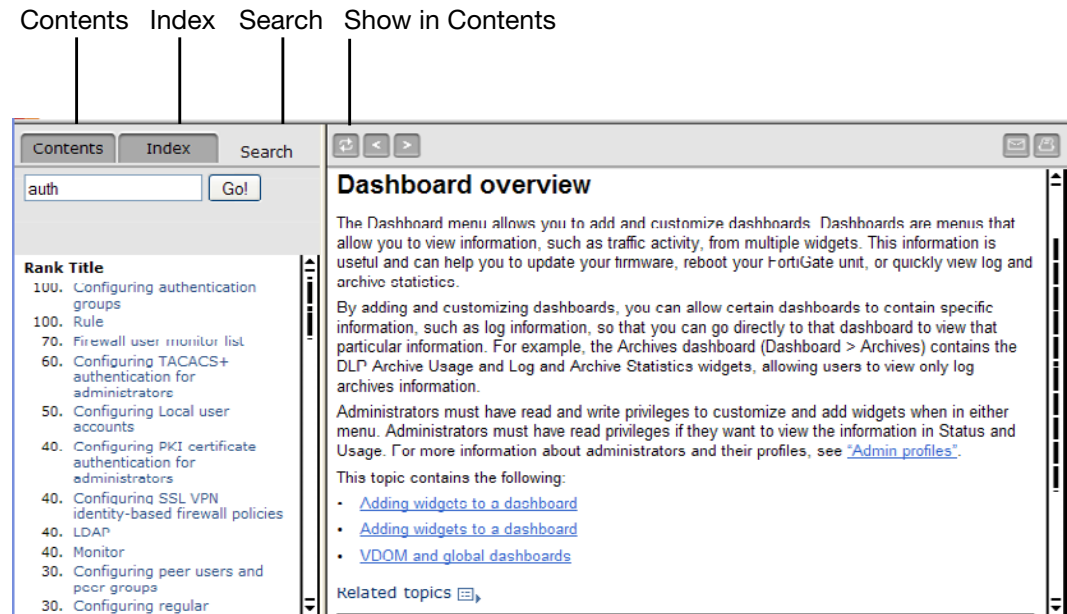
This Online Help button system provides context-sensitive help for the current web-based manager page, as well as access to the online version of the FortiGate Handbook.

Figure 32: A context-sensitive online help page (content pane only)



To view the online help table of contents or index, and to use the search, select *Show Navigation*.

Figure 33: Online help page with navigation pane and content pane



Contents	Display the online help table of contents. The online help is organized in the same way as the web-based manager.
Index	Display the online help index.
Search	Display the online help search.
Show in Contents	Select <i>Show in Contents</i> to display the location of the current help page within the table of contents. If you have used the index, search, or hyperlinks to find information in the online help, the table of contents may not be visible or the table of contents may display where you are within the table of contents.

Online help search tips

- If you search for multiple words, the search finds only those results that contain all of the words that you entered. The search does not find pages that only contain one of the words that you entered.
- The pages found by the search are ranked in order of relevance. The higher the ranking, the more likely the page includes the information a you are searching for. Help pages with the search words in the help page title are ranked highest.
- You can use the asterisk (*) as a wildcard. For example, if you search for **auth*** the search finds help pages containing **auth**, **authenticate**, **authentication**, **authenticates**.

Using the keyboard to navigate in the online help

You can use the keyboard shortcuts listed below to display and find information in the online help.

Key	Function
Alt+1	Display the table of contents.
Alt+2	Display the index.
Alt+3	Display the Search tab.
Alt+4	Go to the previous page.
Alt+5	Go to the next page.
Alt+7	Send an email to Fortinet Technical Documentation at techdoc@fortinet.com if you have comments on or corrections for the online help or any other Fortinet technical documentation product.
Alt+8	Print the current online help page.
Alt+9	Add an entry for this online help page to your browser bookmarks or favorites list, to make it easier to find useful online help pages.

Entering text strings

The configuration of a FortiGate unit is stored as configuration settings in the FortiOS configuration database. To change the configuration you can use the web-based manager or CLI to add, delete or change configuration settings. These configuration changes are stored in the configuration database as you make them.

Individual settings in the configuration database can be text strings, numeric values, selections from a list of allowed options, or on/off (enable/disable) settings.

Entering text strings (names)

Text strings are used to name entities in the configuration. For example, the name of a firewall address, administrative user, and so on. You can enter any character in a FortiGate configuration text string except, to prevent Cross-Site Scripting (XSS) vulnerabilities, text strings in FortiGate configuration names cannot include the following characters:

" (double quote), & (ampersand), ' (single quote), < (less than) and > (greater than)

Most web-based manager text string fields make it easy to add an acceptable number of characters and prevent you from adding the XSS vulnerability characters.

From the CLI, you can also use the `tree` command to view the number of characters that are allowed in a name field. For example, firewall address names can contain up to 64 characters. When you add a firewall address to the web-based manager you are limited to entering 64 characters in the firewall address name field. From the CLI you can enter the following `tree` command to confirm that the firewall address `name` field allows 64 characters.

```
config firewall address
tree
-- [address] --*name (64)
    |- subnet
    |- type
    |- start-ip
```

```
| - end-ip  
| - fqdn (256)  
| - cache-ttl (0,86400)  
| - wildcard  
| - comment (64 xss)  
| - associated-interface (16)  
+- color (0,32)
```

The `tree` command output also shows the number of characters allowed for other firewall address name settings. For example, the fully-qualified domain name (`fqdn`) field can contain up to 256 characters.

Entering numeric values

Numeric values set various sizes, rates, numeric addresses, and other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

Selecting options from a list

If a configuration field can only contain one of a number of selected options, the web-based manager and CLI present you a list of acceptable options and you can select one from the list. No other input is allowed. From the CLI you must spell the selection name correctly.

Enabling or disabling options

If a configuration option can only be on or off (enabled or disabled) the web-based manager presents a check box or other control that can only be enabled or disabled. From the CLI you can set the option to `enable` or `disable`.

Dashboard

The Dashboard menu provides a way to access information about network activity and events, as well as configure basic system settings. FortiOS includes a default dashboard, called Status. You can add more dashboards to contain the content you need at your fingertips.

Each information “chunk” is within a widget. Widgets provide an easy and quick way to view a variety of information, such as statistical information or network activity. There are a selection of widgets to choose from by selecting the *Widgets* option.

Administrators must have read and write privileges for adding and configuring dashboards and widgets.



Your browser must have Java script enabled to view the Dashboard page.

Adding dashboards

Dashboards that you create are automatically added under the default status and usage dashboards. You can add, remove or rename a dashboard, regardless of whether it is default. You can also reset the Dashboard menu to its default settings by selecting *Reset Dashboards*.



If VDOMs are enabled, only the dashboards within Global are available for configuration.

To add a dashboard

- 1 Go to *System > Dashboard > Status*.
- 2 Select *Dashboard*, located at the top left of the page.
- 3 Select *Add Dashboard*.
- 4 Enter a name for the dashboard.
- 5 Select *OK*.

Adding widgets to a dashboard

To add a widget to a dashboard, select *Widget* located at the top left of the dashboard page. Select a widget add it to the dashboard. Select the red X-box to close the window.

Figure 34: A minimized display



In an HA cluster, the information that appears applies to the whole HA cluster, not just the primary FortiGate unit.

System Information widget

The System Information widget status information on the FortiGate unit and provides the access point to update the firmware and backup the configurations.

System Information widget	
Host Name	<p>The name of the FortiGate unit. For details on changing the name, see Changing the FortiGate unit's host name.</p> <p>If the FortiGate unit is in HA mode, this information is not displayed.</p>
Serial Number	<p>The serial number of the FortiGate unit. The serial number is specific to that FortiGate unit and does not change with firmware upgrades.</p>
Operation Mode	<p>The current operating mode of the FortiGate unit. A FortiGate unit can operate in NAT mode or transparent mode. Select <i>Change</i> to switch between NAT and transparent mode. For more information, see Changing the operation mode.</p> <p>If virtual domains are enabled, this field shows the operating mode of the current virtual domain. The Global System Status dashboard does not include this information.</p>
HA Status	<p>The status of high availability within the cluster.</p> <p>Standalone indicates the FortiGate unit is not operating in HA mode.</p> <p>Active-Passive or Active-Active indicate the FortiGate unit is operating in HA mode.</p> <p>Select <i>Configure</i>, to change the HA configuration.</p>
Cluster Name	<p>The name of the HA cluster for this FortiGate unit.</p> <p>The FortiGate unit must be operating in HA mode to display this field.</p>
Cluster Members	<p>The FortiGate units in the HA cluster. Information displayed about each member includes host name, serial number, and whether the FortiGate unit is a primary (master) or subordinate (slave) FortiGate unit in the cluster.</p> <p>The FortiGate unit must be operating in HA mode with virtual domains disabled to display this information.</p>
Virtual Cluster 1 Virtual Cluster 2	<p>The role of each FortiGate unit in virtual cluster 1 and virtual cluster 2.</p> <p>The FortiGate unit must be operating in HA mode with virtual domains enabled to display this information.</p>
System Time	<p>The current date and time. Select <i>Change</i>, to configure the system time. For more information, see Configuring system time.</p>
Firmware Version	<p>The version of the current firmware installed on the FortiGate unit.</p> <p>Select <i>Update</i> to upload a newer or older firmware version. For more information, see Changing the firmware.</p>
System Configuration	<p>The time period of when the configuration file was backed up. Select <i>Backup</i> to back up the current configuration. For more information, see Backing up the configuration.</p> <p>To restore a configuration file, select <i>Restore</i>. For more information, see Restoring your firmware configuration.</p>

Current Administrator	The number of administrators currently logged into the FortiGate unit. Select <i>Details</i> to view more information about each administrator that is currently logged in If you want to changed the current administrator's password, see Changing the currently logged in administrator's password .
Uptime	The time in days, hours, and minutes since the FortiGate unit was started or rebooted.
Virtual Domain	Status of virtual domains on your FortiGate unit. Select <i>Enable</i> or <i>Disable</i> to change the status of virtual domains feature. If you enable or disable virtual domains, your session will be terminated and you will need to log in again.

Changing the FortiGate unit's host name

The host name appears in the *Host Name* row, in the *System Information* widget. The host name also appears at the CLI prompt when you are logged in to the CLI and as the SNMP system name.

The only administrators that can change a FortiGate unit's host name are administrators whose admin profiles permit system configuration write access. If the FortiGate unit is part of an HA cluster, you should use a unique host name to distinguish the FortiGate unit from others in the cluster.

To change the host name on the FortiGate unit, in the *System Information* widget, select *Change* in the *Host Name* row.

Changing the operation mode

FortiGate units and individual VDOMs can operate in NAT or transparent mode. From the *System Information* dashboard widget you can change the operating mode for your FortiGate unit or for a VDOM and perform sufficient network configuration to ensure that you can connect to the web-based manager in the new mode.

NAT mode

In NAT mode (also called NAT mode), the FortiGate unit is visible to the network that it is connected to. All of its interfaces are on different subnets. Each interface that is connected to a network must be configured with an IP address that is valid for that subnetwork. The FortiGate unit functions as a

You would typically use NAT mode when the FortiGate unit is deployed as a gateway between private and public networks (or between any networks). In its default NAT mode configuration, the FortiGate unit functions as a router, routing traffic between its interfaces. Security policies control communications through the FortiGate unit to both the Internet and between internal networks. In NAT mode, the FortiGate unit performs network address translation before IP packets are sent to the destination network.

For example, a company has a FortiGate unit as their interface to the Internet. The FortiGate unit also acts as a router to multiple sub-networks within the company. In this situation the FortiGate unit is set to NAT mode. Using this mode, the FortiGate unit can have a designated port for the Internet, in this example, wan1 with an address of 172.20.120.129, which is the public IP address. The internal network segments are behind the FortiGate unit and invisible to the public access, for example port 2 with an address of 10.10.10.1. The FortiGate unit translates IP addresses passing through it to route the traffic to the correct subnet or the Internet.

Transparent Mode

In transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. To connect the FortiGate unit to your network, all you have to do is configure a management IP address and a default route.

You would typically use the FortiGate unit in transparent mode on a private network behind an existing firewall or behind a router. In transparent mode, the FortiGate unit also functions as a firewall. Security policies control communications through the FortiGate unit to the Internet and internal network. No traffic can pass through the FortiGate unit until you add security policies.

For example, the company has a router or other firewall in place. The network is simple enough that all users are on the same internal network. They need the FortiGate unit to perform application control, antivirus and intrusion protection and similar traffic scanning. In this situation the FortiGate unit is set to transparent mode. The traffic passing through the FortiGate unit does not change the addressing from the router to the internal network. Security policies and UTM profiles define the type of scanning the FortiGate unit performs on traffic entering the network.

To switch from NAT to transparent mode

- 1 From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
- 2 From the *Operation Mode* list, select *Transparent*.
- 3 Enter the *Management IP* address and *Netmask*. This is the IP address to connect to when configuring and maintaining the device.
- 4 Enter the *Default Gateway*.
- 5 Select *OK*.

To change the transparent mode management IP address

- 1 From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
- 2 Enter a new IP address and netmask in the *Management IP/Network* field as required and select *OK*.

Your web browser is disconnected from the web-based manager. To reconnect to the web-based manager browse to the new management IP address.

To switch from transparent to NAT mode

- 1 From the *System Information* dashboard widget select *Change* beside *Operation Mode*.
- 2 From the *Operation Mode* list, select *NAT*.
- 3 Enter a valid IP address and netmask for the network from which you want to manage the FortiGate unit.
- 4 Select the interface to which the *Interface IP/Netmask* settings apply
- 5 Enter the IP address default gateway required to reach other networks from the FortiGate unit. This option address a default route to the static routing table. The gateway setting of this default route is set to the IP address that you enter and the device setting of this default route is set to the interface selected in the *Device* field.
- 6 After the FortiGate unit switches to NAT mode you may need to go to *Router > Static Route* and edit this default route.
- 7 Select *OK*.

Configuring system time

The FortiGate unit's system time can be changed using the *System Information* widget by selecting *Change* in the *System Time* row.

Time Settings page	
System Time	The current system date and time on the FortiGate unit.
Refresh	Update the display of the FortiGate unit's current system date and time.
Time Zone	Select the current system time zone for the FortiGate unit.
Set Time	Select to set the system date and time to the values.
Synchronize with NTP Server	Select to use a Network Time Protocol (NTP) server to automatically set the system date and time. You must specify the server and synchronization interval. FortiGate units use NTP Version 4. For more information about NTP see http://www.ntp.org .
Server	Enter the IP address or domain name of an NTP server. To find an NTP server that you can use, see http://www.ntp.org .
Sync Interval	Specify how often the FortiGate unit should synchronize its time with the NTP server.

Daylight savings time is enabled by default. You can disable daylight savings time using the CLI commands:

```
config system global
    set dst disable
end
```

Changing the firmware



To avoid losing configuration settings you should always back up your configuration before changing the firmware image.

Administrators whose admin profiles permit maintenance read and write access can change the FortiGate unit's firmware. Firmware images can be installed from a number of sources including a local hard disk, a local USB disk, or the FortiGuard Network.

To change the firmware, go to *System > Dashboard > Status > System Information* widget and select the *Update* link on the *Firmware Version* row.

Firmware Upgrade/Downgrade page	
Upgrade From	Select the firmware source from the drop down list of available sources.
Firmware Version	<p>This appears only when selecting <i>FortiGuard Network</i> is selected from the <i>Upgrade From</i> drop-down list. Select a firmware version from the drop-down list.</p> <p>If downgrading the firmware on the FortiGate unit, select the check box beside Allow Firmware Downgrade.</p>
Upgrade File	<p>Browse to the location of the firmware image on your local hard disk.</p> <p>This field is available for local hard disk and USB only.</p>
Allow Firmware Downgrade	<p>Select to confirm the installation of an older firmware image (downgrade).</p> <p>This appears only when selecting <i>FortiGuard Network</i> is selected from the <i>Upgrade From</i> drop-down list.</p>
Upgrade Partition	<p>The number of the partition being updated.</p> <p>This field is available only if your FortiGate unit has more than one firmware partition.</p>
Boot the New Firmware	<p>By default, this is enabled. Select to disable the FortiGate unit's reboot process when installing a firmware image to a partition.</p> <p>This option enables you to install a firmware image to a partition without the FortiGate unit rebooting itself and making the firmware image the default firmware that is currently running.</p>



You need to register your FortiGate unit with Customer Support to access firmware updates for your model. For more information, go to <http://support.fortinet.com> or contact Customer Support.

Backing up the configuration

Administrators can back up the FortiGate unit's configuration file from the *System Information* widget. Select *Backup* in the *System Configuration* row, to back up the firmware configuration file to a local computer, USB disk or to a FortiManager unit.

You should always back up your configuration whenever you make any modifications to the device configuration or performing any firmware updates or changes.

Backup page	
Local PC	Select to back up the configuration file to a local management computer.
FortiManager	<p>Select to back up the configuration file to a FortiManager unit. The Central Management settings must be enabled and a FortiManager unit connected with the FortiGate unit so that the FortiGate unit can send the configuration file to the FortiManager unit.</p> <p>To enable central management, go to <i>System > Admin > Central Management</i>.</p>

USB Disk	Select to back up the configuration file to a USB key that is connected to the FortiGate unit.
Full Config	Select to backup the full VDOM configuration. This appears only when the FortiGate unit has VDOM configuration enabled.
VDOM Config	Select to backup the only the VDOM configuration file. This option backs up only the configuration file within that VDOM. Select the VDOM from the drop-down list, and select <i>Backup</i> .
Encrypt configuration file	Select to enable a password to the configuration file for added security.
Password	Enter the password that will be used to restore the configuration file.
Confirm	Re-enter the password.

Formatting USB

The FortiGate unit enables you to back up the configuration of the device to a USB flash drive. The USB flash drive must be formatted as a FAT16 disk.

To format the USB flash drive, either use the CLI command `exe usb-disk format.` or within Windows at a command prompt, enter the command...

```
"format <drive_letter>: /FS:FAT /V:<drive_label>
```

... where <drive_letter> is the letter of the connected USB flash drive and <drive_label> is the name to give the USB drive.

Remote FortiManager backup and restore options

After successfully connecting to the FortiManager unit from your FortiGate unit, you can back up and restore your configuration to and from the FortiManager unit.

A list of revisions is displayed when restoring the configuration from a remote location. The list allows you to choose the configuration to restore. To use the FortiManager unit as a method of backup and restore of configuration files, you must first configure a connection between the two devices. For more information, see [Central management](#).

Remote FortiGuard backup and restore options

Your FortiGate unit can be remotely managed by a central management server, which is available when you register for the FortiGuard Analysis and Management Service. FortiGuard Analysis and Management Service is a subscription-based service and is purchased by contacting support.

After registering, you can back up or restore your configuration. FortiGuard Analysis and Management Service is useful when administering multiple FortiGate units without having a FortiManager unit. Using this service you can also upgrade the firmware. Upgrading the firmware is available in the *Firmware Upgrade* section of the backup and restore menu.

When restoring the configuration from a remote location, a list of revisions is displayed so that you can choose the configuration file to restore.



The FortiGuard-FortiManager protocol is used when connecting to the FortiGuard Analysis and Management Service. This protocol runs over SSL using IPv4/TCP port 541 and includes the following functions:

- detects FortiGate unit dead or alive status
- detects management service dead or alive status
- notifies the FortiGate units about configuration changes, AV/IPS database update and firewall changes.

Restoring your firmware configuration

Administrators can restore a configuration file that was backed up using the *System Information* widget. If the configuration file was encrypted, you will need the password to restore the configuration file.

Restore	
Local PC	Select to back up the configuration file to a local management computer.
FortiManager	Select to back up the configuration file to a FortiManager unit. The Central Management settings must be enabled and a FortiManager unit connected with the FortiGate unit so that the FortiGate unit can send the configuration file to the FortiManager unit. To enable central management, go to <i>System > Admin > Central Management</i> .
USB Disk	Select to back up the configuration file to a USB key that is connected to the FortiGate unit.
Filename	Select Browse to locate the configuration file
Password	If a password was set when saving the configuration file, enter the password.

Viewing online administrators

The *System Information* widget enables you to view information about the administrators logged into the FortiGate unit. To view logged in administrators, in the *System Information* widget, select *Details*. in the *Current Administrator* row.

Administrators logged in window (System Information widget)	
Lists the administrators that are currently logged into the FortiGate unit.	
Disconnect	To disconnect an administrator, select the check box next to the administrator's name and select <i>Disconnect</i> . This is available only if your admin profile gives you <i>System Configuration</i> write permission. You cannot log off the default "admin" user.
Refresh	Select to update the list.
User Name	The administrator account name.
Type	The type of access: http, https, jsconsole, sshv2.

From	The administrator's IP address. If <i>Type</i> is <i>jsconsole</i> , the value in <i>From</i> is <i>N/A</i> .
Time	The date and time the administrator logged on.

Changing the currently logged in administrator's password

Use the *System Information* widget, to change your password. To do this, select the *Change Password* option in the *Current Administrator* row.

Edit Password	
Administrator	The name of the administrator who is changing their password.
Old Password	Enter your current password.
New Password	Enter the new password.
Confirm Password	Enter the new password again to confirm.

License Information widget

License Information displays the status of your technical support contract and FortiGuard subscriptions. The FortiGate unit updates the license information status indicators automatically when attempting to connect to the FortiGuard Distribution Network (FDN). FortiGuard Subscriptions status indicators are green if the FDN was reachable and the license was valid during the last connection attempt, grey if the FortiGate unit cannot connect to the FDN, and orange if the FDN is reachable but the license has expired.

When a new FortiGate unit is powered on, it automatically searches for FortiGuard services. If the FortiGate unit is configured for central management, it will look for FortiGuard services on the configured FortiManager system. The FortiGate unit sends its serial number to the FortiGuard service provider, which then determines whether the FortiGate unit is registered and has valid contracts for FortiGuard subscriptions and FortiCare support services. If the FortiGate unit is registered and has a valid contract, the License Information is updated.

If the FortiGate unit is not registered, any administrator with the *super_admin* profile sees a reminder message that provides access to a registration form.

When a contract is due to expire within 30 days, any administrator with the *super_admin* profile sees a notification message that provides access to an Add Contract form. Simply enter the new contract number and select *Add*. Fortinet Support also sends contract expiry reminders.

You can optionally disable notification for registration or contract inquiry using the `config system global` command in the CLI. Selecting any of the *Configure* options will take you to the Maintenance page.

License Information widget	
Support Contract	<p>Displays details about your current Fortinet Support contract.</p> <ul style="list-style-type: none"> • If <i>Not Registered</i> appears, select <i>Register</i> to register the FortiGate unit. • If <i>Expired</i> appears, select <i>Renew</i> for information on renewing your technical support contract. Contact your local reseller. • If <i>Registered</i> appears the name of the support that registered this FortiGate unit is also displayed. • You can select <i>Login Now</i> to log into the Fortinet Support account that registered this FortiGate unit. <p>The support contract section also includes information on the number of FortiClient users connecting to the FortiGate unit. It displays the number of FortiClient connections allowed, and the number of users connecting. By selecting the Details link for the number of connections, you can view more information about the connecting user, including IP address, user name and type of operating system the user is connecting with.</p>
FortiGuard Services	Displays the currently installed version of the attack and virus definitions for the various UTM services from FortiGuard. Select <i>Renew</i> to update any of the licenses.
Virtual Domain	<p>Displays the maximum number of virtual domains the FortiGate unit supports with the current license.</p> <p>For high-end models, you can select the <i>Purchase More</i> link to purchase a license key through Fortinet technical support to increase the maximum number of VDOMs.</p>
FortiClient Software	View information about the latest version of FortiClient licenses and users connecting using the software.

Manually updating FortiGuard definitions

You can update the definition files for a number of FortiGuard services from the *License Information* widget.

To update FortiGuard definitions manually

- 1 Download the latest update files from Fortinet support site and copy it to the computer that you use to connect to the web-based manager.
- 2 Log in to the web-based manager and locate the *License Information* widget.
- 3 In the License Information widget, in the *AV Definitions* row, select *Update*.
- 4 Select *Browse* and locate the update file, or type the path and filename.
- 5 Select *OK*.
- 6 Verify the update was successful by locating the License Information widget and viewing the date given in the row.

FortiGate unit Operation widget

The *Unit Operation* widget is an illustrated version of the FortiGate unit's front panel that shows the status of the FortiGate unit's network interfaces. The interface appears green, when the interface is connected. Hover the mouse pointer over the interface to view details about the interface.

The *Unit Operation* widget also is where you reboot or shutdown the FortiGate unit.

Icons around the front panel indicate when the FortiGate unit is connected to a FortiAnalyzer or FortiManager device, or FortiClient installations. Select the icon in the widget to jump to the configuration page for each device. When connected to one of these devices, a green check mark icon appears next to the icon. If the device communication is configured, but the device is unreachable, a red X appears.

System Resources widget

The *System Resources* widget displays basic FortiGate unit resource usage. This widget displays the information for CPU and memory in either real-time or historical data. For FortiGate units with multiple CPUs, you can view the CPU usage as an average of all CPUs or each one individually.

Use the *Refresh* icon when you want to view current system resource information, regardless of whether you are viewing real-time or historical type format.

To change the resource view from real-time to historical, or change the CPU view (for multiple CPU FortiGate units), select the *Edit* icon (visible when you hover the mouse over the widget).

When viewing CPU and memory usage in the web-based manager, only the information for core processes displays. CPU for management processes, is excluded. For example, HTTPS connections to the web-based manager.

Alert Message Console widget

Alert messages help you monitor system events on your FortiGate unit such as firmware changes, network security events, or virus detection events. Each message shows the date and time that the event occurred.

The types of messages can appear in the Alert Message Console include:

System restart	The system restarted. The restart could be due to operator action or power off/on cycling.
System shutdown	An administrator shut down the FortiGate unit from the web-based manager or CLI.
Firmware upgraded by <admin_name>	The named administrator upgraded the firmware to a more recent version on either the active or non-active partition.
Firmware downgraded by <admin_name>	The named administrator downgraded the firmware to an older version on either the active or non-active partition.
FortiGate has reached connection limit for <n> seconds	The antivirus engine was low on memory for the duration of time shown and entered conserve mode. Depending on model and configuration, content can be blocked or can pass unscanned under these conditions.

Found a new FortiAnalyzer Lost the connection to FortiAnalyzer	Shows that the FortiGate unit has either found or lost the connection to a FortiAnalyzer unit.
New firmware is available from FortiGuard	An updated firmware image is available to be downloaded to this FortiGate unit.

You can configure the alert message console settings to control what types of messages are displayed on the console.

To configure the Alert Message Console

- 1 Locate the Alert Message Console widget within the Dashboard menu.
- 2 Select the *Edit* icon in the *Alert Message Console* title bar.
- 3 Select the types of alerts that you do not want to be displayed in the widget.
- 4 Select *OK*.

Log and Archive Statistics widget

The *Log and Archive Statistics* widget displays the activity of what is DLP archiving, network traffic, and security problems including attack attempts, viruses caught, and spam email caught.

The information displayed in the *Log and Archive Statistics* widget is derived from log messages. Various configuration settings are required to collect data, as described below.

Log and Archive Statistics widget	
Since	The date and time when the counts were last reset. Counts are reset when the FortiGate unit reboots, or when you select <i>Reset</i> in the title bar area.

DLP Archive	<p>A summary of the HTTP, HTTPS, MM1, MM3, MM4, MM7, email, FTP IM, and VoIP (also called session control) traffic that has passed through the FortiGate unit, and has been archived by DLP. MM1, MM3, MM4, and MM7 are only available in FortiOS Carrier.</p> <p>This widget also Indicates the average DLP archive bytes per day since the last time it was reset.</p> <p>The <i>Details</i> pages list the last items of the selected type—up to 64 items—and provides links to the FortiAnalyzer unit where the archived traffic is stored. If logging to a FortiAnalyzer unit is not configured, the <i>Details</i> pages provide a link to <i>Log & Report > Log Config > Log Settings</i>.</p> <p>You configure the FortiGate unit to collect DLP archive data for the widget by configuring a DLP sensor to archive its log data.</p> <p>You must also add the profile to a security policy. When the security policy receives sessions for the selected protocols, meta-data is added to the statistics widget.</p> <p>In FortiOS Carrier, you can configure an MMS profile to collect statistics for MM1, MM3, MM4 and MM7 traffic.</p> <p>The Email statistics are based on email POP3, IMAP and SMTP protocols. If your FortiGate unit supports SSL content scanning and inspection, POP3S, IMAPS and SMTPS are also included.</p> <p>The IM statistics are based on the AIM, ICQ, MSN, and Yahoo! protocols and configured by selecting <i>Archive</i> in DLP Sensors for IM DLP rules.</p> <p>The VoIP statistics are based on the SIP, SIMPLE and SCCP session control protocols and configured by selecting <i>Archive</i> in DLP Sensors for Session Control DLP rules.</p>
Log	<p>A summary of traffic, viruses, attacks, spam email messages, and blocked URLs that the FortiGate unit has logged.</p> <p><i>DLP data loss detected</i> displays the number of sessions that have matched DLP sensor profiles. DLP collects meta data about all sessions matched by DLP sensors and records this meta-data in the DLP log. Every time a DLP log message is recorded, the DLP data loss detected number increases. If you are using DLP for summary or full archiving the DLP data loss detected number can get very large. This number may not indicate that data has been lost or leaked.</p>

Viewing DLP archive section of the Log and Archive Statistics widget

From the *Log and Archive Statistics* widget, you can view statistics about HTTP, HTTPS, FTP and IM traffic coming through the FortiGate unit. In FortiOS Carrier, you can view the MM1, MM3, MM4, MM7 email statistics. Select the *Details* link beside each traffic type to view more information.

DLP archive information is viewed from the DLP Archive section of the Log and Archive Statistics widget. You must select *Details* to view the available archive information.

Viewing the Log section of the Log and Archive Statistics widget

From the *Log and Archive Statistics* widget, you can view statistics about the network attacks that the FortiGate unit has stopped, statistics on viruses caught, attacks detected, spam email detected, and URLs blocked. Select the *Details* link beside each attack type to view more information.

CLI Console widget

The *CLI Console* widget enables you to access the CLI without exiting from the web-based manager.

The two controls located on the CLI Console widget title bar are *Customize*, and *Detach*.

- *Detach* moves the CLI Console widget into a pop-up window that you can resize and reposition. Select *Attach*. to move the widget back to the dashboard's page.
- *Customize* enables you to change the appearance of the console by selecting fonts and colors for the text and background.

Session History widget

The *Session History* widget displays the total session activity on the device. Activity displays on a per second basis. Select the *Edit* icon in the title bar (which appears when you hover the mouse over the widget) to change the time period for the widget.

Top Sessions widget

The *Top Sessions* widget polls the FortiGate unit for session information for IPv4 or IPv6 addresses, or both. Rebooting the FortiGate unit will reset the Top Session statistics to zero.

When you select *Details* to view the current sessions list, a list of all sessions currently processed by the FortiGate unit.

Detailed information is available in *System > Monitor > Sessions*. Use the following table to modify the default settings of the Top Sessions widget.

Traffic History widget

The *Traffic History* widget displays the traffic on one selected interface over a specified time period.

Only one interface can be monitored at a time. By default, no interface is monitored. Configure an interface to monitor by selecting the *Edit* icon in the title bar (which appears when you hover the mouse over the widget) and choosing the interface from the drop down menu. All traffic history data is cleared when you select *Apply*.

To expand the information for the widget, select *Enlarge* in the title bar area. The data will appear in a larger, pop up window.

You can modify several default settings for this widget when you select the *Edit* icon in the title bar (which appears when you hover the mouse over the widget).

RAID monitor widget

The *RAID Monitor* widget displays the current state of the RAID array and each RAID disk. This widget does not display unless the FortiGate unit has more than one disk installed, and is not available for FortiOS Carrier.

RAID monitor widget	
Configure	Select to configure the RAID array, or rebuild a degraded array.

Array Status	
Array status icon	<p>Displays the status of the RAID array.</p> <ul style="list-style-type: none"> Green with a check mark shows a healthy RAID array. Yellow triangle shows the array is in a degraded state but it is still functioning. A degraded array is slower than a healthy array. Rebuild the array to fix the degraded state. A wrench shows the array is being rebuilt. <p>Positioning the mouse over the array status icon displays a text message of the status of the array.</p>
Disk status icon	<p>There is one icon for each disk in the array.</p> <ul style="list-style-type: none"> Green with a check mark shows a healthy disk. Red with an X shows the disk has failed and needs attention. <p>Positioning the mouse over the disk status icon displays the status of the disk, and the storage capacity of the disk.</p>
RAID Level	The RAID level of this RAID array. The RAID level is set as part of configuring the RAID array.
Disk Space Usage	
Status bar	The bar shows the percentage of the RAID array that is currently in use.
Used/Free/Total	<p>Displays the amount of RAID array storage that is being used, the amount of storage that is free, and the total storage in the RAID array. The values are in GB.</p> <p><i>Used</i> added to <i>Free</i> should equal <i>Total</i>.</p>
Synchronizing status	<p>Display the percent complete of the RAID array synchronization. Synchronizing may take several hours.</p> <p>When synchronizing the status of the RAID array will indicate synchronizing is happening in the background.</p> <p>Synchronizing progress bar is visible only when the RAID array is synchronizing.</p> <p>You may need to select the refresh icon in the widget title bar to update this progress bar.</p>
Rebuild status	<p>Display the percent complete of the RAID array rebuild. Rebuilding the array may take several hours.</p> <p>While rebuilding the array, it is in a degraded and vulnerable state — any disk failure during a rebuild will result in data loss.</p> <p>A warning is displayed indicating the RAID array is running in reduced reliability mode until the rebuild is completed.</p> <p>You may need to select the refresh icon in the widget title bar to update this progress bar.</p>

RAID disk configuration

The RAID disk is configured from the Disk Configuration page.

Disk Configuration page	
RAID level	<p>Select the level of RAID. Options include:</p> <ul style="list-style-type: none"> • RAID-0 — (striping) better performance, no redundancy • RAID-1 — (mirroring) half the storage capacity, with redundancy • RAID-5 — striping with parity checking, and redundancy <p>Available RAID level options depend on the available number of hard disks. Two or more disks are required for RAID 0 or RAID 1. Three or more disks are required for RAID 5.</p> <p>Changing the RAID level will erase any stored log information on the array, and reboot the FortiGate unit. The FortiGate unit will remain offline while it reconfigures the RAID array. When it reboots, the array will need to synchronize before being fully operational.</p>
Status	<p>The status, or health, of RAID array. This status can be one of:</p> <ul style="list-style-type: none"> • OK — standard status, everything is normal • OK (Background-Synchronizing) (%) — synchronizing the disks after changing RAID level, Synchronizing progress bar shows percent complete • Degraded — One or more of the disks in the array has failed, been removed, or is not working properly. A warning is displayed about the lack of redundancy in this state. Also, a degraded array is slower than a healthy array. Select <i>Rebuild RAID</i> to fix the array. • Degraded (Background-Rebuilding) (%) — The same as degraded, but the RAID array is being rebuilt in the background. The array continues to be in a fragile state until the rebuilding is completed.
Size	<p>The size of the RAID array in gigabytes (GB). The size of the array depends on the RAID level selected, and the number of disks in the array.</p>
Rebuild RAID	<p>Select to rebuild the array after a new disk has been added to the array, or after a disk has been swapped in for a failed disk.</p> <p>If you try to rebuild a RAID array with too few disks you will get a rebuild error. After inserting a functioning disk, the rebuild will start.</p> <p>This button is only available when the RAID array is in a degraded state and has enough disks to be rebuilt.</p> <p>You cannot restart a rebuild once a rebuild is already in progress.</p> <p>Note: If a disk has failed, the number of working disks may not be enough for the RAID level to function. In this case, replace the failed disk with a working disk to rebuild the RAID array.</p>
Disk#	<p>The disk's position in the array. This corresponds to the physical slot of the disk.</p> <p>If a disk is removed from the FortiGate unit, the disk is marked as not a member of the array and its position is retained until a new disk is inserted in that drive bay.</p>

Status	The status of this disk. Options include OK, and unavailable. A disk is unavailable if it is removed or has failed.
Member	Display if the selected disk is part of the RAID array. <ul style="list-style-type: none"> A green icon with a check mark indicates the disk is part of the array. A grey icon with an X indicates the disk is not part of the RAID array. A disk may be displayed as healthy on the dashboard display even when it is not a member in the RAID array. A disk may be available but not used in the RAID array. For example three disks in a RAID 1 array, only two are used.
Capacity	The storage capacity that this drive contributes to the RAID array. The full storage capacity of the disk is used for the RAID array automatically. The total storage capacity of the RAID array depends on the capacity and numbers of the disks, and the RAID level of the array.

Top Application Usage widget

The *Top Application Usage* widget shows the volume of traffic passing through the FortiGate unit classified by application type as either a chart or a table. The chart displays applications in order of use.

From the chart or table display you can:

- View traffic volumes by pausing the mouse pointer over each bar.
- Select an application type on the graph to view information about the source addresses that used the application and the amount of data transferred by sessions from each source address.

Top application usage data collection is started by adding application control lists to security policies. Sessions accepted by security policies (with no application control list applied to that security policy) do not contribute to the data displayed.

Use the following table to modify the default settings for the Top Application Usage widget.

Storage widget

The *Storage* widget displays the status of disks currently installed on your FortiGate unit. The status includes how much space is used and how much free space is available. You can find out more detailed information about a disk's status by going to *System > Config > Disk*. The Storage page displays information regarding the disk's health, RAID events, visual representation of the disk, and configuration of the management of the disk.

P2P Usage widget

The *P2P Usage* widget displays the total bytes and total bandwidth for each supported instant messaging client. These clients are WinNY, BitTorrent, eDonkey, Guntella, and KaZaa. With P2P Usage, you can only modify the default name of the widget.

Per-IP Bandwidth Usage widget

The *Per-IP Bandwidth Usage* widget displays the per-IP address session data. The data, displays each IP address that initiated the traffic (and its current bandwidth consumption), and is similar to the top session widget. Instead of viewing the IP address of the person who initiated the traffic, you can choose to view their name by selecting *Resolve Host Name* in the editing window.

VoIP Usage widget

The *VoIP Usage* widget displays current active VoIP call information (using over SIP and SCCP protocols), which include complete calls, calls that have been dropped, failed or went unanswered.

IM Usage widget

The *IM Usage* widget displays instant messaging clients and their activity that is occurring on your network, including chats, messages, file transfer between clients, and any voice chats. IM Usage provides this information for IM, Yahoo!, AIM, and ICQ.

Network Protocol Usage

The *Network Protocol Usage* widget displays protocol activity over a defined time period and the amount of bandwidth used during the activity.

Basic configurations

Before going ahead and configuring security policies, users, and UTM profiles, you should perform some basic configurations to set up your FortiGate unit.

Changing your administrator password (best practices)

By default, you can log in to the web-based manager by using the admin administrator account and no password. It is highly recommended that you add a password to the admin administrator account. For improved security, you should regularly change the admin administrator account password and the passwords for any other administrator accounts that you add.

To change an administrator's password, go to *System > Admin > Administrators*, edit the administrator account, and then change the password.

For details on selecting a password, and password best practices, see [“Passwords” on page 348](#).



If you forget or lose an administrator account password and cannot log in to the unit, see the Fortinet Knowledge Base article [Recovering a lost FortiGate administrator account password](#).

Changing the web-based manager language

The default language of the web-based manager is English. A selection of localized iterations are available to selected from. For best results, you should select the language that the management computer operating system uses.

To change the language, go to *System > Admin > Settings*. In the *Display Settings* section, select the language you want from the *Language* drop-down list.

Changing administrative access

Through administrative access, an administrator can connect to the FortiGate unit. Access is available through a number of services including HTTPS and SSH. The default configuration allows administrative access to one or more of the unit's interfaces as described in the [QuickStart Guide](#).

To change administrative access

- 1 Go to *System > Network > Interface*.
- 2 Select the interface.
- 3 Select the administrative access type or types for that interface.
- 4 Select OK.

Changing the web-based manager idle timeout

By default, the web-based manager disconnects administrative sessions if no activity occurs for five minutes. This prevents someone from using the web-based manager if the management PC is left unattended.

To change the idle timeout

- 1 Go to *System > Admin > Settings*.
- 2 In the *Administration Settings* section, enter the time in minutes in the *Idle Timeout* field
- 3 Select *Apply*.

Switching VDOMs

When VDOMs are enabled, a menu appears in the left column called *Current VDOM*. This menu displays a drop-down list that lists the configured VDOMs.

To switch to a VDOM using the *Current VDOM* menu, select the VDOM that you want to switch to from the drop-down list. You are automatically redirected to that VDOM.

VDOMs are enabled on the *System Information* Dashboard Widget.

Connecting to the CLI from the web-based manager

You can use the CLI to configure all configuration options available from the web-based manager. Some configuration options are available only from the CLI.

To connect to the CLI console, go to *System > Dashboard > Status*, and in the CLI Console widget select inside the window, and are automatically logged in to the CLI. For more information on using the CLI, see [“Using the CLI” on page 315](#).

Logging out

Select the Logout icon to quit your administrative session. If you only close the browser or leave the web-based manager to surf to another web site, you remain logged in until the idle timeout (default 5 minutes) expires. To change the timeout, see [“Changing the web-based manager idle timeout” on page 313](#).



Using the CLI

The command line interface (CLI) is an alternative configuration tool to the web-based manager.

Both can be used to configure the FortiGate unit. While the configuration, in the web-based manager, a point-and-click method, the CLI, would require typing commands, or upload batches of commands from a text file, like a configuration script.

If you are new to Fortinet products, or if you are new to the CLI, this section can help you to become familiar.

This section includes the topics:

- [Connecting to the CLI](#)
- [Command syntax](#)
- [Sub-commands](#)
- [Permissions](#)
- [Tips](#)

Connecting to the CLI

You can access the CLI in two ways:

- Locally — Connect your computer directly to the FortiGate unit's console port.
- Through the network — Connect your computer through any network attached to one of the FortiGate unit's network ports. The network interface must have enabled Telnet or SSH administrative access if you will connect using an SSH/Telnet client, or HTTP/HTTPS administrative access if you will connect using the *CLI Console* widget in the web-based manager.

Local access is required in some cases.

- If you are installing your FortiGate unit for the first time and it is not yet configured to connect to your network, unless you reconfigure your computer's network settings for a peer connection, you may only be able to connect to the CLI using a local serial console connection. For more information, see [“Connecting to the CLI” on page 336](#).
- Restoring the firmware utilizes a boot interrupt. Network access to the CLI is not available until after the boot process has completed, and therefore local CLI access is the only viable option.

Before you can access the CLI through the network, you usually must enable SSH and/or Telnet on the network interface through which you will access the CLI.

Connecting to the CLI using a local console

Local console connections to the CLI are formed by directly connecting your management computer or console to the FortiGate unit, using its DB-9 or RJ-45 console port. To connect to the local console you need:

- a computer with an available serial communications (COM) port
- the RJ-45-to-DB-9 or null modem cable included in your FortiGate package
- terminal emulation software such as HyperTerminal for Microsoft Windows



The following procedure describes connection using Microsoft HyperTerminal software; steps may vary with other terminal emulators.

To connect to the CLI using a local serial console connection

- 1 Using the null modem or RJ-45-to-DB-9 cable, connect the FortiGate unit's console port to the serial communications (COM) port on your management computer.
- 2 On your management computer, start HyperTerminal.
- 3 For the *Connection Description*, enter a *Name* for the connection, and select *OK*.
- 4 On the *Connect using* drop-down list box, select the communications (COM) port on your management computer you are using to connect to the FortiGate unit.
- 5 Select *OK*.
- 6 Select the following *Port* settings and select *OK*.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 7 Press Enter or Return on your keyboard to connect to the CLI.
- 8 Type a valid administrator account name (such as `admin`) and press Enter.
- 9 Type the password for that administrator account and press Enter. (In its default state, there is no password for the `admin` account.)

The CLI displays the following text:

```
Welcome!
```

```
Type ? to list available commands.
```

You can now enter CLI commands, including configuring access to the CLI through SSH or Telnet. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)”](#) on page 317.

Enabling access to the CLI through the network (SSH or Telnet)

SSH or Telnet access to the CLI is accomplished by connecting your computer to the FortiGate unit using one of its RJ-45 network ports. You can either connect directly, using a peer connection between the two, or through any intermediary network.



If you do not want to use an SSH/Telnet client and you have access to the web-based manager, you can alternatively access the CLI through the network using the *CLI Console* widget in the web-based manager.

You must enable SSH and/or Telnet on the network interface associated with that physical network port. If your computer is not connected directly or through a switch, you must also configure the FortiGate unit with a static route to a router that can forward packets from the FortiGate unit to your computer. You can do this using either a local console connection or the web-based manager.

Requirements

- a computer with an available serial communications (COM) port and RJ-45 port
- terminal emulation software such as HyperTerminal for Microsoft Windows
- the RJ-45-to-DB-9 or null modem cable included in your FortiGate package
- a network cable
- prior configuration of the operating mode, network interface, and static route (for details, see)

To enable SSH or Telnet access to the CLI using a local console connection

- 1 Using the network cable, connect the FortiGate unit's network port either directly to your computer's network port, or to a network through which your computer can reach the FortiGate unit.
- 2 Note the number of the physical network port.
- 3 Using a local console connection, connect and log into the CLI. For details, see [“Connecting to the CLI using a local console” on page 316](#).
- 4 Enter the following command:

```
config system interface
  edit <interface_str>
    set allowaccess <protocols_list>
  next
```

```
end
```

where:

- `<interface_str>` is the name of the network interface associated with the physical network port and containing its number, such as `port1`
- `<protocols_list>` is the complete, space-delimited list of permitted administrative access protocols, such as `https ssh telnet`

For example, to exclude HTTP, HTTPS, SNMP, and PING, and allow only SSH and Telnet administrative access on `port1`:

```
set system interface port1 config allowaccess ssh telnet
```



Telnet is not a secure access method. SSH should be used to access the CLI from the Internet or any other untrusted network.

- 5 To confirm the configuration, enter the command to display the network interface's settings.

```
get system interface <interface_str>
```

The CLI displays the settings, including the allowed administrative access protocols, for the network interfaces.

To connect to the CLI through the network interface, see [“Connecting to the CLI using SSH” on page 318](#) or [“Connecting to the CLI using Telnet” on page 319](#).

Connecting to the CLI using SSH

Once the FortiGate unit is configured to accept SSH connections, you can use an SSH client on your management computer to connect to the CLI.

Secure Shell (SSH) provides both secure authentication and secure communications to the CLI. FortiGate units support 3DES and Blowfish encryption algorithms for SSH.

Before you can connect to the CLI using SSH, you must first configure a network interface to accept SSH connections. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)” on page 317](#). The following procedure uses PuTTY. Steps may vary with other SSH clients.

To connect to the CLI using SSH

- 1 On your management computer, start an SSH client.
- 2 In *Host Name (or IP Address)*, enter the IP address of a network interface on which you have enabled SSH administrative access.
- 3 In *Port*, enter `22`.
- 4 For the *Connection type*, select *SSH*.
- 5 Select *Open*.

The SSH client connects to the FortiGate unit.

The SSH client may display a warning if this is the first time you are connecting to the FortiGate unit and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiGate unit but it used a different IP address or SSH key. If your management computer is directly connected to the FortiGate unit with no network hosts between them, this is normal.

- 6 Click Yes to verify the fingerprint and accept the FortiGate unit's SSH key. You will not be able to log in until you have accepted the key.
The CLI displays a login prompt.
- 7 Type a valid administrator account name (such as `admin`) and press Enter.
- 8 Type the password for this administrator account and press Enter.



If three incorrect login or password attempts occur in a row, you will be disconnected. Wait one minute, then reconnect to attempt the login again.

The FortiGate unit displays a command prompt (its host name followed by a `#`).
You can now enter CLI commands.

Connecting to the CLI using Telnet

Once the FortiGate unit is configured to accept Telnet connections, you can use a Telnet client on your management computer to connect to the CLI.



Telnet is not a secure access method. SSH should be used to access the CLI from the Internet or any other untrusted network.

Before you can connect to the CLI using Telnet, you must first configure a network interface to accept SSH connections. For details, see [“Enabling access to the CLI through the network \(SSH or Telnet\)” on page 317](#).

To connect to the CLI using Telnet

- 1 On your management computer, start a Telnet client.
- 2 Connect to a FortiGate network interface on which you have enabled Telnet.
- 3 Type a valid administrator account name (such as `admin`) and press Enter.
- 4 Type the password for this administrator account and press Enter.



If three incorrect login or password attempts occur in a row, you will be disconnected. Wait one minute, then reconnect to attempt the login again.

The FortiGate unit displays a command prompt (its host name followed by a `#`).
You can now enter CLI commands.

Command syntax

When entering a command, the command line interface (CLI) requires that you use valid syntax, and conform to expected input constraints. It will reject invalid commands.

Fortinet documentation uses the following conventions to describe valid command syntax

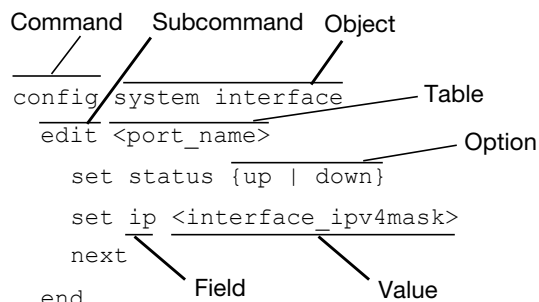
Terminology

Each command line consists of a command word that is usually followed by words for the configuration data or other specific item that the command uses or affects:

```
get system admin
```

To describe the function of each word in the command line, especially if that nature has changed between firmware versions, Fortinet uses terms with the following definitions.

Figure 35: Command syntax terminology



- **command** — A word that begins the command line and indicates an action that the FortiGate unit should perform on a part of the configuration or host on the network, such as `config` or `execute`. Together with other words, such as fields or values, that end when you press the Enter key, it forms a command line. Exceptions include multiline command lines, which can be entered using an escape sequence. (See [“Shortcuts and key commands” on page 327.](#))

Valid command lines must be unambiguous if abbreviated. (See [“Command abbreviation” on page 328.](#)) Optional words or other command line permutations are indicated by syntax notation. (See [“Notation” on page 321.](#))

- **sub-command** — A kind of command that is available only when nested within the scope of another command. After entering a command, its applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into another sub-command. Indentation is used to indicate levels of nested commands. (See [“Indentation” on page 321.](#))

Not all top-level commands have sub-commands. Available sub-commands vary by their containing scope. (See [“Sub-commands” on page 323.](#))

- **object** — A part of the configuration that contains tables and/or fields. Valid command lines must be specific enough to indicate an individual object.
- **table** — A set of fields that is one of possibly multiple similar sets which each have a name or number, such as an administrator account, policy, or network interface. These named or numbered sets are sometimes referenced by other parts of the configuration that use them. (See [“Notation” on page 321.](#))
- **field** — The name of a setting, such as `ip` or `hostname`. Fields in some tables must be configured with values. Failure to configure a required field will result in an invalid object configuration error message, and the FortiGate unit will discard the invalid table.
- **value** — A number, letter, IP address, or other type of input that is usually your configuration setting held by a field. Some commands, however, require multiple input values which may not be named but are simply entered in sequential order in the same command line. Valid input types are indicated by constraint notation. (See [“Notation” on page 321.](#))
- **option** — A kind of value that must be one or more words from of a fixed set of options. (See [“Notation” on page 321.](#))

Indentation

Indentation indicates levels of nested commands, which indicate what other subcommittees are available from within the scope. For example, the `edit` sub-command is available only within a command that affects tables, and the `next` sub-command is available only from within the `edit` sub-command:

```
config system interface
  edit port1
    set status up
  next
end
```

For information about available sub-commands, see [“Sub-commands” on page 323](#).

Notation

Brackets, braces, and pipes are used to denote valid permutations of the syntax. Constraint notations, such as `<address_ipv4>`, indicate which data types or string patterns are acceptable value input.

Table 11: Command syntax notation

Convention	Description
Square brackets []	A non-required word or series of words. For example: [verbose {1 2 3}] indicates that you may either omit or type both the <code>verbose</code> word and its accompanying option, such as <code>verbose 3</code> .

Table 11: Command syntax notation

Angle brackets < >	<p>A word constrained by data type. The angled brackets contain a descriptive name followed by an underscore (_) and suffix that indicates the valid data type. For example, <retries_int>, indicates that you should enter a number of retries, such as 5.</p> <p>Data types include:</p> <ul style="list-style-type: none"> • <xxx_name>: A name referring to another part of the configuration, such as policy_A. • <xxx_index>: An index number referring to another part of the configuration, such as 0 for the first static route. • <xxx_pattern>: A regular expression or word with wild cards that matches possible variations, such as *@example.com to match all email addresses ending in @example.com. • <xxx_fqdn>: A fully qualified domain name (FQDN), such as mail.example.com. • <xxx_email>: An email address, such as admin@example.com. • <xxx_ipv4>: An IPv4 address, such as 192.168.1.99. • <xxx_v4mask>: A dotted decimal IPv4 netmask, such as 255.255.255.0. • <xxx_ipv4mask>: A dotted decimal IPv4 address and netmask separated by a space, such as 192.168.1.99 255.255.255.0. • <xxx_ipv4/mask>: A dotted decimal IPv4 address and CIDR-notation netmask separated by a slash, such as 192.168.1.1/24. • <xxx_ipv4range>: A hyphen (-)-delimited inclusive range of IPv4 addresses, such as 192.168.1.1-192.168.1.255. • <xxx_ipv6>: A colon (:)-delimited hexadecimal IPv6 address, such as 3f2e:6a8b:78a3:0d82:1725:6a2f:0370:6234. • <xxx_v6mask>: An IPv6 netmask, such as /96. • <xxx_ipv6mask>: A dotted decimal IPv6 address and netmask separated by a space. • <xxx_str>: A string of characters that is not another data type, such as P@ssw0rd. Strings containing spaces or special characters must be surrounded in quotes or use escape sequences. See “Special characters” on page 328. • <xxx_int>: An integer number that is not another data type, such as 15 for the number of minutes.
Curly braces { }	<p>A word or series of words that is constrained to a set of options delimited by either vertical bars or spaces. You must enter at least one of the options, unless the set of options is surrounded by square brackets [].</p>
Options delimited by vertical bars	<p>Mutually exclusive options. For example:</p> <pre>{enable disable}</pre> <p>indicates that you must enter either enable or disable, but must not enter both.</p>

Table 11: Command syntax notation

Options delimited by spaces	<p>Non-mutually exclusive options. For example:</p> <pre>{http https ping snmp ssh telnet}</pre> <p>indicates that you may enter all or a subset of those options, in any order, in a space-delimited list, such as:</p> <pre>ping https ssh</pre> <p>Note: To change the options, you must re-type the entire list. For example, to add <code>snmp</code> to the previous example, you would type:</p> <pre>ping https snmp ssh</pre> <p>If the option adds to or subtracts from the existing list of options, instead of replacing it, or if the list is comma-delimited, the exception will be noted.</p>
------------------------------------	---

Sub-commands

Each command line consists of a command word that is usually followed by words for the configuration data or other specific item that the command uses or affects:

```
get system admin
```

Sub-commands are available from within the scope of some commands. When you enter a sub-command level, the command prompt changes to indicate the name of the current command scope. For example, after entering:

```
config system admin
```

the command prompt becomes:

```
(admin) #
```

Applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into another sub-command.

For example, the `edit` sub-command is available only within a command that affects tables; the `next` sub-command is available only from within the `edit` sub-command:

```
config system interface
  edit port1
    set status up
  next
end
```

Sub-command scope is indicated by indentation. See [“Indentation” on page 321](#).

Available sub-commands vary by command. From a command prompt within `config`, two types of sub-commands might become available:

- commands affecting fields

- commands affecting tables

Table 12: Commands for tables

clone <table>	<p>Clone (or make a copy of) a table from the current object.</p> <p>For example, in <code>config firewall policy</code>, you could enter the following command to clone security policy 27 to create security policy 30:</p> <pre>clone 27 to 39</pre> <p>In <code>config antivirus profile</code>, you could enter the following command to clone an antivirus profile named <code>av_pro_1</code> to create a new antivirus profile named <code>av_pro_2</code>:</p> <pre>clone av_pro_1 to av_pro_2</pre> <p><code>clone</code> may not be available for all tables.</p>
delete <table>	<p>Remove a table from the current object.</p> <p>For example, in <code>config system admin</code>, you could delete an administrator account named <code>newadmin</code> by typing <code>delete newadmin</code> and pressing Enter. This deletes <code>newadmin</code> and all its fields, such as <code>newadmin</code>'s first-name and email-address.</p> <p><code>delete</code> is only available within objects containing tables.</p>
edit <table>	<p>Create or edit a table in the current object.</p> <p>For example, in <code>config system admin</code>:</p> <ul style="list-style-type: none"> • edit the settings for the default <code>admin</code> administrator account by typing <code>edit admin</code>. • add a new administrator account with the name <code>newadmin</code> and edit <code>newadmin</code>'s settings by typing <code>edit newadmin</code>. <p><code>edit</code> is an interactive sub-command: further sub-commands are available from within <code>edit</code>.</p> <p><code>edit</code> changes the prompt to reflect the table you are currently editing.</p> <p><code>edit</code> is only available within objects containing tables.</p> <p>In objects such as security policies, <code><table></code> is a sequence number. To create a new entry without the risk of overwriting an existing one, enter <code>edit 0</code>. The CLI initially confirms the creation of entry 0, but assigns the next unused number after you finish editing and enter <code>end</code>.</p>
end	<p>Save the changes to the current object and exit the <code>config</code> command. This returns you to the top-level command prompt.</p>
get	<p>List the configuration of the current object or table.</p> <ul style="list-style-type: none"> • In objects, <code>get</code> lists the table names (if present), or fields and their values. • In a table, <code>get</code> lists the fields and their values. <p>For more information on <code>get</code> commands, see the CLI Reference.</p>

Table 12: Commands for tables

<i>purge</i>	<p>Remove all tables in the current object.</p> <p>For example, in <code>config forensic user</code>, you could type <code>get</code> to see the list of user names, then type <code>purge</code> and then <code>y</code> to confirm that you want to delete all users.</p> <p><code>purge</code> is only available for objects containing tables.</p> <p>Caution: Back up the FortiGate unit before performing a <code>purge</code>. <code>purge</code> cannot be undone. To restore purged tables, the configuration must be restored from a backup.</p> <p>Caution: Do not <code>purge system interface</code> or <code>system admin</code> tables. <code>purge</code> does not provide default tables. This can result in being unable to connect or log in, requiring the FortiGate unit to be formatted and restored.</p>
<i>rename <table> to <table></i>	<p>Rename a table.</p> <p>For example, in <code>config system admin</code>, you could rename <code>admin3</code> to <code>fwadmin</code> by typing <code>rename admin3 to fwadmin</code>.</p> <p><code>rename</code> is only available within objects containing tables.</p>
<i>show</i>	<p>Display changes to the default configuration. Changes are listed in the form of configuration commands.</p>

Example of table commands

From within the `system admin` object, you might enter:

```
edit admin_1
```

The CLI acknowledges the new table, and changes the command prompt to show that you are now within the `admin_1` table:

```
new entry 'admin_1' added
(admin_1)#
```

Table 13: Commands for fields

<i>abort</i>	Exit both the <code>edit</code> and/or <code>config</code> commands without saving the fields.
<i>end</i>	Save the changes made to the current table or object fields, and exit the <code>config</code> command. (To exit without saving, use <code>abort</code> instead.)
<i>get</i>	<p>List the configuration of the current object or table.</p> <ul style="list-style-type: none"> In objects, <code>get</code> lists the table names (if present), or fields and their values. In a table, <code>get</code> lists the fields and their values.
<i>next</i>	<p>Save the changes you have made in the current table's fields, and exit the <code>edit</code> command to the object prompt. (To save and exit completely to the root prompt, use <code>end</code> instead.)</p> <p><code>next</code> is useful when you want to create or edit several tables in the same object, without leaving and re-entering the <code>config</code> command each time.</p> <p><code>next</code> is only available from a table prompt; it is not available from an object prompt.</p>

Table 13: Commands for fields

<code>set <field></code> <code><value></code>	<p>Set a field's value.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, you could type <code>set password newpass</code> to change the password of the <code>admin</code> administrator to <code>newpass</code>.</p> <p>Note: When using <code>set</code> to change a field containing a space-delimited list, type the whole new list. For example, <code>set <field> <new-value></code> will replace the list with the <code><new-value></code> rather than appending <code><new-value></code> to the list.</p>
<code>show</code>	Display changes to the default configuration. Changes are listed in the form of configuration commands.
<code>unset <field></code>	<p>Reset the table or object's fields to default values.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, typing <code>unset password</code> resets the password of the <code>admin</code> administrator account to the default (in this case, no password).</p>

Example of field commands

From within the `admin_1` table, you might enter:

```
set password my1stExamplePassword
```

to assign the value `my1stExamplePassword` to the `password` field. You might then enter the `next` command to save the changes and edit the next administrator's table.

Permissions

Depending on the account that you use to log in to the FortiGate unit, you may not have complete access to all CLI commands.

Access profiles control which CLI commands an administrator account can access.

Access profiles assign either read, write, or no access to each area of the FortiGate software. To view configurations, you must have read access. To make changes, you must have write access.

Unlike other administrator accounts, the administrator account named `admin` exists by default and cannot be deleted. The `admin` administrator account is similar to a root administrator account. This administrator account always has full permission to view and change all FortiGate configuration options, including viewing and changing **all** other administrator accounts. Its name and permissions cannot be changed. It is the only administrator account that can reset another administrator's password without being required to enter that administrator's existing password.



Set a strong password for the `admin` administrator account, and change the password regularly. By default, this administrator account has no password. Failure to maintain the password of the `admin` administrator account could compromise the security of your FortiGate unit.

For complete access to all commands, you must log in with the administrator account named `admin`.

Tips

Basic features and characteristics of the CLI environment provide support and ease of use for many CLI tasks.

Help

To display brief help during command entry, press the question mark (?) key.

- Press the question mark (?) key at the command prompt to display a list of the commands available and a description of each command.
- Type a word or part of a word, then press the question mark (?) key to display a list of valid word completions or subsequent words, and to display a description of each.

Shortcuts and key commands

Table 14: Shortcuts and key commands

Action	Keys
List valid word completions or subsequent words. If multiple words could complete your entry, display all possible completions with helpful descriptions of each.	?
Complete the word with the next available match. Press the key multiple times to cycle through available matches.	Tab
Recall the previous command. Command memory is limited to the current session.	Up arrow, or Ctrl + P
Recall the next command.	Down arrow, or Ctrl + N
Move the cursor left or right within the command line.	Left or Right arrow
Move the cursor to the beginning of the command line.	Ctrl + A
Move the cursor to the end of the command line.	Ctrl + E
Move the cursor backwards one word.	Ctrl + B
Move the cursor forwards one word.	Ctrl + F
Delete the current character.	Ctrl + D
Abort current interactive commands, such as when entering multiple lines. If you are not currently within an interactive command such as <code>config</code> or <code>edit</code> , this closes the CLI connection.	Ctrl + C
Continue typing a command on the next line for a multi-line command. For each line that you want to continue, terminate it with a backslash (\). To complete the command line, terminate it by pressing the spacebar and then the Enter key, without an immediately preceding backslash.	\ then Enter

Command abbreviation

You can abbreviate words in the command line to their smallest number of non-ambiguous characters.

For example, the command `get system status` could be abbreviated to `g sy st`.

Environment variables

The CLI supports the following environment variables. Variable names are case-sensitive.

\$USERFROM	The management access type (<code>ssh</code> , <code>telnet</code> , <code>jsconsole</code> for the <i>CLI Console</i> widget in the web-based manager, and so on) and the IP address of the administrator that configured the item.
\$USERNAME	The account name of the administrator that configured the item.
\$SerialNum	The serial number of the FortiGate unit.

For example, the FortiGate unit's host name can be set to its serial number.

```
config system global
  set hostname $SerialNum
end
```

As another example, you could log in as `admin1`, then configure a restricted secondary administrator account for yourself named `admin2`, whose `first-name` is `admin1` to indicate that it is another of your accounts:

```
config system admin
  edit admin2
    set first-name $USERNAME
```

Special characters

The characters `<`, `>`, `(,)`, `#`, `'`, and `"` are not permitted in most CLI fields. These characters are special characters, sometimes also called reserved characters.

You may be able to enter special character as part of a string's value by using a special command, enclosing it in quotes, or preceding it with an escape sequence — in this case, a backslash (`\`) character.

Table 15: Entering special characters

Character	Keys
<code>?</code>	Ctrl + V then <code>?</code>
Tab	Ctrl + V then Tab
Space (to be interpreted as part of a string value, not to end the string)	Enclose the string in quotation marks: <code>"Security Administrator"</code> . Enclose the string in single quotes: <code>'Security Administrator'</code> . Precede the space with a backslash: <code>Security\ Administrator</code> .
<code>'</code> (to be interpreted as part of a string value, not to end the string)	<code>\'</code>

Table 15: Entering special characters

" (to be interpreted as part of a string value, not to end the string)	\"
\	\\

If you need to add configuration via CLI that requires ? as part of config, you need to input CTRL-V first. If you enter the question mark (?) without first using CTRL-V, the question mark has a different meaning in CLI: it will show available command options in that section.

For example, if you enter ? without CTRL-V:

```
edit "*.xe
token line: Unmatched double quote.
```

If you enter ? with CTRL-V:

```
edit "*.xe?"
new entry '*.xe?' added
```

Using grep to filter get and show command output

In many cases the `get` and `show` (and `diagnose`) commands may produce a large amount of output. If you are looking for specific information in a large `get` or `show` command output you can use the `grep` command to filter the output to only display what you are looking for. The `grep` command is based on the standard UNIX `grep`, used for searching text output based on regular expressions.

Information about how to use `grep` and regular expressions is available on the Internet, just to a search for `grep`. For example, see

<http://www.opengroup.org/onlinepubs/009695399/utilities/grep.html>.

Use the following command to display the MAC address of the FortiGate unit internal interface:

```
get hardware nic internal | grep Current_HWaddr
Current_HWaddr                00:09:0f:cb:c2:75
```

Use the following command to display all TCP sessions in the session list and include the session list line number in the output

```
get system session list | grep -n tcp
```

Use the following command to display all lines in HTTP replacement message commands that contain URL (upper or lower case):

```
show system replacemsg http | grep -i url
```

Language support and regular expressions

Characters such as ñ, é, symbols, and ideographs are sometimes acceptable input. Support varies by the nature of the item being configured. CLI commands, objects, field names, and options must use their exact ASCII characters, but some items with arbitrary names or values may be input using your language of choice.

For example, the host name must not contain special characters, and so the web-based manager and CLI will not accept most symbols and other non-ASCII encoded characters as input when configuring the host name. This means that languages other than English often are not supported. However, some configuration items, such as names and comments, may be able to use the language of your choice.

To use other languages in those cases, you must use the correct encoding.

Input is stored using Unicode UTF-8 encoding, but is not normalized from other encodings into UTF-8 before it is stored. If your input method encodes some characters differently than in UTF-8, your configured items may not display or operate as expected.

Regular expressions are especially impacted. Matching uses the UTF-8 character values. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, matches may not be what you expect.

For example, with Shift-JIS, backslashes (\) could be inadvertently interpreted as yen symbols (¥) and vice versa. A regular expression intended to match HTTP requests containing money values with a yen symbol therefore may not work if the symbol is entered using the wrong encoding.

For best results, you should:

- use UTF-8 encoding, or
- use only the characters whose numerically encoded values are the same in UTF-8, such as the US-ASCII characters that are also encoded using the same values in ISO 8859-1, Windows code page 1252, Shift-JIS and other encodings, or
- for regular expressions that must match HTTP requests, use the same encoding as your HTTP clients



HTTP clients may send requests in encodings other than UTF-8. Encodings usually vary by the client's operating system or input language. If you cannot predict the client's encoding, you may only be able to match any parts of the request that are in English, because regardless of the encoding, the values for English characters tend to be encoded identically. For example, English words may be legible regardless of interpreting a web page as either ISO 8859-1 or as GB2312, whereas simplified Chinese characters might only be legible if the page is interpreted as GB2312.

If you configure your FortiGate unit using other encodings, you may need to switch language settings on your management computer, including for your web browser or Telnet/SSH client. For instructions on how to configure your management computer's operating system language, locale, or input method, see its documentation.



If you choose to configure parts of the FortiGate unit using non-ASCII characters, verify that all systems interacting with the FortiGate unit also support the same encodings. You should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of the web-based manager and your web browser or Telnet/SSH client while you work.


Similarly to input, your web browser or CLI client should usually interpret display output as encoded using UTF-8. If it does not, your configured items may not display correctly in the web-based manager or CLI. Exceptions include items such as regular expressions that you may have configured using other encodings in order to match the encoding of HTTP requests that the FortiGate unit receives.

To enter non-ASCII characters in the CLI Console widget

- 1 On your management computer, start your web browser and go to the URL for the FortiGate unit's web-based manager.
- 2 Configure your web browser to interpret the page as UTF-8 encoded.
- 3 Log in to the FortiGate unit.
- 4 Go to *System > Dashboard > Status*.

- 5 In title bar of the *CLI Console* widget, click *Edit* (the pencil icon).
- 6 Enable *Use external command input box*.
- 7 Select *OK*.
The *Command* field appears below the usual input and display area of the *CLI Console* widget.
- 8 In *Command*, type a command.

Figure 36: Entering encoded characters (CLI Console widget)



- 9 Press Enter.
In the display area, the *CLI Console* widget displays your previous command interpreted into its character code equivalent, such as:



```
edit \743\601\613\743\601\652
```

 and the command's output.

To enter non-ASCII characters in a Telnet/SSH client

- 1 On your management computer, start your Telnet or SSH client.
- 2 Configure your Telnet or SSH client to send and receive characters using UTF-8 encoding.
Support for sending and receiving international characters varies by each Telnet/SSH client. Consult the documentation for your Telnet/SSH client.
- 3 Log in to the FortiGate unit.
- 4 At the command prompt, type your command and press Enter.

Figure 37: Entering encoded characters (PuTTY)



You may need to surround words that use encoded characters with single quotes ('). Depending on your Telnet/SSH client's support for your language's input methods and for sending international characters, you may need to interpret them into character codes before pressing Enter.

For example, you might need to enter:

```
edit '\743\601\613\743\601\652'
```

- 5 The CLI displays your previous command and its output.

Screen paging

You can configure the CLI to, when displaying multiple pages' worth of output, pause after displaying each page's worth of text. When the display pauses, the last line displays `--More--`. You can then either:

- press the spacebar to display the next page.
- type `Q` to truncate the output and return to the command prompt.

This may be useful when displaying lengthy output, such as the list of possible matching commands for command completion, or a long list of settings. Rather than scrolling through or possibly exceeding the buffer of your terminal emulator, you can simply display one page at a time.

To configure the CLI display to pause when the screen is full:

```
config system console
  set output more
end
```

Baud rate

You can change the default baud rate of the local console connection.

To change the baud rate enter the following commands:

```
config system console
  set baudrate {115200 | 19200 | 38400 | 57600 | 9600}
end
```

Editing the configuration file on an external host

You can edit the FortiGate configuration on an external host by first backing up the configuration file to a TFTP server. Then edit the configuration file and restore it to the FortiGate unit.

Editing the configuration on an external host can be time-saving if you have many changes to make, especially if your plain text editor provides advanced features such as batch changes.

To edit the configuration on your computer

- 1 Use `execute backup` to download the configuration file to a TFTP server, such as your management computer.
- 2 Edit the configuration file using a plain text editor that supports Unix-style line endings.



Do not edit the first line. The first line(s) of the configuration file (preceded by a `#` character) contains information about the firmware version and FortiGate model. If you change the model number, the FortiGate unit will reject the configuration file when you attempt to restore it.

- 3 Use `execute restore` to upload the modified configuration file back to the FortiGate unit.

The FortiGate unit downloads the configuration file and checks that the model information is correct. If it is, the FortiGate unit loads the configuration file and checks each command for errors. If a command is invalid, the FortiGate unit ignores the command. If the configuration file is valid, the FortiGate unit restarts and loads the new configuration.

Using Perl regular expressions

Some FortiGate features, such as spam filtering and web content filtering can use either wildcards or Perl regular expressions.

See <http://perl.doc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.

For more information on using Perl expressions see the *UTM* chapter of *The Handbook*.

Differences between regular expression and wildcard pattern matching

In Perl regular expressions, the period (‘.’) character refers to any single character. It is similar to the question mark (‘?’) character in wildcard pattern matching. As a result:

- `fortinet.com` not only matches `example.com` but also matches `exampleacom`, `examplebcom`, `exampleccom` and so on.

To match a special character such as the period (‘.’) and the asterisk (‘*’), regular expressions use the slash (‘\’) escape character. For example:

- To match `example.com`, the regular expression should be `example\.com`.

In Perl regular expressions, the asterisk (‘*’) means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- `exam*.com` matches `examiiii.com` but does not match `eample.com`.

To match any character 0 or more times, use ‘.’ where ‘.’ means any character and the ‘*’ means 0 or more times. For example:

- the wildcard match pattern `exam*.com` is equivalent to the regular expression `exam.*\.com`.

Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression “test” not only matches the word “test” but also matches any word that contains the word “test” such as “atest”, “mytest”, “testimony”, “atestb”. The notation “\b” specifies the word boundary. To match exactly the word “test”, the expression should be `\btest\b`.

Case sensitivity

Regular expression pattern matching is case sensitive in the Web and Spam filters. To make a word or phrase case insensitive, use the regular expression `/i`. For example, `/bad language/i` will block all instances of “bad language” regardless of case.

Table 16: Perl regular expression examples

Expression	Matches
<code>abc</code>	abc (that exact character sequence, but anywhere in the string)
<code>^abc</code>	abc at the beginning of the string

Table 16: Perl regular expression examples

abc\$	abc at the end of the string
a b	either of a and b
^abc abc\$	the string abc at the beginning or at the end of the string
ab{2,4}c	an a followed by two, three or four b's followed by a c
ab{2,}c	an a followed by at least two b's followed by a c
ab*c	an a followed by any number (zero or more) of b's followed by a c
ab+c	an a followed by one or more b's followed by a c
ab?c	an a followed by an optional b followed by a c; that is, either abc or ac
a.c	an a followed by any single character (not newline) followed by a c
a\.c	a.c exactly
[abc]	any one of a, b and c
[Aa]bc	either of Abc and abc
[abc]+	any (nonempty) string of a's, b's and c's (such as a, abba, acbabacaca)
[^abc]+	any (nonempty) string which does not contain any of a, b and c (such as defg)
\d\d	any two decimal digits, such as 42; same as \d{2}
/i	makes the pattern case insensitive. For example, <code>/bad language/i</code> blocks any instance of "bad language" regardless of case.
\w+	a "word": a nonempty sequence of alphanumeric characters and low lines (underscores), such as foo and 12bar8 and foo_1
100\s*mk	the strings 100 and mk optionally separated by any amount of white space (spaces, tabs, newlines)
abc\b	abc when followed by a word boundary (e.g. in abc! but not in abcd)
perl\b	perl when not followed by a word boundary (e.g. in perlert but not in perl stuff)
\x	tells the regular expression parser to ignore white space that is neither backslashed nor within a character class. You can use this to break up your regular expression into (slightly) more readable parts.



Basic setup

The FortiGate unit requires some basic configuration to add it to your network. These basic steps include assigning IP addresses, adding routing and security policies. Until the administrator completes these steps inter-network and internet traffic will not flow through the device.

There are two methods of configuring the FortiGate unit: either the web-based manager or the command line interface (CLI). This chapter will step through both methods to complete the basic configurations to put the device on your network. Use whichever you are most comfortable with.

This chapter also provides guidelines for password and administrator best practices as well as how to upgrade the firmware.

This section includes the topics:

- [Connecting to the FortiGate unit](#)
- [Setup Wizard](#)
- [FortiExplorer](#)
- [Configuring NAT mode](#)
- [Configuring transparent mode](#)
- [Verifying the configuration](#)
- [Additional configuration](#)
- [Passwords](#)
- [Administrators](#)
- [Backing up the configuration](#)
- [Firmware](#)
- [Controlled upgrade](#)

Connecting to the FortiGate unit

To configure, maintain and administer the FortiGate unit, you need to connect to it from a management computer. There are two ways to do this:

- using the web-based manager: a GUI interface that you connect to using a current web browser such as Firefox or Internet Explorer.
- using the command line interface (CLI): a command line interface similar to DOS or UNIX commands that you connect to using SSH or a Telnet terminal.

Connecting to the web-based manager

To connect to the web-based manager, you require:

- a computer with an Ethernet connection
- Microsoft Internet Explorer version 6.0 or higher or any recent version of a common web browser
- an Ethernet cable.

To connect to the web-based manager

- 1 Set the IP address of the management computer to the static IP address 192.168.1.2 with a netmask of 255.255.255.0.
- 2 Using the Ethernet cable, connect the internal or port 1 interface of the FortiGate unit to the computer Ethernet connection.
- 3 Start your browser and enter the address `https://192.168.1.99`. (remember to include the “s” in `https://`).

To support a secure HTTPS authentication method, the FortiGate unit ships with a self-signed security certificate, which is offered to remote clients whenever they initiate a HTTPS connection to the FortiGate unit. When you connect, the FortiGate unit displays two security warnings in a browser.

The first warning prompts you to accept and optionally install the FortiGate unit’s self-signed security certificate. If you do not accept the certificate, the FortiGate unit refuses the connection. If you accept the certificate, the FortiGate login page appears. The credentials entered are encrypted before they are sent to the FortiGate unit. If you choose to accept the certificate permanently, the warning is not displayed again.

Just before the FortiGate login page is displayed, a second warning informs you that the FortiGate certificate distinguished name differs from the original request. This warning occurs because the FortiGate unit redirects the connection. This is an informational message. Select OK to continue logging in.

- 4 Type `admin` in the Name field and select Login.

Connecting to the CLI

The command line interface (CLI) is an alternative method of configuring the FortiGate unit. The CLI compliments the web-based manager in that it not only has the same configuration options, but additional settings not available through the web-based manager.

If you are new to FortiOS or a command line interface configuration tool, see [“Using the CLI” on page 315](#) for an overview of the CLI, how to connect to it, and how to use it.

Setup Wizard

For the FortiGate-50B, 60C and 80C series, FortiOS includes a wizard to step you through the basic configuration of the FortiGate unit. The Setup Wizard will configure your FortiGate unit from factory default settings. If you set your management computer to the default IP address of the FortiGate unit, 192.168.1.99, and connect it to the FortiGate unit, when the device starts it will automatically launch the wizard.

A Wizard button also appears in the web-based manager. Use this button to update the configuration if required. Because the wizard configures from a default setting, it will reset the FortiGate unit to its factory defaults before beginning. The wizard will prompt you to save the existing configuration before proceeding.

FortiExplorer

FortiExplorer is a software tool for easy configuration of a new FortiGate unit, or simple updates to existing FortiGate units on a Microsoft Windows or Mac OS computer. FortiExplorer is included with the FortiGate-60C series of devices, as well as is available from the Fortinet web site.

FortiExplorer uses a USB connection to the FortiGate unit, rather than using a console cable or Ethernet connection. The USB connection does not replace the other options, but adds another option when configuring the FortiGate unit.

Installation

FortiExplorer is available for Microsoft Windows XP, Windows 7, and Mac OS X. The software is available on the Tools and Documentation CD included with your FortiGate unit, or is available for download from the Fortinet web site at http://www.fortinet.com/resource_center/product_demos.html.

Microsoft Windows install

To install FortiClient on Windows

- 1 Extract the ZIP (if downloaded) and double-click the .MSI or .EXE file and follow the instructions on screen. If loading from the CD, select the icon for your version of Windows.
- 2 Connect the USB cable to the FortiGate unit and the management computer.
- 3 For Windows XP, the New Hardware Wizard opens when the cables are connected. Select the option No, not at this time and select Next.
- 4 Select Install the hardware automatically and select Next.
- 5 After a few moments, FortiExplorer will launch.

Apple Macintosh OS X

To install FortiClient on Mac OS X

- 1 Double-click the .dmg file and drag the FortiExplorer program file into the Applications folder.
- 2 Connect the USB cable to the FortiGate unit and the management computer.
- 3 Double-click the FortiExplorer icon to launch the application.

Configuration options

With FortiExplorer, you are provided a number of options on how to configure the FortiGate unit, depending on your level of comfort with various interfaces. The options available are:

- the configuration wizard, which guides you through the basic configuration of IP addresses, passwords and security policies
- the web-based manager, which when chosen, appears within the FortiExplorer window.
- the command line interface (CLI), which when chosen, appears within the FortiExplorer window.

Updating FortiExplorer and firmware

FortiExplorer may be updated from time to time to update and add features, or correct other issues. To ensure you have the most recent FortiExplorer, use the Check for Updates option in FortiExplorer.

To check for updates on Microsoft Windows XP or Windows 7, go to *Help > Check for Updates*.

To check for updates on Apple Macintosh OS X, go to *FortiExplorer > Check for Updates*.

You can also use FortiExplorer to check for new firmware for a FortiGate unit. To check for new firmware, select the FortiGate unit from the *Device* list and select *Check for Update*.

Configuring NAT mode

When configuring NAT mode, you need to define interface addresses and default routes, and simple security policies. You can use the web-based manager or the CLI to configure the FortiGate unit in NAT mode.

Configure the interfaces

When shipped, the FortiGate unit has a default address of 192.168.1.99 and a netmask of 255.255.255.0. for either the Port 1 or Internal interface. You need to configure this and other ports for use on your network.



If you change the IP address of the interface you are connecting to, you must connect through a web browser again using the new address. Browse to https:// followed by the new IP address of the interface. If the new IP address of the interface is on a different subnet, you may have to change the IP address of your computer to the same subnet.

To configure interface for manual addressing - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select an interface and select *Edit*.
- 3 Enter the *IP address* and *netmask* for the interface.
- 4 Select *OK*.

To configure an interface for manual addressing - CLI

```
config system interface
  edit <interface_name>
    set mode static
    set ip <interface_ip> <interface_ipmask>
  end
```

To configure DHCP addressing - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the *Edit* icon for an interface.
- 3 Select *DHCP* and complete the following:

Distance

Enter the administrative distance, between 1 and 255 for the default gateway retrieved from the DHCP server. The administrative distance specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route.

Retrieve default gateway from server	Enable to retrieve a default gateway IP address from the DHCP server.
Override internal DNS	Enable to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page on <i>System > Network > Options</i> . You should also enable Obtain DNS server address automatically in <i>System > Network > Options</i> .

4 Select OK.



For more information on DHCP, see [“DHCP servers and relays” on page 570](#).

To configure DHCP addressing - CLI

```
config system interface
edit <interface_name>
set mode dhcp
set distance <integer>
set defaultgw enable
end
```

To configure PPPoE addressing - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select an interface and select *Edit*.
- 3 Select *PPPoE*, and complete the following:

Username	Enter the username for the PPPoE server. This may have been provided by your Internet Service Provider.
Password	Enter the password for the PPPoE server for the above user name.
Unnumbered IP	Specify the IP address for the interface. If your Internet Service Provider has assigned you a block of IP addresses, use one of these IP addresses. Alternatively, you can use, or borrow, the IP address of a configured interface on the router. You may need to do this to minimize the number of unique IP addresses within your network. If you are borrowing an IP address, remember the interface must be enabled, and the Ethernet cable connected to the FortiGate unit.
Initial Disc Timeout	Initial discovery timeout in seconds. The amount of time to wait before starting to retry a PPPoE discovery. To disable the discovery timeout, set the value to 0.
Initial PADT Timeout	Initial PPPoE Active Discovery Terminate (PADT) timeout in seconds. Use this timeout to shut down the PPPoE session if it is idle for this number of seconds. Your Internet Service Provider must support PADT. To disable the PADT timeout, set the value to 0.

Distance	Enter the administrative distance, between 1 and 255, for the default gateway retrieved from the DHCP server. The distance specifies the relative priority of a route when there are multiple routes to the same destination. A lower distance indicates a more preferred route.
Retrieve default gateway from server	Enable to retrieve a default gateway IP address from the DHCP server. The default gateway is added to the static routing table.
Override internal DNS	Enable to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page on <i>System > Network > Options</i> . On FortiGate-100A units and lower, you should also enable Obtain DNS server address automatically in <i>System > Network > Options</i> .

4 Select *OK*.

To configure PPPoE addressing - CLI

```
config system interface
  edit <interface_name>
    set mode pppoe
    set username <pppoe_username>
    set password <pppoe_password>
    set ipunnumbered <unnumbered_ipv4>
    set disc-retry-timeout <pppoe_retry>
    set padt-retry-timeout <pppoe_retry>
    set distance <integer>
    set defaultgw enable
  end
```

Configure a DNS

A DNS server is a public service that converts symbolic node names to IP addresses. A domain name server (DNS) implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with the computer IP address. This enables you to use readable locations, such as fortinet.com when browsing the Internet.

The FortiGate unit includes default DNS server addresses. However, these should be changed to those provided by your Internet Service Provider. The defaults are DNS proxies and are not as reliable as those from your ISP.

To configure DNS settings - web-based manager

- 1 Go to *System > Network > DNS*.
- 2 Enter the IP address of the primary DNS server.
- 3 Enter the IP address of the secondary DNS server.
- 4 Select *Apply*.



For more information on DNS servers see [“DNS services” on page 574](#).

To configure DNS server settings - CLI

```
config system dns
    set primary <dns_ipv4>
    set secondary <dns_ipv4>
end
```

Add a default route and gateway

A route provides the FortiGate unit with the information it needs to forward a packet to a particular destination. A static route causes packets to be forwarded to a destination other than the default gateway. You define static routes manually. Static routes control traffic exiting the FortiGate unit. You can specify through which interface the packet will leave and to which device the packet should be routed.

In the factory default configuration, entry number 1 in the Static Route list is associated with a destination address of 0.0.0.0/0.0.0.0, which means any/all destinations. This route is called the “static default route”. If no other routes are present in the routing table and a packet needs to be forwarded beyond the FortiGate unit, the factory configured static default route causes the FortiGate unit to forward the packet to the default gateway.

For an initial configuration, you must edit the static default route to specify a different default gateway for the FortiGate unit. This will enable the flow of data through the unit.

To modify the default gateway - web-based manager

- 1 Go to *Router > Static > Static Route*.
- 2 Select the default route and select *Edit*.
- 3 In the *Gateway* field, type the IP address of the next-hop router where outbound traffic is directed.
- 4 If the FortiGate unit reaches the next-hop router through a different interface (compared to the interface that is currently selected in *Device*, select the name of the interface from the *Device* drop-down list.
- 5 Select *OK*.

To modify the default gateway - CLI

```
config router static
    edit <sequence_num>
        set gateway <gateway_address_ipv4>
        set device <interface_name>
    end
```

Add security policies

Security policies enable traffic to flow through the FortiGate interfaces. Security policies define how the FortiGate unit processes the packets in a communication session. For the initial installation, a single security policy that enables all traffic to flow through will enable you to verify your configuration is working. On lower-end units such a default security policy is already in place. For the high-end FortiGate units, you need to add a security policy.

The following steps add two policies that allows all traffic through the FortiGate unit, to enable you to continue testing the configuration on the network. These steps provide a quick way to get traffic flowing through the FortiGate unit. It is a very broad policy and not recommended to keep on the system once initial setup and testing are complete. You will want to add more restrictive security policies to provide better network protection. For more information on security policies, see the [FortiGate Fundamentals](#).

To add an outgoing traffic security policy - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Set the following and select *OK*.

Source Interface/Zone	Select the port connected to the network.
Source Address	All
Destination Interface/Zone	Select the port connected to the Internet.
Destination Address	All
Schedule	always
Service	Any
Action	Accept

To add an outgoing traffic security policy - CLI

```

config firewall policy
  edit <interface_name>
    set srcintf <name_str>
    set srcaddr <name_str>
    set dstintf <name_str>
    set dstaddr <name_str>
    set schedule always
    set service ANY
    set action accept
  end

```

To add an incoming traffic security policy - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Set the following and select *OK*.

Source Interface	Select the port connected to the Internet.
Source Address	All
Destination Interface	Select the port connected to the network.
Destination Address	All
Schedule	always
Service	Any
Action	Accept

To add an incoming traffic security policy - CLI

```

config firewall policy
  edit <interface_name>
    set srcintf <name_str>
    set srcaddr <name_str>

```

```
set dstintf <name_str>
set dstaddr <name_str>
set schedule always
set service ANY
set action accept
end
```

To create an incoming traffic security policy, you use the same commands with the addresses reversed. security policy configuration is the same in NAT and transparent mode.

These policies allow all traffic through. No UTM profiles have been configured or applied. Ensure you create additional security policies to accommodate your network requirements.

Configuring transparent mode

You can then configure the management IP address, default routes, and security policies. You can use the web-based manager or the CLI to configure the FortiGate unit in transparent mode.

Switching to transparent mode

First need to switch to transparent mode.

To switch to transparent mode - web-based manager

- 1 Go to *System > Status*.
- 2 Under *System Information*, select *Change* beside the *Operation Mode*.
- 3 Select *Transparent*.
- 4 Enter the *Management IP/Netmask* address and the *Default Gateway* address.
The default gateway IP address is required to tell the FortiGate unit where to send network traffic to other networks.
- 5 Select *Apply*.

To switch to transparent mode

```
config system settings
set opmode transparent
set manageip <manage_ipv4>
set gateway <gw_ipv4>
end
```

Configure a DNS

A DNS server is a service that converts symbolic node names to IP addresses. A domain name server (DNS) implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with the computer IP address. This enables you to use readable locations, such as fortinet.com when browsing the Internet.

DNS server IP addresses are typically provided by your Internet Service Provider. For further DNS configuration and concepts, see [“DNS services” on page 574](#).

To configure DNS server settings - web-based manager

- 1 Go to *System > Network > DNS*.
- 2 Enter the IP address of the primary DNS server.
- 3 Enter the IP address of the secondary DNS server.
- 4 Select *Apply*.

To configure DNS server settings - CLI

```
config system dns
    set primary <dns_ipv4>
    set secondary <dns_ipv4>
end
```

Add security policies

Security policies enable traffic to flow through the FortiGate interfaces. Security policies define the FortiGate unit process the packets in a communication session. You can configure the security policies to allow only specific traffic, users and specific times when traffic is allowed.

For the initial installation, a single security policy that enables all traffic through will enable you to verify your configuration is working. On lower-end units such a default security policy is already in place. For the higher end FortiGate units, you will need to add a security policy.

The following steps add two policies that allows all traffic through the FortiGate unit, to enable you to continue testing the configuration on the network.

These steps provide a quick way to get traffic flowing through the FortiGate unit. It is a very broad policy and not recommended to keep on the system once initial setup and testing are complete. You will want to add more restrictive security policies to provide better network protection. For more information on security policies, see the [FortiGate Fundamentals](#).

To add an outgoing traffic security policy - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Set the following and select *OK*.

Source Interface/Zone	Select the port connected to the network.
Source Address	All
Destination Interface/Zone	Select the port connected to the Internet.
Destination Address	All
Schedule	always
Service	Any
Action	Accept

To add an outgoing traffic security policy - CLI

```
config firewall policy
    edit <policy_number>
        set srcintf <name_str>
```

```

set srcaddr <name_str>
set dstintf <name_str>
set dstaddr <name_str>
set schedule always
set service ANY
set action accept
end

```

To add an incoming traffic security policy - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Set the following and select *OK*.

Source Interface	Select the port connected to the Internet.
Source Address	All
Destination Interface	Select the port connected to the network.
Destination Address	All
Schedule	always
Service	Any
Action	Accept

To add an incoming traffic security policy - CLI

```

config firewall policy
edit <policy_number>
set srcintf <name_str>
set srcaddr <name_str>
set dstintf <name_str>
set dstaddr <name_str>
set schedule always
set service ANY
set action accept
end

```

To create an incoming traffic security policy, you use the same commands with the addresses reversed.

Security policy configuration is the same in NAT mode and transparent mode.

These policies allow all traffic through. No UTM profiles have been configured or applied. Ensure you create additional security policies to accommodate your network requirements.

Verifying the configuration

Your FortiGate unit is now configured and connected to the network. To verify that the FortiGate unit is connected and configured correctly, use your web browser to browse a web site, or use your email client to send and receive email.

If you cannot browse to the web site or retrieve/send email from your account, review the previous steps to ensure all information was entered correctly and try again.

Remember to verify the security policies. The security policies control the flow of information through the FortiGate unit. If the policies are not set up correctly, or are too restrictive, they can prohibit network traffic flow.

Additional configuration

Once the FortiGate unit is connected and traffic can pass through, several more configuration options are available. While not mandatory, they will help to ensure better control with the firewall.

Setting the time and date

For effective scheduling and logging, the FortiGate system date and time should be accurate. You can either manually set the system date and time or configure the FortiGate unit to automatically keep its time correct by synchronizing with a Network Time Protocol (NTP) server.

To set the date and time - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Under *System Information > System Time*, select *Change*.
- 3 Select your *Time Zone*.
- 4 Select *Set Time* and set the FortiGate system date and time.
- 5 Select *OK*.

Set the time and date - CLI

```
config system global
    set timezone <zone_value>
end
execute date [<date_str>]
execute time [<time_str>]
```



By default, FortiOS has the daylight savings time configuration enabled. The system time must be manually adjusted after daylight saving time ends. To disable DST, in the CLI enter the commands:

```
config system global
    set dst disable
end
```

Using the NTP Server

The Network Time Protocol enables you to keep the FortiGate time in sync with other network systems. By enabling NTP on the FortiGate unit, FortiOS will check with the NTP server you select at the configured intervals. This will also ensure that logs and other time-sensitive settings on the FortiGate unit are correct.



The FortiGate unit maintains its internal clock using the built-in battery. At startup, the time reported by the FortiGate unit will indicate the hardware clock time, which may not be accurate. When using NTP, the system time might change after the FortiGate has successfully obtained the time from a configured NTP server.

For the NTP server, you can identify a specific port/IP address for this self-originating traffic. The configuration is performed in the CLI with the command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
  set ntpsyn enable
  set syncinterval 5
  set source-ip 192.168.4.5
end
```

Configuring FortiGuard

The FDN is a world-wide network of FortiGuard Distribution Servers (FDS). When the FortiGate unit connects to the FDN, it connects to the nearest FDS. To do this, all FortiGate units are programmed with a list of FDS addresses sorted by nearest time zone according to the time zone configured for the FortiGate unit.

Before you can begin receiving updates, you must register your FortiGate unit from the Fortinet web page. After which, you need to configure the FortiGate unit to connect to the FortiGuard Distribution Network (FDN) to update the antivirus, antispyware and IPS attack definitions.

Updating antivirus and IPS signatures

After you have registered your FortiGate unit, you can update antivirus and IPS signatures. The FortiGuard Center enables you to receive push updates, allow push update to a specific IP address, and schedule updates for daily, weekly, or hourly intervals.

To update antivirus definitions and IPS signatures

- 1 Go to *System > Config > FortiGuard*.
- 2 Select the expand arrow for *AntiVirus and IPS Options* to expand the options.
- 3 Select *Update Now* to update the antivirus definitions.

If the connection to the FDN is successful, the web-based manager displays a message similar to the following:

Your update request has been sent. Your database will be updated in a few minutes. Please check your update page for the status of the update.

After a few minutes, if an update is available, the FortiGuard Center Services information on the Dashboard lists new version information for antivirus definitions. The System Status page also displays new dates and version numbers for the antivirus definitions. Messages are recorded to the event log indicating whether or not the update was successful or not.



Updating antivirus definitions can cause a very short disruption in traffic currently being scanned while the FortiGate unit applies the new signature database. Schedule updates when traffic is light, for example overnight, to minimize any disruption.

Passwords

The FortiGate unit ships with a default empty password, that is, there is no password. You will want to apply a password to prevent anyone from logging into the FortiGate unit and changing configuration options.

To change the administrator password - web-based manager

- 1 Go to *System > Admin > Administrators*.
- 2 Select the admin account and select *Change Password*.
- 3 Enter a new password and select *OK*.

Set the admin password - CLI

```
config system admin
  edit admin
    set password <admin_password>
  end
```

Password considerations

When changing the password, consider the following to ensure better security.

- Do not make passwords that are obvious, such as the company name, administrator names or other obvious word or phrase.
- Use numbers in place of letters, for example, `passw0rd`. Alternatively, spell words with extra letters, for example, `password`.
- Administrative passwords can be up to 256 characters.
- Include a mixture of letters, numbers, and upper and lower case.
- Use multiple words together, or possibly even a sentence, for example `keytothehighway`, or with a combination of the above suggestions.
- Use a password generator.
- Change the password regularly and always use a code unique (not a variation of the existing password by adding a “1” to it, for example `password`, `password1`).
- Write the password down and store it in a safe place away from the management computer, in case you forget it.
- Alternatively, ensure at least two people know the password in the event that one person becomes ill, is away on vacation or leaves the company. Alternatively have two different admin logins.

Password policy

The FortiGate unit includes the ability to enforce a password policy for administrator login. with the policy, you can enforce regular changes and specific criteria for a password including:

- minimum length between 8 and 32 characters.
- if the password must contain uppercase (A, B, C) and/or lowercase (a, b, c) characters.
- if the password must contain numbers (1, 2, 3).
- if the password must contain non-alphanumeric characters (!, @, #, \$, %, ^, &, *, ,).
- where the password applies (admin or IPsec or both).

- the duration of the password before a new one must be specified.

To apply a password policy - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 Select *Enable* and configure the settings as required.

To apply a password policy - CLI

```
config system password-policy
set status enable
```

Configure the other settings as required.

Forgotten password?

It happens that the administrator of the FortiGate unit leaves the company and does not have the opportunity to provide the administrative password or forgets. Or you simply forgot the password.

In the event you lose or forget the password, you need to contact Customer Support for the steps required to reset the password. For information on contacting Customer Support, see the Support web site at [web site at https://support.fortinet.com](https://support.fortinet.com).

Administrators

By default, the FortiGate unit has a super administrator called “admin”. This user login cannot be deleted and always has ultimate access over the FortiGate unit. As well you can add administrators for various functions and VDOMs. Each one can have their own username and password and set of access privileges. There are two levels of administrator accounts; regular administrators and system administrators. Regular administrators are administrators with any admin profile other than the default super_admin. System administrators are administrators that are assigned the super_admin profile.

Administrator configuration

To create a new administrator account, go to *System Admin > Administrators* and select *Create New*.

You need to use the default “admin” account, an account with the super_admin admin profile, or an administrator with read-write access control to add new administrator accounts and control their permission levels. If you log in with an administrator account that does not have the super_admin admin profile, the administrators list will show only the administrators for the current virtual domain.



The name of the administrator should not contain the characters <> () # " ' . Using these characters in the administrator account name can result in a cross site scripting (XSS) vulnerability.

Regular (password) authentication for administrators

You can use a password stored on the local FortiGate unit to authenticate an administrator. When you select *Regular* for *Type*, you will see *Local* as the entry in the *Type* column when you view the list of administrators.

If you forget or lose an administrator account password and cannot log in to your FortiGate unit, see the Fortinet Knowledge Base article [Recovering a lost FortiGate administrator account passwords](#).

Management access

Management access defines how administrators are able to log on to the FortiGate unit to perform management tasks such as configuration and maintenance. Methods of access can include local access through the console connection, or remote access over a network or modem interface using various protocols including Telnet and HTTPS.

You can configure management access on any interface in your VDOM. In NAT mode, the interface IP address is used for management access. In transparent mode, you configure a single management IP address that applies to all interfaces in your VDOM that permit management access. The FortiGate unit also uses this IP address to connect to the FDN for virus and attack updates.

The system administrator (admin) can access all VDOMs, and create regular administrator accounts. A regular administrator account can access only the VDOM to which it belongs. The management computer must connect to an interface in that VDOM. It does not matter to which VDOM the interface belongs. In both cases, the management computer must connect to an interface that permits management access and its IP address must be on the same network. Management access can be via HTTP, HTTPS, telnet, or SSH sessions if those services are enabled on the interface. HTTPS and SSH are preferred as they are more secure.

You can allow remote administration of the FortiGate unit. However, allowing remote administration from the Internet could compromise the security of the FortiGate unit. You should avoid this unless it is required for your configuration. To improve the security of a FortiGate unit that allows remote administration from the Internet:

- Use secure administrative user passwords.
- Change these passwords regularly.
- Enable two-factor authentication for administrators.
- Enable secure administrative access to this interface using only HTTPS or SSH.
- Use Trusted Hosts to limit where the remote access can originate from.
- Do not change the system idle timeout from the default value of 5 minutes.

Tightening Security

One point of security breach is at the management computer. Administrators who leave their workstations for a prolonged amount of time while staying logged into the web-based manager or CLI (whether on purpose or not), leave the firewall open to malicious intent.

Passwords

Do not make passwords that are obvious, such as the company name, administrator names or other obvious word or phrase. Administrative passwords can be up to 256 characters. For more information on passwords, see [“Passwords” on page 348](#).

Preventing unwanted login attempts

Setting trusted hosts for an administrators increases limiting what computers an administrator can log in from. When you identify a trusted host, the FortiGate unit will only accept the administrator's login from the configured IP address. Any attempt to log in with the same credentials from any other IP address will be dropped. To ensure the administrator has access from different locations, you can enter up to ten IP addresses. Ideally, this should be kept to a minimum. For higher security, use an IP address with a net mask of 255.255.255.255, and enter an IP address (non-zero) in each of the three default trusted host fields.

Trusted hosts are configured when adding a new administrator by going to *System > Admin > Administrators* in the web-based manager or `config system admin` in the CLI.

The trusted hosts apply to the web-based manager, ping, snmp and the CLI when accessed through Telnet or SSH. CLI access through the console port is not affected.

Also ensure all entries contain actual IP addresses, not the default 0.0.0.0.

Disable admin services

On untrusted networks, turn off the weak administrative services such as TLENET and HTTP. With these services, passwords are passed in the clear, not encrypted.

These services can be disabled by going to *System > Network > Interface* and deselecting the required check boxes.

SSH login time out

When logging into the console using SSH, the default time of inactivity is 120 seconds (2 minutes) to successfully log into the FortiGate unit. To enhance security, you can configure the time to be shorter. Using the CLI, you can change the length of time the command prompt remains idle before the FortiGate unit will log the administrator out. The range can be between 10 and 3600 seconds. To set the logout time enter the following commands:

```
config system global
  set admin-ssh-grace-time <number_of_seconds>
end
```

Administrator lockout

By default, the FortiGate unit includes set number of password retries. That is, the administrator has a maximum of three attempts to log into their account before they are locked out for a set amount of time. The number of attempts can be set to an alternate value.

As well, the default wait time before the administrator can try to enter a password again is 60 seconds. You can also change this to further sway would-be hackers. Both settings are configured only in the CLI

To configure the lockout options use the following commands:

```
config system global
  set admin-lockout-threshold <failed_attempts>
  set admin-lockout-duration <seconds>
end
```

For example, to set the lockout threshold to one attempt and a five minute duration before the administrator can try again to log in enter the commands"

```
config system global
  set admin-lockout-threshold 1
  set admin-lockout-duration 300
end
```

Idle time-out

To avoid the possibility of an administrator walking away from the management computer and leaving it exposed to unauthorized personnel, you can add an idle time-out. That is, if the web-based manager is not used for a specified amount of time, the FortiGate unit will automatically log the user out. To continue their work, they must log in to the device again.

The time-out can be set as high as 480 minutes, or eight hours, although this is not recommend.

To set the idle time out - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 In the *Administration Settings*, enter the amount of time the Administrator login can remain idle in the *Idle Timeout* field.
- 3 Select *Apply*.

To set the idle time out - CLI

```
config system global
  set admintimeout <minutes>
end
```

Administrative ports

You can set the web-based manager access as through HTTP, HTTPS, SSH and Telnet. In these cases, the default ports for these protocols are 80, 443, 22 and 23 respectively. You can change the ports used for network administration to a different, unused port to further limit potential hackers.



Ensure the port you select is not a port you will be using for other applications. For a list of assigned port numbers see <http://www.iana.org/assignments/port-numbers>.

To change the administrative ports - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 In the *Web Administration Ports* section, change the port numbers.
- 3 Select *Apply*.

To change the administrative ports - CLI

```
config system global
  set admin-port <http_port_number>
  set admin-sport <https_port_number>
  set admin-ssh-port <ssh_port_number>
  set admin-telnet-port <telnet_port_number>
end
```

When logging into the FortiGate unit, by default FortiOS will automatically use the default ports. That is, when logging into the FortiGate IP address, you only need to enter the address, for example:

```
https://192.168.1.1
```

When you change the administrative port number, the port number must be added to the url. For example, if the port number for HTTPS access is 2112, the administrator must enter the following address:

```
https://192.168.1.1:2112
```

Disable interfaces

If any of the interfaces on the FortiGate unit are not being used, disable traffic on that interface. This avoids someone plugging in network cables and potentially causing network bypass or loop issues.

To disable an interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the interface from the list and select *Edit*.
- 3 For *Administrative Access*, select *Down*.
- 4 Select *OK*.

To disable an interface - CLI

```
config system interface
  edit <interface_name>
    set status down
  end
```

Change the admin username

The default super administrator user name, admin, is a very standard default administrator name. Leaving this as is, is one half of the key to the FortiGate unit being compromised. The name can be changed.

To do this, you need to create another super user with full access and log in as that user. Then go to *System > Admin > Administrator*, select the *admin* account and select *Edit* to change the user name.

Segregated administrative roles

To minimize the affect of an administrator doing complete harm to the FortiGate configuration and possibly jeopardize the network, create individual administrative roles where none of the administrators have super-admin permissions. For example, and admin solely to create security policies, another for users and groups, another for VPN and so on.

RADIUS authentication for administrators

Remote Authentication and Dial-in User Service (RADIUS) servers provide authentication, authorization, and accounting functions. FortiGate units use the authentication and authorization functions of the RADIUS server. To use the RADIUS server for authentication, you must configure the server before you configure the FortiGate users or user groups that will need it.

If you have configured RADIUS support and a user is required to authenticate using a RADIUS server, the FortiGate unit sends the user's credentials to the RADIUS server for authentication. If the RADIUS server can authenticate the user, the user is successfully authenticated with the FortiGate unit. If the RADIUS server cannot authenticate the user, the FortiGate unit refuses the connection.

If you want to use a RADIUS server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure the FortiGate unit to access the RADIUS server
- create the RADIUS user group
- configure an administrator to authenticate with a RADIUS server.

Configuring LDAP authentication for administrators

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, printers, etc.

If you have configured LDAP support and an administrator is required to authenticate using an LDAP server, the FortiGate unit contacts the LDAP server for authentication. If the LDAP server cannot authenticate the administrator, the FortiGate unit refuses the connection.

If you want to use an LDAP server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure an LDAP server
- create an LDAP user group
- configure an administrator to authenticate with an LDAP server.

To view the LDAP server list, go to *User > Remote > LDAP*.

TACACS+ authentication for administrators

Terminal Access Controller Access-Control System (TACACS+) is a remote authentication protocol that provides access control for routers, network access servers, and other network computing devices via one or more centralized servers.

If you have configured TACACS+ support and an administrator is required to authenticate using a TACACS+ server, the FortiGate unit contacts the TACACS+ server for authentication. If the TACACS+ server cannot authenticate the administrator, the connection is refused by the FortiGate unit.

If you want to use an TACACS+ server to authenticate administrators in your VDOM, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure the FortiGate unit to access the TACACS+ server
- create a TACACS+ user group
- configure an administrator to authenticate with a TACACS+ server.

PKI certificate authentication for administrators

Public Key Infrastructure (PKI) authentication uses a certificate authentication library that takes a list of peers, peer groups, and user groups and returns authentication successful or denied notifications. Users only need a valid certificate for successful authentication; no username or password is necessary.

To use PKI authentication for an administrator, you must configure the authentication before you create the administrator accounts. To do this you need to:

- configure a PKI user
- create a PKI user group
- configure an administrator to authenticate with a PKI certificate.

Administrator profiles

Administer profiles define what the administrator user can do when logged into the FortiGate unit. When you set up an administrator user account, you also assign an administrator profile, which dictates what the administrator user will see. Depending on the nature of the administrator's work, access level or seniority, you can allow them to view and configure as much, or as little, as required.

super_admin profile

The super_admin administrator is the administrative account that the primary administrator should have to log into the FortiGate unit. The profile can not be deleted or modified to ensure there is always a method to administer the FortiGate unit. This user profile has access to all components of FortiOS, including the ability to add and remove other system administrators. For some administrative functions, such as backing up and restoring the configuration using SCP, super_admin access is required.

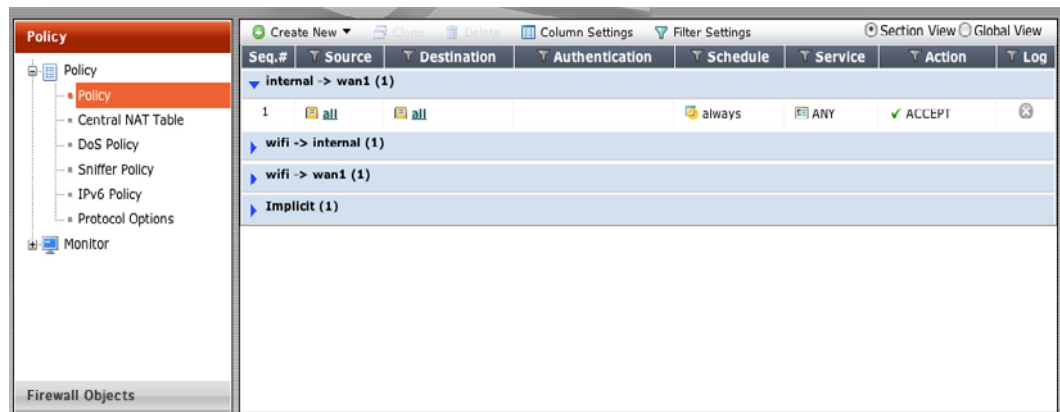


By default, the super_admin user (username is "admin"), does not have a password. Ensure you assign a password. You can also change the name of the account from "admin" to another name for better security.

Creating profiles

To configure administrator profiles go to *System > Admin > Admin Profile*. You can only assign one profile to an administrator user.

On the *New Admin Profile* page, you define the components of FortiOS that will be available to view and/or edit. For example, if you configure a profile so that the administrator can only access the firewall components, when an administrator with that profile logs into the FortiGate unit, they will only be able to view and edit any firewall components including policies, addresses, schedules and any other settings that directly affect security policies.

Figure 38: The view of an administrator with firewall-only access

Global and vdom profiles

By default, when you add a new administrative profile, it is set to have a vdom scope. That is, only the super_admin has a global profile that enables configuration of the entire FortiGate unit.

There may be instances where additional global administrative profiles may be required. To add more global profiles, use the following CLI command to set or change an administrative profile to be global.

```
config system accprofile
  set scope global
  ...
end
```

Once the scope is set, you can enable the read and read/write settings.

Adding administrators

When adding administrators, you are setting up the administrator's user account. An administrator account comprises of an administrator's basic settings as well as their access profile. The access profile is a definition of what the administrator is capable of viewing and editing. For information on administrator profiles, see [“Administrator profiles” on page 355](#).

To add an administrator - web-based manager

- 1 Go to *System > Admin > Administrators*.
- 2 Select *Create New*.
- 3 Enter the administrator name.
- 4 Select the type of account it will be. If you select *Remote*, the FortiGate unit can reference a RADIUS, LDAP or TACAS+ server.
- 5 When selecting *Remote* or *PKI* accounts, select the User Group the account will access.

For information on logging in using remote authentication servers, see the [User Authentication Guide](#). For an example of setting up a user with LDAP, see [“LDAP Admin Access and Authorization” on page 357](#)

- 6 Enter the password for the user.

This may be a temporary password that the administrator can change later. Passwords can be up to 256 characters in length. For more information on passwords, see [“Passwords” on page 348](#).

- 7 Select OK.

To add an administrator - CLI

```
config system admin
  edit <admin_name>
    set password <password>
    set accprofile <profile_name>
  end
```

LDAP Admin Access and Authorization

You can use the LDAP server as a means to add administrative users, saving the time to add users to the FortiGate unit administrator list. After configuring, any user within the selected LDAP group server can automatically log into the FortiGate unit as an administrator. Ensure that the admin profile is the correct level of access, or the users within the LDAP group are the only ones authorized to configure or modify the configuration of the FortiGate unit.

To do this, requires three steps:

- configure the LDAP server
- add the LDAP server to a user group
- configure the administrator account

Configure the LDAP server

First set up the LDAP server as you normally would, and include a group to bind to.

To configure the LDAP server - web-based manager

- 1 Go to *User > Remote > LDAP* and select *Create New*.
- 2 Enter a *Name* for the server.
- 3 Enter the *Server IP* address or name.
- 4 Enter the *Common Name Identifier* and *Distinguished Name*.
- 5 Set the *Bind Type* to *Regular* and enter the *User DN* and *Password*.
- 6 Select OK.

To configure the LDAP server - CLI

```
config user ldap
  edit <ldap_server_name>
    set server <server_ip>
    set cnid cn
    set dn DC=XYZ,DC=COM
    set type regular
    set username CN=Administrator,CN=Users,DC=XYZ,DC=COM
    set password <password>
    set member-attr <group_binding>
  end
```

Add the LDAP server to a user group

Next, create a user group that will include the LDAP server that was created above.

To create a user group - web-based manager

- 1 Go to *User > User Group > User Group* and select *Create New*.
- 2 Enter a *Name* for the group.
- 3 In the section labelled *Remote authentication servers*, select *Add*.
- 4 Select the *Remote Server* from the drop-down list.
- 5 Select *OK*.

To create a user group - CLI

```
config user group
  edit <group_name>
    config match
      edit 1
        set server-name <LDAP_server>
        set group-name <group_name>
      end
    end
  end
```

Configure the administrator account

Now you can create a new administrator, where rather than entering a password, you will use the new user group and the wildcard option for authentication.

To create an administrator - web-based manager

- 1 Go to *System > Admin > Administrators* and select *Create New*.
- 2 In the *Administrator* field, enter the name for the administrator.
- 3 For *Type*, select *Remote*.
- 4 Select the *User Group* created above from the drop-down list.
- 5 Select *Wildcard*.

The Wildcard option allows for LDAP users to connect as this administrator.

- 6 Select an *Admin Profile*.
- 7 Select *OK*.

To create an administrator - CLI

```
config system admin
  edit <admin_name>
    set remote-auth enable
    set accprofile super_admin
    set wildcard enable
    set remote-group ldap
  end
```

Monitoring administrators

You can view the administrators logged in using the *System Information* widget on the Dashboard. On the widget is the *Current Administrator* row that shows the administrator logged in and the total logged in. Selecting *Details* displays the administrators), where they are logging in from and how (CLI, web-based manager) and when they logged in.

You are also able to monitor the activities the administrators perform on the FortiGate unit using the logging of events. Event logs include a number of options to track configuration changes.

To set logging - web-based manager

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select a location to store logs and set the *Minimum log level* to *Information*.
- 3 Select *Apply*.
- 4 Go to *Log&Report > Event Log*.
- 5 Select the following event logs:
 - System activity event
 - Admin event
 - Configuration change event
- 6 Select *Apply*.

To set logging - CLI

```
config log <log_location> (log configuration will vary depending
    on location)
end
config log eventfilter
    set admin enable
    set system enable
    set config enable
end
```

To view the logs go to *Log&Report > Log Access > Event*.

Trusted hosts

Setting trusted hosts for administrators limits what computers an administrator can log in from. When you identify a trusted host, the FortiGate unit will only accept the administrator's login from the configured IP address. Any attempt to log in with the same credentials from any other IP address will be dropped. To ensure the administrator has access from different locations, you can enter up to ten IP addresses. Ideally, this should be kept to a minimum. For higher security, use an IP address with a net mask of 255.255.255.255, and enter an IP address (non-zero) in each of the three default trusted host fields.

Trusted hosts are configured when adding a new administrator by going to *System > Admin > Administrators* in the web-based manager or `config system admin` in the CLI.

The trusted hosts apply to the web-based manager, ping, snmp and the CLI when accessed through Telnet or SSH. CLI access through the console port is not affected.

Also ensure all entries contain actual IP addresses, not the default 0.0.0.0.

General Settings

Go to *System > Admin > Settings* to configure basic settings for administrative access, password policies and displaying additional options in the web-based manager.

Administrative port settings

The Administrative Settings enable you to change the default port configurations for administrative connections to the FortiGate unit for added security. When connecting to the FortiGate unit when the port has changed, the port must be included, such as `https://<ip_address>:<port>`. For example, if you are connecting to the FortiGate unit using port 99, the url would be `https://192.168.1.99:99`.



If you make a change to the default port number for HTTP, HTTPS, Telnet, or SSH, ensure that the port number is unique.

Password policies

Password policies, available by going to *System > Admin > Settings*, enable you to create a password policy that any administrator or user who updates their password, must follow. Using the available options you can define the required length of the password, what it must contain (numbers, upper and lower case, and so on) and an expiry time frame.

The FortiGate unit will warn of any password that is added and does not meet the criteria.

Display options

To minimize clutter on the web-based manager interface, a number of FortiOS features to not appear on the web-based manager. By going to *System > Admin Settings*, you can enable, or if not needed, disable various features. The change takes effect immediately without having to log out or reboot the device.

Backing up the configuration

Once you configure the FortiGate unit and it is working correctly, it is extremely important that you back up the configuration. In some cases, you may need to reset the FortiGate unit to factory defaults, or perform a TFTP upload of the firmware. In these instances, the configuration on the device will be lost.

Always back up the configuration and store it on the management computer or off site. It is also recommended that once the FortiGate is configured, and *any* further changes are made, that you back up the configuration immediately, to ensure you have the most current configuration available.

You have the option to save the configuration file to various locations including the local PC, USB key, FTP and TFTP site. The latter two are configurable through the CLI only.

If you have VDOMs, you can back up the configuration of the entire FortiGate unit, or only a specific VDOM. Note that if you are using FortiManager or the Fortinet Management Services (FAMS), full backups are performed, and the option to backup individual VDOMs will not appear.

To back up the FortiGate configuration - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 On the *System Information* widget, select *Backup* for the *System Configuration*.
- 3 Select to back up to your *Local PC*, *FortiManager* or to a *USB key*.

The *USB Disk* option will be grayed out if no USB drive is inserted in the USB port. The *FortiManager* option will not be available if the FortiGate unit is not being managed by a FortiManager system.

- 4 If VDOMs are enabled, select to backup the entire FortiGate configuration (*Full Config*) or only a specific VDOM configuration (*VDOM Config*).
- 5 If backing up a VDOM configuration, select the VDOM name from the list.
- 6 Select *Encrypt configuration file*.
Encryption must be enabled on the backup file to back up VPN certificates.
- 7 Enter a password and enter it again to confirm it. You will need this password to restore the file.
- 8 Select *Backup*.
- 9 The web browser will prompt you for a location to save the configuration file. The configuration file will have a .conf extension.

To back up the FortiGate configuration - CLI

```
execute backup config management-station <comment>
... or ...
execute backup config usb <backup_filename> [<backup_password>]
... or for FTP, note that port number, username are optional depending on the FTP site...
execute backup config ftp <backup_filename> <ftp_server>
    [<port>] [<user_name>] [<password>]
... or for TFTP ...
execute backup config tftp <backup_filename> <tftp_servers>
    <password>
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config vdom
    edit <vdom_name>
```

It is a good practice to backup the FortiGate configuration after any modification to any of the FortiGate settings. Alternatively, before performing an upgrade to the firmware, ensure you back up the configuration before upgrading. Should anything happen during the upgrade that changes the configuration, you can easily restore the saved configuration.

Backup and restore a configuration file using SCP

You can use secure copy protocol (SCP) to download the configuration file from the FortiGate unit as an alternative method of backing up the configuration file or an individual VDOM configuration file. This is done by enabling SCP for an administrator account and enabling SSH on a port used by the SCP client application to connect to the FortiGate unit. SCP is enabled using the CLI commands:

```
config system global
    set admin-scp enable
end
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config global
    set admin-scp enable
end
config vdom
    edit <vdom_name>
```

Enable SSH access on the interface

SCP uses the SSH protocol to provide secure file transfer. The interface you use for administration must allow SSH access.

To enable SSH - web-based manager:

- 1 Go to *System > Network > Interface*.
- 2 Select the interface you use for administrative access and select *Edit*.
- 3 In the *Administrative Access* section, select *SSH*.
- 4 Select *OK*.

To enable SSH - CLI:

```
config system interface
    edit <interface_name>
        set allowaccess ping https ssh
    end
```



When adding to, or removing a protocol, you must type the entire list again. For example, if you have an access list of HTTPS and SSH, and you want to add PING, typing:

```
set allowaccess ping
```

...only PING will be set. In this case, you must type...

```
set allowaccess https ssh ping
```

Using the SCP client

The FortiGate unit downloads the configuration file as `sys_conf`. Use the following syntax to download the file:

Linux

```
scp admin@<FortiGate_IP>:sys_config <location>
```

Windows

```
pscp admin@<FortiGate_IP>:sys_config <location>
```

These examples show how to download the configuration file from a FortiGate-100A, at IP address 172.20.120.171, using Linux and Windows SCP clients.

Linux client example

To download the configuration file to a local directory called `~/config`, enter the following command:

```
scp admin@172.20.120.171:sys_config ~/config
```

Enter the admin password when prompted.

Windows client example

To download the configuration file to a local directory called `c:\config`, enter the following command in a Command Prompt window:

```
pscp admin@172.20.120.171:sys_config c:\config
```

Enter the admin password when prompted.

SCP public-private key authentication

SCP authenticates itself to the FortiGate unit in the same way as an administrator using SSH accesses the CLI. Instead of using a password, you can configure the SCP client and the FortiGate unit with a public-private key pair.

To configure public-private key authentication

- 1 Create a public-private key pair using a key generator compatible with your SCP client.
- 2 Save the private key to the location on your computer where your SSH keys are stored.

This step depends on your SCP client. The Secure Shell key generator automatically stores the private key.

- 3 Copy the public key to the FortiGate unit using the CLI commands:

```
config system admin
    edit admin
        set ssh-public-key1 "<key-type> <key-value>"
    end
```

<key-type> must be the ssh-dss for a DSA key or ssh-rsa for an RSA key. For the <key-value>, copy the public key data and paste it into the CLI command.

If you are copying the key data from Windows Notepad, copy one line at a time and ensure that you paste each line of key data at the end of the previously pasted data. As well:

- Do not copy the end-of-line characters that appear as small rectangles in Notepad.
- Do not copy the ----- BEGIN SSH2 PUBLIC KEY ----- or Comment: "[2048-bit dsa,...]" lines.
- Do not copy the ----- END SSH2 PUBLIC KEY ----- line.

- 4 Type the closing quotation mark and press Enter.

Your SCP client can now authenticate to the FortiGate unit based on SSH keys rather than the administrator password.

Restoring a configuration using SCP

To restore the configuration using SCP, use the commands:

```
scp <local_file> <admin_user>@<FGT_IP>:fgt_restore_config
```

To use this command/method of restoring the FortiGate configuration, you need to log in as the "admin" administrator.

Restoring a configuration

Should you need to restore a configuration file, use the following steps.

To restore the FortiGate configuration - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 On the *System Information* widget, select *Restore* for the *System Configuration*.
- 3 Select to upload the configuration file to be restored from your *Local PC* or a *USB key*.
The *USB Disk* option will be grayed out if no USB drive is inserted in the USB port.
The *FortiManager* option will not be available if the FortiGate unit is not being managed by a FortiManager system.

- 4 Enter the path and file name of the configuration file, or select *Browse* to locate the file.
- 5 Enter a password if required.
- 6 Select *Restore*.

To back up the FortiGate configuration - CLI

```
execute restore config management-station normal 0
```

... or ...

```
execute restore config usb <filename> [<password>]
```

... or for FTP, note that port number, username are optional depending on the FTP site...

```
execute backup config ftp <backup_filename> <ftp_server> [<port>]  
[<user_name>] [<password>]
```

... or for TFTP ...

```
execute backup config tftp <backup_filename> <tftp_server>  
<password>
```

The FortiGate unit will load the configuration file and restart. Once the restart has completed, verify that the configuration has been restored.

Configuration revisions

The *Configuration Revisions* menu enables you to manage multiple versions of configuration files. Revision control requires either a configured central management server, or FortiGate units with 512 MB or more of memory. The central management server can either be a FortiManager unit or the FortiGuard Analysis & Management Service.

When revision control is enabled on your unit, and configurations have been backed up, a list of saved revisions of those backed-up configurations appears.

Configuration revisions are viewed in *System > Maintenance > Configuration Revision*.

Firmware

Fortinet periodically updates the FortiGate firmware to include new features and address issues. After you have registered your FortiGate unit, you can download firmware updates from the support web site, <http://support.fortinet.com>.

You can also use the instructions in this chapter to revert, to a previous version. The FortiGate unit includes a number of firmware installation options that enables you to test new firmware without disrupting the existing installation, and load it from different locations as required.

Fortinet issues patch releases--maintenance release builds that resolve important issues. Fortinet strongly recommends reviewing the release notes for the patch release, as well as testing and reviewing the patch release before upgrading the firmware. Follow the steps below:

- download and review the release notes for the patch release
- download the patch release
- back up the current configuration
- test the patch release until you are satisfied that it applies to your configuration.

Installing a patch release without reviewing release notes or testing the firmware may result in changes to settings or unexpected issues.

Only FortiGate admin user and administrators whose access profiles contain system read and write privileges can change the FortiGate firmware.

Downloading firmware

Firmware images for all FortiGate units is available on the Fortinet Customer Support web site. You must register your FortiGate unit to access firmware images. Register the FortiGate unit by visiting <http://support.fortinet.com> and select Product Registration.

To download firmware

- 1 Log into the site using your user name and password.
- 2 Go to *Firmware Images > FortiGate*.
- 3 Select the most recent FortiOS version.
- 4 Locate the firmware for your FortiGate unit, right-click the link and select the Download option for your browser.



Always review the [Release Notes](#) for a new firmware release before installing. The [Release Notes](#) can include information that is not available in the regular documentation.

Upgrading the firmware - web-based manager

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.



Always remember to back up your configuration before doing any firmware upgrade or downgrade.

To upgrade the firmware

- 1 Download the firmware image file to your management computer.
- 2 Log into the web-based manager as the admin administrative user.
- 3 Go to *System > Dashboard > Status*.
- 4 Under *System Information > Firmware Version*, select *Update*.
- 5 Type the path and filename of the firmware image file, or select *Browse* and locate the file.
- 6 Select *OK*.

The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, restarts, and displays the FortiGate login. This process takes a few minutes.

Reverting to a previous firmware version

The following procedures revert the FortiGate unit to its factory default configuration and deletes any configuration settings.

Before beginning this procedures, ensure you back up the FortiGate unit configuration.

If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.



Installing firmware replaces the current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.



To use this procedure, you must log in using the admin administrator account, or an administrator account that has system configuration read and write privileges.

To revert to a previous firmware version

- 1 Copy the firmware image file to the management computer.
- 2 Log into the FortiGate web-based manager.
- 3 Go to *System > Dashboard > Status*.
- 4 Under *System Information > Firmware Version*, select *Update*.
- 5 Type the path and filename of the firmware image file, or select *Browse* and locate the file.
- 6 Select *OK*.

The FortiGate unit uploads the firmware image file, reverts to the old firmware version, resets the configuration, restarts, and displays the FortiGate login. This process takes a few minutes.

- 7 Log into the web-based manager.
- 8 Restore your configuration.

For information about restoring your configuration see [“Restoring a configuration” on page 363](#).

Configuration Revision

The *Configuration Revisions* menu enables you to manage multiple versions of configuration files on models that have a 512 flash memory and higher. Revision control requires either a configured central management server, or the local hard drive. The central management server can either be a FortiManager unit or the FortiGuard Analysis and Management Service.

If central management is not configured on your FortiGate unit, a message appears to tell you to do one of the following:

- enable central management (see [Central management](#))
- obtain a valid license.

When revision control is enabled on your FortiGate unit, and configurations have been backed up, a list of saved revisions of those backed-up configurations appears.

Configuration revisions are viewed in *System > Maintenance > Configuration Revision*.

Configuration Revision page.	
Delete	Removes a configuration revision from the list. To remove multiple configurations from within the list, on the Configuration Revision page, in each of the rows of revisions you want removed, select the check box and then select <i>Delete</i> . To remove all configuration revisions from within the list, on the Configuration Revision page, select the check box in the check box column and then select <i>Delete</i> .
Details	View the CLI settings of a configuration revision.
Change Comments	Modifies the description for the configuration revision.
Diff	Select when you want to compare two revisions. You must select two revisions. From the Diff Display window you can view and compare the selected revision to one of: <ul style="list-style-type: none"> the current configuration a selected revision from the displayed list including revision history and templates a specified revision number.
Revert	Restores the previous selected revision.
Upload	Uploads a configuration file to the FortiGate unit, which is then added to the list.
OS Version <firmware_version_build> (appears as sections on the page)	The section of the page that contains the configuration files that belong to the specified FortiOS firmware version and build number. For example, if you have four configuration revisions for 4.0 MR1 (build-178) they appear in the section OS Version 4.00 build178 on the Configuration Revision page.
Revision	An incremental number indicating the order in which the configurations were saved. These may not be consecutive numbers if configurations are deleted.
Date/Time	The date and time this configuration was saved on the FortiGate unit.
Administrator	The administrator account that was used to back up this revision.
Comments	Any relevant information saved with the revision.
Ref.	Displays the number of times the object is referenced to other objects.

Upgrading the firmware - CLI

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date. You can also use the CLI command `execute update-now` to update the antivirus and attack definitions. For more information, see the *FortiGate Administration Guide*.

Before you begin, ensure you have a TFTP server running and accessible to the FortiGate unit.

To upgrade the firmware using the CLI

- 1 Make sure the TFTP server is running.
- 2 Copy the new firmware image file to the root directory of the TFTP server.
- 3 Log into the CLI.
- 4 Make sure the FortiGate unit can connect to the TFTP server.

You can use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is 192.168.1.168:

```
execute ping 192.168.1.168
```

- 5 Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <filename> <tftp_ipv4>
```

Where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is 192.168.1.168, enter:

```
execute restore image tftp image.out 192.168.1.168
```

The FortiGate unit responds with the message:

```
This operation will replace the current firmware version!  
Do you want to continue? (y/n)
```

- 6 Type `y`.

The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.

- 7 Reconnect to the CLI.
- 8 Update antivirus and attack definitions, by entering:

```
execute update-now
```

USB Auto-Install

The USB Auto-Install feature automatically updates the FortiGate configuration file and firmware image file on a system reboot. Also, this feature provides you with an additional backup if you are unable to save your system settings before shutting down or rebooting your FortiGate unit.

You need an unencrypted configuration file for this feature. Also the required files, must be in the root directory of the USB key.



The FortiGate unit will only load a configuration file from a USB key when the FortiGate unit is restarted from a factory reset. This means that with any normal reboot commands, the FortiGate unit will not reload the configuration file.

This was done to ensure that if the USB key was left in the USB port, an older configuration would not be loaded by accident, losing any configuration settings changed after the initial save.

To configure the USB Auto-Install - web-based manager

- 1 Go to *System > Config > Advanced*.
- 2 Select the following:
 - On system restart, automatically update FortiGate configuration file if default file name is available on the USB disk.
 - On system restart, automatically update FortiGate firmware image if default image is available on the USB disk.
- 3 Enter the configuration and image file names or use the default configuration filename (system.conf) and default image name (image.out).
- 4 The default configuration filename should show in the *Default configuration file name* field.
- 5 Select *Apply*.

To configure the USB Auto-Install using the CLI

- 1 Log into the CLI.
- 2 Enter the following command:

```
config system auto-install
    set default-config-file <filename>
    set auto-install-config {enable | disable}
    set default-image-file <filename>
    set auto-install-image {enable | disable}
end
```

Reverting to a previous firmware version

This procedure reverts the FortiGate unit to its factory default configuration and deletes IPS custom signatures, web content lists, email filtering lists, and changes to replacement messages.

Before beginning this procedure, it is recommended that you:

- back up the FortiGate unit system configuration using the command `execute backup config`
- back up the IPS custom signatures using the command `execute backup ipsuserdefsig`
- back up web content and email filtering lists

To use the following procedure, you must have a TFTP server the FortiGate unit can connect to.

To revert to a previous firmware version using the CLI

- 1 Make sure the TFTP server is running
- 2 Copy the firmware image file to the root directory of the TFTP server.
- 3 Log into the FortiGate CLI.
- 4 Make sure the FortiGate unit can connect to the TFTP server execute by using the `execute ping` command.
- 5 Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <name_str> <tftp_ipv4>
```

Where `<name_str>` is the name of the firmware image file and `<tftp_ip4>` is the IP address of the TFTP server. For example, if the firmware image file name is `imagev28.out` and the IP address of the TFTP server is `192.168.1.168`, enter:

```
execute restore image tftp image28.out 192.168.1.168
```

The FortiGate unit responds with this message:

```
This operation will replace the current firmware version!  
Do you want to continue? (y/n)
```

6 Type `y`.

The FortiGate unit uploads the firmware image file. After the file uploads, a message similar to the following appears:

```
Get image from tftp server OK.  
Check image OK.  
This operation will downgrade the current firmware version!  
Do you want to continue? (y/n)
```

7 Type `y`.

The FortiGate unit reverts to the old firmware version, resets the configuration to factory defaults, and restarts. This process takes a few minutes.

8 Reconnect to the CLI.

9 To restore your previous configuration, if needed, use the command:

```
execute restore config <name_str> <tftp_ip4>
```

10 Update antivirus and attack definitions using the command:

```
execute update-now.
```

Installing firmware from a system reboot using the CLI

This procedure installs a firmware image and resets the FortiGate unit to default settings. You can use this procedure to upgrade to a new firmware version, revert to an older firmware version, or re-install the current firmware.

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to DB-9, or null modem cable.

This procedure reverts the FortiGate unit to its factory default configuration.

For this procedure you install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

Before beginning this procedure, ensure you back up the FortiGate unit configuration.

If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.



Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.

To install firmware from a system reboot

- 1 Connect to the CLI using the RJ-45 to DB-9 or null modem cable.
- 2 Make sure the TFTP server is running.
- 3 Copy the new firmware image file to the root directory of the TFTP server.
- 4 Make sure the internal interface is connected to the same network as the TFTP server.
- 5 To confirm the FortiGate unit can connect to the TFTP server, use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is 192.168.1.168:

```
execute ping 192.168.1.168
```

- 6 Enter the following command to restart the FortiGate unit.

```
execute reboot
```

The FortiGate unit responds with the following message:

```
This operation will reboot the system!  
Do you want to continue? (y/n)
```

- 7 Type *y*.

As the FortiGate unit starts, a series of system startup messages appears. When the following messages appears:

```
Press any key to display configuration menu.....
```

Immediately press any key to interrupt the system startup.



You have only 3 seconds to press any key. If you do not press a key soon enough, the FortiGate unit reboots and you must log in and repeat the execute reboot command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.  
[F]: Format boot device.  
[B]: Boot with backup firmware and set as default  
[C]: Configuration and information  
[Q]: Quit menu and continue to boot with default  
firmware.  
[H]: Display this list of options.
```

```
Enter G, F, Q, or H:
```

- 8 Type *G* to get to the new firmware image form the TFTP server.

The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

- 9 Type the address of the TFTP server and press Enter:

The following message appears:

```
Enter Local Address [192.168.1.188]:
```

- 10 Type an IP address the FortiGate unit can use to connect to the TFTP server. The IP address can be any IP address that is valid for the network the interface is connected to. Make sure you do not enter the IP address of another device on this network.

The following message appears:

```
Enter File Name [image.out]:
```

11 Enter the firmware image filename and press Enter.

The TFTP server uploads the firmware image file to the FortiGate unit and a message similar to the following appears:

```
Save as Default firmware/Backup firmware/Run image without
saving: [D/B/R]
```

12 Type D.

The FortiGate unit installs the new firmware image and restarts. The installation might take a few minutes to complete.

Backup and Restore from a USB key

Use a USB key to either backup a configuration file or restore a configuration file. You should always make sure a USB key is properly install before proceeding since the FortiGate unit must recognize that the key is installed in its USB port.



You can only save VPN certificates if you encrypt the file. Make sure the configuration encryption is enabled so you can save the VPN certificates with the configuration file. An encrypted file is ineffective if selected for the USB Auto-Install feature.

To backup configuration using the CLI

- 1 Log into the CLI.
- 2 Enter the following command to backup the configuration files:

```
exec backup config usb <filename>
```
- 3 Enter the following command to check the configuration files are on the key:

```
exec usb-disk list
```

To restore configuration using the CLI

- 1 Log into the CLI.
- 2 Enter the following command to restore the configuration files:

```
exec restore image usb <filename>
```

The FortiGate unit responds with the following message:

```
This operation will replace the current firmware version!
Do you want to continue? (y/n)
```

- 3 Type y.

Testing new firmware before installing

FortiOS enables you to test a new firmware image by installing the firmware image from a system reboot and saving it to system memory. After completing this procedure, the FortiGate unit operates using the new firmware image with the current configuration. This new firmware image is not permanently installed. The next time the FortiGate unit restarts, it operates with the originally installed firmware image using the current configuration. If the new firmware image operates successfully, you can install it permanently using the procedure [“Upgrading the firmware - web-based manager”](#) on [page 365](#).

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to DB-9 or null modem cable. This procedure temporarily installs a new firmware image using your current configuration.

For this procedure you install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

To test the new firmware image

- 1 Connect to the CLI using a RJ-45 to DB-9 or null modem cable.
- 2 Make sure the TFTP server is running.
- 3 Copy the new firmware image file to the root directory of the TFTP server.
- 4 Make sure the FortiGate unit can connect to the TFTP server using the `execute ping` command.
- 5 Enter the following command to restart the FortiGate unit:
`execute reboot`
- 6 As the FortiGate unit reboots, press any key to interrupt the system startup. As the FortiGate unit starts, a series of system startup messages appears. When the following messages appears:
`Press any key to display configuration menu....`
- 7 Immediately press any key to interrupt the system startup.



You have only 3 seconds to press any key. If you do not press a key soon enough, the FortiGate unit reboots and you must login and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default
[C]: Configuration and information
[Q]: Quit menu and continue to boot with default
firmware.
[H]: Display this list of options.
```

Enter G, F, Q, or H:

- 8 Type G to get the new firmware image from the TFTP server.
The following message appears:
Enter TFTP server address [192.168.1.168]:
- 9 Type the address of the TFTP server and press Enter:
The following message appears:
Enter Local Address [192.168.1.188]:
- 10 Type an IP address of the FortiGate unit to connect to the TFTP server.
The IP address must be on the same network as the TFTP server, but make sure you do not use the IP address of another device on the network.
The following message appears:
Enter File Name [image.out]:

- 11 Enter the firmware image file name and press Enter.

The TFTP server uploads the firmware image file to the FortiGate unit and the following appears.

```
Save as Default firmware/Backup firmware/Run image without
saving: [D/B/R]
```

- 12 Type R.

The FortiGate image is installed to system memory and the FortiGate unit starts running the new firmware image, but with its current configuration.

You can test the new firmware image as required. When done testing, you can reboot the FortiGate unit, and the FortiGate unit will resume using the firmware that was running before you installed the test firmware.

Controlled upgrade

Using a controlled upgrade, you can upload a new version of the FortiOS firmware to a separate partition in the FortiGate memory for later upgrade. The FortiGate unit can also be configured so that when it is rebooted, it will automatically load the new firmware (CLI only). Using this option, you can stage a number of FortiGate units to do an upgrade simultaneously to all devices using FortiManager or script.

To load the firmware for later installation - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Under *System Information > Firmware Version*, select *Update*.
- 3 Type the path and filename of the firmware image file, or select *Browse* and locate the file.
- 4 Deselect the *Boot the New Firmware* option
- 5 Select *OK*.

To load the firmware for later installation - CLI

```
execute restore secondary-image {ftp | tftp | usb}
```


To set the FortiGate unit so that when it reboots, the new firmware is loaded, use the CLI command...

```
execute set-next-reboot {primary | secondary}
```

... where {primary | secondary} is the partition with the preloaded firmware.

To trigger the upgrade using the web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Under *System Information > Firmware Version*, select *Details*.
- 3 Select the check box for the new firmware version.
The *Comments* column indicates which firmware version is the current active version.
- 4 Select *Upgrade* icon.



Interfaces

Interfaces, both physical and virtual, enable traffic to flow to and from the internal network, and the Internet and between internal networks. The FortiGate unit has a number of options for setting up interfaces and groupings of subnetworks that can scale to a company's growing requirements.

Physical

FortiGate units have a number of physical ports where you connect Ethernet or optical cables. Depending on the model, they can have anywhere from four to 40 physical ports. Some units have a grouping of ports labelled as internal, providing a built-in switch functionality.

In FortiOS, the port names, as labeled on the FortiGate unit, appear in the web-based manager in the *Unit Operation* the Dashboard. They also appear when you are configuring the interfaces, by going to *System > Network > Interface*. As shown below, the FortiGate-60C has eight interfaces

Figure 39: FortiGate-60C physical interfaces




Figure 40: FortiGate-60C interfaces on the Dashboard

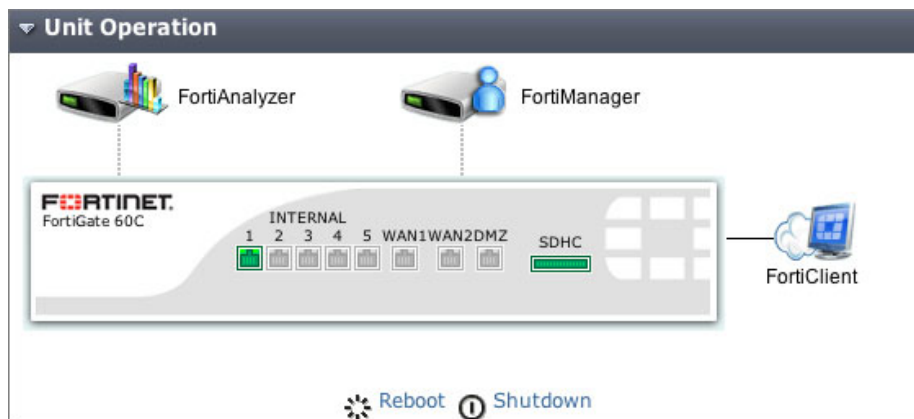


Figure 41: Configuring the FortiGate-60C ports

<input type="checkbox"/>	Name	IP/Netmask	Access	Administrative Status	Link Status	Ref.
<input type="checkbox"/>	dmz	10.10.10.1 / 255.255.255.0	HTTPS,PING	⬆	⬆	0
<input type="checkbox"/>	internal	172.20.120.129 / 255.255.255.0	HTTP,SSH	⬆	⬆	2
<input type="checkbox"/>	wan1	192.168.100.99 / 255.255.255.0	PING	⬆	⬆	2
<input type="checkbox"/>	wan2	192.168.101.99 / 255.255.255.0	PING	⬆	⬆	0

Normally the internal interface is configured as a single interface shared by all physical interface connections - a switch. The switch mode feature has two states - switch mode and interface mode. Switch mode is the default mode with only one interface and one address for the entire internal switch. Interface mode allows you to configure each of the internal switch physical interface connections separately. This enables you to assign different subnets and netmasks to each of the internal physical interface connections.

The larger FortiGate units can also include Advanced Mezzanine Cards (AMC), which can provide additional interfaces (ethernet or optical), with throughput enhancements for more efficient handling of specialized traffic. These interfaces appear in FortiOS as port amc/sw1, amc/sw2 and so on. In the following illustration, the FortiGate-3810A has three AMC cards installed: two single-width (amc/sw1, amc/sw2) and one double-width (amc/dw).

Figure 42: FortiGate-3810A AMC card port naming

<input type="checkbox"/>	Name	IP/Netmask	Access	Administrative Status	Link Status
<input type="checkbox"/>	amc-dw2/1	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-dw2/2	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw1/1	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw1/2	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw1/3	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw1/4	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw2/1	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw2/2	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw2/3	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	amc-sw2/4	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	april	0.0.0.0 / 0.0.0.0		⬆	⬆
<input type="checkbox"/>	port1	10.21.101.101 / 255.255.255.0	HTTPS,PING,SSH	⬆	⬆
<input type="checkbox"/>	port2	192.168.100.99 / 255.255.255.0	PING	⬆	⬆
<input type="checkbox"/>	port3	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port4	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port5	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port6	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port7	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port8	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port9	0.0.0.0 / 0.0.0.0	PING	⬆	⬆
<input type="checkbox"/>	port10	0.0.0.0 / 0.0.0.0	PING	⬆	⬆

•

Interface settings

In *System > Network > Interface*, you configure the interfaces, physical and virtual, for the FortiGate unit. There are different options for configuring interfaces when the FortiGate unit is in NAT mode or transparent mode. On FortiOS Carrier, you can also enable the Gi gatekeeper on each interface for anti-overbilling.

Interface page

Lists all the interfaces that are default and those that you have created. On this page you can view the status of each interface, create a new interface, edit an existing interface, or remove an interface.

Create New	<p>Select to add a new interface, zone or, in transparent mode, port pair. For more information on configuring zones, see “Zones” on page 392. For more information on port pairing in transparent mode, see “Port pairing” on page 269.</p> <p>Depending on the model you can add a VLAN interface, a loopback interface, a IEEE 802.3ad aggregated interface, or a redundant interface.</p> <p>When VDOMs are enabled, you can also add Inter-VDOM links.</p>
Switch Mode	<p>On supported models, select <i>Switch Mode</i> to change between switch mode and interface mode. Interface mode enables you to configure a switch interface to be separate configurable interfaces.</p> <p>On some FortiGate models you can also select <i>Hub Mode</i>. Hub mode is similar to switch mode except that in hub mode the interfaces do not learn the MAC addresses of the devices on the network they are connected to and may also respond quicker to network changes. Normally, you would only select <i>Hub Mode</i> if you are having network performance issues when operating with switch mode.</p> <p>Before switching modes, all configuration settings for the interfaces affected by the change must be set to defaults, that is, no assigned IP addresses or used in any other part of the device. When you select <i>Switch Mode</i> the web-based manager displays the list of affected interfaces.</p>
Show backplane interfaces	<p>Select to make FortiGate-5000 series backplane interfaces visible. Once visible, these interfaces can be configured as regular physical interfaces.</p>
Column Settings	<p>Select to change the columns of information that are displayed on the interface list.</p>
Description (appears as a note icon)	<p>Displays a description for the interface when one has been added. Move your mouse over the icon to view the description.</p>

Name	The names of the physical interfaces on your FortiGate unit. This includes any alias names that have been configured.
	When you combine several interfaces into an aggregate or redundant interface, only the aggregate or redundant interface is listed, not the component interfaces.
	On FortiGate models that support switch mode, the individual interfaces in the switch are not displayed when in switch mode. For more information, see “Switch Mode” on page 382 .
	If you have added VLAN interfaces, they also appear in the name list, below the physical or aggregated interface to which they have been added.
	If you have added loopback interfaces, they also appear in the interface list, below the physical interface to which they have been added. If you have software switch interfaces configured, you will be able to view them. For more information, see “Software switch interfaces” on page 388 .
	If your FortiGate unit supports AMC modules, the interfaces are named amc-sw1/1, amc-dw1/2, and so on.
IP/Netmask	The current IP address/netmask of the interface. In VDOM mode, when VDOMs are not all in NAT or Transparent mode some values may not be available for display and will be displayed as “-”.
Access	The administrative access configuration for the interface.
Administrative Status	Indicates if the interface can be accessed for administrative purposes. If the administrative status is a green arrow, and administrator could connect to the interface using the configured access. If the administrative status is a red arrow, the interface is administratively down and cannot be accessed for administrative purposes.
Link Status	The status of the interface physical connection. Link status can be either up or down. If link status is up (green arrow) the interface is connected to the network and accepting traffic. If link status is down (red arrow) the interface is not connected to the network or there is a problem with the connection. You cannot change link status from the web-based manager. Link status is only displayed for physical interfaces.
MAC	The MAC address of the interface.
Mode	Shows the addressing mode of the interface. The addressing mode can be manual, DHCP, or PPPoE.
MTU	The maximum number of bytes per transmission unit (MTU) for the interface.
Secondary IP	Displays the secondary IP addresses added to the interface.

Type	<p>The type of the interface. Valid types include:</p> <ul style="list-style-type: none"> Physical VLAN Aggregate Redundant VDOM Link Pair Switch Tunnel VAP - a wireless virtual access point (VAP or virtual AP) interface
Virtual Domain	The virtual domain to which the interface belongs. This column is visible when VDOM configuration is enabled.
VLAN ID	The configured VLAN ID for VLAN subinterfaces.
Delete	Removes an interface from the list on the Interface page. Available for interfaces added by selecting <i>Create New</i> . For example, you can remove VLAN, loopback, aggregate, and redundant interfaces. You can only remove an interface if it is not used in another configuration.
Edit	Modifies settings within the configuration of an interface.
View	View the interface's configuration. Displays the number of times the object is referenced to other objects.
Ref.	To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object.

•

Interface configuration and settings

To configure an interface, go to *System > Network > Interface*.

New Interface page	
Provides settings for configuring a new interface.	
Name	Enter a name of the interface. Physical interface names cannot be changed.
Alias	Enter an alternate name for a physical interface on the FortiGate unit. The alias can be a maximum of 25 characters. The alias name will not appear in logs. This field appears when editing an existing physical interface.
Link Status	Indicates whether the interface is connected to a network (link status is <i>Up</i>) or not (link status is <i>Down</i>). This field appears when editing an existing physical interface.
Type	<p>Select the type of interface that you want to add.</p> <p>On some models you can set <i>Type</i> to <i>802.3ad Aggregate</i> or <i>Redundant Interface</i>.</p>

Interface	<p>Select the name of the physical interface to which to add a VLAN interface. Once created, the VLAN interface is listed below its physical interface in the Interface list.</p> <p>You cannot change the physical interface of a VLAN interface except when adding a new VLAN interface.</p> <p>Displayed when <i>Type</i> is set to <i>VLAN</i>.</p>
VLAN ID	<p>Enter the VLAN ID. You cannot change the <i>VLAN ID</i> except when add a new VLAN interface.</p> <p>The VLAN ID can be any number between 1 and 4094 and must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch connected to the VLAN subinterface.</p> <p>Displayed when <i>Type</i> is set to <i>VLAN</i>.</p>
Virtual Domain	<p>Select the virtual domain to add the interface to.</p> <p>Admin accounts with super_admin profile can change the <i>Virtual Domain</i>.</p>
Physical Interface Members	<p>This section has two different forms depending on the interface type:</p> <ul style="list-style-type: none"> • Software switch interface - this section is a display-only field showing the interfaces that belong to the software switch virtual interface. • 802.3ad aggregate or Redundant interface - this section includes available interface and selected interface lists to enable adding or removing interfaces from the interface. For more information, see Redundant interfaces. <p>Select interfaces from this <i>Available Interfaces</i> list and select the right arrow to add an interface to the <i>Selected Interface</i> list.</p>
Addressing mode	<p>Select the addressing mode for the interface.</p> <ul style="list-style-type: none"> • Select <i>Manual</i> and add an <i>IP/Netmask</i> for the interface. If IPv6 configuration is enabled you can add both a IPv4 and an IPv6 IP address. • Select <i>DHCP</i> to get the interface IP address and other network settings from a DHCP server. See DHCP on an interface • Select <i>PPPoE</i> to get the interface IP address and other network settings from a PPPoE server. See PPPoE on an interface.
IP/Netmask	<p>If <i>Addressing Mode</i> is set to <i>Manual</i>, enter an IPv4 address/subnet mask for the interface. FortiGate interfaces cannot have IP addresses on the same subnet.</p>
IPv6 Address	<p>If <i>Addressing Mode</i> is set to <i>Manual</i> and IPv6 support is enabled, enter an IPv6 address/subnet mask for the interface. A single interface can have both an IPv4 and IPv6 address or just one or the other.</p>

Enable one-arm sniffer	Available when editing a physical interface. Select to configure this interface to operate as a one-armed sniffer as part of configuring a FortiGate unit to operate as an IDS appliance by sniffing packets for attacks without actually receiving and otherwise processing the packets. Once the interface is enabled for sniffing you cannot use the interface for other traffic. You must add sniffer policies for the interface to actually sniff packets.
Dedicate this interface to FortiAP connection	Select have the FortiAP connect exclusively to the interface. This option is only available when editing a physical interface, and it has a static IP address.
Reserve IP addresses for FortiAP	Enter the IP address range that will be used for the FortiAP units. When you enter the reserved IP address range, the FortiGate unit automatically creates a DHCP server.
Enable explicit Web Proxy	This option is not available for a VLAN interface selection. Select to enable explicit web proxying on this interface. When enabled, this interface will be displayed on <i>System > Network > Web Proxy</i> under <i>Listen on Interfaces</i> and web traffic on this interface will be proxied according to the Web Proxy settings.
Override Default MTU Value	<p>To change the MTU, select Override default MTU value (1 500) and enter the MTU size based on the addressing mode of the interface</p> <ul style="list-style-type: none"> 68 to 1 500 bytes for static mode 576 to 1 500 bytes for DHCP mode 576 to 1 492 bytes for PPPoE mode larger frame sizes if supported by the FortiGate model <p>Only available on physical interfaces. Virtual interfaces associated with a physical interface inherit the physical interface MTU size.</p> <p>In transparent mode, if you change the MTU of an interface, you must change the MTU of all interfaces to match the new MTU.</p>
Administrative Access	Select the types of administrative access permitted for IPv4 connections to this interface.
HTTPS	Allow secure HTTPS connections to the web-based manager through this interface.
PING	Interface responds to pings. Use this setting to verify your installation and for testing.
HTTP	Allow HTTP connections to the web-based manager through this interface. HTTP connections are not secure and can be intercepted by a third party.
SSH	Allow SSH connections to the CLI through this interface.
SNMP	Allow a remote SNMP manager to request SNMP information by connecting to this interface.
TELNET	Allow Telnet connections to the CLI through this interface. Telnet connections are not secure and can be intercepted by a third party.
FMG-Access	Allow FortiManager authorization automatically during the communication exchange between the FortiManager and FortiGate units. For example, port1 on the FortiGate unit allows the FortiManager unit to connect to it.

IPv6 Administrative Access	Select the types of administrative access permitted for IPv6 connections to this interface. These types are the same as for Administrative Access.
Detect Interface Status for Gateway Load Balancing	Configure interface status detection for the main interface IP address.
Detect Server	Enter the server's IP address.
Detect Protocol	Select the check box beside the protocol that will be detected. for gateway load balancing.
Weight	Enter the load balancing weight.
Spillover Threshold	Enter the spillover threshold number in KBps.
Secondary IP Address	Add additional IPv4 addresses to this interface. Select the Expand Arrow to expand or hide the section.
Comments	Enter a description up to 63 characters to describe the interface.
Administrative Status	<p>Select either <i>Up</i> (green arrow) or <i>Down</i> (red arrow) as the status of this interface.</p> <ul style="list-style-type: none"> <i>Up</i> indicates the interface is active and can accept network traffic. <i>Down</i> indicates the interface is not active and cannot accept traffic.
Gi Gatekeeper (FortiOS Carrier only)	For FortiOS Carrier, enable Gi Gatekeeper to enable the Gi firewall as part of the anti-overbilling configuration. You must also configure <i>Gi Gatekeeper Settings</i> by going to <i>System > Admin > Settings</i> .

Switch Mode

Switch mode enables you to switch a group of related FortiGate interfaces to operate as a multi-port switch with one IP address. Switch mode is available on models with switch hardware; however, it is not available on models running FortiOS Carrier.

Switch mode has two states - switch mode and interface mode. Switch mode enables you to combine multiple interfaces into a single switch with one interface and one address. Interface mode enables you to configure each of the internal switch physical interface connections separately.

Before you are able to change between switch mode and interface mode, all configuration settings for the affected interfaces must be set to defaults. This includes security policies, routing, DNS forwarding, DHCP services, VDOM interface assignments, and routing. If they are not removed, you will not be able to switch modes, and you will see an error message. The web-based manager displays the list of affected interfaces.

To configure a mode, go to *System > Network > Interface*, and select *Switch Mode* at the top of the interface table.

Interface page containing switch mode settings	
Provides settings for switching a group of related FortiGate interfaces to operate as a multi-port switch with one IP address.	
Switch Mode	Only one internal interface is displayed. This is the default mode.

Interface Mode	All internal interfaces on the switch are displayed as individually configurable interfaces.
Hub Mode	On some models you can select <i>Hub Mode</i> . Hub mode is similar to switch mode except that in hub mode the interfaces do not learn the MAC addresses of the devices on the network they are connected to and may also respond quicker to network changes in some circumstances. You should only select <i>Hub Mode</i> if you are having network performance issues when operating with switch mode. The configuration of the FortiGate unit is the same whether in switch mode or hub mode.

Loopback interfaces

A loopback interface is a logical interface that is always up (no physical link dependency) and the attached subnet is always present in the routing table.

The FortiGate's loopback IP address does not depend on one specific external port, and is therefore possible to access it through several physical or VLAN interfaces. Multiple loopback interfaces can be configured in either non-VDOM mode or in each VDOM.

Loopback interfaces still require appropriate firewall policies to allow traffic to and from this type of interface.

A loopback interface can be used with:

- Management access
- BGP (TCP) peering
- PIM RP

Loopback interfaces are a good practice for OSPF. Setting the OSPF router ID the same as loopback IP address troubleshooting OSPF easier, and remembering the management IP addresses (telnet to "router ID").

Dynamic routing protocols can be enabled on loopback interfaces

For blackhole static route, use the blackhole route type instead of the loopback interface.

Redundant interfaces

On some models you can combine two or more physical interfaces to provide link redundancy. This feature enables you to connect to two or more switches to ensure connectivity in the event one physical interface or the equipment on that interface fails.

In a redundant interface, traffic is only going over one interface at any time. This differs from an aggregated interface where traffic is going over all interfaces for distribution of increased bandwidth. This difference means redundant interfaces can have more robust configurations with fewer possible points of failure. This is important in a fully-meshed HA configuration.

An interface is available to be in a redundant interface if:

- it is a physical interface, not a VLAN interface
- it is not already part of an aggregated or redundant interface
- it is in the same VDOM as the redundant interface
- it has no defined IP address
- is not configured for DHCP or PPPoE
- it has no DHCP server or relay configured on it
- it does not have any VLAN subinterfaces

- it is not referenced in any security policy, VIP, or multicast policy
- it is not monitored by HA
- it is not one of the FortiGate-5000 series backplane interfaces

When an interface is included in a redundant interface, it is not listed on the *System > Network > Interface* page. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, or routing.

DHCP on an interface

If you configure an interface to use DHCP, the FortiGate unit automatically broadcasts a DHCP request from the interface. The interface is configured with the IP address and any DNS server addresses and default gateway address that the DHCP server provides.



DHCP IPv6 is similar to DHCP IPv4, however there is:

- no default gateway option defined because a host learns the gateway using router advertisement messages
- there is no WINS servers because it is obsolete.

For more information about DHCP IPv6, see RFC 3315.

Configure DHCP for an interface in *System > Network > Interface* and selecting the interface from the list, and selecting *DHCP* in the *Address Mode*. The table describes the DHCP status information when DHCP is configured for an interface.

Addressing mode section of New Interface page for DHCP information

Status	<p>Displays DHCP status messages as the interface connects to the DHCP server and gets addressing information. Select <i>Status</i> to refresh the addressing mode status message.</p> <p>Status can be one of:</p> <ul style="list-style-type: none"> • initializing - No activity. • connecting - interface attempts to connect to the DHCP server. • connected - interface retrieves an IP address, netmask, and other settings from the DHCP server. • failed - interface was unable to retrieve an IP address and other settings from the DHCP server.
Obtained IP/Netmask	<p>The IP address and netmask leased from the DHCP server.</p> <p>Only displayed if <i>Status</i> is <i>connected</i>.</p>
Renew	<p>Select to renew the DHCP license for this interface.</p> <p>Only displayed if <i>Status</i> is <i>connected</i>.</p>
Expiry Date	<p>The time and date when the leased IP address and netmask is no longer valid for the interface. The IP address is returned to the pool to be allocated to the next user request for an IP address.</p> <p>Only displayed if <i>Status</i> is <i>connected</i>.</p>
Default Gateway	<p>The IP address of the gateway defined by the DHCP server.</p> <p>Only displayed if <i>Status</i> is <i>connected</i>, and if <i>Receive default gateway from server</i> is selected.</p>

Distance	Enter the administrative distance for the default gateway retrieved from the DHCP server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route.
Retrieve default gateway from server	Enable to retrieve a default gateway IP address from the DHCP server. The default gateway is added to the static routing table.
Override internal DNS	Enable to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page. When VDOMs are enabled, you can override the internal DNS only on the management VDOM.

PPPoE on an interface

If you configure the interface to use PPPoE, the FortiGate unit automatically broadcasts a PPPoE request from the interface.

The FortiGate units support many PPPoE RFC features (RFC 2516) including unnumbered IPs, initial discovery timeout and PPPoE Active Discovery Terminate (PADT).

Configure PPPoE on an interface in *System > Network > Interface*. The table describes the PPPoE status information when PPPoE is configured for an interface.

Addressing mode section of New Interface page	
Status	Displays PPPoE status messages as the FortiGate unit connects to the PPPoE server and gets addressing information. Select Status to refresh the addressing mode status message. The status is only displayed if you selected <i>Edit</i> . Status can be any one of the following 4 messages.
initializing	No activity.
connecting	The interface is attempting to connect to the PPPoE server.
connected	The interface retrieves an IP address, netmask, and other settings from the PPPoE server. When the status is connected, PPPoE connection information is displayed.
failed	The interface was unable to retrieve an IP address and other information from the PPPoE server.
Reconnect	Select to reconnect to the PPPoE server. Only displayed if Status is connected.
User Name	The PPPoE account user name.
Password	The PPPoE account password.
Unnumbered IP	Specify the IP address for the interface. If your ISP has assigned you a block of IP addresses, use one of them. Otherwise, this IP address can be the same as the IP address of another interface or can be any IP address.
Initial Disc Timeout	Enter Initial discovery timeout. Enter the time to wait before starting to retry a PPPoE discovery.

Initial PADT timeout	Enter Initial PPPoE Active Discovery Terminate (PADT) timeout in seconds. Use this timeout to shut down the PPPoE session if it is idle for this number of seconds. PADT must be supported by your ISP. Set initial PADT timeout to 0 to disable.
Distance	Enter the administrative distance for the default gateway retrieved from the PPPoE server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route. The default distance for the default gateway is 1.
Retrieve default gateway from server	Enable to retrieve a default gateway IP address from a PPPoE server. The default gateway is added to the static routing table.
Override internal DNS	Enable to replace the DNS server IP addresses on the System DNS page with the DNS addresses retrieved from the PPPoE server. When VDOMs are enabled, you can override the internal DNS only on the management VDOM.

Administrative access

Interfaces, especially the public-facing ports can be potentially accessed by those who you may not want access to the FortiGate unit. When setting up the FortiGate unit, you can set the type of protocol an administrator must use to access the FortiGate unit. The options include:

- HTTPS
- HTTP
- SSH
- TELNET
- PING
- SNMP

You can select as many, or as few, even none, that are accessible by an administrator.

Example

This example adds an IPv4 address 172.20.120.100 to the WAN1 interface as well as the administrative access to HTTPS and SSH. As a good practice, set the administrative access when you are setting the IP address for the port.

To add an IP address on the WAN1 interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the WAN1 interface row and select *Edit*.
- 3 Select the *Addressing Mode of Manual*.
- 4 Enter the IP address for the port of 172.20.120.100/24.
- 5 For *Administrative Access*, select HTTPS and SSH.
- 6 Select *OK*.

To create IP address on the WAN1 interface - CLI

```
config system interface
```



```
edit wan1
  set ip 172.20.120.100/24
  set allowaccess https ssh
end
```



When adding to, or removing a protocol, you must type the entire list again. For example, if you have an access list of HTTPS and SSH, and you want to add PING, typing:

```
set allowaccess ping
```

...only PING will be set. In this case, you must type...

```
set allowaccess https ssh ping
```

Wireless

A wireless interface is similar to a physical interface only it does not include a physical connection. The FortiWiFi units enables you to add multiple wireless interfaces that can be available at the same time (the FortiWiFi-30B can only have one wireless interface). On FortiWiFi units, you can configure the device to be either an access point, or a wireless client. As an access point, the FortiWiFi unit can have up to four separate SSIDs, each on their own subnet for wireless access. In client mode, the FortiWiFi only has one SSID, and is used as a receiver, to enable remote users to connect to the existing network using wireless protocols.

Wireless interfaces also require additional security measures to ensure the signal does not get hijacked and data tampered or stolen.

Interface MTU packet size

You can change the maximum transmission unit (MTU) of the packets that the FortiGate unit transmits to improve network performance. Ideally, the MTU should be the same as the smallest MTU of all the networks between the FortiGate unit and the destination of the packets. If the packets that the FortiGate unit sends are larger than the smallest MTU, they are broken up or fragmented, which slows down transmission. You can easily experiment by lowering the MTU to find an MTU size for optimum network performance.

Interfaces on some models support frames larger than the traditional 1 500 bytes. Contact Fortinet Customer Support for the maximum frame sizes that your FortiGate unit supports.

If you need to enable sending larger frames over a route, you need all Ethernet devices on that route to support that larger frame size, otherwise your larger frames will not be recognized and are dropped.

If you have standard size and larger size frame traffic on the same interface, routing alone cannot route them to different routes based only on frame size. However, you can use VLANs to make sure the larger frame traffic is routed over network devices that support that larger size. VLANs will inherit the MTU size from the parent interface. You will need to configure the VLAN to include both ends of the route as well as all switches and routers along the route.

MTU packet size is changed in the CLI. If you select an MTU size larger than your FortiGate unit supports, an error message will indicate this. In this situation, try a smaller MTU size until the value is supported.



If you change the MTU, you need to reboot the FortiGate unit to update the MTU value of VLAN subinterfaces on the modified interface.

In Transparent mode, if you change the MTU of an interface, you must change the MTU of all interfaces on the FortiGate unit to match the new MTU.

Secondary IP addresses to an interface

If an interface is configured with a manual or static IP address, you can also add secondary static IP addresses to the interface. Adding secondary IP addresses effectively adds multiple IP addresses to the interface. Secondary IP addresses cannot be assigned using DHCP or PPPoE.

All of the IP addresses added to an interface are associated with the single MAC address of the physical interface and all secondary IP addresses are in the same VDOM as the interface that are added to. You configure interface status detection for gateway load balancing separately for each secondary IP addresses. As with all other interface IP addresses, secondary IP addresses cannot be on the same subnet as any other primary or secondary IP address assigned to a FortiGate interface unless they are in separate VDOMs.

To configure a secondary IP, go to *System > Network > Interface*, select *Edit* or *Create New* and select the *Secondary IP Address* check box.

See also

- [Interface configuration and settings](#)

Software switch interfaces

A software switch, or soft switch, is a virtual switch that is implemented at the software, or firmware level, rather than the hardware level. Adding a software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiGate unit.

It can also be useful if you require more hardware ports on for the switch on a FortiGate unit. For example, if your FortiGate unit has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces such as those with FortiWiFi and FortiAP unit.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

There are a few things to consider when setting up a software switch:

- Ensure you create a back up of the configuration.

- Ensure you have at least one port or connection such as the console port to connect to the FortiGate unit. If you accidentally combine too many ports, you will need a way to undo any errors.

The ports that you include must not have any link or relation to any other aspect of the FortiGate unit. For example, DHCP servers, security policies, and so on.

To add a software switch interface

- 1 Go to *System > Network > Interface* and select *Create New*.
- 2 Set *Type* to *Software Switch*.
- 3 Enter a name, select the interfaces to be in the software switch, and configure other standard interface options as required.

Virtual domains

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. A single FortiGate unit is then flexible enough to serve multiple departments of an organization, separate organizations, or to act as the basis for a service provider's managed security service.



Some smaller FortiGate units do not support virtual domains.

VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations. By default, each FortiGate unit has a VDOM named root. This VDOM includes all of the FortiGate physical interfaces, modem, VLAN subinterfaces, zones, security policies, routing settings, and VPN settings.

When a packet enters a VDOM, it is confined to that VDOM. In a VDOM, you can create security policies for connections between Virtual LAN (VLAN) subinterfaces or zones in the VDOM. Packets do not cross the virtual domain border internally. To travel between VDOMs, a packet must pass through a firewall on a physical interface. The packet then arrives at another VDOM on a different interface, but it must pass through another firewall before entering the VDOM. Both VDOMs are on the same FortiGate unit. Inter-VDOMs change this behavior in that they are internal interfaces; however their packets go through all the same security measures as on physical interfaces.

-

Example

This example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First enable Virtual Domains on the FortiGate unit. When you enable VDOMs, the FortiGate unit will log you out.

To enable VDOMs - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 In the *System Information* widget, select *Enable for Virtual Domain*.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all Virtual Domains.

To enable VDOMs - CLI

```
config system global
    set vdom-admin enable
end
```

Next, add the VDOM called accounting.

To add a VDOM - web-based manager

- 1 Go to *System > VDOM > VDOM*, and select *Create New*.
- 2 Enter the VDOM name *accounting*.
- 3 Select *OK*.

To add a VDOM - CLI

```
config vdom
    edit <new_vdom_name>
end
```

With the Virtual Domain created, you can assign a physical interface to it, and assign it an IP address.

To assign physical interface to the accounting Virtual Domain - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the DMZ2 port row and select *Edit*.
- 3 For the *Virtual Domain* drop-down list, select *accounting*.
- 4 Select the *Addressing Mode* of *Manual*.
- 5 Enter the IP address for the port of 10.13.101.100/24.
- 6 Set the *Administrative Access* to *HTTPS* and *SSH*.
- 7 Select *OK*.

To assign physical interface to the accounting Virtual Domain - CLI

```
config global
    config system interface
        edit dmz2
            set vdom accounting
            set ip 10.13.101.100/24
            set allowaccess https ssh
        next
    end
```

Virtual LANs

The term VLAN subinterface correctly implies the VLAN interface is not a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network route that is configured for this VLAN. Without that route, the VLAN will not be connected to the network, and VLAN traffic will not be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

FortiGate unit interfaces cannot have overlapping IP addresses—the IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask. This rule helps prevent a broadcast storm or other similar network problems.

Any FortiGate unit, with or without VDOMs enabled, can have a maximum of 255 interfaces in Transparent operating mode. In NAT/Route operating mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in Transparent operating mode, you need to configure multiple VDOMs with many interfaces on each VDOM.

Example

This example shows how to add a VLAN, `vlan_accounting` on the FortiGate unit internal interface with an IP address of 10.13.101.101.

To add a VLAN - web-based manager

- 1 Go to *System > Network > Interface* and select *Create New*.

The *Type* is by default set to VLAN.

- 2 Enter a name for the VLAN to `vlan_accounting`.
- 3 Select the *Internal* interface.
- 4 Enter the *VLAN ID*.

The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together.

- 5 Select the *Addressing Mode of Manual*.
- 6 Enter the IP address for the port of 10.13.101.101/24.
- 7 Set the *Administrative Access to HTTPS* and *SSH*.
- 8 Select *OK*.

To add a VLAN - CLI


```
config system interface
  edit VLAN_1
    set interface internal
    set type vlan
    set vlanid 100
    set ip 10.13.101.101/24
    set allowaccess https ssh
  next
end
```

Zones

Zones are a group of one or more FortiGate interfaces, both physical and virtual, that you can apply security policies to control inbound and outbound traffic. Grouping interfaces and VLAN subinterfaces into zones simplifies the creation of security policies where a number of network segments can use the same policy settings and protection profiles. When you add a zone, you select the names of the interfaces and VLAN subinterfaces to add to the zone. Each interface still has its own address and routing is still done between interfaces, that is, routing is not affected by zones. Security policies can also be created to control the flow of intra-zone traffic.

For example, in the illustration below, the network includes three separate groups of users representing different entities on the company network. While each group has its own set of port and VLANs, in each area, they can all use the same security policy and protection profiles to access the Internet. Rather than the administrator making nine separate security policies, he can add the required interfaces to a zone, and create three policies, making administration simpler.

Figure 43: Network zones



You can configure policies for connections to and from a zone, but not between interfaces in a zone. Using the above example, you can create a security policy to go between zone 1 and zone 3, but not between WAN2 and WAN1, or WAN1 and DMZ1.

Example

This example explains how to set up a zone on the FortiGate unit to include the Internal interface and a VLAN.

To create a zone - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the arrow on the *Create New* button and select *Zone*.

- 3 Enter a zone name of Zone_1.
- 4 Select the Internal interface and the virtual LAN interface vlan_accounting from the previous section.
- 5 Select OK.

To create a zone - CLI

```
config system zone
  edit Zone_1
    set interface internal VLAN_1
end
```

IPv6

Internet Protocol version 6 (IPv6) is the next-generation version of IP addressing, to eventually replace IPv4. IPv6 was developed because there is a concern that in the near future, the available addresses for the IPv4 infrastructure will be exhausted. The IPv6 infrastructure will supplement, and eventually, replace the IPv4 standard.

Where IPv4 uses 32 bit addressing, IPv6 uses 128 bit addressing, effectively providing trillions upon trillions of unique addresses, whereas IPv4 can have a little over 4 billion. With this larger address space, allocating addresses and routing traffic becomes easier, and network address translation (NAT) becomes virtually unnecessary.

Where IPv4 addresses are written numerals separated by a decimal, the IPv6 address is written with hexadecimal digits separated by a colon. For example, fe80:218:8bff:fe84:4223.

By default, the FortiGate unit is not enabled to use IPv6 addressing. To enable this feature, go to *System > Admin > Settings* and select *IPv6 Support on GUI*. When enabled you can use IPv6 addressing on any of the address-dependant components of the FortiGate unit, including security policies, interface addressing, DNS servers. IPv6 addressing can be configured on the web-based manager and in the CLI.

For further information on IPV6 in FortiOS, see [“Internet Protocol version 6 \(IPv6\)” on page 235](#).

Example

This example adds an IPv6 address 2001:db8:0:1234:0:567:1:1 for the WAN1 interface as well as the administrative access to HTTPS and SSH. As a good practice, set the administrative access when you are setting the IP address for the port.

To add an IP address for the WAN1 interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select WAN1 row and select *Edit*.
- 3 Select the *Addressing Mode of Manual*.
- 4 Enter the *IPv6 Address* of 2001:db8:0:1234:0:567:1:1.
- 5 For *Administrative Access* select *HTTPS* and *SSH*.
- 6 Select OK.

To create IP address for the WAN1 interface - CLI

```
config system interface
  edit wan1
```

```
config ipv6
  set ip6-address 2001:db8:0:1234:0:567:1:1
  set ip6-allowaccess https ssh
end
end
```




Central management

Administering one or two FortiGate units is fairly simple enough, especially when they are in the same room or building. However, if you are administering many FortiGate units that may be located in locations in a large geographical area, or in the world, you will need a more efficient method of maintaining firmware upgrades, configuration changes and updates.

The FortiManager family of appliances supply the tools needed to effectively manage any size Fortinet security infrastructure, from a few devices to thousands of appliances. The appliances provide centralized policy-based provisioning, configuration, and update management. They also offer end-to-end network monitoring for added control. Managers can control administrative access and simplify policy deployment using role-based administration to define user privileges for specific management domains and functions by aggregating collections of Fortinet appliances and agents into independent management domains. By locally hosting security content updates for managed devices and agents, FortiManager appliances minimize Web filtering rating request response time and maximize network protection.

This chapter describes the basics of using FortiManager as an administration tool for multiple FortiGate units. It describes the basics of setting up a FortiGate unit in FortiManager, and some key management features you can use within FortiManager to manage the FortiGate unit. For full details and instructions on FortiManager, see the [FortiManager Administration Guide](#).

This section includes the topics:

- [Adding a FortiGate to FortiManager](#)
- [Configuration through FortiManager](#)
- [Firmware updates](#)
- [FortiGuard](#)
- [Backup and restore configurations](#)
- [Administrative domains](#)

Adding a FortiGate to FortiManager

Before you can use the FortiManager unit to maintain a FortiGate, you need to add it to the FortiManager unit. To do this requires configuration on both the FortiGate and FortiManager. This section describes the basics to configure management using a FortiManager device. For more information on the interaction of FortiManager with the FortiGate unit, see the FortiManager documentation.

FortiGate configuration

These steps ensure that the FortiGate unit will be able to receive updated antivirus and IPS updates, and allow remote management through the FortiManager system. You can add a FortiGate unit whether it is running in either NAT mode or transparent mode. The FortiManager unit provides remote management of a FortiGate unit over TCP port 541.



If you have not already done so, register the FortiGate unit by visiting <http://support.fortinet.com> and select *Product Registration*. By registering your Fortinet unit, you will receive updates to threat detection and prevention databases (Antivirus, Intrusion Detection, etc.) and will also ensure your access to technical support.

You must enable the FortiGate management option so the FortiGate unit can accept management updates to firmware, antivirus signatures and IPS signatures.

To configure the FortiGate unit - web-based manager

- 1 Log in to the FortiGate unit.
- 2 Go to *System > Admin > Settings*.
- 3 Enter the *IP address* for the FortiManager.
- 4 Select *Send Request*.

The FortiManager ID appears in the Trusted FortiManager table, and can now be managed by the FortiManager unit, once you add it to the Device Manager.

As an additional security measure, you can also select *Registration Password* and enter a password to connect to the FortiManager in an upcoming FortiManager release.

To configure the FortiGate unit - CLI

```
config system central-management
  set fmg <ip_address>
end
```

To use the registration password in an upcoming FortiManager release enter:

```
execute central-mgmt register-device <fmg-serial-no><fmg-register-
password><fgt-username><fgt-password>
```

Configuring an SSL connection

With FortiManager 4.0 MR2 Patch 6 and FortiOS 4.0 MR3, you can configure an SSL connection between the two devices, and select the encryption level.

Use the following CLI commands in FortiOS to configure the encryption connection:

```
config central-management setting
  set status enable
  set enc-algorithm {default* | high | low | disable}
end
```

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for high, medium, and low follows openssl definitions:

- **High** - Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
Algorithms are: DHE-RSA-AES256-SHA:AES256-SHA:EDH-RSA-DES-CBC3-SHA:DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA
- **Medium** - Key strengths of 128 bit encryption.
Algorithm are:RC4-SHA:RC4-MD5:RC4-MD

- **Low** - Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites

Algorithms: EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5.

FortiManager configuration

After enabling Central Management and indicating the FortiManager unit that will provide the management of the FortiGate unit, you can add it to the Device Manager in the FortiManager web-based manager.

To add the FortiGate unit to the Device Manager in FortiManager

- 1 Log in to the FortiManager unit.
- 2 Select the Device Manager tab.
- 3 Select *Add Device* from the top tool bar.
- 4 Enter the *IP address* of the FortiGate unit.
- 5 Enter the *Name* of the device.
This can be the model name, or functional name, such as West Building, or Accounting Firewall.
- 6 Enter the remaining information as required.
- 7 Select *OK*.

Configuration through FortiManager

With the FortiManager system, you can monitor and configure multiple FortiGate units from one location and log in. Within the FortiManager system, you can view a FortiGate unit and its web-based manager from the Device Manager. From there you can make the usual configuration updates and changes, without having to log in and out of multiple FortiGate units.

When under control of a FortiManager system, administrators will not be able to configure the FortiGate unit. When trying to change options, the FortiGate unit displays a message that it is configured through FortiManager, and any changes may be reverted.

FortiManager enables you to complete the configuration, by going to the Device Manager, selecting the FortiGate unit and using the same menu structure and pages as you would see in the FortiGate web-based manager. All changes to the FortiGate configuration are stored locally on the FortiManager unit until you synchronize with the FortiGate unit.

Global objects

If you are maintaining a number of FortiGate units within a network, many of the policies and configuration elements will be the same across the corporation. In these instances, the adding and editing of many of the same policies will become a tedious and error-prone activity. With FortiManager global objects, this level of configuration is simplified.

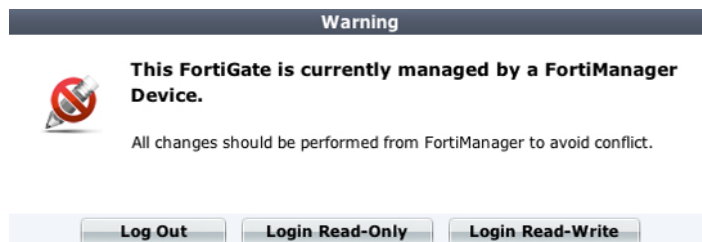
A global object is an object that is not associated specifically with one device or group. Global objects include security policies, a DNS server, VPN, and IP pools.

The Global Objects window is where you can configure global objects and copy the configurations to the FortiManager device database for a selected device or a group of devices. You can also import configurations from the FortiManager device database for a selected device and modify the configuration as required.

When configuring or creating a global policy object the interface, prompts, and fields are the same as creating the same object on a FortiGate unit using the FortiGate web-based manager.

Locking the FortiGate web-based manager

When you use the FortiManager to manager multiple FortiGate units, a local FortiGate unit becomes locked from any configuration using the web-based manager by an administrator. When an administrator logs into the FortiGate unit, the following message appears:



If the administrator selects *Login Read Only*, an icon appears at the top of the web-based manager. All configuration options will only have a *Return* button, rather than the typical *OK*, *Apply* and *Cancel* buttons.

Figure 44: Read-only icon when under FortiManager management



Selecting Login Read-Write, a warning appears that any changes may cause the configuration between FortiManager and the FortiGate unit be become out of sync.

Firmware updates

FortiManager can also be the source where firmware updates are performed for multiple FortiGate units, saving time rather than upgrading each FortiGate unit individually.

The FortiManager unit stores local copies of firmware images by either downloading these images from the Fortinet Distribution Network (FDN) or by accepting firmware images that you upload from your management computer.

If you are using the FortiManager unit to download firmware images from the FDN, FortiManager units first validate device licenses. The FDN validates support contracts and provides a list of currently available firmware images. For devices with valid Fortinet Technical Support contracts, you can download new firmware images from the FDN, including release notes.

After firmware images have been either downloaded from the FDN or imported to the firmware list, you can either schedule or immediately upgrade/downgrade a device or group's firmware.

See the [FortiManager Administration Guide](#) for more information on updating the FortiGate firmware using the FortiManager central management.

FortiGuard

FortiManager can also connect to the FortiGuard Distribution Network to receive push updates for IPS signatures and antivirus definitions. These updates can then be used to update multiple FortiGate units throughout an organization. By the FortiManager as the host for updates, bandwidth use is minimized by downloading to one source instead of many.

To receive IPS and antivirus updates from FortiManager, indicate an alternate IP address on the FortiGate unit.

To configure updates from FortiManager

- 1 Go to *System > Config > FortiGuard*.
- 2 Select *AntiVirus and IPS Options* to expand the options.
- 3 Select the checkbox next to *Use override server address* and enter the IP address of the FortiManager unit.
- 4 Select *Apply*.

Backup and restore configurations

FortiManager stores configuration files for backup and restore purposes. FortiManager also enables you to save revisions of configuration files. Configuration backups occur automatically. Backups occur when the administrator logs out or the administrator login session expires (times out).

FortiManager also enables you to view differences between different configurations to view where changes have been made.

Administrative domains

FortiManager administrative domains enable the `admin` administrator to create groupings of devices for configured administrators to monitor and manage. FortiManager can manage a large number of Fortinet appliances. This enables administrators to maintain managed devices specific to their geographic location or business division. This also includes FortiGate units with multiple configured VDOMs.

Each administrator is tied to an administrative domain (ADOM). When that particular administrator logs in, they see only those devices or VDOMs configured for that administrator and ADOM. The one exception is the `admin` administrator account which can see and maintain all administrative domains and the devices within those domains.

Administrative domains are not enabled by default, and enabling and configuring the domains can only be performed by the `admin` administrator. The maximum number of administrative domains you can add depends on the FortiManager system model.

See the [FortiManager Administration Guide](#) for information on the maximums for each model.



Best practices

The FortiGate unit is installed, and traffic is flowing. With your network sufficiently protected, you can now fine tune the firewall for the best performance and efficiently. This chapter describes configuration options that can ensure your FortiGate unit is running at its best performance.

This section includes the topics on:

- [Hardware](#)
- [Shutting down](#)
- [Performance](#)
- [Firewall](#)
- [Intrusion protection](#)
- [Antivirus](#)
- [Web filtering](#)
- [Email Filtering \(Antispam\)](#)

Hardware

Environmental specifications

Keep the following environmental specifications in mind when installing and setting up your FortiGate unit.

- Operating temperature: 32 to 104°F (0 to 40°C) (temperatures may vary, depending on the FortiGate model)
- If you install the FortiGate unit in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient temperature. Therefore, make sure to install the equipment in an environment compatible with the manufacturer's maximum rated ambient temperature.
- Storage temperature: -13 to 158°F (-25 to 70°C) (temperatures may vary, depending on the FortiGate model)
- Humidity: 5 to 90% non-condensing
- Air flow - For rack installation, make sure that the amount of air flow required for safe operation of the equipment is not compromised.
- For free-standing installation, make sure that the appliance has at least 1.5 in. (3.75 cm) of clearance on each side to allow for adequate air flow and cooling.

This device complies with part FCC Class A, Part 15, UL/CUL, C Tick, CE and VCCI. Operation is subject to the following two conditions:

- This device may not cause harmful interference, and
- This device must accept any interference received, including interference that may cause undesired operation.

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

The equipment compliance with FCC radiation exposure limit set forth for uncontrolled Environment.



Risk of Explosion if battery is replaced by an incorrect type. Dispose of used batteries according to the instructions.



To reduce the risk of fire, use only No. 26 AWG or larger UL Listed or CSA Certified Telecommunication Line Cord.

Grounding

- Ensure the FortiGate unit is connected and properly grounded to a lightning and surge protector. WAN or LAN connections that enter the premises from outside the building should be connected to an Ethernet CAT5 (10/100 Mb/s) surge protector.
- Shielded Twisted Pair (STP) Ethernet cables should be used whenever possible rather than Unshielded Twisted Pair (UTP).
- Do not connect or disconnect cables during lightning activity to avoid damage to the FortiGate unit or personal injury.

Rack mount instructions

Elevated Operating Ambient - If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.

Reduced Air Flow - Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.

Mechanical Loading - Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.

Circuit Overloading - Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

Reliable Earthing - Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips).

Shutting down

Always shut down the FortiGate operating system properly before turning off the power switch to avoid potential hardware problems.

To power off the FortiGate unit - web-based manager

- 1 Go to *System > Status*.
- 2 In the Unit Operation display, select Shutdown.

To power off the FortiGate unit

```
execute shutdown
```

Once completing this step you can safely disconnect the power cables from the power supply.

Performance

- Disable any management features you do not need. If you don't need SSH or SNMP disable them. SSH also provides another possibility for would-be hackers to infiltrate your FortiGate unit.
- Put the most used firewall rules to the top of the interface list.
- Log only necessary traffic. The writing of logs, especially if to an internal hard disk, slows down performance.
- Enable only the required application inspections.
- Keep alert systems to a minimum. If you send logs to a syslog server, you may not need SNMP or email alerts, making for redundant processing.
- Establish scheduled FortiGuard updates at a reasonable rate. Daily every 4-5 hours for most situations, or in more heavy-traffic situations, in the evening when more bandwidth can be available.
- Keep UTM profiles to a minimum. If you do not need a profile on a firewall rule, do not include it.
- Keep VDOMs to a minimum. On low-end FortiGate units, avoid using them if possible.
- Avoid traffic shaping if you need maximum performance. Traffic shaping, by definition, slows down traffic.

Firewall

- Avoid using the *All* selection for the source and destination addresses. Use addresses or address groups.
- Avoid using *Any* for the services.
- Use logging on a policy only when necessary. For example, you may want to log all dropped connections but be aware of the performance impact. However, use this sparingly to sample traffic data rather than have it continually storing log information you may not use.
- Use the comment field to input management data; who requested the rule, who authorized it, etc.
- Avoid FQDN addresses if possible. It can cause a performance impact on DNS queries and security impact from DNS spoofing.
- If possible, avoid port ranges on services for security reasons.
- Use groups whenever possible.
- To ensure that all AV push updates occur, ensure you have an AV profile enabled for UTM in a security policy.

Intrusion protection

- Create and use UTM profiles with specific signatures and anomalies you need per-interface and per-rule.
- Do not use predefined or generic profiles. While convenient to supply immediate protection, you should create profiles to suit your network environment.
- If you do use the default profiles, reduce the IPS signatures/anomalies enabled in the profile to conserve processing time and memory.
- If you are going to enable anomalies, make sure you tune thresholds according to your environment.
- If you need protection, but not audit information, disable the logging option.
- Tune the IP-protocol parameter accordingly.

Antivirus

- Enable only the protocols you need to scan. If you have antivirus scans occurring on the SMTP server, or using FortiMail, it is redundant to have it occur on the FortiGate unit as well.
- Reduce the maximum file size to be scanned. Viruses travel usually in small files of around 1 to 2 megabytes.
- Antivirus scanning within an HA cluster can impact performance.
- Enable grayware scanning on UTM profiles tied to internet browsing.
- Do not quarantine files unless you regularly monitor and review them. This is otherwise a waste of space and impacts performance.
- Use file patterns to avoid scanning where it is not required.
- Enable heuristics from the CLI if high security is required using the command `config antivirus heuristic`.

Web filtering

- Web filtering within an HA cluster impacts performance.
- Always review the DNS settings to ensure the servers are fast.
- Content block may cause performance overhead.
- Local URL filter is faster than FortiGuard web filter, because the filter list is local and the FortiGate unit does not need to go out to the Internet to get the information from a FortiGuard web server.

Email Filtering (Antispam)

- If possible use, a FortiMail unit. The antispam engines are more robust.
- Use fast DNS servers.
- Use specific UTM profiles for the rule that will use antispam.
- DNS checks may cause false positive with HELO DNS lookup.
- Content analysis (banned words) may impose performance overhead.

Security

- Use NTP to synchronize time on the FortiGate and the core network systems such as email servers, web servers and logging services.
- Enable log rules to match corporate policy. For example, log administration authentication events and access to systems from untrusted interfaces.
- Minimize adhoc changes to live systems if possible to minimize interruptions to the network.
- When not possible, create backup configurations and implement sound audit systems using FortiAnalyzer and FortiManager.
- If you only need to allow access to a system on a specific port, limit the access by creating the strictest rule possible.



FortiGuard

FortiGuard is a world-wide network of servers. The FortiGuard Distribution Network (FDN) of servers provides updates to antivirus, antispam and IPS definitions. Worldwide coverage of FortiGuard services is provided by FortiGuard service points. FortiGuard Subscription Services provide comprehensive Unified Threat Management (UTM) security solutions to enable protection against content and network level threats.

Fortinet employs people around the globe monitoring virus, spyware and vulnerability activities. As these various vulnerabilities are found, signatures are created and pushed to the subscribed FortiGate unit. The Global Threat Research Team enables Fortinet to deliver a combination of multi-layered security intelligence and provide true zero-day protection from new and emerging threats. FortiGuard services are continuously updated year round, 24x7x365.

The FortiGuard Network has data centers around the world located in secure, high availability locations that automatically deliver updates to the Fortinet security platforms to and protect the network with the most up-to-date information.

To ensure optimal response and updates, the FortiGate unit will contact a FortiGuard service point closest to the FortiGate installation, using the configured time zone information.

Every FortiGate unit includes a free 30-day FortiGuard trial license. FortiGuard license management is performed by Fortinet servers. The FortiGate unit automatically contacts a FortiGuard service point when enabling FortiGuard services. Contact Fortinet Technical Support to renew a FortiGuard license after the free trial.

This section includes the topics:

- [FortiGuard Services](#)
- [Antivirus and IPS](#)
- [Web filtering](#)
- [Email filtering](#)
- [Security tools](#)
- [Troubleshooting](#)

FortiGuard Services

The FortiGuard services provide a number of services to monitor world-wide activity and provide the best possible security. Services include:

- **Antispam/Web Filtering-** The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools on Fortinet appliances and agents, to detect and block a wide range of spam messages. Updates to the IP reputation and spam signature databases are provided continuously via the global FortiGuard distribution network.
- **Antivirus** -The FortiGuard Antivirus Service provides fully automated updates to ensure protection against the latest content level threats. It employs advanced virus, spyware, and heuristic detection engines to prevent both new and evolving threats and vulnerabilities from gaining access to your network.

- **Intrusion Prevention** - The FortiGuard Intrusion Prevention Service uses a customizable database of more than 4000 known threats to stop attacks that evade conventional firewall defenses. It also provides behavior-based heuristics, enabling the system to recognize threats when no signature has yet been developed. It also provides more than 1000 application identity signatures for complete application control.
- **Web Filtering** - Web Filtering provides Web URL filtering to block access to harmful, inappropriate, and dangerous web sites that may contain phishing/pharming attacks, malware such as spyware, or objectionable content that can expose your organization to legal liability. Based on automatic research tools and targeted research analysis, real-time updates enable you to apply highly-granular policies that filter web access based on 78 web content categories, over 45 million rated web sites, and more than two billion web pages - all continuously updated.

Support Contract and FortiGuard Subscription Services

The *Support Contract* and *FortiGuard Subscription Services* sections are displayed in abbreviated form within the *License Information* widget. A detailed version is available by going to *System > Config > FortiGuard*.

The Support Contract area displays the availability or status of your FortiGate unit's support contract. The status displays can be either *Unreachable*, *Not Registered* or *Valid Contract*.

The FortiGuard Subscription Services area displays detailed information about your FortiGate unit's support contract and FortiGuard subscription services. On this page, you can also manually update the antivirus and IPS engines.

The status icons for each section indicates the state of the subscription service. The icon corresponds to the availability description.

- **Gray (Unreachable)** – the FortiGate unit is not able to connect to service.
- **Orange (Not Registered)** – the FortiGate unit can connect, but not subscribed.
- **Yellow (Expired)** – the FortiGate unit had a valid license that has expired.
- **Green (Valid license)** – the FortiGate unit can connect to FDN and has a registered support contract. If the Status icon is green, the expiry date also appears.

FortiGuard Analysis Service Options

Go to *System > Config > FortiGuard*, and expand the *FortiGuard Analysis & Management Service Options*.

Account ID	Enter the name for the FortiGuard Analysis and Management Service that identifies the account. This is the same account information used when registering for the service.
To launch the service portal, please click here	Select to go directly to the FortiGuard Analysis and Management Service portal web site to view logs or configuration. You can also select this to register your FortiGate unit with the FortiGuard Analysis and Management Service.

Antivirus and IPS

The FortiGuard network is an always updating service. That is, Fortinet employs developers around the clock, monitoring for new and mutating virus and intrusion threats. This includes grayware and signatures for application control. There are two methods of updating the virus and IPS signatures on your FortiGate unit: manually or through push updates.

Antivirus and IPS Options

Go to *System > Config > FortiGuard*, and expand the *Antivirus and IPS Options* section to configure the antivirus and IPS options for connecting and downloading definition files.

Use override server address	Select to configure an override server if you cannot connect to the FDN or if your organization provides updates using their own FortiGuard server.
Allow Push Update	Select to allow updates sent automatically to your FortiGate unit when they are available.
Allow Push Update status icon	The status of the FortiGate unit for receiving push updates: <ul style="list-style-type: none"> Gray (Unreachable) - the FortiGate unit is not able to connect to push update service Yellow (Not Available) - the push update service is not available with current support license Green (Available) - the push update service is allowed.
Use override push IP and Port	Available only if both <i>Use override server address</i> and <i>Allow Push Update</i> are enabled. Enter the IP address and port of the NAT device in front of your FortiGate unit. FDS will connect to this device when attempting to reach the FortiGate unit. The NAT device must be configured to forward the FDS traffic to the FortiGate unit on UDP port 9443.
Schedule Updates	Select this check box to enable updates to be sent to your FortiGate unit at a specific time. For example, to minimize traffic lag times, you can schedule the update to occur on weekends or after work hours. Note that a schedule of once a week means any urgent updates will not be pushed until the scheduled time. However, if there is an urgent update required, select the <i>Update Now</i> button.
Update Now	Select to manually initiate an FDN update.
Submit attack characteristics... (recommended)	Select to help Fortinet maintain and improve IPS signatures. The information sent to the FortiGuard servers when an attack occurs, can be used to keep the database current as variants of attacks evolve.

Manual updates

To manually update the signature definitions file, you need to first go to the Support web site at <https://support.fortinet.com>. Once logged in, select FortiGuard Service Updates from the Download area of the web page. The browser will present you the most current antivirus and IPS signature definitions which you can download.

Once downloaded to your computer, log into the FortiGate unit to load the definition file.

To load the definition file onto the FortiGate unit

- 1 Go to *System > Config > FortiGuard*.
- 2 Select the *Update* link for either *AV Definitions* or *IPS Definitions*.
- 3 Locate the downloaded file and select *OK*.

The upload may take a few minutes to complete.

Automatic updates

The FortiGate unit can be configured to request updates from the FortiGuard Distribution Network. You can configure this to be on a scheduled basis, or with push notifications.

Scheduling updates

Scheduling updates ensures that the virus and IPS definitions are downloaded to your FortiGate unit on a regular basis, ensuring that you do not forget to check for the definition files yourself. As well, by scheduling updates during off-peak hours, such as evenings or weekends, when network usage is minimal, ensures that the network activity will not suffer from the added traffic of downloading the definition files.

If you require the most up-to-date definitions as viruses and intrusions are found in the wild, the FortiGuard Distribution Network can push updates to the FortiGate units as they are developed. This ensures that your network will be protected from any breakouts of a virus within the shortest amount of time, minimizing any damaging effect that can occur. Push updates require that you have registered your FortiGate unit.

Once push updates are enabled, the next time new antivirus or IPS attack definitions are released, the FDN notifies all the FortiGate unit that a new update is available. Within 60 seconds of receiving a push notification, the unit automatically requests the update from the FortiGuard servers.

To enable scheduled updates - web-based manager

- 1 Go to *System > Config > FortiGuard*.
- 2 Click the Expand Arrow for *AntiVirus and IPS Options*.
- 3 Select the *Scheduled Update* check box.
- 4 Select the frequency of the updates and when within that frequency.
- 5 Select *Apply*.

To enable scheduled updates - CLI

```
config system autoupdate schedule
  set status enable
  set frequency {every | daily | weekly}
  set time <hh:mm>
  set day <day_of_week>
end
```

Push updates

Push updates enable you to get immediate updates when new virus or intrusions have been discovered and new signatures are created. This ensures that when the latest signature is available it will be sent to the FortiGate.

When a push notification occurs, the FortiGuard server sends a notice to the FortiGate unit that there is a new signature definition file available. The FortiGate unit then initiates a download of the definition file, similar to the scheduled update.

To ensure maximum security for your network, you should have a scheduled update as well as enable the push update, in case an urgent signature is created, and your cycle of the updates only occurs weekly.

To enable push updates - web-based manager

- 1 Got to *System > Config > FortiGuard*.
- 2 Click the Expand Arrow for *Antivirus and IPS Options*.
- 3 Select *Allow Push Update*.
- 4 Select *Apply*.

To enable push updates - CLI

```
config system autoupdate push-update
    set status enable
end
```


Push IP override

If the FortiGate unit is behind another NAT device (or another FortiGate unit), to ensure it receives the push update notifications, you need to use an override IP address for the notifications. To do this, you create a virtual IP to map to the external port of the NAT device.

Generally speaking, if there are two FortiGate devices as in the diagram below, the following steps need to be completed on the FortiGate NAT device to ensure the FortiGate unit on the internal network receives the updates:

- Add a port forwarding virtual IP to the FortiGate NAT device that connects to the Internet by going to *Firewall Objects > Virtual IP*.
- Add a security policy to the FortiGate NAT device that connects to the Internet that includes the port forwarding virtual IP.
- Configure the FortiGate unit on the internal network with an override push IP and port.

On the FortiGate internal device, the virtual IP is entered as the *Use push override IP* address.

Figure 45: Using a virtual IP for a FortiGate unit behind a NAT device**To enable push update override- web-based manager**

- 1 Got to *System > Config > FortiGuard*.
- 2 Click the Expand Arrow for *Antivirus and IPS Options*.
- 3 Select *Allow Push Update*.
- 4 Select *Use push override IP*.
- 5 Enter the virtual IP address configured on the NAT device.
- 6 Select *Apply*.

To enable push updates - CLI

```
config system autoupdate push-update
  set status enable
  set override enable
  set address <vip_address>
end
```

Web filtering

The multiple FortiGuard data centers around the world hold the entire categorized URL database and receive rating requests from customer FortiGate units typically triggered by browser based URL requests. These rating requests are responded to with the categories stored for specific URLs, the requesting FortiGate unit will then use its own local profile configuration to determine what action is appropriate to the category, that is, to blocking, monitor or permit the request. Fortinet's development team has ensured that providing this powerful filtering capability is as simple as possible to enable.

Further, rating responses can also be cached locally on the FortiGate unit, providing a quicker response time, while easing load on the FortiGuard servers, aiding in a quicker response time for less common URL requests. This is a very effective method for common sites. Search engines and other frequently visited sites for your business can remain cached locally. Other sites less frequently visited, can be cached locally for a determined amount of time. For a site such as Google, the frequency of its access can keep it in the cache, other sites can remain in the cache up to 24 hours, or less depending on the configuration.

By default, the web filtering cache is enabled. The cache includes a time-to-live value, which is the amount of time a url will stay in the cache before expiring. You can change this value to shorten or extend the time between 300 and 86400 seconds.

Web Filtering and Email Filtering Options

Go to *System > Config > FortiGuard*, and expand arrow to view *Web Filtering and Email Filtering* options for setting the size of the caches and ports used.

Web Filter cache TTL	Set the Time To Live value. This is the number of seconds the FortiGate unit will store a blocked IP or URL locally, saving time and network access traffic, checking the FortiGuard server. Once the TTL has expired, the FortiGate unit will contact an FDN server to verify a web address. The TTL must be between 300 and 86400 seconds.
Antispam cache TTL	Set the Time To Live value. This is the number of seconds the FortiGate unit will store a blocked IP or URL locally, saving time and network access traffic, checking the FortiGuard server. Once the TTL has expired, the FortiGate unit will contact an FDN server to verify a web address. The TTL must be between 300 and 86400 seconds.
Port Section	Select the port assignments for contacting the FortiGuard servers. Select the <i>Test Availability</i> button to verify the connection using the selected port.
To have a URL's category rating re-evaluated, please click here.	Select to re-evaluate a URL's category rating on the FortiGuard Web Filter service.

URL verification

If you discover a URL - yours or one you require access to has been incorrectly flagged as an inappropriate site, you can ask the FortiGuard team to re-evaluate the site. To do this, go to *System > Config > FortiGuard*, select the blue arrow for *Web Filtering and Email Filtering Options* and select the link for re-evaluation.

To modify the web filter cache size - web-based manager

- 1 Got to *System > Config > FortiGuard*.
- 2 Click the Expand Arrow for *Web Filtering and Email Filtering Options*.
- 3 Enter the TTL value for the *Web filter cache*.
- 4 Select *Apply*.

To modify the web filter cache size - CLI

```
config system fortiguard
    set webfilter-cache-ttl <integer>
end
```

Further web filtering options can be configured to block specific URLs, and allow others through. These configurations are available through the *UTM > Web Filter* menu.

Email filtering

Similar to web filtering, FortiGuard data centers monitor and update email databases of known spam sources. With FortiGuard antispam enabled, the FortiGate unit verifies incoming email sender address and IPs against the database, and take the necessary action as defined within the antivirus profiles.

Further, spam source IP addresses can also be cached locally on the FortiGate unit, providing a quicker response time, while easing load on the FortiGuard servers, aiding in a quicker response time for less common email address requests.

By default, the antispam cache is enabled. The cache includes a time-to-live value, which is the amount of time an email address will stay in the cache before expiring. You can change this value to shorten or extend the time between 300 and 86400 seconds.

To modify the antispam filter cache size - web-based manager

- 1 Got to *System > Config > FortiGuard*.
- 2 Click the Expand Arrow for *Web Filtering and Email Filtering Options*.
- 3 Enter the TTL value for the *Antispam filter cache*.
- 4 Select *Apply*.

To modify the web filter cache size - CLI

```
config system fortiguard
    set antispam-cache-ttl <integer>
end
```

Further antispam filtering options can be configured to block, allow or quarantine, specific email addresses.

Security tools

The FortiGuard online center provides a number of online security tools that enable you to verify or check ratings of web sites, email addresses as well as check file for viruses. These features are available at <http://www.fortiguard.com>.

URL lookup

By entering a web site address, you can see if it has been rated and what category and classification it is filed as. If you find your web site or a site you commonly go to has been wrongly categorized, use this page to request the site to be re-evaluated.

<http://www.fortiguard.com/webfiltering/webfiltering.html>

IP and signature lookup

The IP and signature lookup enable you to check whether an IP address is blacklisted in the FortiGuard IP reputation database, or whether a URL or email address is in the signature database.

<http://www.fortiguard.com/antispam/antispam.html>

Online virus scanner

If you discover a suspicious file on your machine, or suspect that a program you downloaded from the internet might be malicious you can scan it using the FortiGuard online scanner. The questionable file can be uploaded from your computer to a dedicated server where it will be scanned using FortiClient Antivirus. Only one file of up to 1 MB can be checked at any one time. All files will be forwarded to our research labs for analysis.

http://www.fortiguard.com/antivirus/virus_scanner.html

Malware removal tools

Tools have been developed by FortiGuard Labs to disable and remove the specific malware and related variants. Some tools have been developed to remove specific malware, often tough to remove. A universal cleaning tool, FortiCleanup, is also available for download.

The FortiCleanup is a tool developed to identify and cleanse systems of malicious rootkit files and their associated malware. Rootkits consist of code installed on a system with kernel level privileges, often used to hide malicious files, keylog and thwart detection / security techniques. The aim of this tool is to reduce the effectiveness of such malware by finding and eliminating rootkits. The tool offers a quick memory scan as well as a full system scan. FortiCleanup will not only remove malicious files, but also can cleanse registry entries, kernel module patches, and other tricks commonly used by rootkits - such as SSDT hooks and process enumeration hiding.

A license to use these applications is provided free of charge, courtesy of Fortinet.

http://www.fortiguard.com/antivirus/malware_removal.html

Troubleshooting

If you are not getting FortiGuard web filtering or antispam services, there are a few things to verify communication to the FortiGuard Distribution Network (FDN) is working. Before any troubleshooting, ensure that the FortiGate unit has been registered and you or your company, has subscribed to the FortiGuard services.

Web-based manager verification

The simplest method to check that the FortiGate unit is communicating with the FDN, is to check the License *Information* dashboard widget. Any subscribed services should have a green check mark beside them indicating that connections are successful. Any other icon indicates a problem with the connection, or you are not subscribed to the FortiGuard services.

Figure 46: License Information widget showing FortiGuard availability

License Information		
Support Contract		
Registration	Registered (Login: admin@fortinet.com) [Login Now]	✓
Hardware	8 x 5 support (Expires: 2012-11-26)	✓
Firmware	8 x 5 support (Expires: 2012-11-26)	✓
Enhanced Support	24 x 7 support (Expires: 2012-11-26)	✓
Comprehensive Support	24 x 7 support (Expires: 2012-11-26)	✓
FortiGuard Services		
AntiVirus	Licensed (Expires 2012-11-26)	✓
Intrusion Protection	Licensed (Expires 2012-11-26)	✓
Web Filtering	Not Registered [Configure]	✗
Email Filtering	Not Registered [Configure]	✗
Vulnerability Management	Licensed (Expires 2012-11-26)	✓
Analysis & Management	Expired [Renew]	✗
Virtual Domain		
VDOMs Allowed	10	
FortiClient Software		
FortiClient	Unlicensed [Enter License]	✗
FortiClient Connecting/Allowed 0 / 10		

Alternatively, you can view the FortiGuard connection status by going to *System > Config > FortiGuard*.

Figure 47: FortiGuard availability

Support Contract		
Registration	Registered (Login ID: [Login Now])	✓
Hardware	8 x 5 support (Expires: 2012-11-26)	✓
Firmware	8 x 5 support (Expires: 2012-11-26)	✓
Enhanced Support	24 x 7 support (Expires: 2012-11-26)	✓
Comprehensive Support	24 x 7 support (Expires: 2012-11-26)	✓
FortiGuard Subscription Services		
AntiVirus	Valid License (Expires 2012-11-26)	✓
AV Definitions	14.00000 (Updated 2011-08-24 via Manual Update) [Update]	
AV Engine	4.00382 (Updated 2011-10-28 via Manual Update)	
=====		
Intrusion Protection	Valid License (Expires 2012-11-26)	✓
IPS Definitions	3.00097 (Updated 2011-10-28 via Manual Update) [Update]	
IPS Engine	1.00241 (Updated 2011-10-28 via Manual Update)	
=====		
Web Filtering	Not Registered	✗
=====		
Email Filtering	Not Registered	✗
=====		
Vulnerability Management	Valid License (Expires 2012-11-26)	✓
VCM Plugin	1.00238 (Updated 2011-11-25 via Manual Update) [Update]	
=====		
Analysis & Management Service	Expired [Renew] [Update]	✗
FortiToken Seed Server		
Registration	Reachable (0 Tokens Registered)	✓
=====		
▶ AntiVirus and IPS Options ▶ Web Filtering and Email Filtering Options ▶ FortiGuard Analysis & Management Service Options		

CLI verification

You can also use the CLI to see what FortiGuard servers are available to your FortiGate unit. Use the following CLI command to ping the FDN for a connection:

```
ping guard.fortinet.net
```

You can also use diagnose command to find out what FortiGuard servers are available:

```
diagnose debug rating
```

From this command, you will see output similar to the following:

```
Locale      : english
License     : Contract
Expiration  : Sun Jul 24 20:00:00 2011
Hostname    : service.fortiguard.net

-- Server List (Tue Nov  2 11:12:28 2010) --

IP Weight   RTT Flags  TZ      Packets  Curr Lost Total Lost
69.20.236.180 0    10      -5      77200    0        42
69.20.236.179 0    12      -5      52514    0        34
66.117.56.42  0    32      -5      34390    0        62
80.85.69.38  50   164     0       34430    0       11763
208.91.112.194 81   223 D    -8      42530    0       8129
216.156.209.26 286  241 DI  -8      55602    0      21555
```

An extensive list of servers are available. Should you see a list of three to five available servers, the FortiGuard servers are responding to DNS replies to service.FortiGuard.net, but the INIT requests are not reaching FDS services on the servers.

The rating flags indicate the server status:

Table 17: FortiGuard debug rating flags

D	Indicates the server was found via the DNS lookup of the hostname. If the hostname returns more than one IP address, all of them will be flagged with 'D' and will be used first for INIT requests before falling back to the other servers.
I	Indicates the server to which the last INIT request was sent
F	The server has not responded to requests and is considered to have failed.
T	The server is currently being timed.

The server list is sorted first by weight and then the server with the smallest RTT is put at the top of the list, regardless of weight. When a packet is lost, that is, no response in two seconds, it will be resent to the next server in the list. The top position in the list is selected based on RTT while the other list positions are based on weight.

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a faraway server, the weight is not allowed to dip below a base weight which is calculated as the difference in hours between the FortiGate unit and the server multiplied by 10. The further away the server is, the higher its base weight and the lower in the list it will appear.

Port assignment

FortiGate units contact the FortiGuard Distribution Network (FDN) for the latest list of FDN servers by sending UDP packets with typical source ports of 1027 or 1031, and destination ports of 53 or 8888. The FDN reply packets have a destination port of 1027 or 1031.

If your ISP blocks UDP packets in this port range, the FortiGate unit cannot receive the FDN reply packets. As a result, the FortiGate unit will not receive the complete FDN server list.

You can select a different source port range for the FortiGate unit to use. If your ISP blocks the lower range of UDP ports (around 1024), you can configure your FortiGate unit to use higher-numbered ports, using the CLI command...

```
config system global
    set ip-src-port-range <start port>--<end port>
end
```

...where the <start port> and <end port> are numbers ranging of 1024 to 25000.

For example, you could configure the FortiGate unit to not use ports lower than 2048 or ports higher than the following range:

```
config system global
    set ip-src-port-range 2048-20000
end
```

Trial and error may be required to select the best source port range. You can also contact your ISP to determine the best range to use.

Push updates might be unavailable if:

- there is a NAT device installed between the unit and the FDN
- your unit connects to the Internet using a proxy server.



Monitoring

With network administration, the first step is installing and configuring the FortiGate unit to be the protector of the internal network. Once the system is running efficiently, the next step is to monitor the system and network traffic, to tweak leaks and abusers as well as the overall health of the FortiGate unit(s) that provide that protection.

This chapter discusses the various methods of monitoring both the FortiGate unit and the network traffic through a range of different tools available within FortiOS.

This section includes the topics:

- [Dashboard](#)
- [sFlow](#)
- [Monitor menus](#)
- [Logging](#)
- [Alert email](#)
- [SNMP](#)
- [SNMP get command syntax](#)
- [Fortinet and FortiGate MIB fields](#)

Dashboard

The FortiOS dashboard provides a location to view real-time system information. By default, the dashboard displays the key statistics of the FortiGate unit itself, providing the memory and CPU status, as well as the health of the ports, whether they are up or down and their throughput.

Widgets

Within the dashboard is a number of smaller windows, called widgets, that provide this status information. Beyond what is visible by default, you can add a number of other widgets that display other key traffic information including application use, traffic per IP address, top attacks, traffic history and logging statistics.

You will see when you log into the FortiGate unit, there are two separate dashboards. You can add multiple dashboards to reflect what data you want to monitor, and add the widgets accordingly. Dashboard configuration is only available through the web-based manager. Administrators must have read and write privileges to customize and add widgets when in either menu. Administrators must have read privileges if they want to view the information.

To add a dashboard and widgets

- 1 Go to *System > Dashboard*.
- 2 Select the *Dashboard* menu at the top of the window and select *Add Dashboard*.
- 3 Enter a name such as *Monitoring*.
- 4 Select the *Widget* menu at the top of the window.

5 From the screen, select the type of information you want to add.

6 When done, select the X in the top right of the widget.

Dashboard widgets provide an excellent method to view real-time data about the events occurring on the FortiGate unit and the network. For example, by adding the Network Protocol Usage widget, you can monitor the activity of various protocols over a selected span of time. Based on that information you can add or adjust traffic shaping and/or security policies to control traffic.



You can position widgets within the dashboard frame by clicking and dragging it to a different location.

FortiClient connections

The *License Information* widget includes information for the FortiClient connections. It displays the number of FortiClient connections allowed, and the number of users connecting. By selecting the Details link for the number of connections, you can view more information about the connecting user, including IP address, user name and type of operating system the user is connecting with.

sFlow

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. sFlow is described in <http://www.sflow.org>. FortiOS implements sFlow version 5. sFlow uses packet sampling to monitor network traffic. That is, an sFlow Agent captures packet information at defined intervals and sends them to an sFlow Collector for analysis, providing real-time data analysis. The information sent is only a sampling of the data for minimal impact on network throughput and performance.

The sFlow Agent is embedded in the FortiGate unit. Once configured, the FortiGate unit sends sFlow datagrams of the sampled traffic to the sFlow Collector, also called an sFlow Analyzer. The sFlow Collector receives the datagrams, and provides real-time analysis and graphing to indicate where potential traffic issues are occurring. sFlow Collector software is available from a number of third party software vendors.

sFlow data captures only a sampling of network traffic, not all traffic like the traffic logs on the FortiGate unit. Sampling works by the sFlow Agent looking at traffic packets when they arrive on an interface. A decision is made whether the packet is dropped, and sent on to its destination, or a copy is forwarded to the sFlow Collector. The sample used and its frequency are determined during configuration.



sFlow is not supported on virtual interfaces such as vdom link, ipsec, ssl.<vdom> or gre.

The sFlow datagram sent to the Collector contains the information:

- Packet header (e.g. MAC,IPv4,IPv6,IPX,AppleTalk,TCP,UDP, ICMP)
- Sample process parameters (rate, pool etc.)
- Input/output ports
- Priority (802.1p and TOS)
- VLAN (802.1Q)

- Source/destination prefix
- Next hop address
- Source AS, Source Peer AS
- Destination AS Path
- Communities, local preference
- User IDs (TACACS/RADIUS) for source/destination
- URL associated with source/destination
- Interface statistics (RFC 1573, RFC 2233, and RFC 2358)

sFlow agents can be added to any type of FortiGate interface. sFlow isn't supported on some virtual interfaces such as VDOM link, IPsec, gre, and ssl.<vdom>.

For more information on sFlow, Collector software and sFlow MIBs, visit www.sflow.org.

Configuration

sFlow configuration is available only from the CLI. Configuration requires two steps: enabling the sFlow Agent, and configuring the interface for the sampling information.

Enable sFlow

```
config system sflow
  set collector-ip <ip_address>
  set collector-port <port_number>
end
```

The default port for sFlow is UDP 6343. To configure in VDOM, use the commands:

```
config system vdom-sflow
  set vdom-sflow enable
  set collector-ip <ip_address>
  set collector-port <port_number>
end
```

Configure sFlow agents per interface.

```
config system interface
  edit <interface_name>
    set sflow-sampler enable
    set sample-rate <every_n_packets>
    set sample-direction [tx | rx | both]
    set polling-interval <seconds>
  end
```

Monitor menus

The *Monitor* menus enable you to view session and policy information and other activity occurring on your FortiGate unit. The monitors provide the details of user activity, traffic and policy usage to show live activity. Monitors are available for DHCP, routing, security policies, traffic shaping, a variety of UTM functionality, VPN, user, WiFi controllers and logging.

Logging

FortiOS provides a robust logging environment that enables you to monitor, store and report traffic information and FortiGate events including attempted log ins and hardware status. Depending on your requirements, you can log to a number of different hosts.

To configure logging in the web-based manager, go to *Log & Report > Log Config > Log Setting*.

To configure logging in the CLI use the commands `config log <log_location>`.

For details on configuring logging see the [Logging and Reporting Guide](#) and [FortiAnalyzer Administration Guide](#).

FortiGate memory

Logs are saved to the internal memory by default. Inexpensive yet volatile, for basic event logs or verifying traffic, AV or spam patterns, logging to memory is a simple option. However, because logs are stored in the limited space of the internal memory, only a small amount is available for logs. As such logs can fill up and be overridden with new entries, negating the use of recursive data. This is especially true for traffic logs. Also, should the FortiGate unit be shut down or rebooted, all log information will be lost.

To change the logging options for memory, go to *Log&Report > Log Config > Log Setting*.

FortiGate hard disk

For those FortiGate units with an internal hard disk or SDHC card, you can store logs to this location. Efficient and local, the hard disk provides a convenient storage location. If you choose to store logs in this manner, remember to backup the log data regularly.

Configure log disk settings is performed in the CLI using the commands:

```
config log setting disk
  set status enable
end
```

Further options are available when enabled to configure log file sizes, and uploading/backup events.

As well, note that the write speeds of hard disks compared to the logging of ongoing traffic may cause the dropping of log messages. As such, it is recommended that traffic logging be sent to a FortiAnalyzer or other device meant to handle large volumes of data.

Syslog server

An industry standard for collecting log messages, for off site storage. In the web-based manager, you are able to send logs to a single syslog server, however in the CLI you can configure up to three syslog servers where you can also use multiple configuration options. For example, send traffic logs to one server, antivirus logs to another. The FortiGate unit sends Syslog traffic over UDP port 514. Note that if a secure tunnel is configured for communication to a FortiAnalyzer unit, then Syslog traffic will be sent over an IPsec connection, using UDP 500/4500, protocol IP/50.

To configure a Syslog server in the web-based manager, go to *Log&Report > log Config > Log Setting*. In the CLI use the commands:

```
config log syslogd setting
  set status enable
end
```

Further options are available when enabled to configure a different port, facility and server IP address.

For Syslog traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over. For example, to set the source IP of a Syslog server to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config log syslogd setting
    set status enable
    set source-ip 192.168.4.5
end
```

FortiGuard Analysis and Management service

The FortiGuard Analysis and Management Service is a subscription-based hosted service. With this service, you can have centralized management, logging, and reporting capabilities available in FortiAnalyzer and FortiManager platforms, without any additional hardware to purchase, install or maintain.

This service includes a full range of reporting, analysis and logging, firmware management and configuration revision history. It is hosted within the Fortinet global FortiGuard Network for maximum reliability and performance, and includes reporting, and drill-down analysis widgets makes it easy to develop custom views of network and security events.

The FortiGate unit sends log messages to the FortiGuard Analysis and Management service using TCP port 443. Configuration is available once a user account has been set up and confirmed. To enable the account on the FortiGate unit, go to *System > Maintenance > FortiGuard*, select the blue arrow to expand the option, and enter the account ID.

For FortiGuard Analysis and Management traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of the FortiGuard Analysis and Management server to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config log fortiguard setting
    set status enable
    set source-ip 192.168.4.5
end
```

From the FortiGate unit, you can configure the connection and sending of log messages to be sent over an SSL tunnel to ensure log messages are sent securely. To do this, use the CLI commands to enable the encrypted connection and define the level of encryption.

```
config log fortiguard setting
    set status enable
    set enc-algorithm {default | high | low | disable}
end
```

For more information on each encryption level see [“Configuring an SSL connection” on page 427](#).

FortiAnalyzer

The FortiAnalyzer family of logging, analyzing, and reporting appliances securely aggregate log data from Fortinet devices and other syslog-compatible devices. Using a comprehensive suite of easily-customized reports, users can filter and review records, including traffic, event, virus, attack, Web content, and email data, mining the data to determine your security stance and assure regulatory compliance. FortiAnalyzer also provides advanced security management functions such as quarantined file archiving, event correlation, vulnerability assessments, traffic analysis, and archiving of email, Web access, instant messaging and file transfer content.

The FortiGate unit sends log messages over UDP port 514 or OFTP (TCP 514). If a secure connection has been configured, log traffic is sent over UDP port 500/4500, Protocol IP/50. For more information on configuring a secure connection see [“Sending logs using a secure connection” on page 426](#).

For FortiAnalyzer traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command set `source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of a FortiAnalyzer unit to be on port 3 with an IP of 192.168.21.12, the commands are:

```
config log fortiguard setting
  set status enable
  set source-ip 192.168.21.12
end
```

Sending logs using a secure connection

From the FortiGate unit, you can configure the connection and sending of log messages over an SSL tunnel to ensure log messages are sent securely. To do this, use the CLI commands below to enable the encrypted connection and define the level of encryption.



You must configure the secure tunnel on **both** ends of the tunnel, the FortiGate unit and the FortiAnalyzer unit.

This configuration is for FortiAnalyzer OS version 4.0 MR2 or lower. For version 4.0 MR3, see [“Configuring an SSL connection” on page 427](#).

To configure a secure connection to the FortiAnalyzer unit

On the FortiAnalyzer unit, enter the commands:

```
config log device
  edit <device_name>
    set secure psk
    set psk <name_of_IPSec_tunnel>
    set id <fortigate_device_name_on_the_fortianalyzer>
  end
```

To configure a secure connection on the FortiGate unit

On the FortiGate CLI, enter the commands:

```
config log fortianalyzer setting
  set status enable
  set server <ip_address>
  set local
  set localid <name_of_IPSec_tunnel>
end
```


Configuring an SSL connection

With FortiAnalyzer 4.0 MR3 and FortiOS 4.0 MR3, you can configure an SSL connection between the two devices, and select the encryption level.

Use the CLI commands to configure the encryption connection:

```
config log fortianalyzer setting
  set status enable
  set enc-algorithm {default* | high | low | disable}
end
```



These commands are specific to OS versions 4.0 MR3 and higher. IPSec connections will still be possible between FortiOS 4.0 MR3 and FortiAnalyzer 4.0 MR2 and lower.

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for high, medium, and low follows openssl definitions:

- **High** - Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.
Algorithms are: DHE-RSA-AES256-SHA:AES256-SHA:EDH-RSA-DES-CBC3-SHA:DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA
- **Medium** - Key strengths of 128 bit encryption.
Algorithm are:RC4-SHA:RC4-MD5:RC4-MD
- **Low** - Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites
Algorithms: EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5.

If you want to use an IPSec tunnel to connect to the FortiAnalyzer unit, you need to first disable the enc-algorithm:

```
config log fortianalyzer setting
  set status enable
  set enc-algorithm disable
```

Then set the IPSec encryption:

```
set encrypt enable
set psksecret <preshared_IPSec_tunnel_key>
end
```

Alert email

As an administrator, you want to be certain you can respond quickly to issues occurring on your network or on the FortiGate unit. Alert email provides an efficient and direct method of notifying an administrator of events. By configuring alert messages, you can define the threshold when a problem becomes critical and needs attention. When this threshold is reached, the FortiGate unit will send an email to one or more individuals notifying them of the issue.

In the following example, the FortiGate unit is configured to send email to two administrators (admin1 and admin2) when multiple intrusions are detected every two minutes. The FortiGate unit has its own email address on the mail server.

To configure alert email - web-based manager

- 1 Go to *Log&Report > Log Config > Alert E-mail*.
- 2 Enter the information:

SMTP Server	Enter the address or name of the email server. For example, smtp.example.com.
Email from	fortigate@example.com
Email to	admin1@example.com admin2@example.com
Authentication	Enable authentication if required by the email server.
SMTP User	FortiGate
Password	*****
Interval Time	2

- 3 For the Interval Time, enter 2.
- 4 Select *Intrusion Detected*.
- 5 Select *Apply*.

To configure alert email - CLI

```

config system alert email
    set port 25
    set server smtp.example.com
    set authenticate enable
    set username FortiGate
    set password *****
end
config alertemail setting
    set username fortigate@example.com
    set mailto1 admin1@example.com
    set mailto2 admin2@example.com
    set filter category
    set IPS-logs enable
end

```

SNMP

Simple Network Management Protocol (SNMP) enables you to monitor hardware on your network. You can configure the hardware, such as the FortiGate SNMP agent, to report system information and send traps (alarms or event messages) to SNMP managers. An SNMP manager, or host, is typically a computer running an application that can read the incoming trap and event messages from the agent and send out SNMP queries to the SNMP agents. A FortiManager unit can act as an SNMP manager, or host, to one or more FortiGate units. FortiOS supports SNMP using IPv4 and IPv6 addressing.

By using an SNMP manager, you can access SNMP traps and data from any FortiGate interface or VLAN sub interface configured for SNMP management access. Part of configuring an SNMP manager is to list it as a host in a community on the FortiGate unit it will be monitoring. Otherwise the SNMP monitor will not receive any traps from that FortiGate unit, or be able to query that unit.

The FortiGate SNMP implementation is read-only. SNMP v1, v2c, and v3 compliant SNMP managers have read-only access to FortiGate system information through queries and can receive trap messages from the FortiGate unit.

To monitor FortiGate system information and receive FortiGate traps, you must first compile the Fortinet and FortiGate Management Information Base (MIB) files. A MIB is a text file that describes a list of SNMP data objects that are used by the SNMP manager. These MIBs provide information the SNMP manager needs to interpret the SNMP trap, event, and query messages sent by the FortiGate unit SNMP agent. FortiGate core MIB files are available on the Customer Support web site.

The Fortinet implementation of SNMP includes support for most of RFC 2665 (Ethernet-like MIB) and most of RFC 1213 (MIB II). For more information, see [“Fortinet MIBs” on page 434](#). RFC support for SNMP v3 includes Architecture for SNMP Frameworks (RFC 3411), and partial support of User-based Security Model (RFC 3414).

SNMP traps alert you to events that occur such as an a full log disk or a virus detected. For more information about SNMP traps, see [“Fortinet and FortiGate traps” on page 436](#).

SNMP fields contain information about the FortiGate unit, such as CPU usage percentage or the number of sessions. This information is useful for monitoring the condition of the unit on an ongoing basis and to provide more information when a trap occurs. For more information about SNMP fields, see [“Fortinet and FortiGate MIB fields” on page 441](#).

The FortiGate SNMP v3 implementation includes support for queries, traps, authentication, and privacy. Authentication and encryption are configured in the CLI. See the `system snmp user` command in the [FortiGate CLI Reference](#).

SNMP configuration settings

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections by going to *System > Network > Interface*. Select the interface, and in the *Administrative Access*, select *SNMP*.



When the FortiGate unit is in virtual domain mode, SNMP traps can only be sent on interfaces in the management virtual domain. Traps cannot be sent over other interfaces. IPv6 is supported for SNMP configuration on FortiGate units running FortiOS 4.0 MR3.

To configure SNMP settings, go to *System > Config > SNMP*.

SNMP Agent	Select to enable SNMP communication.
Description	Enter descriptive information about the FortiGate unit. The description can be up to 35 characters.
Location	Enter the physical location of the FortiGate unit. The system location description can be up to 35 characters long.
Contact	Enter the contact information for the person responsible for this FortiGate unit. The contact information can be up to 35 characters.
SNMP v1/v2c section	
To create a new SNMP community, see New SNMP Community page .	
Community Name	The name to identify the community.

Queries	Indicates whether queries protocols (v1 and v2c) are enabled or disabled. A green checkmark indicates queries are enabled; a gray x indicates queries are disabled. If one query is disabled and another one enabled, there will still be a green checkmark.
Traps	Indicates whether trap protocols (v1 and v2c) are enabled or disabled. A green checkmark indicates traps are enabled; a gray x indicates traps are disabled. If one query is disabled and another one enabled, there will still be a green checkmark.
Enable	Select the check box to enable or disable the community.
SNMP v3 section	
To create a new SNMP community, see Create New SNMP V3 User .	
User Name	The name of the SNMPv3 user.
Security Level	The security level of the user.
Notification Host	The IP address or addresses of the host.
Queries	Indicates whether queries are enabled or disabled. A green checkmark indicates queries are enabled; a gray x indicates queries are disabled.
New SNMP Community page	
Community Name	Enter a name to identify the SNMP community.
Hosts (section)	
IP Address	Enter the IP address and Identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit. You can also set the IP address to 0.0.0.0 to so that any SNMP manager can use this SNMP community.
Interface	Optionally select the name of the interface that this SNMP manager uses to connect to the FortiGate unit. You only have to select the interface if the SNMP manager is not on the same subnet as the FortiGate unit. This can occur if the SNMP manager is on the Internet or behind a router. In virtual domain mode, the interface must belong to the management VDOM to be able to pass SNMP traps.
Delete	Removes an SNMP manager from the list within the <i>Hosts</i> section.
Add	Select to add a blank line to the Hosts list. You can add up to eight SNMP managers to a single community.
Queries (section)	
Protocol	The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.

Port	Enter the port number (161 by default) that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the <i>Enable</i> check box to activate queries for each SNMP version. Note: The SNMP client software and the FortiGate unit must use the same port for queries.
Enable	Select to enable that SNMP protocol
Traps (section)	
Protocol	The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.
Local	Enter the remote port numbers (port 162 for each by default) that the FortiGate unit uses to send SNMP v1 or SNMP v2c traps to the SNMP managers in this community. Select the <i>Enable</i> check box to activate traps for each SNMP version. Note: The SNMP client software and the FortiGate unit must use the same port for traps.
Remote	Enter the remote port number (port 162 is default) that the FortiGate unit uses to send SNMP v1 or v2c traps to the SNMP managers in this community. Note: The SNMP client software and the FortiGate unit must use the same port for queries.
Enable	Select to activate traps for each SNMP version.
SNMP Event	Enable each SNMP event for which the FortiGate unit should send traps to the SNMP managers in this community. <i>CPU Overusage</i> traps sensitivity is slightly reduced, by spreading values out over 8 polling cycles. This prevents sharp spikes due to CPU intensive short-term events such as changing a policy. <i>Power Supply Failure</i> event trap is available only on some models. <i>AMC interfaces enter bypass mode</i> event trap is available only on models that support AMC modules.
Enable	Select to enable the SNMP event.
Create New SNMP V3 User	
User Name	Enter the name of the user.
Security Level	Select the type of security level the user will have.
Notification Host	Enter the IP address of the notification host. If you want to add more than one host, after entering the IP address of the first host, select the plus sign to add another host.
Enable Query	Select to enable or disable the query. By default, the query is enabled.
Port	Enter the port number in the field.
Events	Select the SNMP events that will be associated with that user.

Gigabit interfaces

When determining the interface speed of a FortiGate unit with a 10G interface, the IF-MIB.ifSpeed may not return the correct value. IF-MIB.ifSpeed is a 32-bit gauge used to report interface speeds in bits/second, and cannot convert to a 64-bit value. The 32-bit counter wrap the output too fast to be accurate.

In this case, you can use the value ifHighSpeed. It reports interface speeds in megabits/second. This ensures that 10Gb interfaces report the correct value.

SNMP agent

You need to first enter information and enable the FortiGate SNMP Agent. Enter information about the FortiGate unit to identify it so that when your SNMP manager receives traps from the FortiGate unit, you will know which unit sent the information.

To configure the SNMP agent - web-based manager

- 1 Go to *System > Config > SNMP*.
- 2 Select *Enable* for the *SNMP Agent*.
- 3 Enter a descriptive name for the agent.
- 4 Enter the location of the FortiGate unit.
- 5 Enter a contact or administrator for the SNMP Agent or FortiGate unit.
- 6 Select *Apply*.

To configure SNMP agent - CLI

```
config system snmp sysinfo
    set status enable
    set contact-info <contact_information>
    set description <description_of_FortiGate>
    set location <FortiGate_location>
end
```

SNMP community

An SNMP community is a grouping of devices for network administration purposes. Within that SNMP community, devices can communicate by sending and receiving traps and other information. One device can belong to multiple communities, such as one administrator terminal monitoring both a firewall SNMP and a printer SNMP community.

Add SNMP communities to your FortiGate unit so that SNMP managers can connect to view system information and receive SNMP traps.

You can add up to three SNMP communities. Each community can have a different configuration for SNMP queries and traps. Each community can be configured to monitor the FortiGate unit for a different set of events. You can also add the IP addresses of up to 8 SNMP managers to each community.

To add an SNMP v1/v2c community - web-based manager



When the FortiGate unit is in virtual domain mode, SNMP traps can only be sent on interfaces in the management virtual domain. Traps cannot be sent over other interfaces.

- 1 Go to *System > Config > SNMP*.
- 2 In the *SNMP v1/v2c* area, select *Create New*.

- 3 Enter a *Community Name*.
- 4 Enter the IP address and Identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
- 5 Select the interface if the SNMP manager is not on the same subnet as the FortiGate unit.
- 6 Enter the *Port* number that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the *Enable* check box to activate queries for each SNMP version.
- 7 Enter the Local and Remote port numbers that the FortiGate unit uses to send SNMP v1 and SNMP v2c traps to the SNMP managers in this community.
- 8 Select the *Enable* check box to activate traps for each SNMP version.
- 9 Select *OK*.

To add an SNMP v1/v2c community - CLI

```
config system snmp community
  edit <index_number>
    set events <events_list>
    set name <community_name>
    set query-v1-port <port_number>
    set query-v1-status {enable | disable}
    set query-v2c-port <port_number>
    set query-v2c-status {enable | disable}
    set status {enable | disable}
    set trap-v1-lport <port_number>
    set trap-v1-rport <port_number>
    set trap-v1-status {enable | disable}
    set trap-v2c-lport <port_number>
    set trap-v2c-rport <port_number>
    set trap-v2c-status {enable | disable}
  end
```

To add an SNMP v3 community - web-based manager

- 1 Go to System > Config > *SNMP*.
- 2 In the *SNMP v3* area, select *Create New*.
- 3 Enter a *User Name*.
- 4 Select a *Security Level* and associated authorization algorithms.
- 5 Enter the IP address of the *Notification Host* SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
- 6 Enter the *Port* number that the SNMP managers in this community use to receive configuration information from the FortiGate unit. Select the *Enable* check box to activate queries for each SNMP version.
- 7 Select the *Enable* check box to activate traps.
- 8 Select *OK*.

To add an SNMP v3 community - CLI

```
config system snmp user
  edit <index_number>
    set security-level [auth-priv | auth-no-priv | no-auth-no-priv]
  end
```

```
set queries enable
set query-port <port_number>
set notify-hosts <ip_address>
set events <event_selections>
end
```

Enabling on the interface

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections.

To configure SNMP access - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Choose an interface that an SNMP manager connects to and select *Edit*.
- 3 In *Administrative Access*, select *SNMP*.
- 4 Select *OK*.

To configure SNMP access - CLI

```
config system interface
edit <interface_name>
set allowaccess snmp
end
```



When using the `allowaccess` command to add SNMP, you need to also include any other access for the interface. This command will only use what is entered. That is, if you had HTTPS and SSH enabled before, these will be disabled if only the above command is used. In this case, for the `allow access` command, enter `set allowaccess https ssh snmp`.

Fortinet MIBs

The FortiGate SNMP agent supports Fortinet proprietary MIBs as well as standard RFC 1213 and RFC 2665 MIBs. RFC support includes support for the parts of RFC 2665 (Ethernet-like MIB) and the parts of RFC 1213 (MIB II) that apply to FortiGate unit configuration.

There are two MIB files for FortiGate units - the Fortinet MIB, and the FortiGate MIB. The Fortinet MIB contains traps, fields and information that is common to all Fortinet products. The FortiGate MIB contains traps, fields and information that is specific to FortiGate units. Each Fortinet product has its own MIB. If you use other Fortinet products you will need to download their MIB files as well. Both MIB files are used for FortiOS and FortiOS Carrier; there are no additional traps for the Carrier version of the operating system.

The Fortinet MIB and FortiGate MIB along with the two RFC MIBs are listed in tables in this section. You can download the two FortiGate MIB files from Fortinet Customer Support. The Fortinet MIB contains information for Fortinet products in general. the Fortinet FortiGate MIB includes the system information for The FortiGate unit and version of FortiOS. Both files are required for proper SNMP data collection.

To download the MIB files

- 1 Login to the Customer Support web site at support.fortinet.com.
- 2 Go to *Download > Firmware Images*.
- 3 Select *FortiGate > v4.00 > Core MIB*.

- 4 Select and download the Fortinet core MIB file.
- 5 Move up one directory level.
- 6 Select the firmware version, revision and patch (if applicable).
- 7 Select the *MIB* directory.
- 8 Select and download the Fortinet FortiGate MIB file.

Your SNMP manager may already include standard and private MIBs in a compiled database that is ready to use. You must add the Fortinet proprietary MIB to this database to have access to the Fortinet specific information.



There were major changes to the MIB files between v3.0 and v4.0. You need to use the new MIBs for v4.0 or you may mistakenly access the wrong traps and fields.

MIB files are updated for each version of FortiOS. When upgrading the firmware ensure that you updated the Fortinet FortiGate MIB file as well

Table 18: Fortinet MIBs

MIB file name or RFC	Description
FORTINET-CORE-MIB.mib	<p>The Fortinet MIB includes all system configuration information and trap information that is common to all Fortinet products.</p> <p>Your SNMP manager requires this information to monitor FortiGate unit configuration settings and receive traps from the FortiGate SNMP agent. For more information, see “Fortinet and FortiGate traps” on page 436 and “Fortinet and FortiGate MIB fields” on page 441.</p>
FORTINET-FORTIGATE-MIB.mib	<p>The FortiGate MIB includes all system configuration information and trap information that is specific to FortiGate units.</p> <p>Your SNMP manager requires this information to monitor FortiGate configuration settings and receive traps from the FortiGate SNMP agent. FortiManager systems require this MIB to monitor FortiGate units.</p> <p>For more information, see “Fortinet and FortiGate traps” on page 436 and “Fortinet and FortiGate MIB fields” on page 441.</p>
RFC-1213 (MIB II)	<p>The FortiGate SNMP agent supports MIB II groups with these exceptions.</p> <ul style="list-style-type: none"> • No support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP/ICMP/TCP/UDP/etc.) do not accurately capture all FortiGate traffic activity. More accurate information can be obtained from the information reported by the Fortinet MIB.
RFC-2665 (Ethernet-like MIB)	<p>The FortiGate SNMP agent supports Ethernet-like MIB information. FortiGate SNMP does not support for the dot3Tests and dot3Errors groups.</p>

SNMP get command syntax

Normally, to get configuration and status information for a FortiGate unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

...where...

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is `public`.

<address_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself.

The SNMP `get` command gets firmware version running on the FortiGate unit. The community name is `public`. The IP address of the interface configured for SNMP management access is `10.10.10.1`. The firmware version MIB field is `fgSysVersion` and the OID for this MIB field is `1.3.6.1.4.1.12356.101.4.1.1`. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgSysVersion
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.4.1.1
```

Fortinet and FortiGate traps

An SNMP manager can request information from the Fortinet device's SNMP agent, or that agent can send traps when an event occurs. Traps are a method used to inform the SNMP manager that something has happened or changed on the Fortinet device.

To receive FortiGate device SNMP traps, you must load and compile the `FORTINET-CORE-MIB` and `FORTINET-FORTIGATE-MIB` files into your SNMP manager. Traps sent include the trap message as well as the FortiGate unit serial number (`fnSysSerial`) and hostname (`sysName`).

The tables in this section include information about SNMP traps and variables. These tables have been included to help you locate the object identifier number (OID), trap message, and trap description of the Fortinet trap or variable you need.

The name of the table indicates if the trap is found in the Fortinet MIB or the FortiGate MIB. The Trap Message column includes the message included with the trap as well as the SNMP MIB field name to help locate the information about the trap. Traps starting with `fn` such as `fnTrapCpuThreshold` are defined in the Fortinet MIB. Traps starting with `fg` such as `fgTrapAvVirus` are defined in the FortiGate MIB.

The object identifier (OID) is made up of the number at the top of the table with the index added to the end. For example if the OID is `1.3.6.1.4.1.12356.101.2.0` and the index is `4`, the full OID is `1.3.6.1.4.1.12356.101.2.0.4`. The OID and the name of the object are how SNMP managers refer to fields and traps from the Fortinet and FortiGate MIBs.

Indented rows are fields that are part of the message or table associated with the preceding row.

The tables include:

- [Generic Fortinet traps \(OID 1.3.6.1.4.1.12356.101.3.0\)](#)
- [System traps \(OID 1.3.6.1.4.1.12356.1.3.0\)](#)

- FortiGate VPN traps (OID1.3.6.1.4.1.12356.1.3.0)
- FortiGate IPS traps (OID1.3.6.1.4.1.12356.1.3.0)
- FortiGate antivirus traps (OID1.3.6.1.4.1.12356.1.3.0)

- [FortiGate HA traps \(OID1.3.6.1.4.1.12356.1.3.0\)](#)

Table 19: Generic Fortinet traps (OID 1.3.6.1.4.1.12356.101.3.0)

Index	Trap message	Description
.1	ColdStart	Standard traps as described in RFC 1215.
.2	WarmStart	
.3	LinkUp	
.4	LinkDown	

Table 20: System traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.101	CPU usage high (fnTrapCpuThreshold)	CPU usage exceeds 80%. This threshold can be set in the CLI using <code>config system snmp sysinfo, set trap-high-cpu-threshold</code> .
.102	Memory low (fnTrapMemThreshold)	Memory usage exceeds 90%. This threshold can be set in the CLI using <code>config system snmp sysinfo, set trap-low-memory-threshold</code> .
.103	Log disk too full (fnTrapLogDiskThreshold)	Log disk usage has exceeded the configured threshold. Only available on devices with log disks. This threshold can be set in the CLI using <code>config system snmp sysinfo, set trap-log-full-threshold</code> .
.104	Temperature too high (fnTrapTempHigh)	A temperature sensor on the device has exceeded its threshold. Not all devices have thermal sensors. See manual for specifications.
.105	Voltage outside acceptable range (fnTrapVoltageOutOfRange)	Power levels have fluctuated outside of normal levels. Not all devices have voltage monitoring instrumentation.
.106	Power supply failure (fnTrapPowerSupplyFailure)	Power supply failure detected. Not available on all models. Available on some devices which support redundant power supplies.
.201	Interface IP change (fnTrapIpChange)	The IP address for an interface has changed. The trap message includes the name of the interface, the new IP address and the serial number of the Fortinet unit. You can use this trap to track interface IP address changes for interfaces with dynamic IP addresses set using DHCP or PPPoE.
.999	Diagnostic trap (fnTrapTest)	This trap is sent for diagnostic purposes. It has an OID index of .999.

Table 21: FortiGate VPN traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.301	VPN tunnel is up (fgTrapVpnTunUp)	An IPSec VPN tunnel has started.
.302	VPN tunnel down (fgTrapVpnTunDown)	An IPSec VPN tunnel has shut down.

Table 21: FortiGate VPN traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
	Local gateway address (fgVpnTrapLocalGateway)	Address of the local side of the VPN tunnel. This information is associated with both of the VPN tunnel traps. (OID1.3.6.1.4.1.12356.101.12.3.2)
	Remote gateway address (fgVpnTrapRemoteGateway)	Address of remote side of the VPN tunnel. This information is associated with both of the VPN tunnel traps. (OID1.3.6.1.4.1.12356.101.12.3.2)

Table 22: FortiGate IPS traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.503	IPS Signature (fgTrapIpsSignature)	IPS signature detected.
.504	IPS Anomaly (fgTrapIpsAnomaly)	IPS anomaly detected.
.505	IPS Package Update (fgTrapIpsPkgUpdate)	The IPS signature database has been updated.
	(fgIpsTrapSigId)	ID of IPS signature identified in trap. (OID 1.3.6.1.4.1.12356.101.9.3.1)
	(fgIpsTrapSrcIp)	IP Address of the IPS signature trigger. (OID 1.3.6.1.4.1.12356.101.9.3.2)
	(fgIpsTrapSigMsg)	Message associated with IPS event. (OID 1.3.6.1.4.1.12356.101.9.3.3)

Table 23: FortiGate antivirus traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.601	Virus detected (fgTrapAvVirus)	The antivirus engine detected a virus in an infected file from an HTTP or FTP download or from an email message.
.602	Oversize file/email detected (fgTrapAvOversize)	The antivirus scanner detected an oversized file.
.603	Filename block detected (fgTrapAvPattern)	The antivirus scanner blocked a file that matched a known virus pattern.
.604	Fragmented file detected (fgTrapAvFragmented)	The antivirus scanner detected a fragmented file or attachment.
.605	(fgTrapAvEnterConserve)	The AV engine entered conservation mode due to low memory conditions.
.606	(fgTrapAvBypass)	The AV scanner has been bypassed due to conservation mode.
.607	(fgTrapAvOversizePass)	An oversized file has been detected, but has been passed due to configuration.
.608	(fgTrapAvOversizeBlock)	An oversized file has been detected, and has been blocked.
	(fgAvTrapVirName)	The virus name that triggered the event. (OID 1.3.6.1.4.1.12356.101.8.3.1)

Table 24: FortiGate HA traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.401	HA switch (fgTrapHaSwitch)	The specified cluster member has switched from a slave role to a master role.
.402	HA State Change (fgTrapHaStateChange)	The trap sent when the HA cluster member changes its state.

Table 24: FortiGate HA traps (OID1.3.6.1.4.1.12356.1.3.0)

Index	Trap message	Description
.403	HA Heartbeat Failure (fgTrapHaHBFail)	The heartbeat failure count has exceeded the configured threshold.
.404	HA Member Unavailable (fgTrapHaMemberDown)	An HA member becomes unavailable to the cluster.
.405	HA Member Available (fgTrapHaMemberUp)	An HA member becomes available to the cluster.
	(fgHaTrapMemberSerial)	Serial number of an HA cluster member. Used to identify the origin of a trap when a cluster is configured. (OID1.3.6.1.4.1.12356.101.13.3.1)

Fortinet and FortiGate MIB fields

The FortiGate MIB contains fields reporting the current FortiGate unit status information. The following tables list the names of the MIB fields and describe the status information available for each one. You can view more details about the information available from all Fortinet and FortiGate MIB fields by compiling the `FORTINET-CORE-MIB.mib` and `FORTINET-FORTIGATE-MIB.mib` files into your SNMP manager and browsing the MIB fields on your computer.

To help locate a field, the object identifier (OID) number for each table of fields has been included. The OID number for a field is that field's position within the table, starting at 0. For example `fnSysVersion` has an OID of 1.3.6.1.4.1.12356.2.

Fortinet MIB

Table 25: OIDs for the Fortinet-Core-MIB

MIB Field	OID	Description
fnSystem	1.3.6.1.4.1.12356.100.1.1	
fnSysSerial	1.3.6.1.4.1.12356.100.1.1.1	Device serial number. This is the same serial number as given in the ENTITY-MIB tables for the base entity.
fnMgmt	.3.6.1.4.1.12356.100.1.2	
fnMgmtLanguage	1.3.6.1.4.1.12356.100.1.2.1	Language used for administration interfaces.
fnAdmin	1.3.6.1.4.1.12356.100.1.2.100	
fnAdminNumber	1.3.6.1.4.1.12356.100.1.2.100.1	The number of admin accounts in fnAdminTable.
fnAdminTable	1.3.6.1.4.1.12356.100.1.2.100.2	A table of administrator accounts on the device. This table is intended to be extended with platform specific information.
fnAdminEntry	1.3.6.1.4.1.12356.100.1.2.100.2.1	An entry containing information applicable to a particular admin account.

Table 25: OIDs for the Fortinet-Core-MIB

MIB Field	OID	Description
fnAdminIndex	1.3.6.1.4.1.12356.100.1.2.100.2.1.1	An index uniquely defining an administrator account within the fnAdminTable.
fnAdminName	1.3.6.1.4.1.12356.100.1.2.100.2.1.2	The user-name of the specified administrator account.
fnAdminAddrType	1.3.6.1.4.1.12356.100.1.2.100.2.1.3	The type of address stored in fnAdminAddr, in compliance with INET-ADDRESS-MIB
fnAdminAddr	1.3.6.1.4.1.12356.100.1.2.100.2.1.4	The address prefix identifying where the administrator account can be used from, typically an IPv4 address. The address type/format is determined by fnAdminAddrType.
fnAdminMask	1.3.6.1.4.1.12356.100.1.2.100.2.1.5	The address prefix length (or network mask) applied to the fnAdminAddr to determine the subnet or host the administrator can access the device from.
fnTraps	1.3.6.1.4.1.12356.100.1.3	
fnTrapsPrefix	1.3.6.1.4.1.12356.100.1.3.0	
fnTrapCpuThreshold	1.3.6.1.4.1.12356.100.1.3.0.101	Indicates that the CPU usage has exceeded the configured threshold.
fnTrapMemThreshold	1.3.6.1.4.1.12356.100.1.3.0.102	Indicates memory usage has exceeded the configured threshold.
fnTrapLogDiskThreshold	1.3.6.1.4.1.12356.100.1.3.0.103	Log disk usage has exceeded the configured threshold. Only available on devices with log disks.
fnTrapTempHigh	1.3.6.1.4.1.12356.100.1.3.0.104	A temperature sensor on the device has exceeded its threshold. Not all devices have thermal sensors. See manual for specifications.
fnTrapVoltageOutOfRange	1.3.6.1.4.1.12356.100.1.3.0.105	Power levels have fluctuated outside of normal levels. Not all devices have voltage monitoring instrumentation.

Table 25: OIDs for the Fortinet-Core-MIB

MIB Field	OID	Description
fnTrapPowerSupplyFailure	1.3.6.1.4.1.12356.100.1.3.0.106	Power supply failure detected. Not available on all models. Available on some devices which support redundant power supplies. See manual for specifications.
fnTrapIpChange	1.3.6.1.4.1.12356.100.1.3.0.201	Indicates that the IP address of the specified interface has been changed.
fnTrapTest	1.3.6.1.4.1.12356.100.1.3.0.999	Trap sent for diagnostic purposes by an administrator.
fnTrapObjects	1.3.6.1.4.1.12356.100.1.3.1	
fnGenTrapMsg	1.3.6.1.4.1.12356.100.1.3.1.1	Generic message associated with an event. The content will depend on the nature of the trap.
fnMIBConformance	1.3.6.1.4.1.12356.100.10	
fnSystemComplianceGroup	1.3.6.1.4.1.12356.100.10.1	Objects relating to the physical device.
fnMgmtComplianceGroup	1.3.6.1.4.1.12356.100.10.2	Objects relating the management of a device.
fnAdmincomplianceGroup	1.3.6.1.4.1.12356.100.10.3	Administration access control objects.
fnTrapsComplianceGroup	1.3.6.1.4.1.12356.100.10.4	Event notifications.
fnNotifObjectsCompliance Group	1.3.6.1.4.1.12356.100.10.5	Object identifiers used in notifications.
fnMIBCompliance	1.3.6.1.4.1.12356.100.10.100	Object identifiers used in notifications. Objects are required if their containing trap is implemented.

FortiGate MIB

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgModel	1.3.6.1.4.1.12356.101.1	
fgTraps	1.3.6.1.4.1.12356.101.2	
fgTrapPrefix	1.3.6.1.4.1.12356.101.2.0	
fgTrapVpnTunup	1.3.6.1.4.1.12356.101.2.0.301	Indicates that the specified VPN tunnel has been brought up.
fgTrapVpnTunDown	1.3.6.1.4.1.12356.101.2.0.302	The specified VPN tunnel has been brought down.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgTrapHaSwitch	1.3.6.1.4.1.12356.101.2.0.401	The specified cluster member has transitioned from a slave role to a master role.
fgTrapHaStateChange	1.3.6.1.4.1.12356.101.2.0.402	Trap being sent when the HA cluster member changes its state.
fgTrapHaBFail	1.3.6.1.4.1.12356.101.2.0.403	The heartbeat device failure count has exceeded the configured threshold.
fgTrapHaMemberDown	1.3.6.1.4.1.12356.101.2.0.404	The specified device (by serial number) is moving to a down state.
fgTrapHaMemberUp	1.3.6.1.4.1.12356.101.2.0.405	A new cluster member has joined the cluster.
fgTrapIpsSignature	1.3.6.1.4.1.12356.101.2.0.503	An IPS signature has been triggered.
fgTrapIpsAnomaly	1.3.6.1.4.1.12356.101.2.0.504	An IPS anomaly has been detected.
fgTrapIpsPkgUpdate	1.3.6.1.4.1.12356.101.2.0.505	The IPS signature database has been updated.
fgTrapAvVirus	1.3.6.1.4.1.12356.101.2.0.601	A virus has been detected by the antivirus engine.
fgTrapAvOversize	1.3.6.1.4.1.12356.101.2.0.602	An oversized file has been detected by the antivirus engine.
fgTrapAvPattern	1.3.6.1.4.1.12356.101.2.0.603	The antivirus engine has blocked a file because it matched a configured pattern.
fgTrapAvFragmented	1.3.6.1.4.1.12356.101.2.0.604	The antivirus engine has detected a fragmented file.
fgTrapAvEnterConserve	1.3.6.1.4.1.12356.101.2.0.605	The antivirus engine has entered conservation mode due to low memory conditions.
fgTrapAvBypass	1.3.6.1.4.1.12356.101.2.0.606	The antivirus engine has been bypassed due to conservation mode.
fgTrapAvOversizePass	1.3.6.1.4.1.12356.101.2.0.607	An oversized file has been detected, but has been passed due to configuration.
fgTrapAvOversizeBlock	1.3.6.1.4.1.12356.101.2.0.608	An oversized file has been detected and has been blocked.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgTrapFazDisconnect	1.3.6.1.4.1.12356.101.2.0.701	The device has been disconnected from the FortiAnalyzer.
Virtual Domains		
fgVirtualDomain	1.3.6.1.4.1.12356.101.3	
fgVdInfo	1.3.6.1.4.1.12356.101.3.1	
fgVdNumber	1.3.6.1.4.1.12356.101.3.1.1	The number of virtual domains in vdTable.
fgVdMaxVdoms	1.3.6.1.4.1.12356.101.3.1.2	The maximum number of virtual domains allowed on the device as allowed by hardware and/or licensing.
fgVdEnabled	1.3.6.1.4.1.12356.101.3.1.3	Whether virtual domains are enabled on this device.
fgVdTables	1.3.6.1.4.1.12356.101.3.2	
fgVdTable	1.3.6.1.4.1.12356.101.3.2.1	
fgVdEntry	1.3.6.1.4.1.12356.101.3.2.1.1	An entry containing information applicable to a particular virtual domain.
fgVdEntIndex	1.3.6.1.4.1.12356.101.3.2.1.1.1	Internal virtual domain index used to uniquely identify rows in this table. This index is also used by other tables referencing a virtual domain.
fgVdEntName	1.3.6.1.4.1.12356.101.3.2.1.1.2	The name of the virtual domain.
fgVdEntOpMode	1.3.6.1.4.1.12356.101.3.2.1.1.3	Operation mode of the virtual domain (NAT or transparent).
fgVdTpTable	1.3.6.1.4.1.12356.101.3.2.2	A table of virtual domains in transparent operation mode. This table has a dependent relationship with fgVdTable.
fgVdTpEntry	1.3.6.1.4.1.12356.101.3.2.2.1	An entry containing information applicable to a particular virtual domain in transparent mode.
fgVdTpMgmtAddrType	1.3.6.1.4.1.12356.101.3.2.2.1.1	The type of address stored in fgVdTpMgmtAddr, in compliance with INET-ADDRESS-MIB.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgVdTpMgmtAddr	1.3.6.1.4.1.12356.101.3.2.2.1.2	The management IP address of the virtual domain in transparent mode, typically an IPv4 address. The address type/format is determined by fgVdTpMgmtAddrType.
fgVdTpMgmtMask	1.3.6.1.4.1.12356.101.3.2.2.1.3	The address prefix length (or network mask) applied to the fgVdTpMgmtAddr.
System		
fgSystem	1.3.6.1.4.1.12356.101.4	
fgSystemInfo	1.3.6.1.4.1.12356.101.4.1	
fgSysVersion	1.3.6.1.4.1.12356.101.4.1.1	Firmware version.
fgSysMgmtVdom	1.3.6.1.4.1.12356.101.4.1.2	Index that identifies the management virtual domain. This index corresponds to the index used by fgVdTable.
fgSysCpuUsage	1.3.6.1.4.1.12356.101.4.1.3	Current CPU usage (percentage).
fgSysMemUsage	1.3.6.1.4.1.12356.101.4.1.4	Current memory usage (percentage).
fgSysMemCapacity	1.3.6.1.4.1.12356.101.4.1.5	Total physical RAM installed (KB)
fgSysDiskUsage	1.3.6.1.4.1.12356.101.4.1.6	Current hard disk usage (MB), if disk is present.
fgSysDiskCapacity	1.3.6.1.4.1.12356.101.4.1.7	Total hard disk capacity (MB), if disk is present.
fgSysSesCount	1.3.6.1.4.1.12356.101.4.1.8	Number of active sessions on the device.
fgSysLowMemUsage	1.3.6.1.4.1.12356.101.4.1.9	Current lowmem utilization (percentage). Lowmem is memory available for the kernel's own data structures and kernel specific tables. The system can get into a bad state if it runs out of lowmem.
fgSysLowMemCapacity	1.3.6.1.4.1.12356.101.4.1.10	Total lowmem capacity (KB).
Firewall		
fgFirewal	1.3.6.1.4.1.12356.101.5	
fgFwPolicies	1.3.6.1.4.1.12356.101.5.1	
fgFwPolInfo	1.3.6.1.4.1.12356.101.5.1.1	
fgFwPolTables	1.3.6.1.4.1.12356.101.5.1.2	

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgFwPolStatsTable	1.3.6.1.4.1.12356.101.5.1.2.1	Security policy statistics table. This table has a dependent expansion relationship with fgVdTable. Only virtual domains with enabled policies are present in this table.
fgFwPolStatsEntry	1.3.6.1.4.1.12356.101.5.1.2.1.1	Security policy statistics on a virtual domain.
fgFwPolID	1.3.6.1.4.1.12356.101.5.1.2.1.1.1	Security policy ID. Only enabled policies are present in this table. Policy IDs are only unique within a virtual domain.
fgFwPolPktCount	1.3.6.1.4.1.12356.101.5.1.2.1.1.2	Number of packets matched to policy (passed or blocked, depending on policy action). Count is from the time the policy became active.
fgFwPolByteCount	1.3.6.1.4.1.12356.101.5.1.2.1.1.3	Number of bytes in packets matching the policy. See fgFwPolPktCount.
fgFwUsers	1.3.6.1.4.1.12356.101.5.2	
fgFwUserInfo	1.3.6.1.4.1.12356.101.5.2.1	
fgFwUserNumber	1.3.6.1.4.1.12356.101.5.2.1.1	The number of user accounts in fgFwUserTable.
fgFwUserAuthTimeout	1.3.6.1.4.1.12356.101.5.2.1.2	Idle period after which a firewall-authentication user's session is automatically expired.
fgFwUserTables	1.3.6.1.4.1.12356.101.5.2.2	
fgFwUserTable	1.3.6.1.4.1.12356.101.5.2.2.1	A list of local and proxy (Radius server) user accounts for use with firewall user authentication.
fgFwUserEntry	1.3.6.1.4.1.12356.101.5.2.2.1.1	An entry containing information applicable to a particular user account.
fgFwUserIndex	1.3.6.1.4.1.12356.101.5.2.2.1.1.1	An index for uniquely identifying the users in fgFwUserTable.
fgFwUserName	1.3.6.1.4.1.12356.101.5.2.2.1.1.2	User name of the specified account.
fgFwUserAuth	1.3.6.1.4.1.12356.101.5.2.2.1.1.3	Type of authentication the account uses (local, RADIUS, LDAP, etc.).

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgFwUserState	1.3.6.1.4.1.12356.101.5.2.2.1.1.4	Status of the user account (enabled/disabled).
fgFwUserVdom	1.3.6.1.4.1.12356.101.5.2.2.1.1.5	Virtual domain the user account exists in. This index corresponds to the index used in fgVdTable.
FortiManager and Administration		
fgMgmt	1.3.6.1.4.1.12356.101.6	
fgFmTrapPrefix	1.3.6.1.4.1.12356.101.6.0	
fgFmTrapDeployComplete	1.3.6.1.4.1.12356.101.6.0.1000	Indicates when deployment of a new configuration has been completed. Used for verification by FortiManager.
fgFmTrapDeployInProgress	1.3.6.1.4.1.12356.101.6.0.1002	Indicates that a configuration change was not immediate and that the change is currently in progress. Used for verification by FortiManager.
fgFmTrapConfChange	1.3.6.1.4.1.12356.101.6.0.1003	The device configuration has been changed by something other than the managing FortiManager unit.
fgFmTrapIfChange	1.3.6.1.4.1.12356.101.6.0.1004	Trap is sent to the managing FortiManager if an interface IP is changed.
fgAdmin	1.3.6.1.4.1.12356.101.6.1	
fgAdminOptions	1.3.6.1.4.1.12356.101.6.1.1	
fgAdminIdleTimeout	1.3.6.1.4.1.12356.101.6.1.1.1	Idle period after which an administrator is automatically logged out of the system.
fgAdminLcdProtection	1.3.6.1.4.1.12356.101.6.1.1.2	Status of the LCD protection (enabled/disabled).
fgAdminTables	1.3.6.1.4.1.12356.101.6.1.2	
fgAdminTable	1.3.6.1.4.1.12356.101.6.1.2.1	A table of administrator accounts on the device.
fgAdminEntry	1.3.6.1.4.1.12356.101.6.1.2.1.1	An entry containing information applicable to a particular admin account.
fgAdminVdom	1.3.6.1.4.1.12356.101.6.1.2.1.1.1	The virtual domain the administrator belongs to.
fgMgmtTrapObjects	1.3.6.1.4.1.12356.101.6.2	
fgManIfIp	1.3.6.1.4.1.12356.101.6.2.1	IP address of the interface listed in the trap.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgManIfMask	1.3.6.1.4.1.12356.101.6.2.2	Mask of subnet the interface belongs to.
Antivirus		
fgAntivirus	1.3.6.1.4.1.12356.101.8	
fgAvInfo	1.3.6.1.4.1.12356.101.8.1	
fgAvTables	1.3.6.1.4.1.12356.101.8.2	
fgAvStatsTable	1.3.6.1.4.1.12356.101.8.2.1	A table of Antivirus statistics per virtual domain.
fgAvStatsEntry	1.3.6.1.4.1.12356.101.8.2.1.1	Antivirus statistics for a particular virtual domain.
fgAvVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.1	Number of virus transmissions detected in the virtual domain since start up.
fgAvVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.2	Number of virus transmissions blocked in the virtual domain since start up.
fgAvHTTPVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.3	Number of virus transmissions over HTTP detected in the virtual domain since start up.
fgAvHTTPVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.4	Number of virus transmissions over HTTP blocked in the virtual domain since start up.
fgAvSMTPVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.5	Number of virus transmissions over SMTP detected in the virtual domain since start up.
fgAvSMTPVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.6	Number of virus transmissions over SMTP blocked in the virtual domain since start up.
fgAvPOP3VirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.7	Number of virus transmissions over POP3 detected in the virtual domain since start up.
fgAvPOP3VirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.8	Number of virus transmissions over POP3 blocked in the virtual domain since start up.
fgAvIMAPVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.9	Number of virus transmissions over IMAP detected in the virtual domain since start up.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgAvIMAPVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.10	Number of virus transmissions over IMAP blocked in the virtual domain since start up.
fgAvFTPVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.11	Number of virus transmissions over FTP detected in the virtual domain since start up.
fgAvFTPVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.12	Number of virus transmissions over FTP blocked in the virtual domain since start up.
fgAvIMVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.13	Number of virus transmissions over IM protocols detected in the virtual domain since start up.
fgAvIMVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.14	Number of virus transmissions over IM protocols blocked in the virtual domain since start up.
fgAvNNTPVirusDetected	1.3.6.1.4.1.12356.101.8.2.1.1.15	Number of virus transmissions over NNTP detected in the virtual domain since start up.
fgAvNNTPVirusBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.16	Number of virus transmissions over NNTP blocked in the virtual domain since start up.
fgAvOversizedDetected	1.3.6.1.4.1.12356.101.8.2.1.1.17	Number of over-sized file transmissions detected in the virtual domain since start up.
fgAvOversizedBlocked	1.3.6.1.4.1.12356.101.8.2.1.1.18	Number of over-sized file transmissions blocked in the virtual domain since start up.
fgAvTrapObjects	1.3.6.1.4.1.12356.101.8.3	
fgAvTrapVirName	1.3.6.1.4.1.12356.101.8.3.1	Virus name that triggered event.
IPS		
fglps	1.3.6.1.4.1.12356.101.9	
fglpsInfo	1.3.6.1.4.1.12356.101.9.1	
fglpsTables	1.3.6.1.4.1.12356.101.9.2	
fglpsStatsTable	1.3.6.1.4.1.12356.101.9.2.1	A table of IPS/IDS statistics per virtual domain.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fglpsStatsEntry	1.3.6.1.4.1.12356.101.9.2.1.1	IPS/IDS statistics for a particular virtual domain.
fglpsIntrusionDetected	1.3.6.1.4.1.12356.101.9.2.1.1.1	Number of intrusions detected since start up in this virtual domain.
fglpsIntrusionBlocked	1.3.6.1.4.1.12356.101.9.2.1.1.2	Number of intrusions blocked since start up in this virtual domain.
fglpsCritSevDetections	1.3.6.1.4.1.12356.101.9.2.1.1.3	Number of critical severity intrusions detected since start up in this virtual domain.
fglpsHighSevDetections	1.3.6.1.4.1.12356.101.9.2.1.1.4	Number of high severity intrusions detected since start up in this virtual domain.
fglpsMedSevDetections	1.3.6.1.4.1.12356.101.9.2.1.1.5	Number of medium severity intrusions detected since start up in this virtual domain.
fglpsLowSevDetections	1.3.6.1.4.1.12356.101.9.2.1.1.6	Number of low severity intrusions detected since start up in this virtual domain.
fglpsInfoSevDetections	1.3.6.1.4.1.12356.101.9.2.1.1.7	Number of informational severity intrusions detected since start up in this virtual domain.
fglpsSignatureDetections	1.3.6.1.4.1.12356.101.9.2.1.1.8	Number of intrusions detected by signature since start up in this virtual domain.
fglpsAnomalyDetections	1.3.6.1.4.1.12356.101.9.2.1.1.9	Number of intrusions detected as anomalies since start up in this virtual domain.
fglpsTrapObjects	1.3.6.1.4.1.12356.101.9.3	
fglpsTrapSigId	1.3.6.1.4.1.12356.101.9.3.1	ID of IPS signature identified in trap.
fglpsTrapSrcIp	1.3.6.1.4.1.12356.101.9.3.2	Source IP Address of the IPS signature trigger.
fglpsTrapSigMsg	1.3.6.1.4.1.12356.101.9.3.3	Message associated with IPS event.
Application Control		
fgApplications	1.3.6.1.4.1.12356.101.10	
fgWebfilter	1.3.6.1.4.1.12356.101.10.1	
fgWebfilterInfo	1.3.6.1.4.1.12356.101.10.1.1	

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgWebfilterTables	1.3.6.1.4.1.12356.101.10.1.2	
fgWebfilterStatsTable	1.3.6.1.4.1.12356.101.10.1.2.1	A table of Web filter statistics per virtual domain.
fgWebfilterStatsEntry	1.3.6.1.4.1.12356.101.10.1.2.1.1	Web filter statistics for a particular virtual domain.
fgWfHTTPBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.1	Number of HTTP sessions blocked by Web filter since start up.
fgWfHTTPSBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.2	Number of HTTPS sessions blocked by Web filter since start up.
fgWfHTTPURLBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.3	Number of HTTP URLs blocked by Web filter since start up.
fgWfHTTPSURLBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.4	Number of HTTPS URLs blocked by Web filter since start up.
fgWfActiveXBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.5	Number of ActiveX downloads blocked by Web filter since start up.
fgWfCookieBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.6	Number of HTTP Cookies blocked by Web filter since start up.
fgWfAppletBlocked	1.3.6.1.4.1.12356.101.10.1.2.1.1.7	Number of Applets blocked by Web-filter since start up.
fgFortiGuardStatsTable	1.3.6.1.4.1.12356.101.10.1.2.2	A table of FortiGuard statistics per virtual domain.
fgFortiGuardStatsEntry	1.3.6.1.4.1.12356.101.10.1.2.2.1	FortiGuard statistics for a particular virtual domain.
fgFgWfHTTPExamined	1.3.6.1.4.1.12356.101.10.1.2.2.1.1	Number of HTTP requests examined using FortiGuard since start up.
fgFgWfHTTPSExamined	1.3.6.1.4.1.12356.101.10.1.2.2.1.2	Number of HTTPS requests examined using FortiGuard since start up.
fgFgWfHTTPAllowed	1.3.6.1.4.1.12356.101.10.1.2.2.1.3	Number of HTTP requests allowed to proceed using FortiGuard since start up.
fgFgWfHTTPSAllowed	1.3.6.1.4.1.12356.101.10.1.2.2.1.4	Number of HTTPS requests allowed to proceed using FortiGuard since start up.
fgFgWfHTTPBlocked	1.3.6.1.4.1.12356.101.10.1.2.2.1.5	Number of HTTP requests blocked using FortiGuard since start up.
fgFgWfHTTPSBlocked	1.3.6.1.4.1.12356.101.10.1.2.2.1.6	Number of HTTPS requests blocked using FortiGuard since start up.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgFgWfHTTPLogged	1.3.6.1.4.1.12356.101.10.1.2.2.1.7	Number of HTTP requests logged using FortiGuard since start up.
fgFgWfHTTPSLogged	1.3.6.1.4.1.12356.101.10.1.2.2.1.8	Number of HTTPS requests logged using FortiGuard since start up.
fgFgWfHTTPOverridden	1.3.6.1.4.1.12356.101.10.1.2.2.1.9	Number of HTTP requests overridden using FortiGuard since start up.
fgFgWfHTTPSOverridden	1.3.6.1.4.1.12356.101.10.1.2.2.1.10	Number of HTTPS requests overridden using FortiGuard since start up.
fgAppProxyHTTP	1.3.6.1.4.1.12356.101.10.100	
fgApHTTPUpTime	1.3.6.1.4.1.12356.101.10.100.1	HTTP proxy up-time, in seconds.
fgApHTTPMemUsage	1.3.6.1.4.1.12356.101.10.100.2	HTTP proxy memory usage (percentage of system total).
fgApHTTPStatsTable	1.3.6.1.4.1.12356.101.10.100.3	A table of HTTP Proxy statistics per virtual domain.
fgApHTTPStatsEntry	1.3.6.1.4.1.12356.101.10.100.3.1	HTTP Proxy statistics for a particular virtual domain.
fgApHRRPReqProcessed	1.3.6.1.4.1.12356.101.10.100.3.1.1	Number of HTTP requests in this virtual domain processed by the HTTP proxy since start up.
fgApHTTPConnections	1.3.6.1.4.1.12356.101.10.100.4	HTTP proxy current connections.
fgAppProxySMTP	1.3.6.1.4.1.12356.101.10.101	
fgApSMTPUpTime	1.3.6.1.4.1.12356.101.10.101.1	SMTP Proxy up-time, in seconds.
fgAPSMTPMemUsage	1.3.6.1.4.1.12356.101.10.101.2	SMTP Proxy memory utilization (percentage of system total).
fgApSMTPStatsTable	1.3.6.1.4.1.12356.101.10.101.3	A table of SMTP proxy statistics per virtual domain.
fgApSMTPStatsEntry	1.3.6.1.4.1.12356.101.10.101.3.1	SMTP Proxy statistics for a particular virtual domain.
fgApSMTPReqProcessed	1.3.6.1.4.1.12356.101.10.101.3.1.1	Number of requests in this virtual domain processed by the SMTP proxy since start up.
fgApSMTPSpamDetected	1.3.6.1.4.1.12356.101.10.101.3.1.2	Number of spam detected in this virtual domain by the SMTP proxy since start up.
fgApSMTPConnections	1.3.6.1.4.1.12356.101.10.101.4	SMTP proxy current connections.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgAppProxyPOP3	1.3.6.1.4.1.12356.101.10.102	
fgApPOP3UpTime	1.3.6.1.4.1.12356.101.10.102.1	Up time of the POP3 proxy, in seconds.
fgApPOP3MemUsage	1.3.6.1.4.1.12356.101.10.102.2	Memory usage of the POP3 Proxy (percentage of system total).
fgApPOP3StatsTable	1.3.6.1.4.1.12356.101.10.102.3	A table of POP3 proxy statistics per virtual domain.
fgApPOP3StatsEntry	1.3.6.1.4.1.12356.101.10.102.3.1	Proxy pop3 statistics for a particular virtual domain.
fgApPOP3ReqProcessed	1.3.6.1.4.1.12356.101.10.102.3.1.1	Number of requests in this virtual domain processed by the POP3 proxy since start up.
fgApPOP3SpamDetected	1.3.6.1.4.1.12356.101.10.102.3.1.2	Number of spam detected in this virtual domain by the POP3 Proxy since start up.
fgApPOP3Connections	1.3.6.1.4.1.12356.101.10.102.4	POP3 proxy current connections.
fgAppProxyIMAP	1.3.6.1.4.1.12356.101.10.103	
fgApIMAPUpTime	1.3.6.1.4.1.12356.101.10.103.1	Up time of the IMAP proxy, in seconds.
fgApIMAPMemUsage	1.3.6.1.4.1.12356.101.10.103.2	Memory utilization of the IMAP Proxy (as a percentage of the system total).
fgApIMAPStatsTable	1.3.6.1.4.1.12356.101.10.103.3	A table of IMAP proxy statistics per virtual domain.
fgApIMAPStatsEntry	1.3.6.1.4.1.12356.101.10.103.3.1	IMAP Proxy statistics for a particular virtual domain.
fgApIMAPReqProcessed	1.3.6.1.4.1.12356.101.10.103.3.1.1	Number of requests in this virtual domain processed by the IMAP proxy since start up.
fgApIMAPSpamDetected	1.3.6.1.4.1.12356.101.10.103.3.1.2	Number of spam detected in this virtual domain by the IMAP proxy since start up.
fgApIMAPConnections	1.3.6.1.4.1.12356.101.10.103.4	IMAP proxy current connections.
fgAppProxyNNTP	1.3.6.1.4.1.12356.101.10.104	
fgApNNTPUpTime	1.3.6.1.4.1.12356.101.10.104.1	Up time of the NNTP proxy, in seconds.
fgApNNTPMemUsage	1.3.6.1.4.1.12356.101.10.104.2	Memory utilization of the NNTP proxy, as a percentage of the system total.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgApNNTPStatsTable	1.3.6.1.4.1.12356.101.10.104.3	A table of NNTP proxy statistics per virtual domain.
fgApNNTPStatsEntry	1.3.6.1.4.1.12356.101.10.104.3.1	NNTP Proxy statistics for a particular virtual domain.
fgApNNTPReqProcessed	1.3.6.1.4.1.12356.101.10.104.3.1.1	Number of requests in the virtual domain processed by the NNTP proxy since start up.
fgApNNTPConnections	1.3.6.1.4.1.12356.101.10.104.4	NNTP proxy current connections.
fgAppProxyIM	1.3.6.1.4.1.12356.101.10.105	
fgApIMUpTime	1.3.6.1.4.1.12356.101.10.105.1	Up time of the IM proxy, in seconds.
fgApIMMemUsage	1.3.6.1.4.1.12356.101.10.105.2	IM Proxy memory usage, as a percentage of the system total.
fgApIMStatsTable	1.3.6.1.4.1.12356.101.10.105.3	A table of IM proxy statistics per virtual domain.
fgApIMStatsEntry	1.3.6.1.4.1.12356.101.10.105.3.1	IM Proxy statistics for a particular virtual domain.
fgApIMReqProcessed	1.3.6.1.4.1.12356.101.10.105.3.1.1	Number of requests in this virtual domain processed by the IM proxy since start up.
fgAppProxySIP	1.3.6.1.4.1.12356.101.10.106	
fgApSIPUpTime	1.3.6.1.4.1.12356.101.10.106.1	Up time of the SIP Proxy, in seconds.
fgApSIPMemUsage	1.3.6.1.4.1.12356.101.10.106.2	SIP Proxy memory utilization, as a percentage of the system total.
fgApSIPStatsTable	1.3.6.1.4.1.12356.101.10.106.3	A table of SIP proxy statistics per virtual domain.
fgApSIPStatsEntry	1.3.6.1.4.1.12356.101.10.106.3.1	SIP Proxy statistics for a particular virtual domain.
fgApSIPClientReg	1.3.6.1.4.1.12356.101.10.106.3.1.1	Number of client registration requests (Register and Options) in this virtual domain processed by the SIP proxy since start up.
fgApSIPCallHandling	1.3.6.1.4.1.12356.101.10.106.3.1.2	Number of call handling requests (Invite, Ack, Bye, Cancel and Refer) in this virtual domain processed by the SIP proxy since start up.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgApSIPServices	1.3.6.1.4.1.12356.101.10.106.3.1.3	Number of service requests (Subscribe, notify and Message) in this virtual domain processed by the SIP proxy since start up.
fgApSIPOtherReq	1.3.6.1.4.1.12356.101.10.106.3.1.4	Number of other sip requests in this virtual domain processed by the SIP proxy since start up.
fgAppScanUnit	1.3.6.1.4.1.12356.101.10.107	
fgAppSuNumber	1.3.6.1.4.1.12356.101.10.107.1	The number of scan units in the fgAppSuStatsTable.
fgAppSuStatsTable	1.3.6.1.4.1.12356.101.10.107.2	A table of scan unit statistics.
fgAppSuStatsEntry	1.3.6.1.4.1.12356.101.10.107.2.1	Statistics entry for a particular scan unit.
fgAppSuIndex	1.3.6.1.4.1.12356.101.10.107.2.1.1	Index that uniquely identifies a scan unit in the fgAppSuStatsTable.
fgAppSuFileScanned	1.3.6.1.4.1.12356.101.10.107.2.1.2	Number of files scanned by this scan unit.
fgAppVoIP	1.3.6.1.4.1.12356.101.10.108	
fgAppVoIPStatsTable	1.3.6.1.4.1.12356.101.10.108.1	A table of VoIP related statistics per virtual domain.
fgAppVoIPStatsEntry	1.3.6.1.4.1.12356.101.10.108.1.1	VoIP statistics for a particular virtual domain.
fgAppVoIPConn	1.3.6.1.4.1.12356.101.10.108.1.1.1	The current number of VoIP connections on the virtual domain.
fgAppVoIPCallBlocked	1.3.6.1.4.1.12356.101.10.108.1.1.2	Number of VoIP calls blocked (SIP Invites blocked and SCCP calls blocked) in this virtual domain.
fgAppP2P	1.3.6.1.4.1.12356.101.10.109	
fgAppP2PStatsTable	1.3.6.1.4.1.12356.101.10.109.1	A table of P2P protocol related statistics per virtual domain.
fgAppP2PStatsEntry	1.3.6.1.4.1.12356.101.10.109.1.1	P2P statistics for a particular virtual domain.
fgAppP2PConnBlocked	1.3.6.1.4.1.12356.101.10.109.1.1.1	Number of P2P connections blocked in this virtual domain.
fgAppP2PProtoTable	1.3.6.1.4.1.12356.101.10.109.2	A table of peer to peer statistics per virtual domain per protocol. This table has a dependent expansion relationship with fgVdTable.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgAppP2PProtoEntry	1.3.6.1.4.1.12356.101.10.109.2.1	P2P statistics for a particular virtual domain and protocol.
fgAppP2PProtoEntProto	1.3.6.1.4.1.12356.101.10.109.2.1.1	P2P protocol this row of statistics is for, within the specified virtual domain.
fgAppP2PProtoEntBytes	1.3.6.1.4.1.12356.101.10.109.2.1.2	Number of bytes transferred through this virtual domain on this P2P protocol since last reset.
fgAppP2PProtoEntLastReset	1.3.6.1.4.1.12356.101.10.109.2.1.3	Time elapsed since the corresponding fgAppP2PProtoEntBytes was last reset to 0.
fgAppIM	1.3.6.1.4.1.12356.101.10.110	
fgAppIMStatsTable	1.3.6.1.4.1.12356.101.10.110.1	A table of instant messaging statistics per virtual domain.
fgAppIMStatsEntry	1.3.6.1.4.1.12356.101.10.110.1.1	IM statistics for a particular virtual domain.
fgAppIMMessages	1.3.6.1.4.1.12356.101.10.110.1.1.1	Total number of IM messages processed in this virtual domain.
fgAppIMFileTransferred	1.3.6.1.4.1.12356.101.10.110.1.1.2	Number of files transferred through this virtual domain.
fgAppIMFileTxBlocked	1.3.6.1.4.1.12356.101.10.110.1.1.3	Number of blocked file transfers in this virtual domain.
fgAppIMConnBlocked	1.3.6.1.4.1.12356.101.10.110.1.1.4	Number of connections blocked in this virtual domain.
fgAppProxyFTP	1.3.6.1.4.1.12356.101.10.111	
fgApFTPUpTime	1.3.6.1.4.1.12356.101.10.111.1	Up time of the FTP proxy, in seconds.
fgApFTPMemUsage	1.3.6.1.4.1.12356.101.10.111.2	FTP Proxy memory utilization, as a percentage of the system total.
fgApFTPStatsTable	1.3.6.1.4.1.12356.101.10.111.3	A table of FTP proxy statistics per virtual domain.
fgApFTPStatsEntry	1.3.6.1.4.1.12356.101.10.111.3.1	FTP Proxy statistics for a particular virtual domain.
fgApFTPReqProcessed	1.3.6.1.4.1.12356.101.10.111.3.1.1	Number of requests in this virtual domain processed by the FTP proxy since start up.
fgApFTPConnections	1.3.6.1.4.1.12356.101.10.111.4	FTP proxy current connections.
fgAppExplicitProxy	1.3.6.1.4.1.12356.101.10.112	

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgExplicitProxyInfo	1.3.6.1.4.1.12356.101.10.112.1	
fgExplicitProxyUpTime	1.3.6.1.4.1.12356.101.10.112.1.1	Explicit proxy up-time (in seconds).
fgExplicitProxyMemUsage	1.3.6.1.4.1.12356.101.10.112.1.2	Explicit proxy memory usage (percentage of system total).
fgExplicitProxyRequests	1.3.6.1.4.1.12356.101.10.112.1.3	Explicit proxy total number of requests.
fgExplicitProxyStatsTable	1.3.6.1.4.1.12356.101.10.112.2	A table of explicit proxy statistics per virtual domain.
fgExplicitProxyStatsEntry	1.3.6.1.4.1.12356.101.10.112.2.1	Explicit proxy statistics for a particular virtual domain.
fgExplicitProxyUsers	1.3.6.1.4.1.12356.101.10.112.2.1.1	Number of current users in this virtual domain.
fgExplicitProxySessions	1.3.6.1.4.1.12356.101.10.112.2.1.2	Number of current sessions in this virtual domain.
fgExplicitProxyScanStatsTable	1.3.6.1.4.1.12356.101.10.112.3	A table of explicit proxy scan statistics per virtual domain.
fgExplicitProxyScanStatsEntry	1.3.6.1.4.1.12356.101.10.112.3.1	Explicit proxy scan statistics for a particular virtual domain.
fgExplicitProxyScanStatsDisp	1.3.6.1.4.1.12356.101.10.112.3.1.1	Disposition of an scan result.
fgExplicitProxyVirus	1.3.6.1.4.1.12356.101.10.112.3.1.2	Number of viruses in this virtual domain.
fgExplicitProxyBannedWords	1.3.6.1.4.1.12356.101.10.112.3.1.3	Number of elements containing banned words in this virtual domain.
fgExplicitProxyPolicy	1.3.6.1.4.1.12356.101.10.112.3.1.4	Number of elements violating policy (e.g. filename or file type rules) in this virtual domain.
fgExplicitProxyOversized	1.3.6.1.4.1.12356.101.10.112.3.1.5	Number of oversized elements in this virtual domain.
fgExplicitProxyArchNest	1.3.6.1.4.1.12356.101.10.112.3.1.6	Number of too deeply nested archives in this virtual domain.
fgExplicitProxyArchSize	1.3.6.1.4.1.12356.101.10.112.3.1.7	Number of archives that decompress beyond size limit in this virtual domain.
fgExplicitProxyArchEncrypted	1.3.6.1.4.1.12356.101.10.112.3.1.8	Number of encrypted archives in this virtual domain.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgExplicitProxyArchMultiPart	1.3.6.1.4.1.12356.101.10.112.3.1.9	Number of multipart archives in this virtual domain.
fgExplicitProxyArchUnsupported	1.3.6.1.4.1.12356.101.10.112.3.1.10	Number of archives with unsupported (but known) formats in this virtual domain.
fgExplicitProxyArchBomb	1.3.6.1.4.1.12356.101.10.112.3.1.11	Number of archive bombs in this virtual domain.
fgExplicitProxyArchCorrupt	1.3.6.1.4.1.12356.101.10.112.3.1.12	Number of corrupt archives in this virtual domain.
fgExplicitProxyScriptStatsTable	1.3.6.1.4.1.12356.101.10.112.4	A table of explicit proxy script filtering statistics per virtual domain.
fgExplicitProxyScriptStatsEntry	1.3.6.1.4.1.12356.101.10.112.4.1	Explicit proxy scan statistics for a particular virtual domain.
fgExplicitProxyFilteredApplets	1.3.6.1.4.1.12356.101.10.112.4.1.1	Number of applets filtered from files in this virtual domain.
fgExplicitProxyFilteredActiveX	1.3.6.1.4.1.12356.101.10.112.4.1.2	Number of ActiveX scripts filtered from files in this virtual domain.
fgExplicitProxyFilteredJScript	1.3.6.1.4.1.12356.101.10.112.4.1.3	Number of JScript scripts filtered from files in this virtual domain.
fgExplicitProxyFilteredJS	1.3.6.1.4.1.12356.101.10.112.4.1.4	Number of JavaScript scripts filtered from files in this virtual domain.
fgExplicitProxyFilteredVBS	1.3.6.1.4.1.12356.101.10.112.4.1.5	Number of Visual Basic scripts filtered from files in this virtual domain.
fgExplicitProxyFilteredOthScript	1.3.6.1.4.1.12356.101.10.112.4.1.6	Number of other types of scripts filtered from files in this virtual domain.
fgExplicitProxyFilterStatsTable	1.3.6.1.4.1.12356.101.10.112.5	A table of explicit proxy policy enforcement statistics per virtual domain.
fgExplicitProxyFilterStatsEntry	1.3.6.1.4.1.12356.101.10.112.5.1	Explicit proxy scan statistics for a particular virtual domain.
fgExplicitProxyBlockedDLP	1.3.6.1.4.1.12356.101.10.112.5.1.1	Number of elements blocked due to Data Leak Prevention in this virtual domain.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgExplicitProxyBlockedConType	1.3.6.1.4.1.12356.101.10.112.5.1.2	Number of elements blocked due to Content-Type filtering rules in this virtual domain.
fgExplicitProxyExaminedURLs	1.3.6.1.4.1.12356.101.10.112.5.1.3	Number of URLs inspected against filtering rules in this virtual domain.
fgExplicitProxyAllowedURLs	1.3.6.1.4.1.12356.101.10.112.5.1.4	Number of URLs explicitly allowed due to filtering rules in this virtual domain.
fgExplicitProxyBlockedURLs	1.3.6.1.4.1.12356.101.10.112.5.1.5	Number of URLs explicitly blocked due to filtering rules in this virtual domain.
fgExplicitProxyLoggedURLs	1.3.6.1.4.1.12356.101.10.112.5.1.6	Number of URLs logged due to filtering rules in this virtual domain.
fgExplicitProxyOverriddenURLs	1.3.6.1.4.1.12356.101.10.112.5.1.7	Number of URLs access due to overriding filtering rules in this virtual domain.
fgAppWebCache	1.3.6.1.4.1.12356.101.10.113	
fgWebCacheInfo	1.3.6.1.4.1.12356.101.10.113.1	
fgWebCacheRAMLimit	1.3.6.1.4.1.12356.101.10.113.1.1	RAM available for web cache in bytes.
fgWebCacheRAMUsage	1.3.6.1.4.1.12356.101.10.113.1.2	RAM used by web cache in bytes.
fgWebCacheRAMHits	1.3.6.1.4.1.12356.101.10.113.1.3	Number of cache hits in RAM since last reset.
fgWebCacheRAMMisses	1.3.6.1.4.1.12356.101.10.113.1.4	Number of cache misses in RAM since last reset.
fgWebCacheRequests	1.3.6.1.4.1.12356.101.10.113.1.5	Number of cache requests since last reset.
fgWebCacheBypass	1.3.6.1.4.1.12356.101.10.113.1.6	Number of cache bypasses since last reset.
fgWebCacheUpTime	1.3.6.1.4.1.12356.101.10.113.1.7	Web Cache up-time (in seconds).
fgWebCacheDiskStatsTable	1.3.6.1.4.1.12356.101.10.113.2	A table of the Web Cache disk statistics per disk.
fgWebCacheDiskStatsEntry	1.3.6.1.4.1.12356.101.10.113.2.1	The Web Cache disk statistics for a particular disk.
fgWebCacheDisk	1.3.6.1.4.1.12356.101.10.113.2.1.1	The Web Cache Disk index.
fgWebCacheDiskLimit	1.3.6.1.4.1.12356.101.10.113.2.1.2	The about of storage (in bytes) available for the Web Cache on a particular disk.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgWebCacheDiskUsage	1.3.6.1.4.1.12356.101.10.113.2.1.3	The about of storage (in bytes) in use by the Web Cache on a particular disk.
fgWebCacheDiskHits	1.3.6.1.4.1.12356.101.10.113.2.1.4	The number of cache hits on a particular disk.
fgWebCacheDiskMisses	1.3.6.1.4.1.12356.101.10.113.2.1.5	The number of cache misses on a particular disk.
fgAppWanOpt	1.3.6.1.4.1.12356.101.10.114	
fgWanOptInfo	1.3.6.1.4.1.12356.101.10.114.1	
fgMemCacheLimit	1.3.6.1.4.1.12356.101.10.114.1.1	RAM available for mem cache in bytes.
fgMemCacheUsage	1.3.6.1.4.1.12356.101.10.114.1.2	RAM used by mem cache in bytes.
fgMemCacheHits	1.3.6.1.4.1.12356.101.10.114.1.3	Number of hits in mem cache since last reset.
fgMemCacheMisses	1.3.6.1.4.1.12356.101.10.114.1.4	Number of misses in mem cache since last reset.
fgByteCacheRAMLimit	1.3.6.1.4.1.12356.101.10.114.1.5	RAM available for byte cache in bytes.
fgByteCacheRAMUsage	1.3.6.1.4.1.12356.101.10.114.1.6	RAM used by byte cache in bytes.
fgWanOptUpTime	1.3.6.1.4.1.12356.101.10.114.1.7	Wan Optimization up-time (in seconds).
fgWanOptStatsTable	1.3.6.1.4.1.12356.101.10.114.2	A table of WAN optimization statistics per virtual domain.
fgWanOptStatsEntry	1.3.6.1.4.1.12356.101.10.114.2.1	WAN optimization statistics for a particular virtual domain.
fgWanOptTunnels	1.3.6.1.4.1.12356.101.10.114.2.1.1	Number of current tunnels in this virtual domain.
fgWanOptLANBytesIn	1.3.6.1.4.1.12356.101.10.114.2.1.2	Number of bytes received on LAN in last 5 seconds.
fgWanOptLANBytesOut	1.3.6.1.4.1.12356.101.10.114.2.1.3	Number of bytes sent on LAN in last 5 seconds.
fgWanOptWANBytesIn	1.3.6.1.4.1.12356.101.10.114.2.1.4	Number of bytes received on WAN in last 5 seconds.
fgWanOptWANBytesOut	1.3.6.1.4.1.12356.101.10.114.2.1.5	Number of bytes sent on WAN in last 5 seconds.
fgWanOptHistoryStatsTable	1.3.6.1.4.1.12356.101.10.114.3	A table of the WAN optimization history per protocol.
fgWanOptHistoryStatsEntry	1.3.6.1.4.1.12356.101.10.114.3.1	The WAN optimization history for a particular virtual domain, period, and protocol.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgWanOptHistPeriod	1.3.6.1.4.1.12356.101.10.114.3.1.1	WAN optimization table entry period.
fgWanOptProtocol	1.3.6.1.4.1.12356.101.10.114.3.1.2	Internal WAN optimization table entry protocol.
fgWanOptReductionRate	1.3.6.1.4.1.12356.101.10.114.3.1.3	Reduction rate achieved by WAN optimization.
fgWanOptLanTraffic	1.3.6.1.4.1.12356.101.10.114.3.1.4	Number of bytes transferred via LAN.
fgWanOptWanTraffic	1.3.6.1.4.1.12356.101.10.114.3.1.5	Number of bytes transferred via WAN.
fgWanOptTrafficStatsTable	1.3.6.1.4.1.12356.101.10.114.4	A table of the WAN optimization traffic for a particular virtual domain and protocol.
fgWanOptTrafficStatsEntry	1.3.6.1.4.1.12356.101.10.114.4.1	The WAN optimization history for a particular protocol.
fgWanOptLanInTraffic	1.3.6.1.4.1.12356.101.10.114.4.1.1	Amount of traffic received from the LAN by WAN optimization.
fgWanOptLanOutTraffic	1.3.6.1.4.1.12356.101.10.114.4.1.2	Amount of traffic sent to the LAN by WAN optimization.
fgWanOptWanInTraffic	1.3.6.1.4.1.12356.101.10.114.4.1.3	Amount of traffic received from the WAN by WAN optimization.
fgWanOptWanOutTraffic	1.3.6.1.4.1.12356.101.10.114.4.1.4	Amount of traffic sent to the WAN by WAN optimization.
fgWanOptDiskStatsTable	1.3.6.1.4.1.12356.101.10.114.5	A table of the Web Cache disk statistics per disk.
fgWanOptDiskStatsEntry	1.3.6.1.4.1.12356.101.10.114.5.1	The Web Cache disk statistics for a particular disk.
fgWanOptDisk	1.3.6.1.4.1.12356.101.10.114.5.1.1	The Web Cache Disk index.
fgWanOptDiskLimit	1.3.6.1.4.1.12356.101.10.114.5.1.2	The amount of storage (in bytes) available for the Web Cache on a particular disk.
fgWanOptDiskUsage	1.3.6.1.4.1.12356.101.10.114.5.1.3	The amount of storage (in bytes) in use by the Web Cache on a particular disk.
fgWanOptDiskHits	1.3.6.1.4.1.12356.101.10.114.5.1.4	The number of cache hits on a particular disk.
fgWanOptDiskMisses	1.3.6.1.4.1.12356.101.10.114.5.1.5	The number of cache misses on a particular disk.
Protocol and Session Table		
fgInetProto	1.3.6.1.4.1.12356.101.11	
fgInetProtoInfo	1.3.6.1.4.1.12356.101.11.1	

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fglNetProtoTables	1.3.6.1.4.1.12356.101.11.2	
fglpSessTable	1.3.6.1.4.1.12356.101.11.2.1	
fglpSessEntry	1.3.6.1.4.1.12356.101.11.2.1.1	Information on a specific session, including source and destination.
fglpSessIndex	1.3.6.1.4.1.12356.101.11.2.1.1.1	An index value that uniquely identifies an IP session within the fgIpSessTable.
fglpSessProto	1.3.6.1.4.1.12356.101.11.2.1.1.2	The protocol the session is using (IP, TCP, UDP, etc.).
fglpSessFromAddr	1.3.6.1.4.1.12356.101.11.2.1.1.3	Source IP address (IPv4 only) of the session.
fglpSessFromPort	1.3.6.1.4.1.12356.101.11.2.1.1.4	Source port number (UDP and TCP only) of the session.
fglpSessToAddr	1.3.6.1.4.1.12356.101.11.2.1.1.5	Destination IP address (IPv4 only) of the session.
fglpSessToPort	1.3.6.1.4.1.12356.101.11.2.1.1.6	Destination Port number (UDP and TCP only) of the session.
fglpSessExp	1.3.6.1.4.1.12356.101.11.2.1.1.7	Number of seconds remaining before the session expires (if idle).
fglpSessVdom	1.3.6.1.4.1.12356.101.11.2.1.1.8	Virtual domain the session is part of. This index corresponds to the index used by fgVdTable.
fglpSessStatsTable	1.3.6.1.4.1.12356.101.11.2.2	IP session statistics table.
fglpSessStatsEntry	1.3.6.1.4.1.12356.101.11.2.2.1	IP session statistics on a virtual domain.
fglpSessNumber	1.3.6.1.4.1.12356.101.11.2.2.1.1	Current number of sessions on the virtual domain.
VPN		
fgVPN	1.3.6.1.4.1.12356.101.12	
fgVpnInfo	1.3.6.1.4.1.12356.101.12.1	
fgVpnTables	1.3.6.1.4.1.12356.101.12.2	
fgVpnDialupTable	1.3.6.1.4.1.12356.101.12.2.1	Dial-up VPN peers information.
fgVpnDialupEntry	1.3.6.1.4.1.12356.101.12.2.1.1	Dial-up VPN peer info.
fgVpnDialupIndex	1.3.6.1.4.1.12356.101.12.2.1.1.1	An index value that uniquely identifies an VPN dial-up peer within the fgVpnDialupTable.
fgVpnDialupGateway	1.3.6.1.4.1.12356.101.12.2.1.1.2	Remote gateway IP address of the tunnel.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgVpnDialupLifetime	1.3.6.1.4.1.12356.101.12.2.1.1.3	Tunnel life time (seconds) of the tunnel.
fgVpnDialupTimeout	1.3.6.1.4.1.12356.101.12.2.1.1.4	Time before the next key exchange (seconds) of the tunnel.
fgVpnDialupSrcBegin	1.3.6.1.4.1.12356.101.12.2.1.1.5	Remote subnet address of the tunnel.
fgVpnDialupSrcEnd	1.3.6.1.4.1.12356.101.12.2.1.1.6	Remote subnet mask of the tunnel.
fgVpnDialupDstAddr	1.3.6.1.4.1.12356.101.12.2.1.1.7	Local subnet address of the tunnel.
fgVpnDialupVdom	1.3.6.1.4.1.12356.101.12.2.1.1.8	Virtual domain tunnel is part of. This index corresponds to the index used by fgVdTable.
fgVpnDialupInOctets	1.3.6.1.4.1.12356.101.12.2.1.1.9	Number of bytes received on tunnel since instantiation.
fgVpnDialupOutOctets	1.3.6.1.4.1.12356.101.12.2.1.1.10	Number of bytes sent on tunnel since instantiation.
fgVpnTunTable	1.3.6.1.4.1.12356.101.12.2.2	Table of non-dial-up VPN tunnels.
fgVpnTunEntry	1.3.6.1.4.1.12356.101.12.2.2.1	Tunnel VPN peer info.
fgVpnTunEntIndex	1.3.6.1.4.1.12356.101.12.2.2.1.1	An index value that uniquely identifies a VPN tunnel within the fgVpnTunTable.
fgVpnTunEntPhase1Name	1.3.6.1.4.1.12356.101.12.2.2.1.2	Descriptive name of phase1 configuration for the tunnel.
fgVpnTunEntPhase2Name	1.3.6.1.4.1.12356.101.12.2.2.1.3	Descriptive name of phase2 configuration for the tunnel.
fgVpnTunEntRemGwIpl	1.3.6.1.4.1.12356.101.12.2.2.1.4	IP of remote gateway used by the tunnel.
fgVpnTunEntRemGwyPort	1.3.6.1.4.1.12356.101.12.2.2.1.5	Port of remote gateway used by tunnel, if UDP.
fgVpnTunEntLocGwIpl	1.3.6.1.4.1.12356.101.12.2.2.1.6	IP of local gateway used by the tunnel.
fgVpnTunEntLocGwyPort	1.3.6.1.4.1.12356.101.12.2.2.1.7	Port of local gateway used by tunnel, if UDP.
fgVpnTunEntSelectorSrcBeginIp	1.3.6.1.4.1.12356.101.12.2.2.1.8	Beginning of address range of source selector.
fgVpnTunEntSelectorSrcEndIp	1.3.6.1.4.1.12356.101.12.2.2.1.9	End of address range of source selector.
fgVpnTunEntSelectorSrcPort	1.3.6.1.4.1.12356.101.12.2.2.1.10	Source selector port.
fgVpnTunEntSelectorDstBeginIp	1.3.6.1.4.1.12356.101.12.2.2.1.11	Beginning of address range of destination selector.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgVpnTunEntSelectorDstEndIp	1.3.6.1.4.1.12356.101.12.2.2.1.12	End of address range of destination selector.
fgVpnTunEntSelectorDstPort	1.3.6.1.4.1.12356.101.12.2.2.1.13	Destination selector port.
fgVpnTunEntSelectorProto	1.3.6.1.4.1.12356.101.12.2.2.1.14	Protocol number for selector.
fgVpnTunEntLifeSecs	1.3.6.1.4.1.12356.101.12.2.2.1.15	Lifetime of tunnel in seconds, if time based lifetime used.
fgVpnTunEntLifeBytes	1.3.6.1.4.1.12356.101.12.2.2.1.16	Lifetime of tunnel in bytes, if byte transfer based lifetime used.
fgVpnTunEntTimeout	1.3.6.1.4.1.12356.101.12.2.2.1.17	Timeout of tunnel in seconds.
fgVpnTunEntInOctets	1.3.6.1.4.1.12356.101.12.2.2.1.18	Number of bytes received on tunnel.
fgVpnTunEntOutOctets	1.3.6.1.4.1.12356.101.12.2.2.1.19	Number of bytes sent out on tunnel.
fgVpnTunEntStatus	1.3.6.1.4.1.12356.101.12.2.2.1.20	Current status of tunnel (up or down).
fgVpnTunEntVdom	1.3.6.1.4.1.12356.101.12.2.2.1.21	Virtual domain the tunnel is part of. This index corresponds to the index used by fgVdTable.
fgVpnSslStatsTable	1.3.6.1.4.1.12356.101.12.2.3	SSL VPN statistics table.
fgVpnSslStatsEntry	1.3.6.1.4.1.12356.101.12.2.3.1	SSL VPN statistics for a given virtual domain.
fgVpnSslState	1.3.6.1.4.1.12356.101.12.2.3.1.1	Whether SSL-VPN is enabled on this virtual domain.
fgVpnSslStatsLoginUsers	1.3.6.1.4.1.12356.101.12.2.3.1.2	The current number of users logged in through SSL-VPN tunnels in the virtual domain.
fgVpnSslStatsMaxUsers	1.3.6.1.4.1.12356.101.12.2.3.1.3	The maximum number of total users that can be logged in at any one time on the virtual domain.
fgVpnSslStatsActiveWebSessions	1.3.6.1.4.1.12356.101.12.2.3.1.4	The current number of active SSL web sessions in the virtual domain.
fgVpnSslStatsMaxWebSessions	1.3.6.1.4.1.12356.101.12.2.3.1.5	The maximum number of active SSL web sessions at any one time within the virtual domain.
fgVpnSslStatsActiveTunnels	1.3.6.1.4.1.12356.101.12.2.3.1.6	The current number of active SSL tunnels in the virtual domain.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgVpnSslStatsMaxTunnels	1.3.6.1.4.1.12356.101.12.2.3.1.7	The maximum number of active SSL tunnels at any one time in the virtual domain.
fgVpnSslTunnelTable	1.3.6.1.4.1.12356.101.12.2.4	A list of active SSL VPN tunnel entries.
fgVpnSslTunnelEntry	1.3.6.1.4.1.12356.101.12.2.4.1	An SSL VPN tunnel entry containing connection information and traffic statistics.
fgVpnSslTunnelIndex	1.3.6.1.4.1.12356.101.12.2.4.1.1	An index value that uniquely identifies an active SSL VPN tunnel within the fgVpnSslTunnelTable.
fgVpnSslTunnelVdom	1.3.6.1.4.1.12356.101.12.2.4.1.2	The index of the virtual domain this tunnel belongs to. This index corresponds to the index used by fgVdTable.
fgVpnSslTunnelUserName	1.3.6.1.4.1.12356.101.12.2.4.1.3	The user name used to authenticate the tunnel.
fgVpnSslTunnelSrcIp	1.3.6.1.4.1.12356.101.12.2.4.1.4	The source IP address of this tunnel.
fgVpnSslTunnelIp	1.3.6.1.4.1.12356.101.12.2.4.1.5	The connection IP address of this tunnel.
fgVpnSslTunnelUpTime	1.3.6.1.4.1.12356.101.12.2.4.1.6	The up-time of this tunnel in seconds.
fgVpnSslTunnelBytesIn	1.3.6.1.4.1.12356.101.12.2.4.1.7	The number of incoming bytes of L2 traffic through this tunnel since it was established.
fgVpnSslTunnelBytesOut	1.3.6.1.4.1.12356.101.12.2.4.1.8	The number of outgoing bytes of L2 traffic through this tunnel since it was established.
fgVpnTrapObjects	1.3.6.1.4.1.12356.101.12.3	
fgVpnTrapLocalGateway	1.3.6.1.4.1.12356.101.12.3.2	Local gateway IP address. Used in VPN related traps.
fgVpnTrapRemoteGateway	1.3.6.1.4.1.12356.101.12.3.3	Remote gateway IP address. Used in VPN related traps.
High Availability		
fgHighAvailability	1.3.6.1.4.1.12356.101.13	
fgHaInfo	1.3.6.1.4.1.12356.101.13.1	
fgHaSystemMode	1.3.6.1.4.1.12356.101.13.1.1	High-availability mode (Standalone, A-A or A-P).

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgHaGroupId	1.3.6.1.4.1.12356.101.13.1.2	HA cluster group ID device is configured for.
fgHaPriority	1.3.6.1.4.1.12356.101.13.1.3	HA clustering priority of the device (default = 127).
fgHaOverride	1.3.6.1.4.1.12356.101.13.1.4	Status of a master override flag.
fgHaAutoSync	1.3.6.1.4.1.12356.101.13.1.5	Configuration of an automatic configuration synchronization (enabled or disabled).
fgHaSchedule	1.3.6.1.4.1.12356.101.13.1.6	Load-balancing schedule of cluster (in A-A mode).
fgHaGroupName	1.3.6.1.4.1.12356.101.13.1.7	Ha cluster group name.
fgHaTables	1.3.6.1.4.1.12356.101.13.2	
fgHaStatsTable	1.3.6.1.4.1.12356.101.13.2.1	Some useful statistics for all members of a cluster. This table is also available in standalone mode.
fgHaStatsEntry	1.3.6.1.4.1.12356.101.13.2.1.1	Statistics for a particular HA cluster's unit.
fgHaStatsIndex	1.3.6.1.4.1.12356.101.13.2.1.1.1	An index value that uniquely identifies an unit in the HA Cluster.
fgHaStatsSerial	1.3.6.1.4.1.12356.101.13.2.1.1.2	Serial number of the HA cluster member for this row.
fgHaStatsCpuUsage	1.3.6.1.4.1.12356.101.13.2.1.1.3	CPU usage of the specified cluster member (percentage).
fgHaStatsMemUsage	1.3.6.1.4.1.12356.101.13.2.1.1.4	Memory usage of the specified cluster member (percentage).
fgHaStatsNetUsage	1.3.6.1.4.1.12356.101.13.2.1.1.5	Network bandwidth usage of specified cluster member (kbps).
fgHaStatsSesCount	1.3.6.1.4.1.12356.101.13.2.1.1.6	Current session count of specified cluster member.
fgHaStatsPktCount	1.3.6.1.4.1.12356.101.13.2.1.1.7	Number of packets processed by the specified cluster member since start up.
fgHaStatsByteCount	1.3.6.1.4.1.12356.101.13.2.1.1.8	Number of bytes processed by the specified cluster member since start up.
fgHaStatsIdsCount	1.3.6.1.4.1.12356.101.13.2.1.1.9	Number of IDS/IPS events triggered on the specified cluster member since start up.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgHaStatsAvCount	1.3.6.1.4.1.12356.101.13.2.1.1.10	Number of anti-virus events triggered on the specified cluster member since start up.
fgHaStatsHostname	1.3.6.1.4.1.12356.101.13.2.1.1.11	Host name of the specified cluster member.
fgHaTrapObjects	1.3.6.1.4.1.12356.101.13.3	
fgHaTrapMemberSerial	1.3.6.1.4.1.12356.101.13.3.1	Serial number of an HA cluster member. Used to identify the origin of a trap when a cluster is configured.
fgFmTrapGroup	1.3.6.1.4.1.12356.101.100.1	Traps are intended for use in conjunction with a FortiManager.
fgFmTrapObjectGroup	1.3.6.1.4.1.12356.101.100.2	These objects support the traps in the fgFmTrapGroup.
fgAdminObjectGroup	1.3.6.1.4.1.12356.101.100.3	Objects pertaining to administration of the device.
fgSystemObjectGroup	1.3.6.1.4.1.12356.101.100.4	Objects pertaining to the system status of the device.
fgSoftwareObjectGroup	1.3.6.1.4.1.12356.101.100.5	Objects pertaining to software running on the device.
fgHwSensorsObjectGroup	1.3.6.1.4.1.12356.101.100.6	Object pertaining to hardware sensors on the device.
fgHighAvailabilityObjectGroup	1.3.6.1.4.1.12356.101.100.7	Objects pertaining to High Availability clustering of FortiGate devices.
fgVpnObjectGroup	1.3.6.1.4.1.12356.101.100.8	Objects pertaining to Virtual Private Networking on FortiGate devices.
fgFirewallObjectGroup	1.3.6.1.4.1.12356.101.100.9	Objects pertaining to Firewall functionality on FortiGate devices.
fgAppServicesObjectGroup	1.3.6.1.4.1.12356.101.100.10	Objects pertaining to application proxy and filtering services on FortiGate devices.
fgAntivirusObjectGroup	1.3.6.1.4.1.12356.101.100.11	Objects pertaining to Antivirus services on FortiGate devices.
fgIntrusionPrevObjectGroup	1.3.6.1.4.1.12356.101.100.12	Objects pertaining to Intrusion Detection and Prevention services on FortiGate devices.

Table 26: OIDs for the Fortinet-FortiGate-MIB

MIB Field	OID	Description
fgWebFilterObjectGroup	1.3.6.1.4.1.12356.101.100.13	Objects pertaining to FortiGate and FortiGuard based Web Filtering services on FortiGate devices.
fgVirtualDomainObjectGroup	1.3.6.1.4.1.12356.101.100.14	Objects pertaining to Virtual Firewall Domain services on FortiGate devices.
fgAdministrationObjectGroup	1.3.6.1.4.1.12356.101.100.15	Objects pertaining to the administration of FortiGate device.
fgIntfObjectGroup	1.3.6.1.4.1.12356.101.100.16	Objects pertaining to the interface table of FortiGate device.
fgProcessorsObjectGroup	1.3.6.1.4.1.12356.101.100.17	Objects pertaining to the processors table of FortiGate device.
fgNotificationGropu	1.3.6.1.4.1.12356.101.100.18	Notifications that can be generated from a FortiGate device.
fgObsoleteNotificationsGroup	1.3.6.1.4.1.12356.101.100.19	Notifications that have been deprecated, but may still be generated by older models.
fgMIBCompliance	1.3.6.1.4.1.12356.101.100.100	Model and feature specific.



Multicast forwarding

Multicasting (also called IP multicasting) consists of using a single multicast source to send data to many receivers. Multicasting can be used to send data to many receivers simultaneously while conserving bandwidth and reducing network traffic. Multicasting can be used for one-way delivery of media streams to multiple receivers and for one-way data transmission for news feeds, financial information, and so on. Also RIPv2 uses multicasting to share routing table information.

A FortiGate unit can operate as a Protocol Independent Multicast (PIM) version 2 router. FortiGate units support PIM sparse mode (RFC 4601) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate unit interface is connected. Multicast routing is not supported in transparent mode (TP mode).



To support PIM communications, the sending/receiving applications and all connecting PIM routers in between must be enabled with PIM version 2. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations. To enable source-to-destination packet delivery, either sparse mode or dense mode must be enabled on the PIM-router interfaces. Sparse mode routers cannot send multicast messages to dense mode routers. In addition, if a FortiGate unit is located between a source and a PIM router, two PIM routers, or is connected directly to a receiver, you must create a security policy manually to pass encapsulated (multicast) packets or decapsulated data (IP traffic) between the source and destination.

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Boot Strap Router (BSR), and if sparse mode is enabled, a number of Rendezvous Points (RPs) and Designated Routers (DRs). When PIM is enabled on a FortiGate unit, the FortiGate unit can perform any of these functions at any time as configured.

Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast group(s) that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.



When a FortiGate unit interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP's distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that is associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly connected network segments. The PIM routers exchange state refresh messages to update their distribution trees. FortiGate units store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that do not have downstream receivers—PIM routers that do not manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

FortiGate units operating in NAT mode can also be configured as multicast routers. You can configure a FortiGate unit to be a Protocol Independent Multicast (PIM) router operating in Sparse Mode (SM) or Dense Mode (DM).

Multicast IP addresses

Multicast uses the Class D address space. The 224.0.0.0 to 239.255.255.255 IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets. Table 27 lists reserved multicast address ranges and describes what they are reserved for:

Table 27: Reserved Multicast address ranges

Reserved Address Range	Use	Notes
224.0.0.0 to 224.0.0.255	Used for network protocols on local networks. For more information, see RFC 1700.	In this range, packets are not forwarded by the router but remain on the local network. They have a Time to Live (TTL) of 1. These addresses are used for communicating routing information.
224.0.1.0 to 238.255.255.255	Global addresses used for multicasting data between organizations and across the Internet. For more information, see RFC 1700.	Some of these addresses are reserved, for example, 224.0.1.1 is used for Network Time Protocol (NTP).
239.0.0.0 to 239.255.255.255	Limited scope addresses used for local groups and organizations. For more information, see RFC 2365.	Routers are configured with filters to prevent multicasts to these addresses from leaving the local system.

PIM Support

A FortiGate unit can be configured to support PIM using the `config router multicast` CLI command. When PIM is enabled, the FortiGate unit allocates memory to manage mapping information. The FortiGate unit communicates with neighboring PIM routers to acquire mapping information and if required, processes the multicast traffic associated with specific multicast groups.



The end-user multicast client-server applications must be installed and configured to initiate Internet connections and handle broadband content such as audio/video information.

Client applications send multicast data by registering IP traffic with a PIM-enabled router. An end-user could type in a class D multicast group address, an alias for the multicast group address, or a call-conference number to initiate the session.

Rather than sending multiple copies of generated IP traffic to more than one specific IP destination address, PIM-enabled routers encapsulate the data and use the one multicast group address to forward multicast packets to multiple destinations. Because one destination address is used, a single stream of data can be sent. Client applications receive multicast data by requesting that the traffic destined for a certain multicast group address be delivered to them — end-users may use phone books, a menu of ongoing or future sessions, or some other method through a user interface to select the address of interest.

A class D address in the 224.0.0.0 to 239.255.255.255 range may be used as a multicast group address, subject to the rules assigned by the Internet Assigned Numbers Authority (IANA). All class D addresses must be assigned in advance. Because there is no way to determine in advance if a certain multicast group address is in use, collisions may occur (to resolve this problem, end-users may switch to a different multicast address).

To configure a PIM domain

- 1 If you will be using sparse mode, determine appropriate paths for multicast packets.
- 2 Make a note of the interfaces that will be PIM-enabled. These interfaces may run a unicast routing protocol.
- 3 If you will be using sparse mode and want multicast packets to be handled by specific (static) RPs, record the IP addresses of the PIM-enabled interfaces on those RPs.
- 4 Enable PIM version 2 on all participating routers between the source and receivers. On FortiGate units, use the `config router multicast` command to set global operating parameters.
- 5 Configure the PIM routers that have good connections throughout the PIM domain to be candidate BSRs.
- 6 If sparse mode is enabled, configure one or more of the PIM routers to be candidate RPs.
- 7 If required, adjust the default settings of PIM-enabled interface(s).

Multicast forwarding and FortiGate units

In both transparent mode and NAT mode you can configure FortiGate units to forward multicast traffic.

For a FortiGate unit to forward multicast traffic you must add FortiGate multicast security policies. Basic multicast security policies accept any multicast packets at one FortiGate interface and forward the packets out another FortiGate interface. You can also use multicast security policies to be selective about the multicast traffic that is accepted based on source and destination address, and to perform NAT on multicast packets.

In the example shown in [Figure 48](#), a multicast source on the Marketing network with IP address 192.168.5.18 sends multicast packets to the members of network 239.168.4.0. At the FortiGate unit, the source IP address for multicast packets originating from workstation 192.168.5.18 is translated to 192.168.18.10. In this example, the FortiGate unit is not acting as a multicast router.

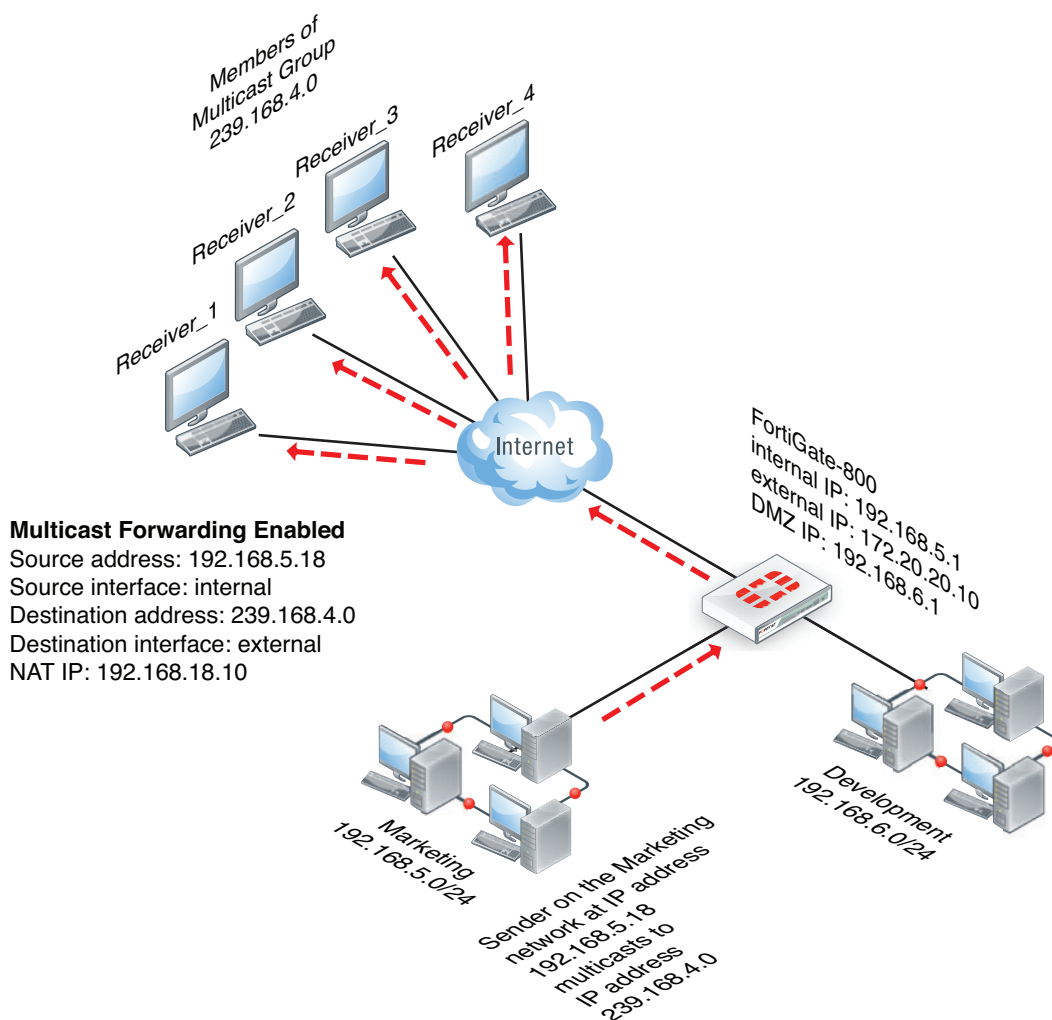
Multicast forwarding and RIPv2

RIPv2 uses multicast to share routing table information. If your FortiGate unit is installed on a network that includes RIPv2 routers, you must configure the FortiGate unit to forward multicast packets so that RIPv2 devices can share routing data through the FortiGate unit. No special FortiGate configuration is required to share RIPv2 data, you can simply use the information in the following sections to configure the FortiGate unit to forward multicast packets.



RIPv1 uses broadcasting to share routing table information. To allow RIPv1 packets through a FortiGate unit you can add standard security policies. Security policies to accept RIPv1 packets can use the ANY predefined firewall service or the RIP predefined firewall service.

Figure 48: Example multicast network including a FortiGate unit that forwards multicast packets



Configuring FortiGate multicast forwarding

You configure FortiGate multicast forwarding from the Command Line Interface (CLI). Two steps are required:

- [Adding multicast security policies](#)
- [Enabling multicast forwarding](#)

This second step is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

Adding multicast security policies

You need to add security policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. You add multicast security policies from the CLI using the `config firewall multicast-policy` command. As with unicast security policies, you specify the source and destination interfaces and optionally the allowed address ranges for the source and destination addresses of the packets.

You can also use multicast security policies to configure source NAT and destination NAT for multicast packets. For full details on the `config firewall multicast-policy` command, see the [FortiGate CLI Reference](#).

Keep the following in mind when configuring multicast security policies:

- The matched forwarded (outgoing) IP multicast source IP address is changed to the configured IP address.
- Source and Destination interfaces are optional. If left blank, then the multicast will be forwarded to ALL interfaces.
- Source and Destination addresses are optional. If left un set, then it will mean ALL addresses.
- The `nat` keyword is optional. Use it when source address translation is needed.

Enabling multicast forwarding

Multicast forwarding is disabled by default. In NAT mode you must use the `multicast-forward` keyword of the `system settings` CLI command to enable multicast forwarding. When `multicast-forward` is enabled, the FortiGate unit forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header will be reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add security policies to actually allow multicast packets through the FortiGate. In our example, the security policy allows multicast packets received by the internal interface to exit to the external interface.



Enabling multicast forwarding is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

Enter the following CLI command to enable multicast forwarding:

```
config system settings
    set multicast-forward enable
end
```

If multicast forwarding is disabled and the FortiGate unit drops packets that have multicast source or destination addresses.

You can also use the `multicast-ttl-notchange` keyword of the `system settings` command so that the FortiGate unit does not increase the TTL value for forwarded multicast packets. You should use this option only if packets are expiring before reaching the multicast router.

```
config system settings
    set multicast-ttl-notchange enable
end
```

In transparent mode, the FortiGate unit does not forward frames with multicast destination addresses. Multicast traffic such as the one used by routing protocols or streaming media may need to traverse the FortiGate unit, and should not be interfere with the communication. To avoid any issues during transmission, you can set up multicast security policies. These types of security policies can only be enabled using the CLI.



The CLI parameter `multicast-skip-policy` must be disabled when using multicast security policies. To disable enter the command

```
config system settings
    set multicast-skip-policy disable
end
```

In this simple example, no check is performed on the source or destination interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces on the forwarding domain, except the incoming interface.

To enable the multicast policy

```
config firewall multicast-policy
    edit 1
        set action accept
    end
```

In this example, the multicast policy only applies to the source port of WAN1 and the destination port of Internal.

To enable the restrictive multicast policy

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set dstintf internal
        set action accept
    end
```

In this example, packets are allowed to flow from WAN1 to Internal, and sourced by the address 172.20.120.129.

To enable the restrictive multicast policy

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set srcaddr 172.20.120.129 255.255.255.255
        set dstintf internal
        set action accept
    end
```

This example shows how to configure the multicast security policy required for the configuration shown in [Figure 48 on page 475](#). This policy accepts multicast packets that are sent from a PC with IP address 192.168.5.18 to destination address range 239.168.4.0. The policy allows the multicast packets to enter the internal interface and then exit the external interface. When the packets leave the external interface their source address is translated to 192.168.18.10

```
config firewall multicast-policy
    edit 5
        set srcaddr 192.168.5.18 255.255.255.255
```

```
set srcintf internal
set destaddr 239.168.4.0 255.255.255.0
set dstintf external
set nat 192.168.18.10
end
```

This example shows how to configure a multicast security policy so that the FortiGate unit forwards multicast packets from a multicast Server with an IP 10.10.10.10 is broadcasting to address 225.1.1.1. This Server is on the network connected to the FortiGate DMZ interface.

```
config firewall multicast-policy
edit 1
set srcintf DMZ
set srcaddr 10.10.10.10 255.255.255.255
set dstintf Internal
set dstaddr 225.1.1.1 255.255.255.255
set action accept
edit 2
set action deny
end
```

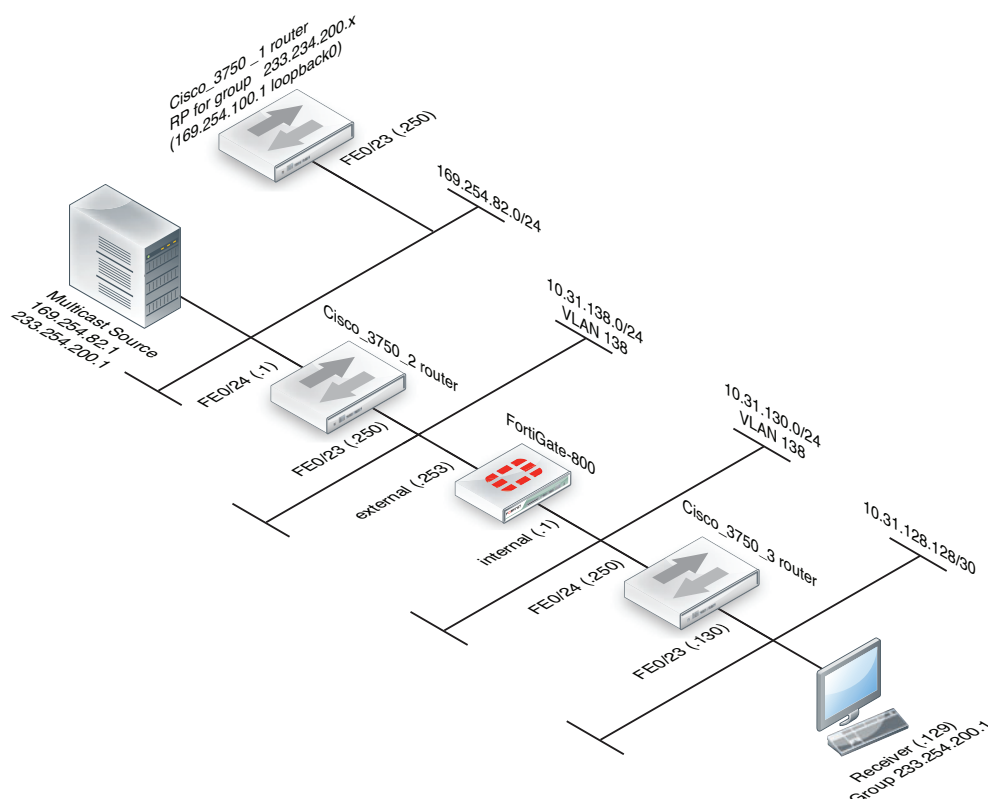
Multicast routing examples

This section contains the following multicast routing configuration examples and information:

- [Example FortiGate PIM-SM configuration using a static RP](#)
- [FortiGate PIM-SM debugging examples](#)
- [Example multicast destination NAT \(DNAT\) configuration](#)
- [Example PIM configuration that uses BSR to find the RP](#)

Example FortiGate PIM-SM configuration using a static RP

The example Protocol Independent Multicast Sparse Mode (PIM-SM) configuration shown in [Figure 49](#) has been tested for multicast interoperability using PIM-SM between Cisco 3750 switches running 12.2 and a FortiGate-800 running FortiOS v3.0 MR5 patch 1. In this configuration, the receiver receives the multicast stream when it joins the group 233.254.200.1.

Figure 49: Example FortiGate PIM-SM topology

The configuration uses a statically configured rendezvous point (RP) which resides on the Cisco_3750_1. Using a bootstrap router (BSR) was not tested in this example. See [“Example PIM configuration that uses BSR to find the RP” on page 491](#) for an example that uses a BSR.

Configuration steps

The following procedures show how to configure the multicast configuration settings for the devices in the example configuration.

- [Cisco_3750_1 router configuration](#)
- [Cisco_3750_2 router configuration](#)
- [To configure the FortiGate-800 unit](#)
- [Cisco_3750_3 router configuration](#)

Cisco_3750_1 router configuration

```
version 12.2
!
hostname Cisco-3750-1
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
```

```

spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface Loopback0
 ip address 169.254.100.1 255.255.255.255
!
interface FastEthernet1/0/23
 switchport access vlan 182
 switchport mode access
!
interface FastEthernet1/0/24
 switchport access vlan 172
 switchport mode access
!
interface Vlan172
 ip address 10.31.138.1 255.255.255.0
 ip pim sparse-mode
 ip igmp query-interval 125
 ip mroute-cache distributed
!
interface Vlan182
 ip address 169.254.82.250 255.255.255.0
 ip pim sparse-mode
 ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 169.254.82.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
ip access-list standard Source-RP
 permit 233.254.200.0 0.0.0.255

```

Cisco_3750_2 router configuration

```

version 12.2
!
hostname Cisco-3750-2
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
 switchport access vlan 138
 switchport mode access
!

```

```

interface FastEthernet1/0/24
    switchport access vlan 182
    switchport mode access
!
interface Vlan138
    ip address 10.31.138.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
interface Vlan182
    ip address 169.254.82.1 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.138.253
ip route 169.254.100.1 255.255.255.255 169.254.82.250
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255

```

To configure the FortiGate-800 unit

1 Configure the internal and external interfaces.

```

config system interface
    edit internal
        set vdom root
        set ip 10.31.130.1 255.255.255.0
        set allowaccess ping https
        set type physical
    next
    edit external
        set vdom root
        set ip 10.31.138.253 255.255.255.0
        set allowaccess ping
        set type physical
    end
end

```

2 Add a firewall address for the RP.

```

config firewall address
    edit RP
        set subnet 169.254.100.1/32
    end

```

3 Add standard security policies to allow traffic to reach the RP.

```

config firewall policy
    edit 1
        set srcintf internal
        set dstintf external
        set srcaddr all
        set dstaddr RP
    end

```

```
        set action accept
        set schedule always
        set service ANY
    next
    edit 2
        set srcintf external
        set dstintf internal
        set srcaddr RP
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    end
```

4 Add the multicast security policy.

```
config firewall multicast-policy
    edit 1
        set dstaddr 233.254.200.0 255.255.255.0
        set dstintf internal
        set srcaddr 169.254.82.0 255.255.255.0
        set srcintf external
    end
```

5 Add an access list.

```
config router access-list
    edit Source-RP
        config rule
            edit 1
                set prefix 233.254.200.0 255.255.255.0
                set exact-match disable
            next
        end
```

6 Add some static routes.

```
config router static
    edit 1
        set device internal
        set gateway 10.31.130.250
    next
    edit 2
        set device external
        set dst 169.254.0.0 255.255.0.0
        set gateway 10.31.138.250
    next
```

7 Configure multicast routing.

```
config router multicast
    config interface
        edit internal
            set pim-mode sparse-mode
            config igmp
                set version 2
            end
        next
        edit external
            set pim-mode sparse-mode
```



```
        config igmp
        set version 2
    end
    next
end
set multicast-routing enable
config pim-sm-global
config rp-address
edit 1
    set ip-address 169.254.100.1
    set group Source-RP
next
```

Cisco_3750_3 router configuration

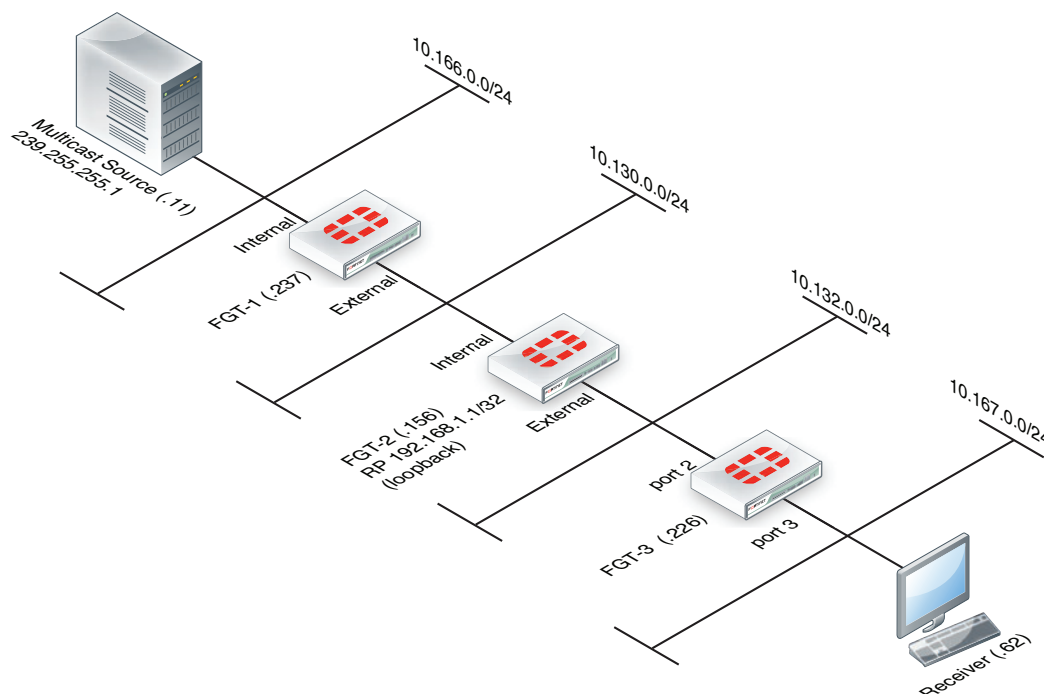
```
version 12.2
!
hostname Cisco-3750-3
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
    switchport access vlan 128
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 130
    switchport mode access
!
interface Vlan128
    ip address 10.31.128.130 255.255.255.252
    ip pim sparse-mode
    ip mroute-cache distributed
!
interface Vlan130
    ip address 10.31.130.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.130.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
```

```
permit 233.254.200.0 0.0.0.255
```

FortiGate PIM-SM debugging examples

Using the example topology shown in [Figure 50](#) you can trace the multicast streams and states within the three FortiGate units (FGT-1, FGT-2, and FGT-3) using the debug commands described in this section. The command output in this section is taken from FortiGate unit when the multicast stream is flowing correctly from source to receiver.

Figure 50: PIM-SM debugging topology



Checking that the receiver has joined the required group

From the last hop router, FGT-3, you can use the following command to check that the receiver has correctly joined the required group.

```
FGT-3 # get router info multicast igmp groups
IGMP Connected Group Membership
Group Address      Interface      Uptime    Expires    Last
Reporter
239.255.255.1     port3         00:31:15  00:04:02
10.167.0.62
```

Only 1 receiver is displayed for a particular group, this is the device that responded to the IGMP query request from the FGT-3. If a receiver is active the expire time should drop to approximately 2 minutes before being refreshed.

Checking the PIM-SM neighbors

Next the PIM-SM neighbors should be checked. A PIM router becomes a neighbor when the PIM router receives a PIM hello. Use the following command to display the PIM-SM neighbors of FGT-3.

```
FGT-3 # get router info multicast pim sparse-mode neighbour
Neighbor          Interface      Uptime/Expires    Ver    DR
```

```
Address Priority/Mode
10.132.0.156      port2      01:57:12/00:01:33 v2      1 /
```

Checking that the PIM router can reach the RP

The rendezvous point (RP) must be reachable for the PIM router (FGT-3) to be able to send the *,G join to request the stream. This can be checked for FGT-3 using the following command:

```
FGT-3 # get router info multicast pim sparse-mode rp-mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
RP: 192.168.1.1
Uptime: 07:23:00
```

Viewing the multicast routing table (FGT-3)

The FGT-3 unicast routing table can be used to determine the path taken to reach the RP at 192.168.1.1. You can then check the stream state entries using the following commands:

```
FGT-3 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

(*,*,RP) Entries	This state may be reached by general joins for all groups served by a specified RP.
(*,G) Entries	State that maintains the RP tree for a given group.
(S,G) Entries	State that maintains a source-specific tree for source S and group G.
(S,G,rpt) Entries	State that maintains source-specific information about source S on the RP tree for G. For example, if a source is being received on the source-specific tree, it will normally have been pruned off the RP tree.
FCR	The FCR state entries are for tracking the sources in the <*, G> when <S, G> is not available for any reason, the stream would typically be flowing when this state exists.

Breaking down each entry in detail:

```
(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: JOINED
Local:
    port3
Joined:
Asserted:
FCR:
```

The RP will always be listed in a *, G entry, the RPF neighbor and interface index will also be shown. In this topology these are the same in all downstream PIM routers. The state is active so the upstream state is joined.

In this case FGT-3 is the last hop router so the IGMP join is received locally on port3. There is no PIM outgoing interface listed for this entry as it is used for the upstream PIM join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.132.0.156
RPF idx: port2
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
Asserted:
Outgoing:
port3
```

This is the entry for the SPT, no RP IS listed. The S, G stream will be forwarded out of the stated outgoing interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: NOT PRUNED
Local:
Pruned:
Outgoing:
```

The above S, G, RPT state is created for all streams that have both a S, G and a *, G entry on the router. This is not pruned in this case because of the topology, the RP and source are reachable over the same interface.

Although not seen in this scenario, assert states may be seen when multiple PIM routers exist on the same LAN which can lead to more than one upstream router having a valid forwarding state. Assert messages are used to elect a single forwarder from the upstream devices.

Viewing the PIM next-hop table

The PIM next-hop table is also very useful for checking the various states, it can be used to quickly identify the states of multiple multicast streams

```
FGT-3 # get router info multicast pim sparse-mode next-hop
Flags: N = New, R = RP, S = Source, U = Unreachable
Destination      Type  Nexthop  Nexthop      Nexthop  Metric Pref
Refcnt
              Num      Addr      Ifindex
-----
10.166.0.11      ..S.  1        10.132.0.156  9 21     110    3
192.168.1.1      .R..  1        10.132.0.156  9 111    110    2
```

Viewing the PIM multicast forwarding table

Also you can check the multicast forwarding table showing the ingress and egress ports of the multicast stream.

```

FGT-3 # get router info multicast table

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL threshold)

(10.166.0.11, 239.255.255.1), uptime 04:02:55, stat expires
00:02:25
Owner PIM-SM, Flags: TF
  Incoming interface: port2
  Outgoing interface list:
    port3 (TTL threshold 1)

```

Viewing the kernel forwarding table

Also the kernel forwarding table can be verified, however this should give similar information to the above command:

```

FGT-3 # diag ip multicast mroute
grp=239.255.255.1 src=10.166.0.11 intf=9 flags=(0x10000000)[ ]
status=resolved
  last_assert=2615136 bytes=1192116 pkt=14538 wrong_if=0
num_ifs=1
  index(ttl)=[6(1),]

```

Viewing the multicast routing table (FGT-2)

If you check the output on FGT-2 there are some small differences:

```

FGT-2 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
  Local:
  Joined:
    external
  Asserted:
FCR:

```

The *,G entry now has a joined interface rather than local because it has received a PIM join from FGT-3 rather than a local IGMP join.

```

(10.166.0.11, 239.255.255.1)
RPF nbr: 10.130.0.237
RPF idx: internal
SPT bit: 1
Upstream State: JOINED

```

```

Local:
Joined:
    external
Asserted:
Outgoing:
    external

```

The *S, G* entry shows that we have received a join on the external interface and the stream is being forwarded out of this interface.

```

(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: PRUNED
Local:
Pruned:
Outgoing:
    External

```

The *S, G, RPT* is different from FGT-3 because FGT-2 is the RP, it has pruned back the SPT for the RP to the first hop router.

Viewing the multicast routing table (FGT-1)

FGT-1 again has some differences with regard to the PIM-SM states, there is no **, G* entry because it is not in the path of a receiver and the RP.

```

FGT-1_master # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

```

Below the *S, G* is the SPT termination because this FortiGate unit is the first hop router, the RPF neighbor always shows as 0.0.0.0 because the source is local to this device. Both the joined and outgoing fields show as external because the PIM join and the stream is egressing on this interface.

```

(10.166.0.11, 239.255.255.1)
RPF nbr: 0.0.0.0
RPF idx: None
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
    external
Asserted:
Outgoing:
    external

```

The stream has been pruned back from the RP because the end-to-end SPT is flowing, there is no requirement for the stream to be sent to the RP in this case.

```

(10.166.0.11, 239.255.255.1, rpt)
RP: 0.0.0.0
RPF nbr: 10.130.0.156

```

```

RPF idx: external
Upstream State: RPT NOT JOINED
Local:
Pruned:
Outgoing:

```

Example multicast destination NAT (DNAT) configuration

The example topology shown in [Figure 51](#) and described below shows how to configure destination NAT (DNAT) for two multicast streams. Both of these streams originate from the same source IP address, which is 10.166.0.11. The example configuration keeps the streams separate by creating 2 multicast NAT policies.

In this example the FortiGate units in [Figure 51](#) have the following roles:

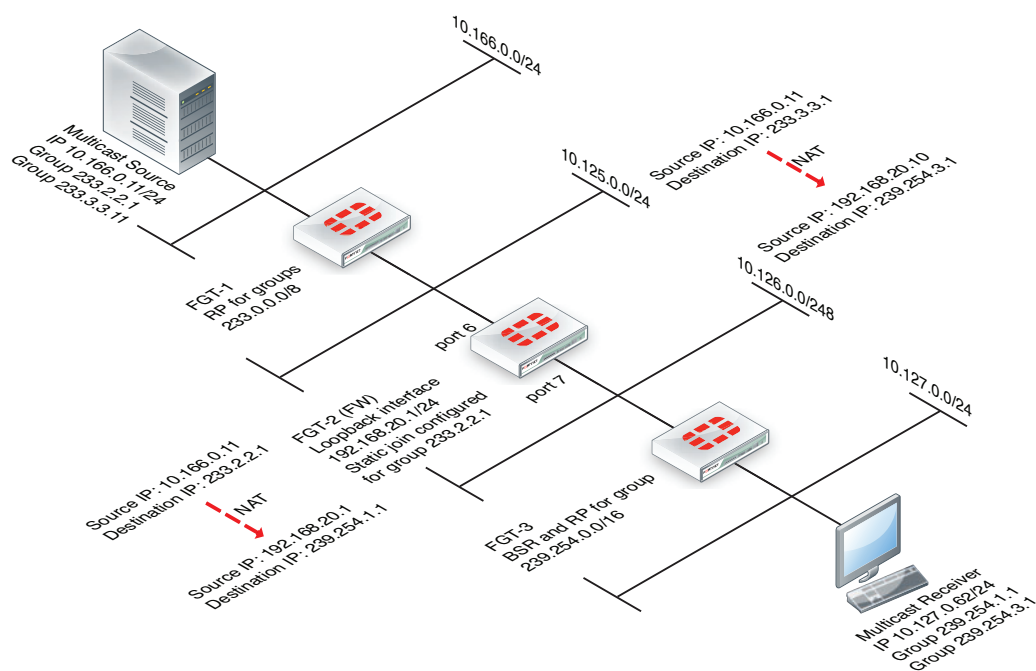
- FGT-1 is the RP for dirty networks, 233.0.0.0/8.
- FGT-2 performs all firewall and DNAT translations.
- FGT-3 is the RP for the clean networks, 239.254.0.0/16.
- FGT-1 and FGT-3 are functioning as PM enabled routers and could be replaced can be any PIM enabled router.

This example only describes the configuration of FGT-2.

FGT-2 performs NAT so that the receivers connected to FGT-3 receive the following translated multicast streams.

- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.2.2.1; FGT-3 translates the source and destination IPs to 192.168.20.1 and 239.254.1.1
- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.3.3.1; FGT-3 translates the source and destination IPs to 192.168.20.10 and 239.254.3.1

Figure 51: Example multicast DNAT topology



To configure FGT-2 for DNAT multicast

- 1 Add a loopback interface. In the example, the loopback interface is named loopback.

```
config system interface
  edit loopback
    set vdom root
    set ip 192.168.20.1 255.255.255.0
    set type loopback
  next
end
```

- 2 Add PIM and add a unicast routing protocol to the loopback interface as if it was a normal routed interface. Also add static joins to the loopback interface for any groups to be translated.

```
config router multicast
  config interface
    edit loopback
      set pim-mode sparse-mode
      config join-group
        edit 233.2.2.1
        next
        edit 233.3.3.1
        next
      end
    next
  end
```

- 3 In this example, to add firewall multicast policies, different source IP addresses are required so you must first add an IP pool:

```
config firewall ippool
  edit Multicast_source
    set endip 192.168.20.20
    set interface port6
    set startip 192.168.20.10
  next
end
```

- 4 Add the translation security policies.

Policy 2, which is the source NAT policy, uses the actual IP address of port6. Policy 1, the DNAT policy, uses an address from the IP pool.

```
config firewall multicast-policy
  edit 1
    set dnat 239.254.3.1
    set dstaddr 233.3.3.1 255.255.255.255
    set dstintf loopback
    set nat 192.168.20.10
    set srcaddr 10.166.0.11 255.255.255.255
    set srcintf port6
  next
```



```
edit 2
  set dnat 239.254.1.1
  set dstaddr 233.2.2.1 255.255.255.255
  set dstintf loopback
  set nat 192.168.20.1
  set srcaddr 10.166.0.11 255.255.255.255
  set srcintf port6
next
```

- 5 Add a firewall multicast policy to forward the stream from the loopback interface to the physical outbound interface.

This example is an any/any policy that makes sure traffic accepted by the other multicast policies can exit the FortiGate unit.

```
config firewall multicast-policy
  edit 3
    set dstintf port7
    set srcintf loopback
  next
```

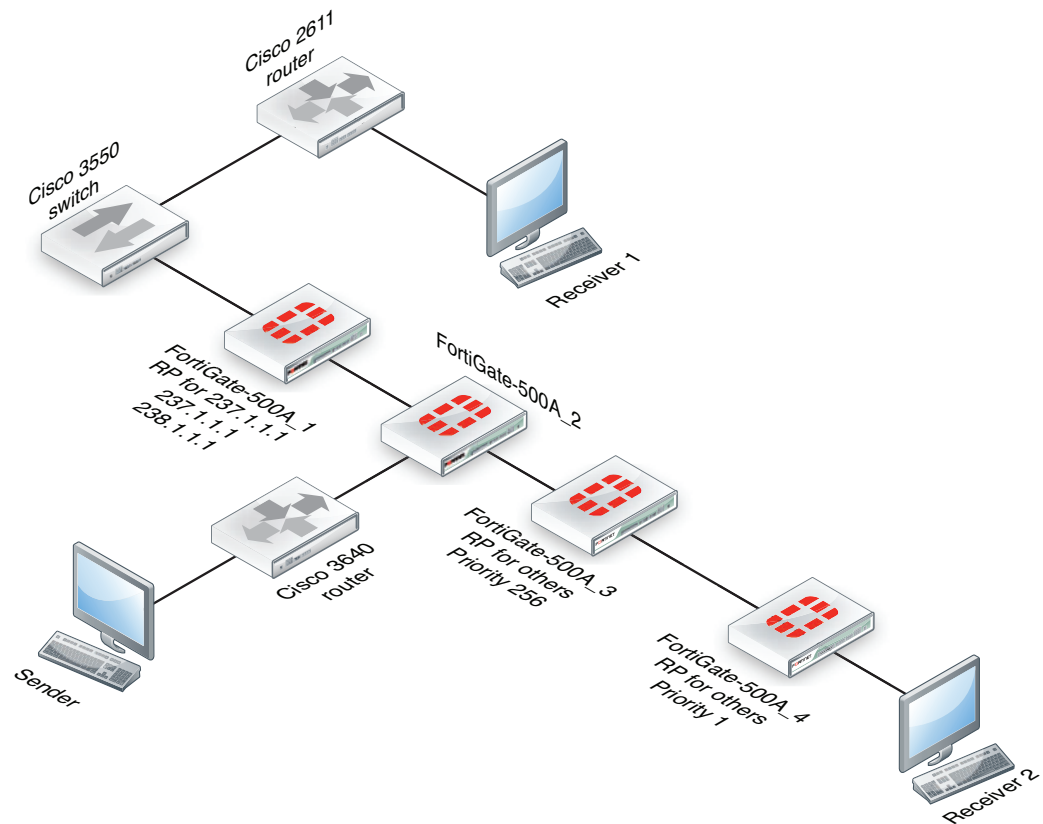
Example PIM configuration that uses BSR to find the RP

This example shows how to configure a multicast routing network for a network consisting of four FortiGate-500A units (FortiGate-500A_1 to FortiGate-500A_4, see [Figure 52](#)). A multicast sender is connected to FortiGate-500A_2. FortiGate-500A_2 forwards multicast packets in two directions to reach Receiver 1 and Receiver 2.

The configuration uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface `lo0` must join the 236.1.1.1 group (source).

This example describes:

- [Commands used in this example](#)
- [Configuration steps](#)
- [Example debug commands](#)

Figure 52: PIM network topology using BSR to find the RP

Commands used in this example

This example uses CLI commands for the following configuration settings:

- Adding a loopback interface (lo0)
- Defining the multicast routing
- Adding the NAT multicast policy

Adding a loopback interface (lo0)

Where required, the following command is used to define a loopback interface named lo0.

```

config system interface
  edit lo0
    set vdom root
    set ip 1.4.50.4 255.255.255.255
    set allowaccess ping https ssh snmp http telnet
    set type loopback
  next
end

```

Defining the multicast routing

In this example, the following command syntax is used to define multicast routing. The example uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface lo0 must join the 236.1.1.1 group (source).

```

config router multicast
  config interface
    edit port6
      set pim-mode sparse-mode
    next
    edit port1
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      config join-group
        edit 236.1.1.1
        next
      end
      set rp-candidate-priority 1
    next
  end
set multicast-routing enable
config pim-sm-global
  set bsr-allow-quick-refresh enable
  set bsr-candidate enable
  set bsr-interface lo0
  set bsr-priority 200
end
end

```

Adding the NAT multicast policy

In this example, the incoming multicast policy does the address translation. The NAT address should be the same as the IP address of the of loopback interface. The DNAT address is the translated address, which should be a new group.

```

config firewall multicast-policy
  edit 1
    set dstintf port6
    set srcintf lo0
  next
  edit 2
    set dnat 238.1.1.1
    set dstintf lo0
    set nat 1.4.50.4
    set srcintf port1
  next

```

Configuration steps

In this sample, FortiGate-500A_1 is the RP for the group 228.1.1.1, 237.1.1.1, 238.1.1.1, and FortiGate-500A_4 is the RP for the other group which has a priority of 1. OSPF is used in this example to distribute routes including the loopback interface. All firewalls have full mesh security policies to allow any to any.

- In the FortiGate-500A_1 configuration, the NAT policy translates source address 236.1.1.1 to 237.1.1.1
- In the FortiGate-500A_4, configuration, the NAT policy translates source 236.1.1.1 to 238.1.1.1
- Source 236.1.1.1 is injected into network as well.

The following procedures include the CLI commands for configuring each of the FortiGate units in the example configuration.

To configure FortiGate-500A_1

1 Configure multicast routing.

```
config router multicast
  config interface
    edit port5
      set pim-mode sparse-mode
    next
    edit port4
      set pim-mode sparse-mode
    next
    edit lan
      set pim-mode sparse-mode
    next
    edit port1
      set pim-mode sparse-mode
    next
    edit lo999
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      set rp-candidate-group 1
    next
  end
set multicast-routing enable
  config pim-sm-global
    set bsr-candidate enable
    set bsr-interface lo0
  end
end
```

2 Add multicast security policies.

```
config firewall multicast-policy
  edit 1
    set dstintf port5
    set srcintf port4
  next
  edit 2
    set dstintf port4
    set srcintf port5
  next
  edit 3
  next
end
```

3 Add router access lists.

```
config router access-list
  edit 1
    config rule
      edit 1
        set prefix 228.1.1.1 255.255.255.255
        set exact-match enable
      next
      edit 2
        set prefix 237.1.1.1 255.255.255.255
        set exact-match enable
      next
      edit 3
        set prefix 238.1.1.1 255.255.255.255
        set exact-match enable
      next
    end
  next
end
```

To configure FortiGate-500A_2**1 Configure multicast routing.**

```
config router multicast
  config interface
    edit "lan"
      set pim-mode sparse-mode
    next
    edit "port5"
      set pim-mode sparse-mode
    next
    edit "port2"
      set pim-mode sparse-mode
    next
    edit "port4"
      set pim-mode sparse-mode
    next
    edit "lo_5"
      set pim-mode sparse-mode
      config join-group
        edit 236.1.1.1
        next
      end
    next
  end
  set multicast-routing enable
end
```

2 Add multicast security policies.

```
config firewall multicast-policy
  edit 1
    set dstintf lan
    set srcintf port5
  next
end
```

```
edit 2
    set dstintf port5
    set srcintf lan
next
edit 4
    set dstintf lan
    set srcintf port2
next
edit 5
    set dstintf port2
    set srcintf lan
next
edit 7
    set dstintf port1
    set srcintf port2
next
edit 8
    set dstintf port2
    set srcintf port1
next
edit 9
    set dstintf port5
    set srcintf port2
next
edit 10
    set dstintf port2
    set srcintf port5
next
edit 11
    set dnat 237.1.1.1
    set dstintf lo_5
    set nat 5.5.5.5
    set srcintf port2
next
edit 12
    set dstintf lan
    set srcintf lo_5
next
edit 13
    set dstintf port1
    set srcintf lo_5
next
edit 14
    set dstintf port5
    set srcintf lo_5
next
edit 15
    set dstintf port2
    set srcintf lo_5
next
edit 16
next
end
```

To configure FortiGate-500A_3**1 Configure multicast routing.**

```
config router multicast
  config interface
    edit port5
      set pim-mode sparse-mode
    next
    edit port6
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      set rp-candidate-priority 255
    next
    edit lan
      set pim-mode sparse-mode
    next
  end
set multicast-routing enable
config pim-sm-global
  set bsr-candidate enable
  set bsr-interface lo0
end
end
```

2 Add multicast security policies.

```
config firewall multicast-policy
  edit 1
    set dstintf port5
    set srcintf port6
  next
  edit 2
    set dstintf port6
    set srcintf port5
  next
  edit 3
    set dstintf port6
    set srcintf lan
  next
  edit 4
    set dstintf lan
    set srcintf port6
  next
  edit 5
    set dstintf port5
    set srcintf lan
  next
  edit 6
    set dstintf lan
    set srcintf port5
  next
end
```

To configure FortiGate-500A_4**1 Configure multicast routing.**

```
config router multicast
  config interface
    edit port6
      set pim-mode sparse-mode
    next
    edit lan
      set pim-mode sparse-mode
    next
    edit port1
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      config join-group
        edit 236.1.1.1
        next
      end
      set rp-candidate-priority 1
    next
  end
set multicast-routing enable
config pim-sm-global
  set bsr-allow-quick-refresh enable
  set bsr-candidate enable
  set bsr-interface lo0
  set bsr-priority 1
end
end
```

2 Add multicast security policies.

```
config firewall policy
  edit 1
    set srcintf lan
    set dstintf port6
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
  edit 2
    set srcintf port6
    set dstintf lan
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
end
```



```
edit 3
    set srcintf port1
    set dstintf port6
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 4
    set srcintf port6
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 5
    set srcintf port1
    set dstintf lan
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 6
    set srcintf lan
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 7
    set srcintf port1
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 8
    set srcintf port6
    set dstintf lo0
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
```

```

edit 9
    set srcintf port1
    set dstintf lo0
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
edit 10
    set srcintf lan
    set dstintf lo0
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
next
end

```

Example debug commands

You can use the following CLI commands to view information about and status of the multicast configuration. This section includes `get` and `diagnose` commands and some sample output.

```

get router info multicast pim sparse-mode table 236.1.1.1
get router info multicast pim sparse-mode neighbour

```

Neighbor Address	Interface	Uptime/Expires	Ver	DR Priority/
83.97.1.2	port6	02:22:01/00:01:44	v2	1 / DR

```

diagnose ip multicast mroute
    grp=236.1.1.1 src=19.2.1.1 intf=7 flags=(0x10000000) [ ]
status=resolved
    last_assert=171963 bytes=1766104 pkt=1718 wrong_if=1
num_ifs=2
    index(ttl)=[6(1),10(1),]
grp=236.1.1.1 src=1.4.50.4 intf=10 flags=(0x10000000) [ ]
status=resolved
    last_assert=834864 bytes=4416 pkt=138 wrong_if=0 num_ifs=2
    index(ttl)=[7(1),6(1),]
grp=238.1.1.1 src=1.4.50.4 intf=10 flags=(0x10000000) [ ]
status=resolved
    last_assert=834864 bytes=1765076 pkt=1717 wrong_if=0
num_ifs=1
    index(ttl)=[7(1),]

get router info multicast igmp groups
    IGMP Connected Group Membership
Group Address    Interface    Uptime    Expires    Last
Reporter

```

```

236.1.1.1      lan      00:45:48 00:03:21 10.4.1.1
236.1.1.1      lo0      02:19:31 00:03:23 1.4.50.4

get router info multicast pim sparse-mode interface
  Address      Interface VIFindex Ver/   Nbr    DR    DR
                Mode    Count  Prior
10.4.1.2      lan      2      v2/S   0      1     10.4.1.2
83.97.1.1     port6    0      v2/S   1      1     83.97.1.2
1.4.50.4      lo0      3      v2/S   0      1     1.4.50.4

get router info multicast pim sparse-mode rp-mapping
  PIM Group-to-RP Mappings
This system is the Bootstrap Router (v2)
Group(s): 224.0.0.0/4
  RP: 1.4.50.4
    Info source: 1.4.50.4, via bootstrap, priority 1
    Uptime: 02:20:32, expires: 00:01:58
  RP: 1.4.50.3
    Info source: 1.4.50.3, via bootstrap, priority 255
    Uptime: 02:20:07, expires: 00:02:24
Group(s): 228.1.1.1/32
  RP: 1.4.50.1
    Info source: 1.4.50.1, via bootstrap, priority 192
    Uptime: 02:18:24, expires: 00:02:06
Group(s): 237.1.1.1/32
  RP: 1.4.50.1
    Info source: 1.4.50.1, via bootstrap, priority 192
    Uptime: 02:18:24, expires: 00:02:06
Group(s): 238.1.1.1/32
  RP: 1.4.50.1
    Info source: 1.4.50.1, via bootstrap, priority 192
    Uptime: 02:18:24, expires: 00:02:06

get router info multicast pim sparse-mode bsr-info
  PIMv2 Bootstrap information
This system is the Bootstrap Router (BSR)
BSR address: 1.4.50.4
Uptime:      02:23:08, BSR Priority: 1, Hash mask length: 10
Next bootstrap message in 00:00:18
Role: Candidate BSR
State: Elected BSR

Candidate RP: 1.4.50.4(lo0)
  Advertisement interval 60 seconds
  Next Cand_RP_advertisement in 00:00:54

```




Virtual LANs

Virtual Local Area Networks (VLANs) multiply the capabilities of your FortiGate unit, and can also provide added network security. Virtual LANs (VLANs) use ID tags to logically separate devices on a network into smaller broadcast domains. These smaller domains forward packets only to devices that are part of that VLAN domain. This reduces traffic and increases network security.

A Local Area Network (LAN) is a group of connected computers and devices that are arranged into network broadcast domains. A LAN broadcast domain includes all the computers that receive a packet broadcast from any computer in that broadcast domain. A switch will automatically forward the packets to all of its ports; in contrast, routers do not automatically forward network broadcast packets. This means routers separate broadcast domains. If a network has only switches and no routers, that network is considered one broadcast domain, no matter how large or small it is. Smaller broadcast domains are more efficient because fewer devices receive unnecessary packets. They are more secure as well because a hacker reading traffic on the network will have access to only a small portion of the network instead of the entire network's traffic.

Virtual LANs (VLANs) use ID tags to logically separate a LAN into smaller broadcast domains. Each VLAN is its own broadcast domain. Smaller broadcast domains reduce traffic and increase network security. The IEEE 802.1Q standard defines VLANs. All layer-2 and layer-3 devices along a route must be 802.1Q-compliant to support VLANs along that route. For more information, see [“VLAN switching and routing” on page 504](#) and [“VLAN layer-3 routing” on page 507](#).

VLANs reduce the size of the broadcast domains by only forwarding packets to interfaces that are part of that VLAN or part of a VLAN trunk link. Trunk links form switch-to-switch or switch-to-router connections, and forward traffic for all VLANs. This enables a VLAN to include devices that are part of the same broadcast domain, but physically distant from each other.

VLAN ID tags consist of a 4-byte frame extension that switches and routers apply to every packet sent and received in the VLAN. Workstations and desktop computers, which are commonly originators or destinations of network traffic, are not an active part of the VLAN process. All the VLAN tagging and tag removal is done after the packet has left the computer. For more information, see [“VLAN ID rules” on page 504](#).

Any FortiGate unit without VDOMs enabled can have a maximum of 255 interfaces in transparent operating mode. The same is true for any single VDOM. In NAT mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in transparent operating mode, you need to configure multiple VDOMs that enable you to divide the total number of interfaces over all the VDOMs.

One example of an application of VLANs is a company's accounting department. Accounting computers may be located at both main and branch offices. However, accounting computers need to communicate with each other frequently and require increased security. VLANs allow the accounting network traffic to be sent only to accounting computers and to connect accounting computers in different locations as if they were on the same physical subnet.



This guide uses the term “packet” to refer to both layer-2 frames and layer-3 packets.

VLAN ID rules

Layer-2 switches and layer-3 devices add VLAN ID tags to the traffic as it arrives and remove them before they deliver the traffic to its final destination. Devices such as PCs and servers on the network do not require any special configuration for VLANs. Twelve bits of the 4-byte VLAN tag are reserved for the VLAN ID number. Valid VLAN ID numbers are from 1 to 4094, while 0 is used for high priority frames, and 4095 is reserved.

On a layer-2 switch, you can have only one VLAN subinterface per physical interface, unless that interface is configured as a trunk link. Trunk links can transport traffic for multiple VLANs to other parts of the network.

On a FortiGate unit, you can add multiple VLANs to the same physical interface. However, VLAN subinterfaces added to the same physical interface cannot have the same VLAN ID or have IP addresses on the same subnet. You can add VLAN subinterfaces with the same VLAN ID to different physical interfaces.

Creating VLAN subinterfaces with the same VLAN ID does not create any internal connection between them. For example a VLAN ID of 300 on port1 and VLAN ID of 300 on port2 are allowed, but they are not connected. Their relationship is the same as between any two FortiGate network interfaces.

VLAN switching and routing

VLAN switching takes place on the OSI model layer-2, just like other network switching. VLAN routing takes place on the OSI model layer-3. The difference between them is that during VLAN switching, VLAN packets are simply forwarded to their destination. This is different from VLAN routing where devices can open the VLAN packets and change their VLAN ID tags to route the packets to a new destination.

VLAN layer-2 switching

Ethernet switches are layer-2 devices, and generally are 802.1Q compliant. Layer 2 refers to the second layer of the seven layer Open Systems Interconnect (OSI) basic networking model; the Data Link layer. FortiGate units act as layer-2 switches or bridges when they are in transparent mode. The units simply tag and forward the VLAN traffic or receive and remove the tags from the packets. A layer-2 device does not inspect incoming packets or change their contents; it only adds or removes tags and routes the packet.

A VLAN can have any number of physical interfaces assigned to it. Multiple VLANs can be assigned to the same physical interface. Typically two or more physical interfaces are assigned to a VLAN, one for incoming and one for outgoing traffic. Multiple VLANs can be configured on one FortiGate unit, including trunk links.

Layer-2 VLAN example

To better understand VLAN operation, this example shows what happens to a data frame on a network that uses VLANs.


The network topology consists of two 8-port switches that are configured to support VLANs on a network. Both switches are connected through port 8 using an 802.1Q trunk link. Subnet 1 is connected to switch A, and subnet 2 is connected to switch B. The ports on the switches are configured as follows.

Table 28: How ports and VLANs are used on Switch A and B

Switch	Ports	VLAN
A	1 - 4	100
A	5 - 7	200
A & B	8	Trunk link
B	4 - 5	100
B	6	200


In this example, switch A is connected to the Branch Office and switch B to the Main Office.

- 1 A computer on port 1 of switch A sends a data frame over the network.




- 2 Switch A tags the data frame with a VLAN 100 ID tag upon arrival because port 1 is part of VLAN 100.
- 3 Switch A forwards the tagged data frame to the other VLAN 100 ports — ports 2 through 4. Switch A also forwards the data frame to the 802.1Q trunk link (port 8) so other parts of the network that may contain VLAN 100 groups will receive VLAN 100 traffic.

This data frame is not forwarded to the other ports on switch A because they are not part of VLAN 100. This increases security and decreases network traffic.



- 4 Switch B receives the data frame over the trunk link (port 8).
- 5 Because there are VLAN 100 ports on switch B (ports 4 and 5), the data frame is forwarded to those ports. As with switch A, the data frame is not delivered to VLAN 200.

If there were no VLAN 100 ports on switch B, the switch would not forward the data frame and it would stop there.



- 6 The switch removes the VLAN 100 ID tag before it forwards the data frame to an end destination.

The sending and receiving computers are not aware of any VLAN tagging on the data frames that are being transmitted. When any computer receives that data frame, it appears as a normal data frame.

VLAN layer-3 routing

Routers are layer-3 devices. Layer 3 refers to the third layer of the OSI networking model, the Network layer. FortiGate units in NAT mode act as layer-3 devices. As with layer 2, FortiGate units acting as layer-3 devices are 802.1Q-compliant.

The main difference between layer-2 and layer-3 devices is how they process VLAN tags. Layer-2 switches just add, read and remove the tags. They do not alter the tags or do any other high-level actions. Layer-3 routers not only add, read and remove tags but also analyze the data frame and its contents. This analysis allows layer-3 routers to change the VLAN tag if it is appropriate and send the data frame out on a different VLAN.

In a layer-3 environment, the 802.1Q-compliant router receives the data frame and assigns a VLAN ID. The router then forwards the data frame to other members of the same VLAN broadcast domain. The broadcast domain can include local ports, layer-2 devices and layer-3 devices such as routers and firewalls. When a layer-3 device receives the data frame, the device removes the VLAN tag and examines its contents to decide what to do with the data frame. The layer-3 device considers:

- source and destination addresses
- protocol
- port number.


The data frame may be forwarded to another VLAN, sent to a regular non-VLAN-tagged network or just forwarded to the same VLAN as a layer-2 switch would do. Or, the data frame may be discarded if the proper security policy has been configured to do so.

Layer-3 VLAN example

In this example, switch A is connected to the Branch Office subnet, the same as subnet 1 in the layer-2 example. In the Main Office subnet, VLAN 300 is on port 5 of switch B. The FortiGate unit is connected to switch B on port 1 and the trunk link connects the FortiGate unit's port 3 to switch A. The other ports on switch B are unassigned.


This example explains how traffic can change VLANs originating on VLAN 100 and arriving at a destination on VLAN 300. Layer-2 switches alone cannot accomplish this, but a layer-3 router can.

- 1 The VLAN 100 computer at the Branch Office sends the data frame to switch A, where the VLAN 100 tag is added.



- 2 Switch A forwards the tagged data frame to the FortiGate unit over the 802.1Q trunk link, and to the VLAN 100 interfaces on Switch A.


Up to this point everything is the same as in the layer-2 example.




- 3 The FortiGate unit removes the VLAN 100 tag, and inspects the content of the data frame. The FortiGate unit uses the content to select the correct security policy and routing options.

- 4 The FortiGate unit's security policy allows the data frame to go to VLAN 300 in this example. The data frame will be sent to all VLAN 300 interfaces, but in the example there is only port 1 on the FortiGate unit. Before the data frame leaves, the FortiGate unit adds the VLAN ID 300 tag to the data frame.

This is the step that layer 2 cannot do. Only layer 3 can retag a data frame as a different VLAN.



- 5 Switch B receives the data frame, and removes the VLAN ID 300 tag, because this is the last hop, and forwards the data frame to the computer on port 5.



In this example, a data frame arrived at the FortiGate unit tagged as VLAN 100. After checking its content, the FortiGate unit retagged the data frame for VLAN 300. It is this change from VLAN 100 to VLAN 300 that requires a layer-3 routing device, in this case the FortiGate unit. Layer-2 switches cannot perform this change.

VLANs in NAT mode

In NAT mode the FortiGate unit functions as a layer-3 device. In this mode, the FortiGate unit controls the flow of packets between VLANs, but can also remove VLAN tags from incoming VLAN packets. The FortiGate unit can also forward untagged packets to other networks, such as the Internet.

In NAT mode, the FortiGate unit supports VLAN trunk links with IEEE 802.1Q-compliant switches, or routers. The trunk link transports VLAN-tagged packets between physical subnets or networks. When you add VLAN sub-interfaces to the FortiGate unit physical interfaces, the VLANs have IDs that match the VLAN IDs of packets on the trunk link. The FortiGate unit directs packets with VLAN IDs to sub-interfaces with matching IDs.

You can define VLAN sub-interfaces on all FortiGate physical interfaces. However, if multiple virtual domains are configured on the FortiGate unit, you will have access to only the physical interfaces on your virtual domain. The FortiGate unit can tag packets leaving on a VLAN subinterface. It can also remove VLAN tags from incoming packets and add a different VLAN tag to outgoing packets.

Normally in VLAN configurations, the FortiGate unit's internal interface is connected to a VLAN trunk, and the external interface connects to an Internet router that is not configured for VLANs. In this configuration the FortiGate unit can apply different policies for traffic on each VLAN interface connected to the internal interface, which results in less network traffic and better security.

Adding VLAN subinterfaces

A VLAN subinterface, also called a VLAN, is a virtual interface on a physical interface. The subinterface allows routing of VLAN tagged packets using that physical interface, but it is separate from any other traffic on the physical interface.

Adding a VLAN subinterface includes configuring:

- Physical interface
- IP address and netmask
- VLAN ID
- VDOM

Physical interface

The term VLAN subinterface correctly implies the VLAN interface is not a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network router that is configured for this VLAN. Without that router, the VLAN will not be connected to the network, and VLAN traffic will not be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

When you are working with interfaces on your FortiGate unit, use the *Column Settings* on the Interface display to make sure the information you need is displayed. When working with VLANs, it is useful to position the *VLAN ID* column close to the IP address. If you are working with VDOMs, including the *Virtual Domain* column as well will help you troubleshoot problems more quickly.

To view the Interface display, go to *System > Network > Interface*.

IP address and netmask

FortiGate unit interfaces cannot have overlapping IP addresses. The IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask pair. This rule helps prevent a broadcast storm or other similar network problems.



If you are unable to change your existing configurations to prevent IP overlap, enter the CLI command `config system global and set ip-overlap enable` to allow IP address overlap. If you enter this command, multiple VLAN interfaces can have an IP address that is part of a subnet used by another interface. This command is recommended for advanced users only.

VLAN ID

The VLAN ID is part of the VLAN tag added to the packets by VLAN switches and routers. The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together. VLAN ID 0 is used only for high priority frames, and 4095 is reserved.

All devices along a route must support the VLAN ID of the traffic along that route. Otherwise, the traffic will be discarded before reaching its destination. For example, if your computer is part of VLAN_100 and a co-worker on a different floor of your building is also on the same VLAN_100, you can communicate with each other over VLAN_100, only if all the switches and routers support VLANs and are configured to pass along VLAN_100 traffic properly. Otherwise, any traffic you send your co-worker will be blocked or not delivered.

VDOM

If VDOMs are enabled, each VLAN subinterface must belong to a VDOM. This rule also applies for physical interfaces.



Interface-related CLI commands require a VDOM to be specified, regardless of whether the FortiGate unit has VDOMs enabled.

VLAN subinterfaces on separate VDOMs cannot communicate directly with each other. In this situation, the VLAN traffic must exit the FortiGate unit and re-enter the unit again, passing through firewalls in both directions. This situation is the same for physical interfaces.

A VLAN subinterface can belong to a different VDOM than the physical interface it is part of. This is because the traffic on the VLAN is handled separately from the other traffic on that interface. This is one of the main strengths of VLANs.

The following procedure will add a VLAN subinterface called `VLAN_100` to the FortiGate internal interface with a VLAN ID of 100. It will have an IP address and netmask of `172.100.1.1/255.255.255.0`, and allow HTTPS, PING, and TELNET administrative access. Note that in the CLI, you must enter “`set type vlan`” before setting the `vlanid`, and that the `allowaccess` protocols are lower case.

To add a VLAN subinterface in NAT mode - web-based manager

- 1 If *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
- 2 Go to *System > Network > Interface*.
- 3 Select *Create New* to add a VLAN subinterface.
- 4 Enter the following:

VLAN Name	VLAN_100
Type	VLAN
Interface	internal
VLAN ID	100
Addressing Mode	Manual
IP/Netmask	172.100.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET

- 5 Select *OK*.

To view the new VLAN subinterface, select the expand arrow next to the parent physical interface (the internal interface). This will expand the display to show all VLAN subinterfaces on this physical interface. If there is no expand arrow displayed, there are no subinterfaces configured on that physical interface.

For each VLAN, the list displays the name of the VLAN, and, depending on column settings, its IP address, the Administrative access you selected for it, the VLAN ID number, and which VDOM it belongs to if VDOMs are enabled.

To add a VLAN subinterface in NAT mode - CLI

```
config system interface
  edit VLAN_100
    set interface internal
    set type vlan
    set vlanid 100
    set ip 172.100.1.1 255.255.255.0
    set allowaccess https ping telnet
  end
```

Configuring security policies and routing

Once you have created a VLAN subinterface on the FortiGate unit, you need to configure security policies and routing for that VLAN. Without these, the FortiGate unit will not pass VLAN traffic to its intended destination. Security policies direct traffic through the FortiGate unit between interfaces. Routing directs traffic across the network.

Configuring security policies

Security policies permit communication between the FortiGate unit's network interfaces based on source and destination IP addresses. Interfaces that communicate with the VLAN interface need security policies to permit traffic to pass between them and the VLAN interface.

Each VLAN needs a security policy for each of the following connections the VLAN will be using:

- from this VLAN to an external network
- from an external network to this VLAN
- from this VLAN to another VLAN in the same virtual domain on the FortiGate unit
- from another VLAN to this VLAN in the same virtual domain on the FortiGate unit.

The packets on each VLAN are subject to antivirus scans and other UTM measures as they pass through the FortiGate unit.

Configuring routing

As a minimum, you need to configure a default static route to a gateway with access to an external network for outbound packets. In more complex cases, you will have to configure different static or dynamic routes based on packet source and destination addresses.

As with firewalls, you need to configure routes for VLAN traffic. VLANs need routing and a gateway configured to send and receive packets outside their local subnet just as physical interfaces do. The type of routing you configure, static or dynamic, will depend on the routing used by the subnet and interfaces you are connecting to. Dynamic routing can be routing information protocol (RIP), border gateway protocol (BGP), open shortest path first (OSPF), or multicast.

If you enable SSH, PING, TELNET, HTTPS and HTTP on the VLAN, you can use those protocols to troubleshoot your routing and test that it is properly configured. Enabling logging on the interfaces and using CLI diagnose commands such as `diagnose sniff packet <interface_name>` can also help locate any possible configuration or hardware issues.


Example VLAN configuration in NAT mode

In this example two different internal VLAN networks share one interface on the FortiGate unit, and share the connection to the Internet. This example shows that two networks can have separate traffic streams while sharing a single interface. This configuration could apply to two departments in a single company, or to different companies.

There are two different internal network VLANs in this example. VLAN_100 is on the 10.1.1.0/255.255.255.0 subnet, and VLAN_200 is on the 10.1.2.0/255.255.255.0 subnet. These VLANs are connected to the VLAN switch, such as a Cisco 2950 Catalyst switch.

The FortiGate internal interface connects to the VLAN switch through an 802.1Q trunk. The internal interface has an IP address of 192.168.110.126 and is configured with two VLAN subinterfaces (VLAN_100 and VLAN_200). The external interface has an IP address of 172.16.21.2 and connects to the Internet. The external interface has no VLAN subinterfaces on it.

Figure 53: FortiGate unit with VLANs in NAT mode



When the VLAN switch receives packets from VLAN_100 and VLAN_200, it applies VLAN ID tags and forwards the packets of each VLAN both to local ports and to the FortiGate unit across the trunk link. The FortiGate unit has policies that allow traffic to flow between the VLANs, and from the VLANs to the external network.

This section describes how to configure a FortiGate-800 unit and a Cisco Catalyst 2950 switch for this example network topology. The Cisco configuration commands used in this section are IOS commands.

It is assumed that both the FortiGate-800 and the Cisco 2950 switch are installed and connected and that basic configuration has been completed. On the switch, you will need to be able to access the CLI to enter commands. Refer to the manual for your FortiGate model as well as the manual for the switch you select for more information.

It is also assumed that no VDOMs are enabled.

General configuration steps

The following steps provide an overview of configuring and testing the hardware used in this example. For best results in this configuration, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Configure the FortiGate unit
 - Configure the external interface
 - Add two VLAN subinterfaces to the internal network interface
 - Add firewall addresses and address ranges for the internal and external networks
 - Add security policies to allow:
 - the VLAN networks to access each other
 - the VLAN networks to access the external network.
- 2 Configure the VLAN switch

Configure the FortiGate unit

Configuring the FortiGate unit includes:

- [Configure the external interface](#)
- [Add VLAN subinterfaces](#)
- [Add the firewall addresses](#)
- [Add the security policies](#)

Configure the external interface

The FortiGate unit's external interface will provide access to the Internet for all internal networks, including the two VLANs.

To configure the external interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Edit* for the external interface.
- 3 Enter the following information and select *OK*:

Addressing mode	Manual
IP/Netmask	172.16.21.2/255.255.255.0

To configure the external interface - CLI

```
config system interface
  edit external
    set mode static
    set ip 172.16.21.2 255.255.255.0
  end
```

Add VLAN subinterfaces

This step creates the VLANs on the FortiGate unit internal physical interface. The IP address of the internal interface does not matter to us, as long as it does not overlap with the subnets of the VLAN subinterfaces we are configuring on it.

The rest of this example shows how to configure the VLAN behavior on the FortiGate unit, configure the switches to direct VLAN traffic the same as the FortiGate unit, and test that the configuration is correct.

Adding VLAN subinterfaces can be completed through the web-based manager, or the CLI.

To add VLAN subinterfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Name	VLAN_100
Interface	internal
VLAN ID	100
Addressing mode	Manual
IP/Netmask	10.1.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Name	VLAN_200
Interface	internal
VLAN ID	200
Addressing mode	Manual
IP/Netmask	10.1.2.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET

To add VLAN subinterfaces - CLI

```
config system interface
  edit VLAN_100
    set vdom root
    set interface internal
    set type vlan
    set vlanid 100
    set mode static
    set ip 10.1.1.1 255.255.255.0
    set allowaccess https ping telnet
  next
  edit VLAN_200
    set vdom root
```

```

set interface internal
set type vlan
set vlanid 200
set mode static
set ip 10.1.2.1 255.255.255.0
set allowaccess https ping telnet
end

```

Add the firewall addresses

You need to define the addresses of the VLAN subnets for use in security policies. The FortiGate unit provides one default address, “all”, that you can use when a security policy applies to all addresses as a source or destination of a packet. However, using “all” is less secure and should be avoided when possible.

In this example, the “_Net” part of the address name indicates a range of addresses instead of a unique address. When choosing firewall address names, use informative and unique names.

To add the firewall addresses - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Address Name	VLAN_100_Net
Type	Subnet / IP Range
Subnet / IP Range	10.1.1.0/255.255.255.0

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Address Name	VLAN_200_Net
Type	Subnet / IP Range
Subnet / IP Range	10.1.2.0/255.255.255.0

To add the firewall addresses - CLI

```

config firewall address
edit VLAN_100_Net
set type ipmask
set subnet 10.1.1.0 255.255.255.0
next
edit VLAN_200_Net
set type ipmask
set subnet 10.1.2.0 255.255.255.0
end

```

Add the security policies

Once you have assigned addresses to the VLANs, you need to configure security policies for them to allow valid packets to pass from one VLAN to another and to the Internet.



You can customize the Security Policy display by including some or all columns, and customize the column order onscreen. Due to this feature, security policy screenshots may not appear the same as on your screen.

If you do not want to allow all services on a VLAN, you can create a security policy for each service you want to allow. This example allows all services.

To add the security policies - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100
Source Address	VLAN_100_Net
Destination Interface/Zone	VLAN_200
Destination Address	VLAN_200_Net
Schedule	Always
Service	ANY
Action	ACCEPT
Enable NAT	Enable

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_200
Source Address	VLAN_200_Net
Destination Interface/Zone	VLAN_100
Destination Address	VLAN_100_Net
Schedule	Always
Service	ANY
Action	ACCEPT
Enable NAT	Enable

- 6 Select *Create New*.
- 7 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100
Source Address	VLAN_100_Net

Destination Interface/Zone	external
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT
Enable NAT	Enable

- 8 Select *Create New*.
- 9 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_200
Source Address	VLAN_200_Net
Destination Interface/Zone	external
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT
Enable NAT	Enable

To add the security policies - CLI

```

config firewall policy
  edit 1
    set srcintf VLAN_100
    set srcaddr VLAN_100_Net
    set dstintf VLAN_200
    set dstaddr VLAN_200_Net
    set schedule always
    set service ANY
    set action accept
    set nat enable
    set status enable
  next
  edit 2
    set srcintf VLAN_200
    set srcaddr VLAN_200_Net
    set dstintf VLAN_100
    set dstaddr VLAN_100_Net
    set schedule always
    set service ANY
    set action accept
    set nat enable
    set status enable
  next
  edit 3
    set srcintf VLAN_100
    set srcaddr VLAN_100_Net

```

```
        set dstintf external
        set dstaddr all
        set schedule always
        set service ANY
        set action accept
        set nat enable
        set status enable
    next
    edit 4
        set srcintf VLAN_200
        set srcaddr VLAN_200_Net
        set dstintf external
        set dstaddr all
        set schedule always
        set service ANY
        set action accept
        set nat enable
        set status enable
    end
```

Configure the VLAN switch

On the Cisco Catalyst 2950 Catalyst VLAN switch, you need to define VLANs 100 and 200 in the VLAN database, and then add a configuration file to define the VLAN subinterfaces and the 802.1Q trunk interface.

One method to configure a Cisco switch is to connect over a serial connection to the console port on the switch, and enter the commands at the CLI. Another method is to designate one interface on the switch as the management interface and use a web browser to connect to the switch's graphical interface. For details on connecting and configuring your Cisco switch, refer to the installation and configuration manuals for the switch.

The switch used in this example is a Cisco Catalyst 2950 switch. The commands used are IOS commands. Refer to the switch manual for help with these commands.

To configure the VLAN subinterfaces and the trunk interfaces

Add this file to the Cisco switch:

```
!
interface FastEthernet0/3
    switchport access vlan 100
!
interface FastEthernet0/9
    switchport access vlan 200
!
interface FastEthernet0/24
    switchport trunk encapsulation dot1q
    switchport mode trunk
!
```

The switch has the configuration:

Port 0/3	VLAN ID 100
Port 0/9	VLAN ID 200
Port 0/24	802.1Q trunk



To complete the setup, configure devices on VLAN_100 and VLAN_200 with default gateways. The default gateway for VLAN_100 is the FortiGate VLAN_100 subinterface. The default gateway for VLAN_200 is the FortiGate VLAN_200 subinterface.

Test the configuration

Use diagnostic commands, such as `tracert`, to test traffic routed through the FortiGate unit and the Cisco switch.

Testing traffic from VLAN_100 to VLAN_200

In this example, a route is traced between the two internal networks. The route target is a host on VLAN_200.

Access a command prompt on a Windows computer on the VLAN_100 network, and enter the following command:

```
C:\>tracert 10.1.2.2
Tracing route to 10.1.2.2 over a maximum of 30 hops:
  1  <10 ms  <10 ms  <10 ms  10.1.1.1
  2  <10 ms  <10 ms  <10 ms  10.1.2.2
Trace complete.
```

Testing traffic from VLAN_200 to the external network

In this example, a route is traced from an internal network to the external network. The route target is the external network interface of the FortiGate-800 unit.

From VLAN_200, access a command prompt and enter this command:

```
C:\>tracert 172.16.21.2
Tracing route to 172.16.21.2 over a maximum of 30 hops:
  1  <10 ms  <10 ms  <10 ms  10.1.2.1
  2  <10 ms  <10 ms  <10 ms  172.16.21.2
Trace complete.
```

See also

VLANs in transparent mode

In transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering and intrusion protection to traffic. There are some limitations in transparent mode in that you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

VLANs and transparent mode

You can insert the FortiGate unit operating in transparent mode into the VLAN trunk without making changes to your network. In a typical configuration, the FortiGate unit internal interface accepts VLAN packets on a VLAN trunk from a VLAN switch or router connected to internal network VLANs. The FortiGate external interface forwards VLAN-tagged packets through another VLAN trunk to an external VLAN switch or router and on to external networks such as the Internet. You can configure the unit to apply different policies for traffic on each VLAN in the trunk.

To pass VLAN traffic through the FortiGate unit, you add two VLAN subinterfaces with the same VLAN ID, one to the internal interface and the other to the external interface. You then create a security policy to permit packets to flow from the internal VLAN interface to the external VLAN interface. If required, you create another security policy to permit packets to flow from the external VLAN interface to the internal VLAN interface. Typically in transparent mode, you do not permit packets to move between different VLANs. Network protection features, such as spam filtering, web filtering and anti-virus scanning, are applied through the UTM profiles specified in each security policy, enabling very detailed control over traffic.

When the FortiGate unit receives a VLAN-tagged packet at a physical interface, it directs the packet to the VLAN subinterface with the matching VLAN ID. The VLAN tag is removed from the packet, and the FortiGate unit then applies security policies using the same method it uses for non-VLAN packets. If the packet exits the FortiGate unit through a VLAN subinterface, the VLAN ID for that subinterface is added to the packet and the packet is sent to the corresponding physical interface. For a configuration example, see [“Example of VLANs in transparent mode” on page 525](#).

There are two essential steps to configure your FortiGate unit to work with VLANs in transparent mode:

- [Add VLAN subinterfaces](#)
- [Create security policies](#)

You can also configure the protection profiles that manage antivirus scanning, web filtering and spam filtering. For more information on UTM profiles, see [“UTM Guide” on page 877](#).

Add VLAN subinterfaces

The VLAN ID of each VLAN subinterface must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch. The VLAN ID can be any number between 1 and 4094, with 0 being used only for high priority frames and 4095 being reserved. You add VLAN subinterfaces to the physical interface that receives VLAN-tagged packets.

For this example, we are creating a VLAN called internal_v225 on the internal interface, with a VLAN ID of 225. Administrative access is enabled for HTTPS and SSH. VDOMs are not enabled.

To add VLAN subinterfaces in transparent mode - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*.

Name	internal_v225
Type	VLAN
Interface	internal

VLAN ID	225
Ping Server	not enabled
Administrative Access	Enable HTTPS, and SSH. These are very secure access methods.
Description	VLAN 225 on internal interface

The FortiGate unit adds the new subinterface to the interface that you selected.

Repeat steps 2 and 3 to add additional VLANs. You will need to change the *VLAN ID*, *Name*, and possibly *Interface* when adding additional VLANs.

To add VLAN subinterfaces in transparent mode - CLI

```
config system interface
  edit internal_v225
    set interface internal
    set vlanid 225
    set allowaccess HTTPS SSH
    set description "VLAN 225 on internal interface"
    set vdom root
  end
```

Create security policies

In transparent mode, the FortiGate unit performs antivirus and antispam scanning on each VLAN's packets as they pass through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another.

To add security policies for VLAN subinterfaces - web based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New* to add firewall addresses that match the source and destination IP addresses of VLAN packets.
- 3 Go to *Policy > Policy > Policy*.
- 4 Select *Create New*.
- 5 From the Source Interface/Zone list, select the VLAN interface where packets enter the unit.
- 6 From the Destination Interface/Zone list, select the VLAN interface where packets exit the unit.
- 7 Select the Source and Destination Address names that you added in step 2.
- 8 Select *Protection Profile*, and select the profile from the list.
- 9 Configure other settings as required.
- 10 Select *OK*.

To add security policies for VLAN subinterfaces - CLI

```
config firewall address
  edit incoming_VLAN_address
    set associated-interface <incoming_VLAN_interface>
    set type ipmask
    set subnet <IPv4_address_mask>
```

```
next
edit outgoing_VLAN_address
    set associated-interface <outgoing_VLAN_interface>
    set type ipmask
    set subnet <IPv4_address_mask>
next
end
config firewall policy
edit <unused_policy_number>
    set srcintf <incoming_VLAN_interface>
    set srcaddr incoming_VLAN_address
    set destintf <outgoing_VLAN_interface>
    set destaddr outgoing_VLAN_address
    set service <protocol_to_allow_on_VLAN>
    set action ACCEPT
    set profile-status enable
    set profile <selected_profile>
next
end
end
```


Example of VLANs in transparent mode

In this example, the FortiGate unit is operating in transparent mode and is configured with two VLANs: one with an ID of 100 and the other with ID 200. The internal and external physical interfaces each have two VLAN subinterfaces, one for VLAN_100 and one for VLAN_200.

The IP range for the internal VLAN_100 network is 10.100.0.0/255.255.0.0, and for the internal VLAN_200 network is 10.200.0.0/255.255.0.0.

The internal networks are connected to a Cisco 2950 VLAN switch, which combines traffic from the two VLANs onto one the FortiGate unit internal interface. The VLAN traffic leaves the FortiGate unit on the external network interface, goes on to the VLAN switch, and on to the Internet. When the FortiGate unit receives a tagged packet, it directs it from the incoming VLAN subinterface to the outgoing VLAN subinterface for that VLAN.

This section describes how to configure a FortiGate-800 unit, Cisco switch, and Cisco router in the network topology shown in [Figure 180](#).

Figure 54: VLAN transparent network topology

General configuration steps

The following steps summarize the configuration for this example. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Configure the FortiGate unit which includes
 - Adding VLAN subinterfaces
 - Adding the security policies
- 2 Configure the Cisco switch and router

Configure the FortiGate unit

The FortiGate unit must be configured with the VLAN subinterfaces and the proper security policies to enable traffic to flow through the FortiGate unit.

Add VLAN subinterfaces

For each VLAN, you need to create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID.

To add VLAN subinterfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Create New*.

- 3 Enter the following information and select *OK*:

Name	VLAN_100_int
Interface	internal
VLAN ID	100

- 4 Select *Create New*.

- 5 Enter the following information and select *OK*:

Name	VLAN_100_ext
Interface	external
VLAN ID	100

- 6 Select *Create New*.

- 7 Enter the following information and select *OK*:

Name	VLAN_200_int
Interface	internal
VLAN ID	200

- 8 Select *Create New*.

- 9 Enter the following information and select *OK*:

Name	VLAN_200_ext
Interface	external
VLAN ID	200

To add VLAN subinterfaces - CLI

```
config system interface
  edit VLAN_100_int
    set status down
    set type vlan
    set interface internal
    set vlanid 100
  next
  edit VLAN_100_ext
    set status down
    set type vlan
    set interface external
    set vlanid 100
  next
  edit VLAN_200_int
    set status down
    set type vlan
    set interface internal
    set vlanid 200
  next
```

```

edit VLAN_200_ext
  set status down
  set type vlan
  set interface external
  set vlanid 200
end

```

Add the security policies

Security policies allow packets to travel between the VLAN_100_int interface and the VLAN_100_ext interface. Two policies are required; one for each direction of traffic. The same is required between the VLAN_200_int interface and the VLAN_200_ext interface, for a total of four required security policies.

To add the security policies - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_int
Source Address	all
Destination Interface/Zone	VLAN_100_ext
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_ext
Source Address	all
Destination Interface/Zone	VLAN_100_int
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

- 6 Go to *Policy > Policy > Policy*.
- 7 Select *Create New*.

8 Enter the following information and select OK:

Source Interface/Zone	VLAN_200_int
Source Address	all
Destination Interface/Zone	VLAN_200_ext
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT
Enable NAT	enable

9 Select *Create New*.

10 Enter the following information and select OK:

Source Interface/Zone	VLAN_200_ext
Source Address	all
Destination Interface/Zone	VLAN_200_int
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

To add the security policies - CLI

```
config firewall policy
edit 1
set srcintf VLAN_100_int
set srcaddr all
set dstintf VLAN_100_ext
set dstaddr all
set action accept
set schedule always
set service ANY
next
edit 2
set srcintf VLAN_100_ext
set srcaddr all
set dstintf VLAN_100_int
set dstaddr all
set action accept
set schedule always
set service ANY
next
edit 3
set srcintf VLAN_200_int
set srcaddr all
set dstintf VLAN_200_ext
```

```

        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 4
        set srcintf VLAN_200_ext
        set srcaddr all
        set dstintf VLAN_200_int
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    end

```

Configure the Cisco switch and router

This example includes configuration for the Cisco Catalyst 2900 ethernet switch, and for the Cisco Multiservice 2620 ethernet router. If you have access to a different VLAN enabled switch or VLAN router you can use them instead, however their configuration is not included in this document.

Configure the Cisco switch

On the VLAN switch, you need to define VLAN_100 and VLAN_200 in the VLAN database and then add a configuration file to define the VLAN subinterfaces and the 802.1Q trunk interface.

Add this file to the Cisco switch:

```

interface FastEthernet0/3
    switchport access vlan 100
!
interface FastEthernet0/9
    switchport access vlan 200
!
interface FastEthernet0/24
    switchport trunk encapsulation dot1q
    switchport mode trunk
!

```

The switch has the following configuration:

Port 0/3	VLAN ID 100
Port 0/9	VLAN ID 200
Port 0/24	802.1Q trunk

Configure the Cisco router

You need to add a configuration file to the Cisco Multiservice 2620 ethernet router. The file defines the VLAN subinterfaces and the 802.1Q trunk interface on the router. The 802.1Q trunk is the physical interface on the router.

The IP address for each VLAN on the router is the gateway for that VLAN. For example, all devices on the internal VLAN_100 network will have 10.100.0.1 as their gateway.

Add this file to the Cisco router:

```
!  
interface FastEthernet0/0  
!  
interface FastEthernet0/0.1  
  encapsulation dot1Q 100  
  ip address 10.100.0.1 255.255.255.0  
!  
interface FastEthernet0/0.2  
  encapsulation dot1Q 200  
  ip address 10.200.0.1 255.255.255.0  
!
```

The router has the following configuration:

Port 0/0.1	VLAN ID 100
Port 0/0.2	VLAN ID 200
Port 0/0	802.1Q trunk

Test the configuration

Use diagnostic network commands such as `tracert` (tracert) and `ping` to test traffic routed through the network.

Testing traffic from VLAN_100 to VLAN_200

In this example, a route is traced between the two internal networks. The route target is a host on VLAN_200. The Windows `tracert` command `tracert` is used.

From VLAN_100, access a Windows command prompt and enter this command:

```
C:\>tracert 10.1.2.2  
Tracing route to 10.1.2.2 over a maximum of 30 hops:  
  1  <10 ms  <10 ms  <10 ms  10.1.1.1  
  2  <10 ms  <10 ms  <10 ms  10.1.2.2  
Trace complete.
```

Troubleshooting VLAN issues

Several problems can occur with your VLANs. Since VLANs are interfaces with IP addresses, they behave as interfaces and can have similar problems that you can diagnose with tools such as `ping`, `tracert`, packet sniffing, and `diag debug`.

Asymmetric routing

You might discover unexpectedly that hosts on some networks are unable to reach certain other networks. This occurs when request and response packets follow different paths. If the FortiGate unit recognizes the response packets, but not the requests, it blocks the packets as invalid. Also, if the FortiGate unit recognizes the same packets repeated on multiple interfaces, it blocks the session as a potential attack.

This is asymmetric routing. By default, the FortiGate unit blocks packets or drops the session when this happens. You can configure the FortiGate unit to permit asymmetric routing by using the following CLI commands:

```
config vdom
  edit <vdom_name>
    config system settings
      set asymroute enable
    end
  end
```

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem. If this solves your blocked traffic issue, you know that asymmetric routing is the cause. But allowing asymmetric routing is not the best solution, because it reduces the security of your network.

For a long-term solution, it is better to change your routing configuration or change how your FortiGate unit connects to your network. The [Asymmetric Routing and Other FortiGate Layer-2 Installation Issues](#) technical note provides detailed examples of asymmetric routing situations and possible solutions.



If you enable asymmetric routing, antivirus and intrusion prevention systems will not be effective. Your FortiGate unit will be unaware of connections and treat each packet individually. It will become a stateless firewall.

Layer-2 and Arp traffic

By default, FortiGate units do not pass layer-2 traffic. If there are layer-2 protocols such as IPX, PPTP or L2TP in use on your network, you need to configure your FortiGate unit interfaces to pass these protocols without blocking. Another type of layer-2 traffic is ARP traffic. For more information on ARP traffic, see [“ARP traffic” on page 532](#).

You can allow these layer-2 protocols using the CLI command:

```
config vdom
  edit <vdom_name>
    config system interface
      edit <name_str>
        set l2forward enable
      end
    end
```

where `<name_str>` is the name of an interface.

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem. If you enable layer-2 traffic, you may experience a problem if packets are allowed to repeatedly loop through the network. This repeated looping, very similar to a broadcast storm, occurs when you have more than one layer-2 path to a destination. Traffic may overflow and bring your network to a halt. You can break the loop by enabling Spanning Tree Protocol (STP) on your network's switches and routers. For more information, see [“STP forwarding” on page 1262](#).

ARP traffic

Address Resolution Protocol (ARP) packets are vital to communication on a network, and ARP support is enabled on FortiGate unit interfaces by default. Normally you want ARP packets to pass through the FortiGate unit, especially if it is sitting between a client and a server or between a client and a router.

ARP traffic can cause problems, especially in transparent mode where ARP packets arriving on one interface are sent to all other interfaces including VLAN subinterfaces. Some layer-2 switches become unstable when they detect the same MAC address originating on more than one switch interface or from more than one VLAN. This instability can occur if the layer-2 switch does not maintain separate MAC address tables for each VLAN. Unstable switches may reset and cause network traffic to slow down considerably.

Multiple VDOMs solution

By default, physical interfaces are in the root domain. If you do not configure any of your VLANs in the root VDOM, it will not matter how many interfaces are in the root VDOM.

The multiple VDOMs solution is to configure multiple VDOMs on the FortiGate unit, one for each VLAN. In this solution, you configure one inbound and one outbound VLAN interface in each VDOM. ARP packets are not forwarded between VDOMs. This configuration limits the VLANs in a VDOM and correspondingly reduces the administration needed per VDOM.

As a result of this configuration, the switches do not receive multiple ARP packets with duplicate MACs. Instead, the switches receive ARP packets with different VLAN IDs and different MACs. Your switches are stable.

However, you should **not** use the multiple VDOMs solution under any of the following conditions:

- you have more VLANs than licensed VDOMs
- you do not have enough physical interfaces

Instead, use one of two possible solutions, depending on which operation mode you are using:

- In NAT mode, you can use the `vlan forward` CLI command.
- In transparent mode, you can use the `forward-domain` CLI command. But you still need to be careful in some rare configurations.

Vlanforward solution

If you are using NAT mode, the solution is to use the `vlanforward` CLI command for the interface in question. By default, this command is enabled and will forward VLAN traffic to all VLANs on this interface. When disabled, each VLAN on this physical interface can send traffic only to the same VLAN. There is no "cross-talk" between VLANs, and ARP packets are forced to take one path along the network which prevents the multiple paths problem.

In the following example, `vlanforward` is disabled on `port1`. All VLANs configured on `port1` will be separate and will not forward any traffic to each other.

```
config system interface
  edit port1
    set vlanforward disable
end
```

Forward-domain solution

If you are using transparent mode, the solution is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic. It is like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0. The many benefits of this solution include reduced administration, the need for fewer physical interfaces, and the availability of more flexible network solutions.

In the following example, forward-domain collision group 340 includes VLAN 340 traffic on port1 and untagged traffic on port 2. Forward-domain collision group 341 includes VLAN 341 traffic on port 1 and untagged traffic on port 3. All other interfaces are part of forward-domain collision group 0 by default. This configuration separates VLANs 340 and 341 from each other on port 1, and prevents the ARP packet problems from before.

Use these CLI commands:

```
config system interface
  edit port1
  next
  edit port2
    set forward_domain 340
  next
  edit port3
    set forward_domain 341
  next
  edit port1-340
    set forward_domain 340
    set interface port1
    set vlanid 340
  next
  edit port1-341
    set forward_domain 341
    set interface port1
    set vlanid 341
end
```

You may experience connection issues with layer-2 traffic, such as ping, if your network configuration has:

- packets going through the FortiGate unit in transparent mode more than once
- more than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled.

Now IPS and AV is applied the first time packets go through the FortiGate unit, but not on subsequent passes. Only applying IPS and AV to this first pass fixes the network layer-2 related connection issues.

NetBIOS

Computers running Microsoft Windows operating systems that are connected through a network rely on a WINS server to resolve host names to IP addresses. The hosts communicate with the WINS server by using the NetBIOS protocol.

To support this type of network, you need to enable the forwarding of NetBIOS requests to a WINS server. The following example will forward NetBIOS requests on the internal interface for the WINS server located at an IP address of 192.168.111.222.

```
config system interface
  edit internal
    set netbios_forward enable
    set wins-ip 192.168.111.222
  end
```

These commands apply only in NAT mode. If VDOMs are enabled, these commands are per VDOM. You must set them for each VDOM that has the problem.

STP forwarding

The FortiGate unit does not participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, reducing available bandwidth to nothing.

If you use your FortiGate unit in a network topology that relies on STP for network loop protection, you need to make changes to your FortiGate configuration. Otherwise, STP recognizes your FortiGate unit as a blocked link and forwards the data to another path. By default, your FortiGate unit blocks STP as well as other non-IP protocol traffic.

Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the external interface:

```
config system interface
  edit external
    set l2forward enable
    set stpforward enable
  end
```

By substituting different commands for `stpforward enable`, you can also allow layer-2 protocols such as IPX, PPTP or L2TP to be used on the network. For more information, see [“Layer-2 and Arp traffic” on page 532](#).

Too many VLAN interfaces

Any virtual domain can have a maximum of 255 interfaces in transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. NAT mode supports from 255 to 8192 depending on the FortiGate model. This total number of interfaces includes VLANs, other virtual interfaces, and physical interfaces.

Your FortiGate unit may allow you to configure more interfaces than this. However, if you configure more than 255 interfaces, your system will become unstable and, over time, will not work properly. As all interfaces are used, they will overflow the routing table that stores the interface information, and connections will fail. When you try to add more interfaces, an error message will state that the maximum limit has already been reached.

If you see this error message, chances are you already have too many VLANs on your system and your routing has become unstable. To verify, delete a VLAN and try to add it back. If you have too many, you will not be able to add it back on to the system. In this case, you will need to remove enough interfaces (including VLANs) so that the total number of interfaces drops to 255 or less. After doing this, you should also reboot your FortiGate unit to clean up its memory and buffers, or you will continue to experience unstable behavior.

To configure more than 255 interfaces on your FortiGate unit in transparent mode, you have to configure multiple VDOMs, each with many VLANs. However, if you want to create more than the default 10 VDOMs (or a maximum of 2550 interfaces), you must buy a license for additional VDOMs. Only FortiGate models 3000 and higher support more than 10 VDOMs.

With these extra licenses, you can configure up to 500 VDOMs, with each VDOM containing up to 255 VLANs in transparent mode. This is a theoretical maximum of over 127 500 interfaces. However, system resources will quickly get used up before reaching that theoretical maximum. To achieve the maximum number of VDOMs, you need to have top-end hardware with the most resources possible.

In NAT mode, if you have a top-end model, the maximum interfaces per VDOM can be as high as 8192, enough for all the VLANs in your configuration.



Your FortiGate unit has limited resources, such as CPU load and memory, that are divided between all configured VDOMs. When running 250 or more VDOMs, you may need to monitor the system resources to ensure there is enough to support the configured traffic processing.



PPTP and L2TP

A virtual private network (VPN) is a way to use a public network, such as the Internet, as a vehicle to provide remote offices or individual users with secure access to private networks. FortiOS supports the Point-to-Point Tunneling Protocol (PPTP), which enables interoperability between FortiGate units and Windows or Linux PPTP clients. Because FortiGate units support industry standard PPTP VPN technologies, you can configure a PPTP VPN between a FortiGate unit and most third-party PPTP VPN peers.

This section describes how to configure PPTP and L2TP VPNs as well as PPTP passthrough.

This section includes the topics:

- [How PPTP VPNs work](#)
- [FortiGate unit as a PPTP server](#)
- [Configuring the FortiGate unit for PPTP VPN](#)
- [Configuring the FortiGate unit for PPTP pass through](#)
- [Testing PPTP VPN connections](#)
- [Logging VPN events](#)
- [Configuring L2TP VPNs](#)
- [L2TP configuration overview](#)

How PPTP VPNs work

The Point-to-Point Tunneling Protocol enables you to create a VPN between a remote client and your internal network. Because it is a Microsoft Windows standard, PPTP does not require third-party software on the client computer. As long as the ISP supports PPTP on its servers, you can create a secure connection by making relatively simple configuration changes to the client computer and the FortiGate unit.

PPTP uses Point-to-Point protocol (PPP) authentication protocols so that standard PPP software can operate on tunneled PPP links. PPTP packages data in PPP packets and then encapsulates the PPP packets within IP packets for transmission through a VPN tunnel.

When the FortiGate unit acts as a PPTP server, a PPTP session and tunnel is created as soon as the PPTP client connects to the FortiGate unit. More than one PPTP session can be supported on the same tunnel. FortiGate units support PAP, CHAP, and plain text authentication. PPTP clients are authenticated as members of a user group.

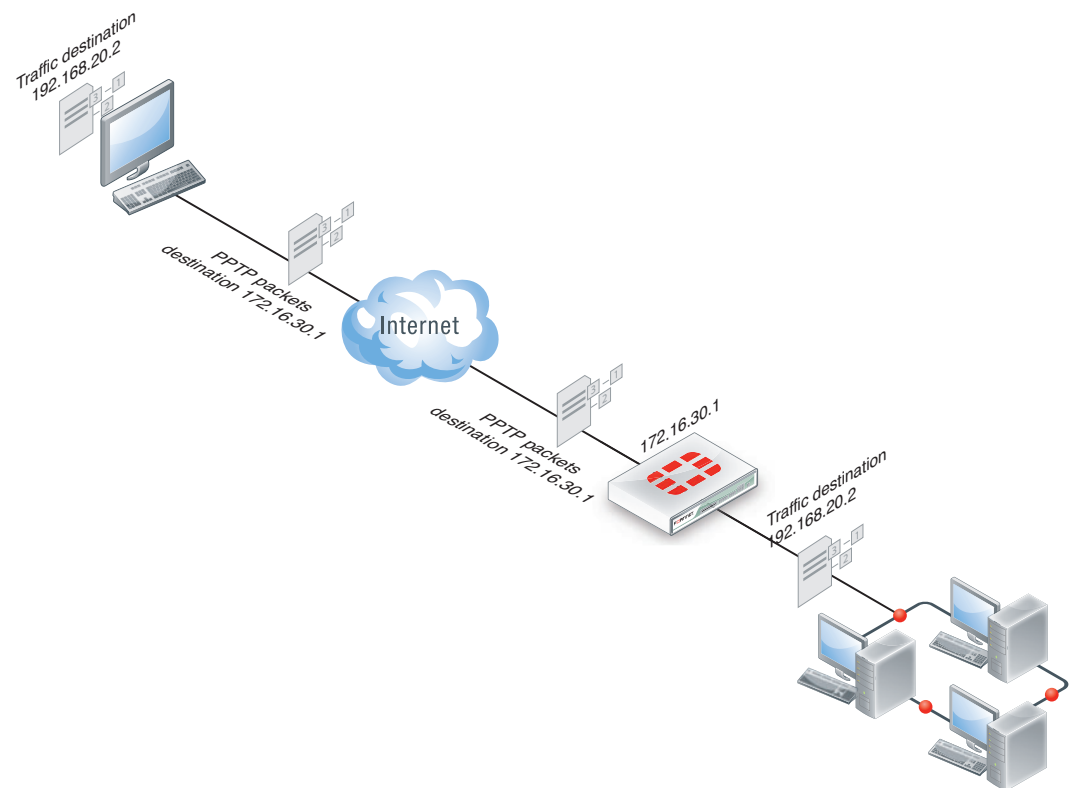
Traffic from one PPTP peer is encrypted using PPP before it is encapsulated using Generic Routing Encapsulation (GRE) and routed to the other PPTP peer through an ISP network. PPP packets from the remote client are addressed to a computer on the private network behind the FortiGate unit. PPTP packets from the remote client are addressed to the public interface of the FortiGate unit. See [Figure 55 on page 538](#).



PPTP control channel messages are not authenticated, and their integrity is not protected. Furthermore, encapsulated PPP packets are not cryptographically protected and may be read or modified unless appropriate encryption software such as Secure Shell (SSH) or Secure File Transfer Protocol (SFTP) is used to transfer data after the tunnel has been established.

As an alternative, you can use encryption software such as Microsoft Point-to-Point Encryption (MPPE) to secure the channel. MPPE is built into Microsoft Windows clients and can be installed on Linux clients. FortiGate units support MPPE.

Figure 55: Packet encapsulation



In [Figure 55](#), traffic from the remote client is addressed to a computer on the network behind the FortiGate unit. When the PPTP tunnel is established, packets from the remote client are encapsulated and addressed to the FortiGate unit. The FortiGate unit forwards disassembled packets to the computer on the internal network.

When the remote PPTP client connects, the FortiGate unit assigns an IP address from a reserved range of IP addresses to the client PPTP interface. The PPTP client uses the assigned IP address as its source address for the duration of the connection.

When the FortiGate unit receives a PPTP packet, the unit disassembles the PPTP packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.




PPTP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate PPTP clients. All PPTP clients are challenged when a connection attempt is made.

FortiGate unit as a PPTP server

In the most common Internet scenario, the PPTP client connects to an ISP that offers PPP connections with dynamically-assigned IP addresses. The ISP forwards PPTP packets to the Internet, where they are routed to the FortiGate unit.

Figure 56: FortiGate unit as a PPTP server



If the FortiGate unit will act as a PPTP server, there are a number of steps to complete:

- Configure user authentication for PPTP clients.
- Enable PPTP.
- Specify the range of addresses that are assigned to PPTP clients when connecting
- Configure the security policy.

Configuring user authentication for PPTP clients

To enable authentication for PPTP clients, you must create user accounts and a user group to identify the PPTP clients that need access to the network behind the FortiGate unit. Within the user group, you must add a user for each PPTP client.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS, LDAP, or TACACS+ server. If password protection will be provided through a RADIUS, LDAP, or TACACS+ server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

This example creates a basic user/password combination.

Configuring a user account

To add a local user - web-based manager

- 1 Go to *User > User > User* and select *Create New*.
- 2 Enter a *User Name*.
- 3 Enter a *Password* for the user. The password should be at least six characters.
- 4 Select *OK*.

To add a local user - CLI

```
config user local
  edit <username>
    set type password
    set passwd <password>
  end
```

Configuring a user group

To ease configuration, create user groups that contain users in similar categories or departments.

To create a user group - web-based manager

- 1 Go to *User > User Group > User Group* and select *Create New*.
- 2 Enter a *Name* for the group.
- 3 Select the *Type of Firewall*.
- 4 From the *Available Users* list, select the required users and select the right-facing arrow to add them to the *Members* list.
- 5 Select *OK*.

To create a user group - CLI

```
config user group
  edit <group_name>
    set group-type firewall
    set members <user_names>
  end
```

Enabling PPTP and specifying the PPTP IP address range

The PPTP address range specifies the range of addresses reserved for remote PPTP clients. When a PPTP client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the PPTP client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the PPTP client appear to be part of the internal network.

PPTP requires two IP addresses, one for each end of the tunnel. The PPTP address range is the range of addresses reserved for remote PPTP clients. When the remote PPTP client establishes a connection, the FortiGate unit assigns an IP address from the reserved range of IP addresses to the client PPTP interface or retrieves the assigned IP address from the PPTP user group. If you use the PPTP user group, you must also define the FortiGate end of the tunnel by entering the IP address of the unit in *Local IP* (web-based manager) or `local-ip` (CLI). The PPTP client uses the assigned IP address as its source address for the duration of the connection.

PPTP configuration is only available through the CLI. In the example below, PPTP is enabled with the use of an IP range of 192.168.1.1 to 192.168.1.10 for addressing.



The start and end IPs in the PPTP address range must be in the same 24-bit subnet, for example, 192.168.1.1 - 192.168.1.254.

```
config vpn pptp
  set status enable
  set ip-mode range
  set eip 192.168.1.10
  set sip 192.168.1.1
end
```

In this example, PPTP is enabled with the use of a user group for addressing, where the IP address of the PPTP server is 192.168.1.2 and the user group is `hr_admin`.

```
config vpn pptp
  set status enable
  set ip-mode range
  set local-ip 192.168.2.1
  set usrgrp hr_admin
end
```

Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the PPTP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

To configure the firewall for the PPTP tunnel - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.

2 Complete the following and select OK:

Source Interface/Zone	The FortiGate interface connected to the Internet.
Source Address	Select the name that corresponds to the range of addresses that you reserved for PPTP clients
Destination Interface/Zone	The FortiGate interface connected to the internal network.
Destination Address	Select the name that corresponds to the IP addresses behind the FortiGate unit
Schedule	always
Service	ANY
Action	ACCEPT



Do not select identity-based policy, as this will cause the PPTP access to fail. Authentication is configured in the PPTP configuration setup.

To configure the firewall for the PPTP tunnel - CLI

```
config firewall policy
  edit 1
    set srcintf <interface to internet>
    set dstintf <interface to internal network>
    set srcaddr <reserved_range>
    set dstaddr <internal_addresses>
    set action accept
    set schedule always
    set service ANY
  end
```

Configuring the FortiGate unit for PPTP VPN

To arrange for PPTP packets to pass through the FortiGate unit to an external PPTP server, perform the following tasks in the order given:

- Configure user authentication for PPTP clients.
- Enable PPTP on the FortiGate unit and specify the range of addresses that can be assigned to PPTP clients when they connect.
- Configure PPTP pass through on the FortiGate unit.

Configuring the FortiGate unit for PPTP pass through

To forward PPTP packets to a PPTP server on the network behind the FortiGate unit, you need to perform the following configuration tasks on the FortiGate unit:

- Define a virtual IP address that points to the PPTP server.

- Create a security policy that allows incoming PPTP packets to pass through to the PPTP server.



The address range is the external (public) ip address range which requires access to the internal PPTP server through the FortiGate virtual port-forwarding firewall. IP addresses used in this document are fictional and follow the technical documentation guidelines specific to Fortinet. Real external IP addresses are not used.

Configuring a virtual IP address

The virtual IP address will be the address of the PPTP server host.

To define a virtual IP for PPTP pass through - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.
- 3 Enter the name of the VIP, for example, `PPTP_Server`.
- 4 Select the *External Interface* where the packets will be received for the PPTP server.
- 5 Enter the *External IP Address* for the VIP.
- 6 Select *Port Forwarding*.
- 7 Set the *Protocol* to *TCP*.
- 8 Enter the *External Service Port* of `1723`, the default for PPTP.
- 9 Enter the *Map to Port* to `1723`.
- 10 Select *OK*.

To define a virtual IP for PPTP pass through - web-based manager

```
config firewall vip
  edit PPTP_Server
    set extintf <interface>
    set extip <ip_address>
    set portforward enable
    set protocol tcp
    set extport 1723
    set mappedport 1723
end
```

Configuring a port-forwarding security policy

To create a port-forwarding security policy for PPTP pass through you must first create an address range reserved for the PPTP clients.

To create an address range - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.
- 2 Enter a *Name* for the range, for example, `External_PPTP`.
- 3 Select a *Type* of *Subnet/IP Range*.
- 4 Enter the IP address range.
- 5 Select the *Interface* to the Internet.
- 6 Select *OK*.

To create an address range - CLI

```

config firewall address
  edit External_PPTP
    set iprange <ip_range>
    set start-ip <ip_address>
    set end-ip <ip_address>
    set associated-interface <internet_interface>
  end

```

With the address set, you can add the security policy.

To add the security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following and select *OK*:

Source Interface/Zone	The FortiGate interface connected to the Internet.
Source Address	Select the address range created in the previous step.
Destination Interface/Zone	The FortiGate interface connected to the PPTP server.
Destination Address	Select the VIP address created in the previous steps.
Schedule	always
Service	PPTP
Action	ACCEPT

To add the security policy - CLI

```

config firewall policy
  edit <policy_number>
    set srcintf <interface to internet>
    set dstintf <interface to PPTP server>
    set srcaddr <address_range>
    set dstaddr <PPTP_server_address>
    set action accept
    set schedule always
    set service PPTP
  end

```

Testing PPTP VPN connections

To confirm that a PPTP VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The PPTP VPN tunnel initializes when the dialup client attempts to connect.

Logging VPN events

PPTP VPN, activity is logged when enabling VPN logging. The FortiGate unit connection events and tunnel status (up/down) are logged.

To log VPN events

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Enable the storage of log messages to one or more locations.
- 3 Select *L2TP/PPTP/PPPoE* service event.
- 4 Select *Apply*.

To view event logs

- 1 Go to *Log&Report > Log & Archive Access > Event Log*.
- 2 If the option is available from the Log Type list, select the log file from disk or memory.

Configuring L2TP VPNs

This section describes how to configure a FortiGate unit to establish a Layer Two Tunneling Protocol (L2TP) tunnel with a remote dialup client. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly.

According to RFC 2661, an Access Concentrator (LAC) can establish an L2TP tunnel with an L2TP Network Server (LNS). In a typical scenario, the LAC is managed by an ISP and located on the ISP premises; the LNS is the gateway to a private network. When a remote dialup client connects to the Internet through the ISP, the ISP uses a local database to establish the identity of the caller and determine whether the caller needs access to an LNS through an L2TP tunnel. If the services registered to the caller indicate that an L2TP connection to the LNS is required, the ISP LAC attempts to establish an L2TP tunnel with the LNS.

A FortiGate unit can be configured to act as an LNS. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly, bypassing any LAC managed by an ISP. The ISP must configure its network access server to forward L2TP traffic from the remote client to the FortiGate unit directly whenever the remote client requires an L2TP connection to the FortiGate unit.

When the FortiGate unit acts as an LNS, an L2TP session and tunnel is created as soon as the remote client connects to the FortiGate unit. The FortiGate unit assigns an IP address to the client from a reserved range of IP addresses. The remote client uses the assigned IP address as its source address for the duration of the connection.

More than one L2TP session can be supported on the same tunnel. FortiGate units can be configured to authenticate remote clients using a plain text user name and password, or authentication can be forwarded to an external RADIUS or LDAP server. L2TP clients are authenticated as members of a user group.




FortiGate units support L2TP with Microsoft Point-to-Point Encryption (MPPE) encryption only. Later implementations of Microsoft L2TP for Windows use IPsec and require certificates for authentication and encryption. If you want to use Microsoft L2TP with IPsec to connect to a FortiGate unit, the IPsec and certificate elements must be disabled on the remote client

Traffic from the remote client must be encrypted using MPPE before it is encapsulated and routed to the FortiGate unit. Packets originating at the remote client are addressed to a computer on the private network behind the FortiGate unit. Encapsulated packets are addressed to the public interface of the FortiGate unit. See [Figure 57](#).

When the FortiGate unit receives an L2TP packet, the unit disassembles the packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.


Figure 57: L2TP encapsulation



Fortinet units cannot deliver non-IP traffic such as Frame Relay or ATM frames encapsulated in L2TP packets — FortiGate units support the IPv4 and IPv6 addressing schemes only.

Network topology

The remote client connects to an ISP that determines whether the client requires an L2TP connection to the FortiGate unit. If an L2TP connection is required, the connection request is forwarded to the FortiGate unit directly.

Figure 58: Example L2TP configuration

L2TP infrastructure requirements

- The FortiGate unit must be operating in NAT mode and have a static public IP address.
- The ISP must configure its network access server to forward L2TP traffic from remote clients to the FortiGate unit directly.
- The remote client must not generate non-IP traffic (Frame Relay or ATM frames).
- The remote client includes L2TP support with MPPE encryption. If the remote client includes Microsoft L2TP with IPsec, the IPsec and certificate components must be disabled.

L2TP configuration overview

To configure a FortiGate unit to act as an LNS, you perform the following tasks on the FortiGate unit:

- Create an L2TP user group containing one user for each remote client.
- Enable L2TP on the FortiGate unit and specify the range of addresses that can be assigned to remote clients when they connect.
- Define firewall source and destination addresses to indicate where packets transported through the L2TP tunnel will originate and be delivered.
- Create the security policy and define the scope of permitted services between the source and destination addresses.
- Configure the remote clients.

Authenticating L2TP clients

L2TP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate L2TP clients. All L2TP clients are challenged when a connection attempt is made.

To enable authentication, you must create user accounts and a user group to identify the L2TP clients that need access to the network behind the FortiGate unit.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS or LDAP server. If password protection will be provided through a RADIUS or LDAP server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

Enabling L2TP and specifying an address range

The L2TP address range specifies the range of addresses reserved for remote clients. When a remote client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the remote client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the remote client appear to be part of the internal network.

To enable L2TP and specify the L2TP address range, use the `config vpn l2tp` CLI command.

The following example shows how to enable L2TP and set the L2TP address range using a starting address of 192.168.10.80 and an ending address of 192.168.10.100 for an existing group of L2TP users named `L2TP_users`:

```
config vpn l2tp
  set sip 192.168.10.80
  set eip 192.168.10.100
  set status enable
  set usrgroup L2TP_users
end
```

Defining firewall source and destination addresses

Before you define the security policy, you must define the source and destination addresses of packets that are to be transported through the L2TP tunnel:

- For the source address, enter the range of addresses that you reserved for remote L2TP clients (for example 192.168.10.[80-100]).
- For the destination address, enter the IP addresses of the computers that the L2TP clients need to access on the private network behind the FortiGate unit (for example, 172.16.5.0/24 for a subnet, or 172.16.5.1 for a server or host, or 192.168.10.[10-15] for an IP address range).

To define the firewall source address

- 1 Go to *Firewall Objects > Address* and select *Create New*.
- 2 In the *Address Name* field, type a name that represents the range of addresses that you reserved for remote clients (for example, `Ext_L2TPrange`).
- 3 In *Type*, select *Subnet / IP Range*.
- 4 In the *Subnet / IP Range* field, type the corresponding IP address range.

- 5 In *Interface*, select the FortiGate interface that connects to the clients.
This is usually the interface that connects to the Internet.
- 6 Select *OK*.

To define the firewall destination address

- 1 Go to *Firewall Objects > Address* and select *Create New*.
- 2 In the *Address Name* field, type a name that represents a range of IP addresses on the network behind the FortiGate unit (for example, *Int_L2TPaccess*).
- 3 In *Type*, select *Subnet / IP Range*.
- 4 In the *Subnet / IP Range* field, type the corresponding IP address range.
- 5 In *Interface*, select the FortiGate interface that connects to the network behind the FortiGate unit.
- 6 Select *OK*.

Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the L2TP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

To define the traffic and services permitted inside the L2TP tunnel

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter these settings in particular:

Source Interface/Zone	Select the FortiGate interface to the Internet.
Source Address	Select the name that corresponds to the range of addresses that you reserved for L2TP clients (for example, <i>Ext_L2TPrange</i>).
Destination Interface/Zone	Select the FortiGate interface to the internal (private) network.
Destination Address	Select the name that corresponds to the IP addresses behind the FortiGate unit (for example, <i>Int_L2TPaccess</i>).
Service	Select ANY, or if selected services are required instead, select the service group that you defined previously.
Action	Select ACCEPT.

- 3 You may enable NAT, a protection profile, and/or event logging, or select *Enable Identity Based Policy* to add authentication or shape traffic. For more information on identity based policies, see the [Firewall](#) chapter of the handbook.
- 4 Select *OK*.

Configuring a Linux client

The following procedure outlines how to install L2TP client software and run an L2TP tunnel on a Linux computer. Obtain an L2TP client package that meets your requirements (for example, *rp-l2tp*). If needed to encrypt traffic, obtain L2TP client software that supports encryption using MPPE.

To establish an L2TP tunnel with a FortiGate unit that has been set up to accept L2TP connections, you can obtain and install the client software following these guidelines:

- 1 If encryption is required but MPPE support is not already present in the kernel, download and install an MPPE kernel module and reboot your computer.
- 2 Download and install the L2TP client package.
- 3 Configure an L2TP connection to run the L2TP program.
- 4 Configure routes to determine whether all or some of your network traffic will be sent through the tunnel. You must define a route to the remote network over the L2TP link and a host route to the FortiGate unit.
- 5 Run `l2tpd` to start the tunnel.

Follow the software supplier's documentation to complete the steps.

To configure the system, you need to know the public IP address of the FortiGate unit, and the user name and password that has been set up on the FortiGate unit to authenticate L2TP clients. Contact the FortiGate administrator if required to obtain this information.

Monitoring L2TP sessions

You can display a list of all active sessions and view activity by port number. By default, port 1701 is used for L2TP VPN-related communications. If required, active sessions can be stopped from this view. Use the Top Sessions Dashboard Widget.

Testing L2TP VPN connections

To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

Logging L2TP VPN events

You can configure the FortiGate unit to log VPN events. For L2TP VPNs, connection events and tunnel status (up/down) are logged.

To log VPN events - web-based manager

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Enable the storage of log messages to one or more locations.
- 3 Select *Enable*, and then select *L2TP/PPTP/PPPoE service event*.
- 4 Select *Apply*.

To log VPN events - CLI

```
config log memory setting
    set diskfull overright
    set status enable
end
config log eventfilter
    set ppp
end
```



Session helpers

The FortiOS firewall can analyze most TCP/IP protocol traffic by comparing packet header information to security policies. This comparison determines whether to accept or deny the packet and the session that the packet belongs to.

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. But the packets that carry the actual conversation can use a variety of UDP protocols with a variety of source and destination port numbers. The information about the protocols and port numbers used for a SIP call is contained in the body of the SIP TCP control packets. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and adjust the firewall to allow those protocols to send packets through the firewall.

This section includes the topics:

- [Viewing the session helper configuration](#)
- [Changing the session helper configuration](#)
- [DCE-RPC session helper \(dcerpc\)](#)
- [DNS session helpers \(dns-tcp and dns-udp\)](#)
- [File transfer protocol \(FTP\) session helper \(ftp\)](#)
- [H.245 session helpers \(h245I and h245O\)](#)
- [H.323 and RAS session helpers \(h323 and ras\)](#)
- [Media Gateway Controller Protocol \(MGCP\) session helper \(mgcp\)](#)
- [ONC-RPC portmapper session helper \(pmap\)](#)
- [PPTP session helper for PPTP traffic \(pptp\)](#)
- [Remote shell session helper \(rsh\)](#)
- [Real-Time Streaming Protocol \(RTSP\) session helper \(rtsp\)](#)
- [Session Initiation Protocol \(SIP\) session helper \(sip\)](#)
- [Trivial File Transfer Protocol \(TFTP\) session helper \(tftp\)](#)
- [Oracle TNS listener session helper \(tns\)](#)

Viewing the session helper configuration

You can view the session helpers enabled on your FortiGate unit in the CLI using the commands below. The following output shows the first two session helpers. The number of session helpers can vary to around 20.

```
show system session-helper
config system session-helper
edit 1
    set name pptp
    set port 1723
    set protocol 6
end
next
    set name h323
    set port 1720
    set protocol 6
next
end
.
```

The configuration for each session helper includes the name of the session helper and the port and protocol number on which the session helper listens for sessions. Session helpers listed on protocol number 6 (TCP) or 17 (UDP). For a complete list of protocol numbers see: [Assigned Internet Protocol Numbers](#).

For example, the output above shows that FortiOS listens for PPTP packets on TCP port 1723 and H.323 packets on port TCP port 1720.

If a session helper listens on more than one port or protocol the more than one entry for the session helper appears in the `config system session-helper` list. For example, the pmap session helper appears twice because it listens on TCP port 111 and UDP port 111. The rsh session helper appears twice because it listens on TCP ports 514 and 512.

Changing the session helper configuration

Normally you will not need to change the configuration of the session helpers. However in some cases you may need to change the protocol or port the session helper listens on.

Changing the protocol or port that a session helper listens on

Most session helpers are configured to listen for their sessions on the port and protocol that they typically use. If your FortiGate unit receives sessions that should be handled by a session helper on a non-standard port or protocol you can use the following procedure to change the port and protocol used by a session helper. The following example shows how to change the port that the pmap session helper listens on for Sun RPC portmapper TCP sessions. By default pmap listens on TCP port 111.

To change the port that the pmap session helper listens on to TCP port 112

- 1 Confirm that the TCP pmap session helper entry is 11 in the session-helper list:

```
show system session-helper 11
config system session-helper
edit 11
    set name pmap
    set port 111
    set protocol 6
next
end
```

- 2 Enter the following command to change the TCP port to 112.

```
config system session-helper
edit 11
set port 112
end
```

- 3 The pmap session helper also listens on UDP port 111. Confirm that the UDP pmap session helper entry is 12 in the session-helper list:

```
show system session-helper 12
config system session-helper
edit 12
set name pmap
set port 111
set protocol 17
next
end
```

- 4 Enter the following command to change the UDP port to 112.

```
config system session-helper
edit 12
set port 112
end
end
```

Use the following command to set the h323 session helper to listen for ports on the UDP protocol.

To change the protocol that the h323 session helper listens on

- 1 Confirm that the h323 session helper entry is 2 in the session-helper list:

```
show system session-helper 2
config system session-helper
edit 2
set name h323
set port 1720
set protocol 6
next
end
```

- 2 Enter the following command to change the protocol to UDP.

```
config system session-helper
edit 2
set protocol 17
end
end
```

If a session helper listens on more than one port or protocol, then multiple entries for the session helper must be added to the session helper list, one for each port and protocol combination. For example, the rtsp session helper listens on TCP ports 554, 7070, and 8554 so there are three rtsp entries in the session-helper list. If your FortiGate unit receives rtsp packets on a different TCP port (for example, 6677) you can use the following command to configure the rtsp session helper to listen on TCP port 6677.

To configure a session helper to listen on a new port and protocol

```
config system session-helper
edit 0
    set name rtsp
    set port 6677
    set protocol 6
end
```

Disabling a session helper

In some cases you may need to disable a session helper. Disabling a session helper just means removing it from the session-helper list so that the session helper is not listening on a port. You can completely disable a session helper by deleting all of its entries from the session helper list. If there are multiple entries for a session helper on the list you can delete one of the entries to prevent the session helper from listening on that port.

To disable the mgcp session helper from listening on UDP port 2427

- 1 Enter the following command to find the mgcp session helper entry that listens on UDP port 2427:

```
show system session-helper
.
.
.
edit 19
    set name mgcp
    set port 2427
    set protocol 17
next
.
.
.
```

- 2 Enter the following command to delete session-helper list entry number 19 to disable the mgcp session helper from listening on UDP port 2427:

```
config system session-helper
delete 19
```

By default the mgcp session helper listens on UDP ports 2427 and 2727. The previous procedure shows how to disable the mgcp protocol from listening on port 2427. The following procedure completely disables the mgcp session helper by also disabling it from listening on UDP port 2727.

To completely disable the mgcp session helper

- 1 Enter the following command to find the mgcp session helper entry that listens on UDP port 2727:

```
show system session-helper
.
.
.
edit 20
    set name mgcp
    set port 2727
    set protocol 17
```



```
next
.
```

- 2 Enter the following command to delete session-helper list entry number 20 to disable the mgcp session helper from listening on UDP port 2727:

```
config system session-helper
delete 20
```

DCE-RPC session helper (dcerpc)

Distributed Computing Environment Remote Procedure Call (DCE-RPC) provides a way for a program running on one host to call procedures in a program running on another host. DCE-RPC (also called MS RPC for Microsoft RPC) is similar to ONC-RPC. Because of the large number of RPC services, for example, MAPI, the transport address of an RPC service is dynamically negotiated based on the service program's universal unique identifier (UUID). The Endpoint Mapper (EPM) binding protocol in FortiOS maps the specific UUID to a transport address.

To accept DCE-RPC sessions you must add a security policy with service set to any or to the DCE-RPC pre-defined service (which listens on TCP and UDP ports 135). The dcerpc session helper also listens on TCP and UDP ports 135.

The session allows FortiOS to handle DCE-RPC dynamic transport address negotiation and to ensure UUID-based security policy enforcement. You can define a security policy to permit all RPC requests or to permit by specific UUID number.

In addition, because a TCP segment in a DCE-RPC stream might be fragmented, it might not include an intact RPC PDU. This fragmentation occurs in the RPC layer; so FortiOS does not support parsing fragmented packets.

DNS session helpers (dns-tcp and dns-udp)

FortiOS includes two DNS session helpers, dns-tcp, a session helper for DNS over TCP, and dns-udp, a session helper for DNS over UDP. The DNS session helpers monitor DNS query and reply packets and close sessions if the DNS flag indicates the packet is a reply message.

To accept DNS sessions you must add a security policy with service set to any or to the DNS pre-defined service (which listens on TCP and UDP ports 53). The dns-udp session helper also listens on UDP port 53. By default the dns-tcp session helper is disabled. If needed you can use the following command to enable the dns-tcp session helper to listen for DNS sessions on TCP port 53:

```
config system session-helper
edit 0
set name dns-tcp
set port 53
set protocol 6
end
```

File transfer protocol (FTP) session helper (ftp)

The FTP session helper monitors PORT, PASV and 227 commands and NATs the IP addresses and port numbers in the body of the FTP packets and opens ports on the FortiGate unit as required.

To accept FTP sessions you must add a security policy with service set to any or to the FTP, FTP_Put, and FTP_GET pre-defined services (which all listen on TCP port 21).

H.245 session helpers (h245I and h245O)

H.245 is a control channel protocol used for H.323 and other similar communication sessions. H.245 sessions transmit non-telephone signals. H.245 sessions carry information needed for multimedia communication, such as encryption, flow control jitter management and others.

FortiOS includes two H.245 sessions helpers, h245I which is for H.245 call in and h245O which is for H.245 call out sessions. There is no standard port for H.245. By default the H.245 sessions helpers are disabled. You can enable them as you would any other session helper. When you enable them, you should specify the port and protocol on which the FortiGate unit receives H.245 sessions.

H.323 and RAS session helpers (h323 and ras)

The H.323 session helper supports secure H.323 voice over IP (VoIP) sessions between terminal endpoints such as IP phones and multimedia devices. In H.323 VoIP networks, gatekeeper devices manage call registration, admission, and call status for VoIP calls. The FortiOS h323 session helper supports gatekeepers installed on two different networks or on the same network.

To accept H.323 sessions you must add a security policy with service set to any or to the H323 pre-defined service (which listens on TCP port numbers 1720 and 1503 and on UDP port number 1719). The h323 session helper listens on TCP port 1720.

The ras session helper is used with the h323 session helper for H.323 Registration, Admission, and Status (RAS) services. The ras session helper listens on UDP port 1719.

Alternate H.323 gatekeepers

The h323 session helper supports using H.323 alternate gatekeepers. All the H.323 end points must register with a gatekeeper through the Registration, Admission, and Status (RAS) protocol before they make calls. During the registration process, the primary gatekeeper sends Gatekeeper Confirm (GCF) and Registration Confirm (RCF) messages to the H.323 end points that contain the list of available alternate gatekeepers.

The alternate gatekeeper provides redundancy and scalability for the H.323 end points. If the primary gatekeeper fails the H.323 end points that have registered with that gatekeeper are automatically registered with the alternate gatekeeper. To use the H.323 alternate gatekeeper, you need to configure security policies that allow H.323 end points to reach the alternate gatekeeper.

Media Gateway Controller Protocol (MGCP) session helper (mgcp)

The Media Gateway Control Protocol (MGCP) is a text-based application layer protocol used for VoIP call setup and control. MGCP uses a master-slave call control architecture in which the media gateway controller uses a call agent to maintain call control intelligence, while the media gateways perform the instructions of the call agent.

To accept MGCP sessions you must add a security policy with service set to any or to the MGCP pre-defined service (which listens on UDP port numbers 2427 and 2727). The h323 session helper also listens on UDP port numbers 2427 and 2727.

The MGCP session helper does the following:

- VoIP signalling payload inspection. The payload of the incoming VoIP signalling packet is inspected and malformed packets are blocked.
- Signaling packet body inspection. The payload of the incoming MGCP signaling packet is inspected according to RFC 3435. Malformed packets are blocked.
- Stateful processing of MGCP sessions. State machines are invoked to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- MGCP Network Address Translation (NAT). Embedded IP addresses and ports in packet bodies is properly translated based on current routing information and network topology, and is replaced with the translated IP address and port number, if necessary.
- Manages pinholes for VoIP traffic. To keep the VoIP network secure, the IP address and port information used for media or signalling is identified by the session helper, and pinholes are dynamically created and closed during call setup.

ONC-RPC portmapper session helper (pmap)

Open Network Computing Remote Procedure Call (ONC-RPC) is a widely deployed remote procedure call system. Also called Sun RPC, ONC-RPC allows a program running on one host to call a program running on another. The transport address of an ONC-RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

To accept ONC-RPC sessions you must add a security policy with service set to any or to the ONC-RPC pre-defined service (which listens on TCP and UDP port number 111). The RPC portmapper session helper (called pmap) handles the dynamic transport address negotiation mechanisms of ONC-RPC.

PPTP session helper for PPTP traffic (pptp)

The PPTP session help supports port address translation (PAT) for PPTP traffic. PPTP provides IP security at the Network Layer. PPTP consists of a control session and a data tunnel. The control session runs over TCP and helps in establishing and disconnecting the data tunnel. The data tunnel handles encapsulated Point-to-Point Protocol (PPP) packets carried over IP.

To accept PPTP sessions that pass through the FortiGate unit you must add a security policy with service set to any or to the PPTP pre-defined service (which listens on IP port 47 and TCP port 1723). The pptp session helper listens on TCP port 1723.

PPTP uses TCP port 1723 for control sessions and Generic Routing Encapsulation (GRE) (IP protocol 47) for tunneling the encapsulated PPP data. The GRE traffic carries no port number, making it difficult to distinguish between two clients with the same public IP address. PPTP uses the source IP address and the Call ID field in the GRE header to identify a tunnel. When multiple clients sharing the same IP address establish tunnels with the same PPTP server, they may get the same Call ID. The call ID value can be translated in both the control message and the data traffic, but only when the client is in a private network and the server is in a public network.

PPTP clients can either directly connect to the Internet or dial into a network access server to reach the Internet. A FortiGate unit that protects PPTP clients can translate the clients' private IP addresses to a pool of public IP addresses using NAT port translation (NAT-PT). Because the GRE traffic carries no port number for address translation, the pptp session helper treats the Call ID field as a port number as a way of distinguishing multiple clients.

After the PPTP establishing a TCP connection with the PPTP server, the client sends a start control connection request message to establish a control connection. The server replies with a start control connection reply message. The client then sends a request to establish a call and sends an outgoing call request message. FortiOS assigns a Call ID (bytes 12-13 of the control message) that is unique to each PPTP tunnel. The server replies with an outgoing call reply message that carries its own Call ID in bytes 12-13 and the client's call ID in bytes 14-15. The pptp session helper parses the control connection messages for the Call ID to identify the call to which a specific PPP packet belongs. The session helper also identifies an outgoing call request message using the control message type field (bytes 8-9) with the value 7. When the session helper receives this message, it parses the control message for the call ID field (bytes 12-13). FortiOS translates the call ID so that it is unique across multiple calls from the same translated client IP. After receiving outgoing call response message, the session helper holds this message and opens a port that accepts GRE traffic that the PPTP server sends. An outgoing call request message contains the following parts:

- The protocol used for the outgoing call request message (usually GRE)
- Source IP address (PPTP server IP)
- Destination IP address (translated client IP)
- Destination port number (translated client call ID)

The session helper identifies an outgoing call reply message using the control message type field (bytes 8-9) with the value 8. The session helper parses these control messages for the call ID field (bytes 12-13) and the client's call ID (bytes 14-15). The session helper then uses the client's call ID value to find the mapping created for the other direction, and then opens a pinhole to accept the GRE traffic that the client sends.

An outgoing call reply message contains the following parts:

- Protocol used for the outgoing call reply message (usually GRE)
- Source IP address (PPTP client IP)
- Destination IP address (PPTP server IP)
- Destination port number (PPTP server Call ID)

Each port that the session opens creates a session for data traffic arriving in that direction. The session helper opens the following two data sessions for each tunnel:

- Traffic from the PPTP client to the server, using the server's call ID as the destination port
- Traffic from the PPTP server to the client, using the client's translated call ID as the destination port

The default timeout value of the control connection is 30 minutes. The session helper closes the pinhole when the data session exceeds the timeout value or is idle for an extended period.

Remote shell session helper (rsh)

Using the remote shell program (RSH), authenticated users can run shell commands on remote hosts. RSH sessions most often use TCP port 514. To accept RSH sessions you must add a security policy with service set to any or to the RSH pre-defined service (which listens on TCP port number 514).

FortiOS automatically invokes the rsh session helper to process all RSH sessions on TCP port 514. The rsh session helper opens ports required for the RSH service to operate through a FortiGate unit running NAT or transparent and supports port translation of RSH traffic.

Real-Time Streaming Protocol (RTSP) session helper (rtsp)

The Real-Time Streaming Protocol (RTSP) is an application layer protocol often used by SIP to control the delivery of multiple synchronized multimedia streams, for example, related audio and video streams. Although RTSP is capable of delivering the data streams itself it is usually used like a network remote control for multimedia servers. The protocol is intended for selecting delivery channels (like UDP, multicast UDP, and TCP) and for selecting a delivery mechanism based on the Real-Time Protocol (RTP). RTSP may also use the SIP Session Description Protocol (SDP) as a means of providing information to clients for aggregate control of a presentation consisting of streams from one or more servers, and non-aggregate control of a presentation consisting of multiple streams from a single server.

To accept RTSP sessions you must add a security policy with service set to any or to the RTSP pre-defined service (which listens on TCP ports 554, 770, and 8554 and on UDP port 554). The rtsp session helper listens on TCP ports 554, 770, and 8554.

The rtsp session help is required because RTSP uses dynamically assigned port numbers that are communicated in the packet body when end points establish a control connection. The session helper keeps track of the port numbers and opens pinholes as required. In Network Address Translation (NAT) mode, the session helper translates IP addresses and port numbers as necessary.

In a typical RTSP session the client starts the session (for example, when the user selects the Play button on a media player application) and establishes a TCP connection to the RTSP server on port 554. The client then sends an OPTIONS message to find out what audio and video features the server supports. The server responds to the OPTIONS message by specifying the name and version of the server, and a session identifier, for example, 24256-1.

The client then sends the DESCRIBE message with the URL of the actual media file the client wants to play. The server responds to the DESCRIBE message with a description of the media in the form of SDP code. The client then sends the SETUP message, which specifies the transport mechanisms acceptable to the client for streamed media, for example RTP/RTCP or RDT, and the ports on which it receives the media.

In a NAT configuration the rtsp session helper keeps track of these ports and addresses translates them as necessary. The server responds to the SETUP message and selects one of the transport protocols. When both client and server agree on a mechanism for media transport the client sends the PLAY message, and the server begins streaming the media.

Session Initiation Protocol (SIP) session helper (sip)

The sip session helper is described in [“The SIP session helper” on page 2536](#).

Trivial File Transfer Protocol (TFTP) session helper (tftp)

To accept TFTP sessions you must add a security policy with service set to any or to the TFTP pre-defined service (which listens on UDP port number 69). The TFTP session helper also listens on UTP port number 69.

TFTP initiates transfers on UDP port 69, but the actual data transfer ports are selected by the server and client during initialization of the connection. The tftp session helper reads the transfer ports selected by the TFTP client and server during negotiation and opens these ports on the firewall so that the TFTP data transfer can be completed. When the transfer is complete the tftp session helper closes the open ports.

Oracle TNS listener session helper (tns)

The Oracle Transparent Network Substrate (TNS) listener listens on port TCP port 1521 for network requests to be passed to a database instance. The Oracle TNS listener session helper (tns) listens for TNS sessions on TCP port 1521. TNS is a foundation technology built into the Oracle Net foundation layer and used by SQLNET.



Advanced concepts

This chapter provides configuration concepts and techniques to enhance your network security.

This section includes the topics:

- [Dual internet connections](#)
- [Advanced concepts Single firewall vs. multiple virtual domains](#)
- [Modem](#)
- [DHCP servers and relays](#)
- [Assigning IP address by MAC address](#)
- [DNS services](#)
- [Dynamic DNS](#)
- [Aggregate Interfaces](#)
- [IP addresses for self-originated traffic](#)
- [Administration for schools](#)
- [Tag management](#)
- [Software switch](#)
- [Replacement messages list](#)
- [Disk](#)
- [CLI Scripts](#)
- [Rejecting PING requests](#)
- [Opening TCP 113](#)
- [Obfuscate HTTP headers](#)


Dual internet connections

Dual internet connection, dual WAN, or redundant internet connection refers to using two FortiGate interfaces to connect to the Internet. Dual internet connections can be used in three ways:

- redundant interfaces, should one interface go down, the second automatically becomes the main internet connection
- for load sharing to ensure better throughput.
- a combination of redundancy and load sharing.

Redundant interfaces

Redundant interfaces, ensures that should your internet access be no longer available through a certain port, the FortiGate unit will use an alternate port to connect to the Internet.

Figure 59: Configuring redundant interfaces

In this scenario, two interfaces, WAN1 and WAN2 are connected to the Internet using two different ISPs. WAN1 is the primary connection. In an event of a failure of WAN1, WAN2 automatically becomes the connection to the Internet. For this configuration to function correctly, you need to configure three specific settings:

- configure a ping server to determine when the primary interface (WAN1) is down and when the connection returns
- configure a default route for each interface.
- configure security policies to allow traffic through each interface to the internal network.

Ping server

Adding a ping server is required for routing fail over traffic. A ping server will confirm the connectivity of the device's interface

To add a ping server - web-based manager

- 1 Go to *Router > Static > Settings* and select *Create New*.
- 2 Select the *Interface* that will send ping requests.
- 3 For the *Ping Server* field, enter the IP address of a server that the FortiGate unit will send ping requests to. This is typically a next hop router or gateway device.
- 4 Select the *Detect Protocol* type.
- 5 For the *Ping Interval* field, enter the number of seconds to send ping requests.
- 6 For the *Failover Threshold*, enter the number of lost pings is acceptable before the port is determined to be down.
- 7 Select *OK*.

To add a ping server - CLI

```
config router gwdetect
edit wan1
set server <ISP_IP_address>
set failtime <failure_count>
set interval <seconds>
end
```

Routing

You need to configure a default route for each interface and indicate which route is preferred by specifying the distance. The lower distance is declared active and placed higher in the routing table.



When you have dual WAN interfaces that are configured to provide fail over, you might not be able to connect to the backup WAN interface because the FortiGate unit may not route traffic (even responses) out of the backup interface. The FortiGate unit performs a reverse path lookup to prevent spoofed traffic. If no entry can be found in the routing table which sends the return traffic out the same interface, then the incoming traffic is dropped.

To configure the routing of the two interfaces - web-based manager

- 1 Go to *Router > Static > Static Route* and select *Create New*.
- 2 Set the *Destination IP/Mask* to the address and netmask to 0.0.0.0/0.0.0.0.
- 3 Select the *Device* to the primary connection, *WAN1*.
- 4 Enter the *Gateway* address.
- 5 Select *Advanced*.
- 6 Set the *Distance* to 10.
- 7 Select *OK*.
- 8 Repeat steps 1 through 7 setting the *Device* to *WAN2* and a *Distance* of 20.

To configure the routing of the two interfaces - CLI

```
config router static
edit 1
set dst 0.0.0.0 0.0.0.0
set device WAN1
set gateway 0.0.0.0 0.0.0.0
set distance 10
next
edit 1
set dst <ISP_Address>
set device WAN2
set gateway <gateway_address>
set distance 20
next
end
```

Security policies

When creating security policies, you need to configure duplicate policies to ensure that after traffic fails over WAN1, regular traffic will be allowed to pass through WAN2 as it did with WAN1. This ensures that fail-over will occur with minimal affect to users. For more information on creating security policies see the [Firewall Guide](#).

Load sharing

Load sharing enables you to use both connections to the internet at the same time, but do not provide fail over support. When configuring for load sharing, you need to ensure routing is configured for both external ports, for example, WAN1 and WAN2, have static routes with the same distance and priority.

Further configuration can be done using Equal Cost Multiple Path (ECMP). For more information on ECMP and load sharing, see the [Advanced Routing Guide](#).

Link redundancy and load sharing

In this scenario, both links are available to distribute Internet traffic over both links. Should one of the interfaces fail, the FortiGate unit will continue to send traffic over the other active interface. Configuration is similar to the [Redundant interfaces](#) configuration, with the main difference being that the configured routes should have equal distance settings.

This means both routes will remain active in the routing table. To make one interface the preferred interface, use a default policy route to indicate the interface that is preferred for accessing the Internet. If traffic matches the security policy, the policy overrides all entries in the routing table, including connected routes. You may need to add a specific policy routes that override these default policy routes.

To redirect traffic over the secondary interface, create policy routes to direct some traffic onto it rather than the primary interface. When adding the policy route, only define the outgoing interface and leave the gateway blank. This ensures that the policy route will not be active when the link is down.

Single firewall vs. multiple virtual domains

A typical FortiGate setup, with a small to mid-range appliance, enables you to include a number of subnets on your network using the available ports and switch interfaces. This can potentially provide a means of having three or more mini networks for the various groups in a company. Within this infrastructure, multiple network administrators have access to the FortiGate to maintain security policies.

However, the FortiGate unit may not have enough interfaces to match the number of departments in the organization. If the FortiGate unit is running in transparent mode however, there is only one interface, and multiple network branches through the FortiGate are not possible.

A FortiGate unit with Virtual Domains (VDOMs) enabled, provides a means to provide the same functionality in transparent mode as a FortiGate in NAT mode. VDOMs are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs can provide separate security policies and, in NAT mode, completely separate configurations for routing and VPN services for each connected network. For administration, an administrator can be assigned to each VDOM, minimizing the possibility of error or fouling network communications.


By default, your FortiGate unit supports a maximum of 10 VDOMs. For FortiGate models 3000 and higher, you can purchase a license key to increase the number of VDOMs to 25, 50, 100 or 250.




The FortiGate-20C and 30B and FortiWifi-20C and 30B do not support VDOMs.

Single firewall vs. vdoms

When VDOMs are not enabled, and the FortiGate unit is in transparent mode, all the interfaces on your unit become broadcast interfaces. The problem is there are no interfaces free for additional network segments.



A FortiGate with three interfaces means only limited network segments are possible without purchasing more FortiGate devices.



With multiple VDOMs you can have one of them configured in transparent mode, and the rest in NAT mode. In this configuration, you have an available transparent mode FortiGate unit you can drop into your network for troubleshooting, and you also have the standard.

This example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First enable Virtual Domains on the FortiGate unit.

To enable VDOMs - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 In the *System Information* widget, select *Enable for Virtual Domain*.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all Virtual Domains.

To enable VDOMs - CLI

```
config system global
    set vdom-admin enable
end
```

Next, add the VDOM called accounting.

To add a VDOM - web-based manager

- 1 Go to *System > VDOM > VDOM*, and select *Create New*.
- 2 Enter the VDOM name `accounting`.
- 3 Select *OK*.

To add a VDOM - CLI

```
config vdom
    edit <new_vdom_name>
end
```

With the Virtual Domain created, you can assign a physical interface to it, and assign it an IP address.

To assign physical interface to the accounting Virtual Domain - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the DMZ2 port row and select *Edit*.
- 3 For the *Virtual Domain* drop-down list, select *accounting*.
- 4 Select the *Addressing Mode* of *Manual*.
- 5 Enter the IP address for the port of 10.13.101.100/24.
- 6 Set the *Administrative Access* to *HTTPS* and *SSH*.
- 7 Select *OK*.

To assign physical interface to the accounting Virtual Domain - CLI

```
config global
  config system interface
    edit dmz2
      set vdom accounting
      set ip 10.13.101.100/24
      set allowaccess https ssh
    next
  end
```

Modem

FortiGate units support the use of wireless, 3G and 4G modems connected using the USB port or, if available, the express card slot. Modem access provides either primary or secondary (redundant) access to the Internet. For FortiGate units that do not include an internal modem (those units with an “M” designation), the modem interface will not appear in the web-based manager until enabled in the CLI. To enable the modem interface enter the CLI commands:

```
config system modem
  set status enable
end
```

Once enabled, modem options become available by going to *System > Network > Interface*.



The modem interface is only available when the FortiGate unit is in NAT mode.

Configuring the modem settings is a matter of entering the ISP phone number, user name and password. Depending on the modem, additional information may need to be supplied such as product identifiers, and initialization strings.

The FortiGate unit includes a number of common modems within its internal database. You can view these by selecting the *Configure Modem* link on the *Modem Settings* page. If your modem is not on the list, select *Create New* to add the information. This information is stored on the device, and will remain after a reboot.

Fortinet has an online database of modem models and configuration settings through FortiGuard. A subscription to the FortiGuard services is not required to access the information. As models are added, you can select the *Configure Modem* link and select *Update Now* to download new configurations.

USB modem port

Each USB modem has a specific dial-out ttyusb port. This will be indicated with the documentation for your modem. To enable the correct USB port, use the CLI commands:

```
config system modem
  set wireless-port {ttyusb0 | ttyusb1 | ttyusb2}
end
```

To test the port, use the diagnose command:

```
diagnose sys modem com /ttyusb1
```

The ttyusb1 will be the value of your USB port selected. The response will be:

```
Serial port: /dev/ttyusb1
Press Ctrl+W to exit.
```

If the port does not respond the output will be:

```
Can not open modem device '/dev/ttyusb1' : Broken pipe
```

Modes

The FortiGate unit allows for two modes of operation for the modem; stand alone and redundant. In stand alone mode, the modem connects to a dialup ISP account to provide the connection to the Internet. In redundant mode, the modem acts as a backup method of connecting to the Internet, should the primary port for this function fails.

Configuring either stand alone or redundant modes are very similar. The primary difference is the selection of the interface that the modem will replace in the event of it failing, and the configuration of a PING server to monitor the chosen interface.

Configuring stand alone mode

Configuring stand alone mode is a matter of configuring the modem information and the dialing mode. The dial mode is either *Always Connect* or *Dial on Demand*. Selecting *Always Connect* ensures that once the modem has connected, it remains connected to the ISP. Selecting *Dial on Demand*, the modem only calls the ISP if packets are routed to the modem interface. Once sent, the modem will disconnect after a specified amount of time.

To configure standalone mode as needed - web-based manager

- 1 Go to *System > Network > Modem*.
- 2 Select the *Mode of Standalone*.
- 3 Select the *Dial Mode of Dial on Demand*.
- 4 Enter the *Idle Timeout* of 2 minutes.
- 5 Select the number of redials the modem attempts if connection fails to 5.
- 6 Select *Apply*.

To configure standalone mode as needed- CLI

```
config system modem
  set mode standalone
  set auto-dial enable
  set idle-timer 2
  set redial 5
end
```

Configuring redundant mode

Redundant mode provides a backup to an interface, typically to the Internet. If that interface fails or disconnects, the modem automatically dials the configured phone number(s). Once connected, the FortiGate unit routes all traffic to the modem interface until the monitored interface is up again. The FortiGate unit pings the connection to determine when it is back online.

For the FortiGate to verify when the interface is back up, you need to configure a Ping server for that interface. You will also need to configure security policies between the modem interface and the other interfaces of the FortiGate unit to ensure traffic flow.

To configure redundant mode as needed - web-based manager

- 1 Go to *System > Network > Modem*.
- 2 Select the *Mode of Redundant*.
- 3 Select the interface the modem takes over from if it fails.
- 4 Select the *Dial Mode of Dial on Demand*.
- 5 Enter the *Idle Timeout* of 2 minutes.
- 6 Select the number of redials the modem attempts if connection fails to 5.
- 7 Select *Apply*.

To configure standalone mode as needed- CLI

```
config system modem
  set mode redundant
  set interface wan1
  set auto-dial enable
  set idle-timer 2
  set redial 5
end
```

Ping server

Adding a ping server is required for routing fail over traffic. A ping server will confirm the connectivity of the device's interface. You can only configure the ping server in the CLI.

To add a ping server - CLI

```
config router gwdetect
  edit wan1
    set server <ISP_IP_address>
    set failtime <failure_count>
    set interval <seconds>
  end
```

Additional modem configuration

The CLI provides additional configuration options when setting up the modem options including adding multiple ISP dialing and initialization options and routing. For more information, see the [CLI Reference](#).

Modem interface routing

The modem interface can be used in FortiOS as a dedicated interface. Once enabled and configured, you can use it in security policies and define static and dynamic routing. Within the CLI commands for the modem, you can configure the distance and priority of routes involving the modem interface. The CLI commands are:

```
config syssetm modem
    set distance <route_distance>
    set priority <priority_value>
end
```

For more information on the routing configuration in the CLI, see the [CLI Reference](#). For more information on routing and configuring routing, see the [Advanced Routing](#) Guide.

DHCP servers and relays

A DHCP server provides an address to a client on the network, when requested, from a defined address range.



DHCP server options are not available in transparent mode.

An interface cannot provide both a server and a relay for connections of the same type (regular or IPsec). However, you can configure a Regular DHCP server on an interface only if the interface is a physical interface with a static IP address. You can configure an IPsec DHCP server on an interface that has either a static or a dynamic IP address.

You can configure one or more DHCP servers on any FortiGate interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

If an interface is connected to multiple networks via routers, you can add a DHCP server for each network. The IP range of each DHCP server must match the network address range. The routers must be configured for DHCP relay.

You can configure a FortiGate interface as a DHCP relay. The interface forwards DHCP requests from DHCP clients to an external DHCP server and returns the responses to the DHCP clients. The DHCP server must have appropriate routing so that its response packets to the DHCP clients arrive at the unit.

DHCP Server configuration

To add a DHCP server, go to *System > Network > DHCP Server*, select *Create New* and complete the following:

Interface Name	Select an interface from the drop-down list.
Mode	Select the type of DHCP server.
Enable	Select to enable the DHCP server.

Type	Select the type of <i>DHCP</i> server. You cannot configure a regular DHCP server on an interface that has a dynamic IP address.
DHCP Server IP	Enter the IP address for the relay DHCP server. This appears only when <i>Mode</i> is <i>Relay</i> .
IP Range	Enter the start and end for the range of IP addresses that this DHCP server assigns to DHCP clients.
Network Mask	Enter the netmask of the addresses that the DHCP server assigns.
Default Gateway	Enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
DNS Service	Select to use either a specific DNS server or the system's DNS settings. You can add multiple DNS servers by selecting the plus sign (+) beside <i>DNS Server 1</i> . For more information see DNS services and DNS server .
DNS Server 0	Enter the DNS server.
DNS Server 1	Enter the second DNS server. If you need to add more DNS servers, select the plus sign (+).
IP Reservation	Select to match an IP address from the DHCP server to a specific client or device using its MAC address. In a typical situation, an IP address is assigned ad hoc to a client, and that assignment times out after a specific time of inactivity from the client, known as the lease time. To ensure a client or device always has the same IP address, that is, there is no lease time, use IP reservation.
Add from DHCP Client List	If the client is currently connected and using an IP address from the DHCP server, you can select this option to select the client from the list.
Advanced section of the New DHCP Service page	
Domain	Enter the domain that the DHCP server assigns to clients.
Lease Time	Set the length of time an IP address remains assigned to a client. Once the lease expires, the address is released for allocation to the next client request for an IP address. To set the lease to never expire, select <i>Unlimited</i> .

IP Assignment Mode	<p>Configure how the IP addresses for an IPsec DHCP server are assigned to dialup IPsec VPN users. These options are available when the DHCP server type is <i>IPsec</i>. Select:</p> <ul style="list-style-type: none"> <i>Server IP Range</i> - The IPsec DHCP server will assign the IP addresses as specified in <i>IP Range</i>, and <i>Exclude Ranges</i>. <i>User-group defined method</i> - The IP addresses will be assigned by a user group used to authenticate the user. The user group is used to authenticate XAUTH users. <p>When <i>User-group defined method</i> is selected, the <i>IP Range</i> fields are greyed out, and the <i>Exclude Ranges</i> table and controls are not visible.</p>
WINS Server 0 WINS Server 1	Add the IP addresses of one or two WINS servers that the DHCP server assigns to DHCP clients.
Options	<p>When adding a DHCP server, you have the ability to include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you may need to configure a FortiGate DHCP server that gives out a separate option as well as an IP address. For example, an environment that needs to support PXE boot with Windows images.</p> <p>The option numbers and codes are specific to the particular application. The documentation for the application will indicate the values to use. Option codes are represented in a option value/HEX value pairs. The option is a value 1 and 255. You can add up to three DHCP code/option pairs per server.</p>
Exclude Ranges	Enter a range of IP addresses from the IP range that should not be assigned. This option is only available when the DHCP type is <i>IPsec</i> , and the <i>IP Assignment Mode</i> is <i>Server IP range</i> .
Match VCI	.Select when connecting a FortiAP unit to the FortiGate. In the field that appears when selected, enter the FortiAP model number as the Vendor Class Identifier (VCI).

Service

On FortiGate-50 and FortiGate-60 series units, a DHCP server is configured, by default on the Internal interface:

IP Range	192.168.1.110 to 192.168.1.210
Netmask	255.255.255.0
Default gateway	192.168.1.99
Lease time	7 days
DNS Server 1	192.168.1.99

These settings are appropriate for the default Internal interface IP address of 192.168.1.99. If you change this address to a different network, you need to change the DHCP server settings to match.

Reserving IP addresses for specific clients

Within the DHCP pool of addresses, you can ensure certain computers will always have the same address. This can be to ensure certain users always have an IP address when connecting to the network, or if you want a device that connects occasionally to have the same address for monitoring its activity or use.

In the example below, the IP address 172.20.120.129 will be matched to MAC address 00:1f:5c:b8:03:57. This configuration is now only available in the CLI.

To configure IP reservation - CLI

```
config system dhcp reserved-address
  edit 1
    set ip 172.20.120.129
    set mac 00:1f:5c:b8:03:57
  end
```

Alternatively, after the FortiGate unit assigns an address, you can go to *System > Monitor > DHCP Monitor*, locate the particular user. Select the check box for the user and select *Add to Reserved*.

DHCP options

When adding a DHCP server, you have the ability to include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you may need to configure a FortiGate DHCP server that gives out a separate option as well as an IP address. For example, an environment that needs to support PXE boot with Windows images.

The option numbers and codes are specific to the particular application. The documentation for the application will indicate the values to use. Option codes are represented in a option value/HEX value pairs. The option is a value 1 and 255.

You can add up to three DHCP code/option pairs per DHCP server.

To configure option 252 with value <http://192.168.1.1/wpad.dat> - web-based manager

- 1 Go to *System > Network > DHCP Server* and select *Create New*.
- 2 Select a *Mode of Server*.
- 3 Select the blue arrow to expand the *Advanced* options.
- 4 Select *Options*.
- 5 Enter a *Code* of 252.
- 6 Enter the *Options* of
687474703a2f2f3139322e3136382e312e312f777061642e646174.

In the CLI, use the commands:

```
config system dhcp server
  edit <server_entry_number>
    set option1 252
    687474703a2f2f3139322e3136382e312e312f777061642e646174
  end
```

For detailed information about DHCP options, see [RFC 2132](#), DHCP Options and BOOTP Vendor Extensions.

DHCP Monitor

To view information about DHCP server connections, go to *System > Monitor > DHCP Monitor*. On this page, you can also add IP address to the reserved IP address list.

Assigning IP address by MAC address

To prevent users in the from changing their IP addresses and causing IP address conflicts or unauthorized use of IP addresses, you can bind an IP address to a specific MAC address using DHCP.

Use the CLI to reserve an IP address for a particular client identified by its device MAC address and type of connection. The DHCP server then always assigns the reserved IP address to the client. The number of reserved addresses that you can define ranges from 10 to 200 depending on the FortiGate model.

In the example below, the IP address 10.10.10.55 for User1 is assigned to MAC address 00:09:0F:30:CA:4F.

To assign an IP address to a specific MAC address

```
config system dhcp reserved-address
  edit User1
    set ip 10.10.10.55
    set mac 00:09:0F:30:CA:4F
    set type regular
  end
```

DNS services

A DNS server is a public service that converts symbolic node names to IP addresses. A Domain Name System (DNS) server implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with the computer IP address. This enables you to use readable locations, such as fortinet.com when browsing the Internet. FortiOS supports DNS configuration for both IPv4 and IPv6 addressing.

The FortiGate unit includes default DNS server addresses. However, these should be changed to those provided by your Internet Service Provider. The defaults are DNS proxies and are not as reliable as those from your ISP.

Within FortiOS, there are two DNS configuration options; each provide a specific service, and can work together to provide a complete DNS solution.

DNS queries

Basic DNS queries are configured on interfaces that connect to the Internet. When a web site is requested, for example, the FortiGate unit will look to the configured DNS servers to provide the IP address to know which server to contact to complete the transaction.

DNS server addresses are configured by going to *System > Network > DNS*. Here you specify the DNS server addresses. Typically, these addresses are supplied by your ISP. An additional option is available if you have local Microsoft domains on the network, by entering a domain name in the *Local Domain Name* field.

In a situation where all three fields are configured, the FortiGate unit will first look to the local domain. If no match is found, a request is sent to the external DNS servers.

If virtual domains are enabled, you create a DNS database in each VDOM. All of the interfaces in a VDOM share the DNS database in that VDOM.

Additional DNS CLI configuration

Further options are available from the CLI with the command `config system dns`. Within this command you can set the following commands:

- `dns-cache-limit` - enables you to set how many DNS entries are stored in the cache. Entries that remain in the cache provide a quicker response to requests than going out to the Internet to get the same information.
- `dns-cache-ttl` - enables you to set how long entries remain in the cache in seconds, between 60 and 86,400 (24 hours).
- `cache-notfound-responses` - when enabled, any DNS requests that are returned with NOTFOUND can be stored in the cache.
- `source-ip` - enables you to define a dedicated IP address for communications with the DNS server.

DNS server

You can also create local DNS servers for your network. Depending on your requirements, you can manually maintain your entries (master DNS server), or use it as a jumping point, where the server refers to an outside source (slave DNS server). A local master DNS server works similarly to the DNS server addresses configured in *System > Network > DNS*, but all entries must be added manually. This enables you to add a local DNS server to include specific URL/IP address combinations.



The DNS server options are not visible in the web-based manager by default. To enable the server, go to *System > Admin > Settings* and select *DNS Database*.

While a master DNS server is an easy method of including regularly used addresses to save on going to an outside DNS server, it is not recommended to make it the authoritative DNS server. IP addresses may change, and maintaining any type of list can quickly become labor-intensive.

A FortiGate master DNS server is best set for local services. For example, if your company has a web server on the DMZ that is accessed by internal employees as well as external users, such as customers or remote users. In this situation, the internal users when accessing the site would send a request for `website.example.com`, that would go out to the DNS server on the web, to return an IP address or virtual IP. With an internal DNS, the same site request is resolved internally to the internal web server IP address, minimizing inbound/outbound traffic and access time.

As a slave, DNS server, the FortiGate server refers to an external or alternate source as way to obtain the url/IP combination. This useful if there is a master DNS server for a large company where a list is maintained. Satellite offices can then connect to the master DNS server to obtain the correct addressing.



The DNS server entries does not allow CNAME entries, as per [rfc 1912](#), section 2.4.

To configure a master DNS server - web-based manager

- 1 Go to *System > Network > DNS Server*, and select *Create New*.
- 2 Select the *Type of Master*.

- 3 Select the *View* as *Shadow*.
The view is the accessibility of the DNS server. Selecting *Public*, external users can access, or use, the DNS server. Selecting *Shadow*, only internal users can use it.
- 4 Enter the DNS *Zone*, for example, *WebServer*.
- 5 Enter the domain name for the zone, for example *example.com*.
- 6 Enter the hostname of the DNS server, for example, *Corporate*.
- 7 Enter the contact address for the administrator, for example, *admin@example.com*.
- 8 Set *Authoritative* to *Disable*.
- 9 Select *OK*.
- 10 Enter the DNS entries for the server by selecting *Create New*.
- 11 Select the *Type*, for example, *Address (A)*.
- 12 Enter the *Hostname*, for example *web.example.com*.
- 13 Enter the remaining information, which varies depending on the *Type* selected.
- 14 Select *OK*.

To configure a DNS server - CLI

```

config system dns-database
  edit WebServer
    set domain example.com
    set type master
    set view shadow
    set ttl 86400
    set primary-name corporate
    set contact admin@example.com
    set authoritative disable
    config dns-entry
      edit 1
        set hostname web.example.com
        set type A
        set ip 192.168.21.12
        set status enable
      end
    end
  end
end

```

Recursive DNS

You can set an option to ensure these types of DNS server is not the authoritative server. When configured, the FortiGate unit will check its internal DNS server (Master or Slave). If the request cannot be fulfilled, it will look to the external DNS servers. This is known as a split DNS configuration.

You can also have the FortiGate unit look to an internal server should the Master or Slave not fulfill the request by using the CLI commands:

```

config system dns-database
  edit example.com
    ...
    set view shadow
  end
end

```

For this behavior to work completely, for the external port, you must set the DNS query for the external interface to be recursive. This option is configured in the CLI only.

To set the DNS query

```
config system dns-server
    edit wan1
        set mode recursive
    end
```

Dynamic DNS

If your ISP changes the your external IP address on a regular basis, and you have a static domain name, you can configure the external interface to use a dynamic DNS service to ensure external users and/or customers can always connect to your company firewall.

To configure dynamic DNS in the web-based manager, go to *System > Network > DNS*, select *Use DDNS*, and enter the relevant information for the interface communicating to the server, and which server to use, and relevant information.

To configure dynamic DNS in the CLI use the commands below. Within the CLI you can configure a DDNS for each interface. Only the first configured port appears in the web-based manager. Additional commands vary with the DDNS server you select.

```
config system ddns
    edit <instance_value>
        set monitor-interface <external_interface>
        set ddns-server <ddns_server_selection>
    end
```

Aggregate Interfaces

Link aggregation (IEEE 802.3ad) enables you to bind two or more physical interfaces together to form an aggregated (combined) link. This new link has the bandwidth of all the links combined. If a link in the group fails, traffic is transferred automatically to the remaining interfaces with the only noticeable effect being a reduced bandwidth.

This is similar to redundant interfaces with the major difference being that a redundant interface group only uses one link at a time, where an aggregate link group uses the total bandwidth of the functioning links in the group, up to eight.

Support of the IEEE standard 802.3ad for link aggregation is available on some models.

An interface is available to be an aggregate interface if:

- it is a physical interface, not a VLAN interface or subinterface
- it is not already part of an aggregate or redundant interface
- it is in the same VDOM as the aggregated interface. Aggregate ports cannot span multiple VDOMs.
- it does not have an IP address and is not configured for DHCP or PPPoE
- it is not referenced in any security policy, VIP, IP Pool or multicast policy
- it is not an HA heartbeat interface
- it is not one of the FortiGate-5000 series backplane interfaces

Some models of FortiGate units do not support aggregate interfaces. In this case, the aggregate option is not an option in the web-based manager or CLI. As well, you cannot create aggregate interfaces from the interfaces in a switch port.

To see if a port is being used or has other dependencies, use the following diagnose command:

```
diagnose sys checkused system.interface.name <interface_name>
```

When an interface is included in an aggregate interface, it is not listed on the *System > Network > Interface* page. Interfaces will still appear in the CLI, although configuration for those interfaces will not take affect. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, IP pools, or routing.

Example

This example creates an aggregate interface on a FortiGate-3810A using ports 4-6 with an internal IP address of 10.13.101.100, as well as the administrative access to HTTPS and SSH.

To create an aggregate interface - web-based manager

- 1 Go to *System > Network > Interface* and select *Create New*.
- 2 Enter the Name as *Aggregate*.
- 3 For the *Type*, select *802.3ad Aggregate*.
If this option does not appear, your FortiGate unit does not support aggregate interfaces.
- 4 In the *Available Interfaces* list, select port 4, 5 and 6 and move it to the *Selected Interfaces* list.
- 5 Select the *Addressing Mode of Manual*.
- 6 Enter the IP address for the port of 10.13.101.100/24.
- 7 For *Administrative Access* select HTTPS and SSH.
- 8 Select *OK*.

To create aggregate interface - CLI

```
config system interface
  edit Aggregate
    set type aggregate
    set member port4 port5 port6
    set vdom root
    set ip 172.20.120.100/24
    set allowaccess https ssh
  end
```

IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP

- DNS
- Authorization requests such as RADIUS
- FSSO

Configuration of these services is performed in the CLI. In each instance, there is a command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
    set ntpsyn enable
    set syncinterval 5
    set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

Administration for schools

For system administrator in the school system it is particularly difficult to maintain a network and access to the Internet. There are potential legal liabilities if content is not properly filtered and children are allowed to view pornography and other non-productive and potentially dangerous content. For a school, too much filtering is better than too little. This section describes some basic practices administrators can employ to help maintain control without being too draconian for access to the internet.

Security policies

The default security policies in FortiOS allow all traffic on all ports and all IP addresses. Not the most secure. While applying UTM profiles can help to block viruses, detect attacks and prevent spam, this doesn't provide a solid overall security option. The best approach is a layered approach; the first layer being the security policy.

When creating outbound security policies, you need to know the answer to the question "What are the students allowed to do?" The answer is surf the web, connect to FTP sites, send/receive email, and so on.

Once you know what the students need to do, you can research the software used and determine the ports the applications use. For example, if the students only require web surfing, then there are only two ports (80 - HTTP and 443 - HTTPS) needed to complete their tasks. Setting the security policies to only allow traffic through two ports (rather than all 65,000), this will significantly lower any possible exploits. By restricting the ports to known services, means stopping the use of proxy servers, as many of them operate on a non-standard port to hide their traffic from URL filtering or HTTP inspection.

DNS

Students should not be allowed to use whatever DNS they want. This opens another port for them to use and potentially smuggle traffic on. The best approach is to point to an internal DNS server and only allow those devices out on port 53. It's the same approach one would use for SMTP. Only allow the mail server to use port 25 since nothing else should be sending email.

If there is no internal DNS server, then the list of allowed DNS servers they can use should be restrictive. One possible exploit would be for them to set up their own DNS server at home that serves different IPs for known hosts, such as having Google.com sent back the IP for playboy.com.

Encrypted traffic (HTTPS)

Generally speaking, students should not be allowed to access encrypted web sites. Encrypted traffic cannot be sniffed, and therefore, cannot be monitored. HTTPS traffic should only be allowed when necessary. Most web sites a student needs to access are HTTP, not HTTPS. Due to the nature of HTTPS protocol, and the fact that encryption is an inherent security risk to your network, its use should be restricted.

Adding a security policy that encompasses a list of allowed secure sites will ensure that any HTTPS sites that are required are the only sites a student can go to.

FTP

For the most part, students should not be using FTP. FTP is not HTTP or HTTPS so you cannot use URL filtering to restrict where they go. This can be controlled with destination IPs in the security policy. With a policy that specifically outlines which FTP addresses are allowed, all other will be blocked.

Example security policies

Given these requirements, an example set of security policies could look like the following illustration. In a large setup, all the IPs for the students are treated by one of these four policies.

Figure 60: Simple security policy setup

<input type="checkbox"/>	Seq. No.	ID	Source	Destination	Schedule	Service	Action	Status
<input type="checkbox"/>	1	2	Student PCs	Allowed Websites	always	HTTPS	✓	✓
<input type="checkbox"/>	2	3	Student PCs	all	always	HTTP	✓	✓
<input type="checkbox"/>	3	4	Student PCs	Allowed DNS	always	DNS	✓	✓
<input type="checkbox"/>	4	5	Student PCs	Allowed FTP	always	FTP	✓	✓
<input type="checkbox"/>	5		all	all	always	ANY	✗	Implicit

The last policy in the list, included by default, is a deny policy. This adds to the potential of error that could end up allowing unwanted traffic to pass. The deny policy ensures that any traffic making it to this point is stopped. It can also help in further troubleshooting by viewing the logs for denied traffic.

With these policies in place, even before packet inspection occurs, the FortiGate, and the network are fairly secure. Should any of the UTM profiles fail, there is still a basic level of security.

UTM Profiles

In FortiOS 4.0 MR2, the protection profiles have been broken into individual profiles. Each UTM feature is now its own component, which can make setting up network security easier.

Antivirus profiles

Antivirus screening should be enabled for any service you have enabled in the security policies. In the case above, HTTP, FTP, as well as POP3 and SMTP (assuming there is email access for students). There is not a virus scan option for HTTPS, because the content is encrypted. Generally speaking, most of the network traffic will be students surfing the web.

To configure antivirus profiles in the web-based manager, go to *UTM Profiles > Antivirus > Profile*, or use the CLI commands under `config antivirus profile`.

Web filtering

The actual filtering of URLs - sites and content - should be performed by FortiGuard. It is easier and web sites are constantly being monitored, and new ones reviewed and added to the FortiGuard databases every day. The FortiGuard categories provide an extensive list of offensive, and non-productive sites.

As well, there are additional settings to include in a web filtering profile to best contain a student's web browsing.

- Web URL filtering should be enabled to set up exemptions for web sites that are blocked or reasons other than category filtering. It also prevents the use of IP addresses to get around web filtering.
- Block invalid URLs - HTTPS only. This option inspects the HTTPS certificate and looks at the URL to ensure it's valid. It is common for proxy sites to create an HTTPS certificate with a garbage URL. If the site is legitimate, it should be set up correctly. If the site approach to security is to ignore it, then their security policy puts your network at risk and the site should be blocked.

Web filtering options are configured in the web-based manager by going to *UTM Profiles > Web filter > Profile*, or in the CLI under `config webfilter profile`.

Advanced options

There are a few Advanced options to consider for a web filtering profile:

- Enable *Provide details for blocked HTTP 4xx and 5xx errors*. Under normal circumstances there are exploits that can be used with 400 and 500 series messages to access the web site. While most students probably won't know how to do this, there is no harm in being cautious. It only takes one.
- Enable *Rate Images by URL*. This option only works with Google images. It examines the URL that the images is stored at to get a rating on it, then blocks or allows the image based on the rating of the originating URL. It does not inspect the image contents. Most image search engines to a prefect and pass the images directly to the browser.
- Enable *Block HTTP redirects by rating*. An HTTP redirect is one method of getting around ratings. Go to one web site that has an allowed rating, and it redirects to another web site that may want blocked.

Categories and Classifications

For the selection of what FortiGuard categories and classifications that should be blocked, that is purely based on the school system and its Internet information policy.

Email Filtering

Other than specific teacher-led email inboxes, there is no reason why a student should be able to access, read or send personal email. Ports for POP3, SMTP and IMAP should not be opened in a security policies.

IPS

The intrusion protection profiles should be used to ensure the student PCs are not vulnerable to attacks, nor do you want students making attacks. As well, IPS can do more than simple vulnerability scans. With a FortiGuard subscription, IPS signatures are pushed to the FortiGate unit. New signatures are released constantly for various intrusions as they are discovered.

FortiOS includes a number of predefined IPS sensors that you can enable by default. Selecting the `all_default` signature is a good place to start as it includes the major signatures.

To configure IPS sensors in the web-based manager, go to *UTM Profiles > Intrusion Protection > IPS Sensor*, on the CLI use commands under `config ips sensor`.

Application control

Application control uses IPS signatures to limit the use of instant messaging and peer-to-peer applications which can lead to possible infections on a student's PC. FortiOS includes a number of pre-defined application categories. To configure and maintain application control profiles in the web-based manager, go to *UTM Profiles > Application Control > Application Control List*. In the CLI use commands under `config application list`.

Some applications to consider include proxies, botnets, toolbars and P2P applications.

Logging

Turn on all logging - every option in this section should be enabled. This is not where you decide what you are going to log. It is simply defining what the UTM profiles can log.

Logging everything is a way to monitor traffic on the network, see what student's are utilizing the most, and locate any potential holes in your security plan. As well, keeping this information may help to prove negligence later in necessary.

Tag management

Tag management provide a method of categorizing, or labelling objects within FortiOS using keywords. You can give the following elements a "tag", similar to a keyword:

- IPS signature
- application signature
- security policy
- firewall address

Tagging is way to organize the various elements, especially if you have a large number of addresses, security policies to manage and keep track of. Tagging enables you to break these elements into groups, but each element can belong to more than one group. Tags help you find elements which have something in common, be it a group, user or location. This is very similar to tagging found on photo sharing sites.

To use tagging, you need to enable it for 1U FortiGate units. It is enabled by default on all 2U FortiGate units and blades.

To enable tagging - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 Select *Object Tagging and Coloring*.
- 3 Select *Apply*.

To enable tagging - CLI

```
config system settings
  set gui-object-tags
end
```

Adding and removing tags

You add and remove tags when you create the various elements. For example, when adding a firewall address, a section below the Interface selection enables you to add tags for that element, such as the department, region, or really, anything to help identify the element. When editing, applied tags appear as well.

Figure 61: Adding tags to a new address.

The screenshot shows the 'New Address' configuration window. It contains the following fields and options:

- Address Name:** User_1
- Color:** [Change]
- Type:** Subnet / IP Range (dropdown)
- Subnet / IP Range:** 172.20.120.12
- Interface:** dmz2 (dropdown)
- Tags:**
 - Applied tags:** accounting (dropdown)
 - Add tags:** west coast (text input with a plus icon)

At the bottom right, there are 'OK' and 'Cancel' buttons.

To remove a tag, in the element, click the tag in the Applied Tags list.

Reviewing tags

Tags can be reviewed in one location by going to *System > Config > Tag Management*. In this screen, all tags used appear. The visual size of the tag name indicates the usage; the bigger the size, the more it is used. By hovering over the keyword, a fly out indicates how many times it has been used.

To see where it was used, click the keyword. An *Object Usage* window displays all the reference categories where the keyword was used, and the number of times. Selecting the expand arrow further details its use.

Further, for security policies for example, you can select the *View* icon and see the details of the particular element. If need be, select the *Edit* icon to modify the element.

Figure 62: Viewing the address information for a tagged object

Tagging guidelines

Given the ease that tags can be added to elements in FortiOS, it makes sense to jump right in and begin applying tags to elements and object. However, this type of methodology will lead to problems down the road as new elements are added.

A methodology should be considered and developed before applying tags. This doesn't mean you need to develop an entire thesaurus or reference guide for all possibilities of tags. However, taking some time to develop a methodology for the keywords you intend to use will benefit later when new security policies, addresses, and so on are added. Some things to consider when developing a tag list:

- the hierarchy used for the organization such as region, city location, building location
- department names and if short forms or long forms are used
- will acronyms be used or terms spelled out.
- how granular will the tagging be

As tags are added, previously used tags appear so there is an opportunity to use previously used tags. However, you want to avoid a situation where both accounting and acct are both options. This is also important if there are multiple administrators in different locations to ensure consistency.

At any time, you can change or even remove tags. It is best to do a bit of planning ahead of time to avoid unnecessary work later on.

Software switch

A software switch, or soft switch, is a virtual switch that is implemented at the software, or firmware level, rather than the hardware level. Adding a software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiGate unit.

It can also be useful if you require more hardware ports on for the switch on a FortiGate unit. For example, if your FortiGate unit has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces such as those with FortiWiFi and FortiAP unit.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

There are a few things to consider when setting up a software switch:

- Ensure you create a back up of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiGate unit. If you accidentally combine too many ports, you will need a way to undo any errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiGate unit. For example, DHCP servers, security policies, and so on.

To create a software switch - web-based manager

- 1 Go to *System > Network > Interface* and select *Create New*.
- 2 Enter a *Name* for the switch.
- 3 Set *Type* to *Software Switch*.
- 4 Select the interfaces to add to the switch.
- 5 Enter an *IP address*.
- 6 Select *OK*.

To create a software switch - CLI

```
config system switch-interface
  edit <switch-name>
    set type switch
    set member <interface_list>
  end
config system interface
  edit <switch_name>
    set ip <ip_address>
    set allowaccess https ssh ping
  end
```

Soft switch example

For this example, the wireless interface (WiFi) needs to be on the same subnet as the DMZ1 interface to facilitate wireless syncing from an iPhone and a local computer. The syncing between two subnets is problematic. By putting both interfaces on the same subnet the syncing will work. The software switch will accomplish this.



In this example, the soft switch includes a wireless interface. Remember to configure any wireless security before proceeding. If you leave this interface open without any password or other security, it leaves open access to not only the wireless interface, but any other interfaces and devices connected within the software switch.

Clear the interfaces and back up the configuration

First, ensure that the interfaces are not being used with any other security policy or other use on the FortiGate unit. Check the WiFi and DMZ1 ports to ensure DHCP is not enabled on the interface and there are no other dependencies with these interfaces.

Next, save the current configuration, in the event something doesn't work, recovery can be quick.

To back up the configuration

Go to *System > Dashboard > Status*.

- 1 In the *System Information* widget, select *Backup* in the *System Configuration* row.
- 2 Select *Backup* from the backup window.
- 3 Select the location to save the configuration file.

Merge the interfaces

The plan is to merge the WiFi port and DMZ1 port. This will create a software switch with a name of "synchro" with an IP address of 10.10.21.12. The steps will create the switch, add the IP and then set the administrative access for HTTPS, SSH and Ping.

To merge the interfaces - web-based manager

- 1 Go to *System > Network > Interface* and select *Create New*.
- 2 Enter the *Name* of *synchro*.
- 3 Select the *Type of Software Switch*.
- 4 From the *Available Interfaces* list, select *DMZ1* and select the *Right arrow* to move it to the *Selected Interfaces* list.
- 5 Repeat the above step for the *WiFi* interface.
- 6 Enter the *IP/Netmask* of *10.10.21.12/255.255.255.0*.
- 7 For the *Administrative Access*, select *HTTPS*, *PING* and *SSH*.
- 8 Select *OK*.

To merge the interfaces - CLI

```
config system switch-interface
  edit synchro
    set type switch
    set member dmz1 wifi
  end
config system interface
```



```
edit synchro
set ip 10.10.21.12
set allowaccess https ssh ping
end
```

Final steps

With the switch set up, you can now add security policies, DHCP servers and any other configuration that you would normally do to configure interfaces on the FortiGate unit.

Replacement messages list

The replacement message list is in *System > Config > Replacement Message*.

The replacement messages list enables you to view and customize replacement messages. Use the expand arrow beside each type to display the replacement messages for that category. Select the *Edit* icon beside each replacement message to customize that message for your requirements.

If you are viewing the replacement messages list in a VDOM, any messages that have been customized for that VDOM are displayed with a Reset icon that you can use to reset the replacement message to the global version.

For connections requiring authentication, the FortiGate unit uses HTTP to send an authentication disclaimer page for the user to accept before a security policy is in effect. Therefore, the user must initiate HTTP traffic first in order to trigger the authentication disclaimer page. Once the disclaimer is accepted, the user can send whatever traffic is allowed by the security policy.

Replacement message images

You can add images to replacement messages to:

- disclaimer pages
- login pages
- declined disclaimer pages
- login failed page
- login challenge pages
- keepalive pages

Image embedding is also available to the endpoint NAC download portal and recommendation portal replacement messages, as well as HTTP replacement messages.

Supported image formats are GIF, JPEG, TIFF and PNG. The maximum file size supported is 6000 bytes.

Adding images to replacement messages

To upload an image for use in a message

- 1 Go to *System > Config > Replacement Message*.
- 2 Select *Manage Images* at the top of the page.
- 3 Select *Create New*.
- 4 Enter a *Name* for the image.
- 5 Select the *Content Type*.
- 6 Select *Browse* to locate the file and select *OK*.

The image that you include in a replacement message, must have the following html:

```
<img src=%%IMAGE: <config_image_name>%% size=<bytes> >
```

For example:

```
<img src=%%IMAGE: logo_hq%% size=4272>
```

Modifying replacement messages

Replacement messages can be modified to include a message or content that suits your organization.

Use the expand arrows to view the replacement message list for a given category. Messages are in HTML format. For descriptions of the replacement message tags, see [Replacement message tags](#).

To change a replacement message, go to *System > Config > Replacement Message* and expand the replacement message category to access the replacement message that you want to modify. Select the message and select *Edit*.

Within the message edit screen, the Allowed Formats line indicates whether the content can be HTML code or simple text. indicates what type of format the replacement message is in, Text or HTML. For example, the HTTP virus replacement message's format is HTTP and the Email file block replacement message is Text. The Size indicates the maximum number of characters allowed in the message.

Replacement message tags

Replacement messages can include replacement message tags, or variables. When users receive the message, the message tag is replaced with content relevant to the message. The table lists the replacement message tags that you can use.

Table 29: Replacement message tags

Tag	Description
%%AUTH_LOGOUT%%	The URL that will immediately delete the current policy and close the session. Used on the auth-keepalive page.
%%AUTH_REDIR_URL%%	The auth-keepalive page can prompt the user to open a new window which links to this tag.
%%CATEGORY%%	The name of the content category of the web site.
%%DEST_IP%%	The IP address of the request destination from which a virus was received. For email this is the IP address of the email server that sent the email containing the virus. For HTTP this is the IP address of web page that sent the virus.
%%DURATION%% (FortiOS Carrier only)	The amount of time in the reporting period. This is user defined in the protection profile.
%%EMAIL_FROM%%	The email address of the sender of the message from which the file was removed.
%%EMAIL_TO%%	The email address of the intended receiver of the message from which the file was removed.
%%FAILED_MESSAGE%%	The failed to login message displayed on the auth-login-failed page.

Table 29: Replacement message tags (Continued)

Tag	Description
%%FILE%%	The name of a file that has been removed from a content stream. This could be a file that contained a virus or was blocked by antivirus file blocking. %%FILE%% can be used in virus and file block messages.
%%FORTIGUARD_WF%%	The FortiGuard - Web Filtering logo.
%%FORTINET%%	The Fortinet logo.
%%LINK%%	The link to the FortiClient Host Security installs download for the Endpoint Control feature.
%%HTTP_ERR_CODE%%	The HTTP error code. "404" for example.
%%HTTP_ERR_DESC%%	The HTTP error description.
%%KEEPAALIVEURL%% (FortiOS Carrier only)	auth-keepalive-page automatically connects to this URL every %%TIMEOUT%% seconds to renew the connection policy.
%%MMS_SENDER%% (FortiOS Carrier only)	Senders MSISDN from message header.
%%MMS_RECIPIENT%% (FortiOS Carrier only)	Recipients MSISDN from message header.
%%MMS_SUBJECT%% (FortiOS Carrier only)	MMS Subject line to help with message identity.
%%MMS_HASH_CHECKSUM%%	Value derived from hash calculation - will only be shown on duplicate message alerts.
%%MMS_THRESH%%	Mass MMS alert threshold that triggered this alert.
%%NIDSEVENT%%	The IPS attack message. %%NIDSEVENT%% is added to alert email intrusion messages.
%%NUM_MSG%% (FortiOS Carrier only)	The number of time the device tried to send the message with banned content within the reporting period.
%%OVERRIDE%%	The link to the FortiGuard Web Filtering override form. This is visible only if the user belongs to a group that is permitted to create FortiGuard web filtering overrides.
%%OVRD_FORM%%	The FortiGuard web filter block override form. This tag must be present in the FortiGuard Web Filtering override form and should not be used in other replacement messages.
%%PROTOCOL%%	The protocol (http, ftp, pop3, imap, or smtp) in which a virus was detected. %%PROTOCOL%% is added to alert email virus messages.
%%QUARFILENAME%%	The name of a file that has been removed from a content stream and added to the quarantine. This could be a file that contained a virus or was blocked by antivirus file blocking. %%QUARFILENAME%% can be used in virus and file block messages. Quarantining is only available on FortiGate units with a local disk.

Table 29: Replacement message tags (Continued)

Tag	Description
%%QUOTA_INFO%%	Display information about the traffic shaping quota setting that is blocking the user. Used in traffic quota control replacement messages.
%%QUESTION%%	Authentication challenge question on auth-challenge page. Prompt to enter username and password on auth-login page.
%%SERVICE%%	The name of the web filtering service.
%%SOURCE_IP%%	The IP address of the request originator who would have received the blocked file. For email this is the IP address of the user's computer that attempted to download the message from which the file was removed.
%%TIMEOUT%%	Configured number of seconds between authentication keepalive connections. Used on the auth-keepalive page.
%%URL%%	The URL of a web page. This can be a web page that is blocked by web filter content or URL blocking. %%URL%% can also be used in http virus and file block messages to be the URL of the web page from which a user attempted to download a file that is blocked.
%%VIRUS%%	The name of a virus that was found in a file by the antivirus system. %%VIRUS%% can be used in virus messages

Mail replacement messages

The FortiGate unit sends the mail replacement messages to email clients using IMAP, POP3, or SMTP when an event occurs such as antivirus blocking a file attached to an email that contains a virus. Email replacement messages are text messages.

If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to IMAPS, POP3S, and SMTPS email messages.

Table 30: Mail replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus message	<i>Virus Scan</i> (any email protocol within an antivirus profile)	If a match is detected, the infected file from the email message is deleted and replaced with the message.
File block message	<i>File Filter</i> (file filter list selected within the antivirus profile)	If a match is detected, the incoming file is blocked and triggers this message.
Oversized file message	<i>Oversized File/Email</i> set to <i>Block</i> (within protocol options list)	If a match is detected, the file is removed and replaced with this message.
Fragmented email	<i>Allow Fragmented Emails</i> (an exception: this is not enabled)	If a match is detected, the fragmented email is blocked and this replacement message replaces the first fragment of the fragmented email.

Table 30: Mail replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Data leak prevention message	A rule is set to <i>Block</i> (DLP sensor)	If a match is detected, the FortiGate unit blocks messages and this replacement message is sent to the sender.
Subject of data leak prevention message	<i>Block, Ban, Ban Sender, Quarantine IP address, and Quarantine interface</i> (DLP sensor)	This replacement message is added to the subject field of all email messages when a match is found.
Banned by data leak prevention message	A rule is set to <i>Ban</i> (DLP sensor)	Replaces a blocked email message and replaces any additional email messages that the banned user sends until they are removed from the banned user list.
Sender banned by data leak prevention message	A rule set to <i>Ban Sender</i> (DLP sensor)	Replaces a blocked email message with this message and replaces any additional email messages that the banned user sends until the user is removed from the banned user list.
Virus message (splice mode)	Splice mode	If the antivirus system detects a virus in an SMTP email message, then the FortiGate FortiGate unit aborts the SMTP session and returns a 554 SMTP error message to the sender that is included in the message.
File block message (splice mode)	Splice mode	If the antivirus file filter deleted a file from an SMTP email message, the FortiGate unit aborts the SMTP session and returns a 554 SMTP error message to the sender that is included in the message.
Oversized file message (splice mode)	Splice mode AND <i>Oversized File/Email</i> is set to <i>Block</i> (protocol option list)	If the FortiGate unit blocks an oversized SMTP email message, the FortiGate unit aborts the SMTP session and returns a 554 SMTP error message to the sender that is included in the message.

HTTP replacement messages

The FortiGate unit sends the HTTP replacement messages listed in the following table to web browsers using the HTTP protocol when an event occurs such as antivirus blocking a file that contains a virus in an HTTP session. HTTP replacement messages are HTML pages.

If the FortiGate unit supports SSL content scanning and inspection, and if under HTTPS in the protocol option list has Enable Deep Scan enabled, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

Table 31: HTTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus message	<i>Virus Scan</i> for HTTP or HTTPS (antivirus profile)	Displays a web page in the client's browser when an entry in the selected file filter list matches HTTP GET. The FortiGate unit blocks the file being downloaded using HTTP GET.
Infection cache message	Client comforting (web filter profile)	This message is triggered only after the blocked URL is attempted for a second time.
File block message	<i>File Filter</i> for HTTP or HTTPS (antivirus profile)	Displays in the client's browser when an entry in the selected file filter list matches HTTP GET and the FortiGate unit blocks that file that is being downloaded.
Oversized file message	<i>Oversized File/Email set to Block</i> for HTTP or HTTPS (protocol options list)	The FortiGate unit blocks an oversized file that is being downloaded that uses a HTTP GET and replaces the file with this web page that is displayed by the client browser.
Data leak prevention message	A rule is set to <i>Block</i> (DLP sensor)	If the FortiGate unit blocks a page/file that the user loads using HTTP GET with this web page, the message appears. This message appears if the FortiGate unit also blocks the user that is sending information using HTTP POST.
Banned by data leak prevention message	A rule is set to <i>Ban</i> (DLP sensor)	This message replaces a blocked web page or file. This message also replaces any additional web pages or files that the banned user attempts to access until the user is removed from the banned user list.
Banned word message	Banned word's score (web filter profile)	If the banned word's score exceeds the threshold set in the web filter profile, the page is blocked and the blocked page is replaced with this message.
Content-type block message	N/A	Email headers include information about content types such as image for pictures, and so on. If a specific content-type is blocked, the blocked message is replaced with this web page message.
URL block message	Web URL Filtering (web filter profile)	This message replaces a web page that the FortiGate unit blocked.

Table 31: HTTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Archive block message	A DLP sensor has, in a sensor filter, <i>Archive</i> set to <i>Full</i> or <i>Summary</i>	This message displays in the client's browser when archive HTTP downloads are blocked by the FortiGate unit.
Web Filter error message	<i>Allow Websites When a Rating Error Occurs</i> option in <i>Advanced Filter</i> , in a web filter profile	This message displays in the clients browser when there are HTTP web filter errors.
Client block	<i>File Filter</i> for HTTP or HTTPS (antivirus profile)	This message displays in the client's browser when a file that is being uploaded by an HTTP POST is blocked by the FortiGate unit.
Client anti-virus	<i>Virus Scan</i> for HTTP or HTTPS (antivirus profile)	This message displays in a client's browser when an infected file that is being uploaded using FTP PUT is detected by the FortiGate unit.
Client filesize	<i>Oversized File/Email</i> set to <i>Block</i> for HTTP or HTTPS (protocol options list)	If an oversized file is being uploaded using FTP PUT, and the FortiGate unit blocks the file that is being uploaded, this message replaces the web page.
Client banned word	Web Content filtering (this is in the CLI)	This message displays in a client's browser when the FortiGate unit blocks a web page that is being uploaded with an HTTP PUT and contains content that matches an entry in the Web Content Filter list.
Client archive block	set archive-log {corrupted encrypted mailbomb multipart nested unhandled} (config antivirus profile, under config setting in the CLI)	This message displays when a user is attempting to upload a banned archived file.
POST block	<i>HTTP POST Action</i> is set to <i>Block</i> (profile)	This message displays this web page when the FortiGate unit blocks a HTTP POST.
Invalid certificate block	When <code>block</code> is set in: smtps-client-cert-request ftps-client-cert-request https-client-cert-request (CLI)	This message displays when a request for a certificate is determined by the FortiGate unit to be invalid and blocks the certificate.

Web Proxy replacement messages

The FortiGate unit sends Web Proxy replacement messages listed in the table below when a web proxy event occurs that is detected and matches the web proxy configuration. These replacement messages are web pages that appear within your web browser.

The following web proxy replacement messages require an identity-based security policy so that the web proxy is successful. You can also enable FTP-over-HTTP by selecting the *FTP* option in *System > Network > Explicit Proxy*.

Table 32: Web Proxy replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Web proxy access denied	Web Proxy (default action set to Deny)	<p>This message displays when both of the following are true, as well as when there is no web proxy policy defined:</p> <ul style="list-style-type: none"> no web proxy policy is defined OR no existing policy matches the incoming request default action is set to Deny (<i>System > Network > Explicit Proxy</i>) <p>Note: The default action is ignored when there is at least one web policy defined.</p>
Web proxy login challenge	N/A	This replacement message is triggered by a log in, and is always sent to the client's browser with it is triggered; however, some browsers (Internet Explorer and Firefox) are unable to display this replacement message.
Web proxy login fail	N/A	If a user name and password authentication combination is entered, and is accepted as incorrect, this replacement message appears.
Web proxy authorization fail	N/A	<p>If a username and password is entered and is correct, this message appears. However, if the following is true, this message also appears:</p> <ul style="list-style-type: none"> The user is not allowed to view the request resources, (for example, in an Fortinet Single Sign On Agent setup and the authentication passes), and the username and password combo is correct, but the user group does not match a user group defined in the security policy.

Table 32: Web Proxy replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Web proxy HTTP error	N/A	This message is triggered whenever there is a web proxy HTTP error. This message forwards the actual servers' error message and a web proxy internal error message, for example, error 404: web page is not found.
Web proxy user-limit (CLI only)	user-limit (config system replacemsg webproxy)	This message is triggered when a web proxy user has met the threshold that is defined in global resources or vdom resources.

FTP Proxy replacement message

The FortiGate unit sends the FTP proxy replacement message whenever a user access the FTP proxy. The replacement message is a banner-message that contains only text.

FTP replacement messages

The FortiGate unit sends the FTP replacement messages listed in the table below to FTP clients when an event occurs such as antivirus blocking a file that contains a virus in an FTP session. FTP replacement messages are text messages.

Table 33: FTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus message	Virus Scan for FTP (antivirus profile)	If a match is detected and the infected file is deleted when being downloaded using FTP, the FortiGate unit sends this message to the FTP client.
Blocked message	File Filter for FTP (antivirus profile)	If a match is detected and a file that is being downloaded uses FTP, and the FortiGate unit blocks the download, the FortiGate unit sends this message to the FTP client.
Oversized message	<i>Oversized File/Email</i> set to <i>Block</i> for FTP (antivirus profile)	If an oversize file that is being downloaded using FTP is blocked, the FortiGate unit sends this message to the FTP client.
DLP message	A rule set to <i>Block</i> (DLP sensor)	This replacement message replaces a blocked FTP download.

Table 33: FTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
DLP ban message	A rule set to <i>Ban</i> (blocks an FTP session)	If a match is detected, and is using protocols such as FTP PUT and FTP GET, this replacement message displays. This message is displayed whenever the banned user attempts to access, until the user is removed from the banned user list.
Archive block	A DLP rule has <i>Archived</i> enabled, <i>Action</i> is set to <i>Block</i> , and <i>All-FTP</i> (file transfers) is also selected.	If a match is detected, and a file is being transferred, this replacement message displays. This message is displayed whenever a file is being transferred over FTP, and that DLP rule has archiving enabled as well as <i>Action</i> set to <i>Block</i> .

NNTP replacement messages

The FortiGate unit sends the NNTP replacement messages listed in the following table to NNTP clients when an event occurs such as antivirus blocking a file attached to an NNTP message that contains a virus. NNTP replacement messages are text messages.

Table 34: NNTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus message	<i>Virus Scan</i> for NNTP (antivirus profile)	If a match is detected, and an infected file is attached to an NNTP message, and the FortiGate unit deletes the NNTP message, the FortiGate unit sends the replacement message to the client.
Blocked message	<i>File Filter</i> for NNTP (antivirus profile)	This message is sent to the client when a file attached to an NNTP message is blocked by the FortiGate unit.
Oversized message	<i>Oversized File/Email</i> set to <i>Block</i> for NNTP (protocol options list)	If the FortiGate unit removes an oversized file from an NNTP message, this message is sent instead.
Data Leak prevention message	A rule set to <i>Block</i> (DLP sensor)	If the FortiGate unit blocks an NNTP message, the message is replaced with this replacement message.

Table 34: NNTP replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Subject of data leak prevention message	<i>Block</i> , <i>Ban</i> , <i>Quarantine IP address</i> , and <i>Quarantine interface</i> (DLP sensor)	This message is added to the subject field of all NNTP messages.
Banned by data leak prevention message	A rule set to <i>Ban</i> (DLP sensor)	If a match is detected, this message replaces a blocked NNTP message. This message also replaces any additional NNTP messages that the banned user sends until the user is removed from the banned user list.

Alert Mail replacement messages

The FortiGate unit adds the alert mail replacement messages listed in the following table to alert email messages sent to administrators.



If you enable the option *Send alert email for logs based on severity*, whether or not replacement messages are sent by alert email depends on how you set the alert email in *Minimum log level*.

Table 35: Alert mail replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus message	<i>Virus detected</i> (alert email message) AND <i>Virus Scan</i> (antivirus profile)	If a match is detected, this message displays. Note: Both options/settings must be enabled for this replacement message to appear.
Block message	<i>Virus detected</i> (alert email messages) AND <i>File Filter</i> (antivirus profile)	This message displays when if the FortiGate unit blocks a file. Note: Both options/settings must be enabled for this replacement message to appear.
Intrusion message	<i>Intrusion detected</i> (alert email message) AND IPS sensor or DoS sensor is enabled	If a match is detected, as well as an attack, this message displays. Note: Both options/settings must be enabled for this replacement message to appear.
Critical event message	<i>Send alert email for logs based on severity</i> AND <i>Minimum log level</i> set to <i>Alert</i> or <i>Emergency</i> (alert email message)	Whenever a critical level event log message is generated, this message is sent; however, unless you configure an alert email message with both of the said options enabled, this message does not appear.

Table 35: Alert mail replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Disk full message	<i>Disk Usage</i> (alert email message)	If the disk usage reaches the percentage configured in the alert email notification settings, this replacement message displays.

Spam replacement messages

The FortiGate unit adds the Spam replacement messages listed in the following table to SMTP server responses if the email message is identified as spam and the spam action is discard. If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to SMTPS server responses.

Table 36: Spam replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Email IP	<i>IP address BWL check</i> (for any email protocol within an email filter profile)	If a match is detected to the last hop IP address, then this message is added.
DNSBL/ORD BL	<i>spamrbl (CLI)</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies the email message as spam, this message is added.
HELO/EHLO domain	<i>HELO DNS lookup</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies an email message as spam, this message is added. Note: <i>HELO DNS lookup</i> is not available for SMTPS.
Email address	<i>E-mail Address BWL check</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies an email message as spam, this message is added.
Mime header	<i>spamhdr check (CLI)</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies an email message as spa, this message is added.
Returned email domain	<i>Return e-mail DNS check</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies an email message as spam, this message is added.

Table 36: Spam replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Banned word	<i>Banned word check</i> (for any email protocol within an email filter profile)	If the FortiGate unit identifies an email message as spam, this message is added.
Spam submission message	Any email filtering option for any email protocol within an email filter profile	If the FortiGate unit identifies an email message as spam, it adds this message. Note: Email Filtering adds this message to all email tagged as spam. The message describes a button that the recipient of the message can select to submit the email signatures to the FortiGuard Antispam service if the email was incorrectly tagged as spam (a false positive).

Administration replacement message

If you enter the following CLI command the FortiGate unit displays the *Administration Login Disclaimer* whenever an administrator logs into the FortiGate unit's web-based manager or CLI.

```
config system global
    set access-banner enable
end
```

The web-based manager administrator login disclaimer contains the text of the Login Disclaimer replacement message as well as Accept and Decline buttons. The administrator must select accept to login.

Authentication replacement messages

The FortiGate unit uses the text of the authentication replacement messages listed in [Authentication replacement messages](#) for various user authentication HTML pages that are displayed when a user is required to authenticate because a security policy includes at least one identity-based policy that requires firewall users to authenticate.

These replacement message pages are for authentication using HTTP and HTTPS. You cannot customize the firewall authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

Users see the authentication login page when they use a VPN or a security policy that requires authentication. You can customize this page in the same way as you modify other replacement messages.

There are some unique requirements for these replacement messages:

- The login page must be an HTML page containing a form with ACTION="/" and METHOD="POST"

- The form must contain the following hidden controls:
 - `<INPUT TYPE="hidden" NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%">`
 - `<INPUT TYPE="hidden" NAME="%%STATEID%%" VALUE="%%STATEVAL%%">`
 - `<INPUT TYPE="hidden" NAME="%%REDIRID%%" VALUE="%%PROTURI%%">`
- The form must contain the following visible controls:
 - `<INPUT TYPE="text" NAME="%%USERNAMEID%%" size=25>`
 - `<INPUT TYPE="password" NAME="%%PASSWORDID%%" size=25>`

Example

The following is an example of a simple authentication page that meets the requirements listed above.

```
<HTML><HEAD><TITLE>Firewall Authentication</TITLE></HEAD>
<BODY><H4>You must authenticate to use this service.</H4>

<FORM ACTION="/" method="post">
<INPUT NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%" TYPE="hidden">

<TABLE ALIGN="center" BGCOLOR="#00cccc" BORDER="0"
CELLPADDING="15" CELLSPACING="0" WIDTH="320"><TBODY>

<TR><TH>Username:</TH>
<TD><INPUT NAME="%%USERNAMEID%%" SIZE="25" TYPE="text"> </TD></TR>

<TR><TH>Password:</TH>
<TD><INPUT NAME="%%PASSWORDID%%" SIZE="25" TYPE="password">
</TD></TR>

<TR><TD COLSPAN="2" ALIGN="center" BGCOLOR="#00cccc">
<INPUT NAME="%%STATEID%%" VALUE="%%STATEVAL%%" TYPE="hidden">
<INPUT NAME="%%REDIRID%%" VALUE="%%PROTURI%%" TYPE="hidden">
<INPUT VALUE="Continue" TYPE="submit"> </TD></TR>

</TBODY></TABLE></FORM></BODY></HTML>
```

Table 37: Authentication replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Disclaimer page	<i>Enable Disclaimer and Redirect URL to</i> (identity-based security policy)	After a firewall user authenticates with the FortiGate unit using HTTP or HTTPS, this message, which is a disclaimer page, displays. Note: The CLI includes <code>auth-disclaimer-page-1</code> , <code>auth-disclaimer-page-2</code> , and <code>auth-disclaimer-page-3</code> that you can use to increase the size of the authentication disclaimer page replacement message.

Table 37: Authentication replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Declined disclaimer page	N/A	When a firewall user selects the button on the Disclaimer page to decline access through the FortiGate unit, the <i>Declined disclaimer page</i> (replacement message) is displayed.
Login page	N/A	The HTML page displayed for firewall users who are required to authenticate using HTTP or HTTPS before connecting through the FortiGate FortiGate unit.
Login failed page	N/A	The HTML page displayed if firewall users enter an incorrect user name and password combination.
Success message	N/A	The page displays when a user authenticates for a Telnet session.
Login challenge page	N/A	<p>The HTML page displayed if firewall users are required to answer a question to complete authentication. The page displays the question and includes a field in which to type the answer. This feature is supported by RADIUS and uses the generic RADIUS challenge-access auth response. Usually, challenge-access responses contain a Reply-Message attribute that contains a message for the user (for example, "Please enter new PIN"). This message is displayed on the login challenge page. The user enters a response that is sent back to the RADIUS server to be verified.</p> <p>The Login challenge page is most often used with RSA RADIUS server for RSA SecurID authentication. The login challenge appears when the server needs the user to enter a new PIN. You can customize the replacement message to ask the user for a SecurID PIN.</p>

Table 37: Authentication replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Keepalive page	<pre>config system global set auth-keepalive enable end</pre>	The HTML page displayed with firewall authentication keepalive is enabled using the following command: Authentication keepalive keeps authenticated firewall sessions from ending when the authentication timeout ends. Go to <i>User > Options</i> to set the <i>Authentication Timeout</i> .
FortiToken page	<i>Two-factor authentication</i> is enabled	The message displays when the token password code is required as part of a user's login credentials when that user has two-factor authentication enabled.
Email token page	<i>Two-factor authentication</i> is enabled and <i>Email to</i> is also enabled	The message displays when the token password code has been emailed to a user and is required as part of a user's login credentials when that user has two-factor authentication enabled.
SMS token page	<i>Two-factor authentication</i> is enabled and <i>SMS</i> is also enabled	The message displays when the token password code has been sent to a user's mobile phone, and that code must be included in a user's login credentials when that user has two-factor authentication enabled.

Captive Portal Default replacement messages

The Captive Portal Default replacement messages are used for wireless authentication only. You must have a VAP interface with the security set as captive portal to trigger these replacement messages.

Table 38: Captive Portal Default replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Disclaimer page	VAP interface that has captive portal set	<i>The message is a disclaimer agreement; if the user who is trying to access a web site that is not under the control of the network access provider and is given a choice to either agree to the terms and continue or not gain access to the site.</i>
Declined disclaimer page	VAP interface that has captive portal set	<i>Appears when the user who did not agree to the terms in the Disclaimer page message.</i>
Login page	VAP interface that has captive portal set	<i>The message is an authentication page.</i>

Table 38: Captive Portal Default replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Login failed page	VAP interface that has captive portal set	<i>The message that appears when the user has failed to log in.</i>

FortiGuard Web Filtering replacement messages

The FortiGate unit sends the FortiGuard Web Filtering replacement messages listed in the table to web browsers using the HTTP protocol when FortiGuard web filtering blocks a URL, provides details about blocked HTTP 4xx and 5xx errors, and for FortiGuard overrides. FortiGuard Web Filtering replacement messages are HTTP pages.

If the FortiGate unit supports SSL content scanning and inspection and if *Protocol Recognition > HTTPS Content Filtering Mode* is set to Deep Scan in the antivirus profile, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

Table 39: FortiGuard Web Filtering replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
URL block message	<code>config ftgd-wf</code> <code>set options</code> (under <code>config webfilter profile</code>) for HTTP or HTTPS	This message replaces a web page that is blocked by the FortiGate unit.
HTTP error message	<i>Provide details for blocked HTTP 4xx and 5xx errors</i> for HTTP or HTTPS (web filter profile)	This message replaces a web page that is blocked.
FortiGuard Web Filtering override form	<code>ovrd-perm</code> in <code>config webfilter profile</code>	<p>If FortiGuard Web Filtering blocks a web page in this category, this message displays a web page.</p> <p>By using this web page, users can authenticate to get access to the page. Overrides are configured within the CLI using the <code>ovrd-perm</code> value in the <code>config webfilter profile</code> command.</p> <p>Note: Do not remove the <code>%%OVRD_FORM%%</code> tag. This tag provides the form used to initiate an override if FortiGuard Web Filtering blocks access to a web page.</p>

Table 39: FortiGuard Web Filtering replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
FortiGuard Web Filtering quota expired message	<i>Enforced Quota</i> (web filter profile)	This message is added when a match is detected regarding FortiGuard quota.
FortiGuard Webfiltering warning portal message	<i>Warning</i> (action) for a filter in a web filter profile	This message replaces the blocked web page. The user can either select <i>Proceed</i> or <i>Cancel</i> , to proceed to the web site or cancel and return back to the previous web site. If they select <i>Proceed</i> , and filter category <i>Require additional Authentication to proceed</i> is enabled in the web filter profile, the user is required to log in.

IM and P2P replacement messages

The FortiGate unit sends the IM and P2P replacement messages listed in [Table 40](#) to IM and P2P clients using AIM, ICQ, MSN, or Yahoo! Messenger when an event occurs such as antivirus blocking a file attached to an email that contains a virus. IM and P2P replacement messages are text messages.

Table 40: IM and P2P replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
File block message	<i>File Filter</i> for IM (application control list)	If a file that matches an entry in the selected file filter list for IM is deleted, this message replaces the file.
File name block message	<i>File Filter</i> antivirus for IM (application control list)	If a file name matches an entry in the selected file filter list, and is deleted, this message replaces it.
Virus message	<i>Virus Scan</i> for IM (application control list)	If an infected file is deleted, this message replaces the file.
Oversized file message	<i>Oversized File/Email</i> set to <i>Block</i> for IM (protocol options list)	If an oversized file is removed, this message replaces it.
Data leak prevention message	A rule set to <i>Block</i> (DLP sensor)	If a blocked IM or P2P message is blocked, this message replaces it.

Table 40: IM and P2P replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Banned by data leak prevention message	A rule set to <i>Ban</i> (DLP sensor)	If an IM or P2P message is blocked, this message replaces it with this message. This message also replaces any additional messages that the banned user sends until they are removed from the banned user list.
Voice chat block message	<i>Block Audio</i> for AIM, ICQ, MSN, or Yahoo! (application control list)	If a match is detected, this message displays.
Video chat block message	<i>set block-video</i> enable in either AIM, ICQ, MSN, or Yahoo application list entries.	If a match is detected, this message displays.
Photo share block message	<i>block-photo</i> for MSN or Yahoo! (CLI) (application control list)	If a match is detected, this message displays.

Endpoint NAC replacement messages

The FortiGate unit sends one of the following messages to non-compliant users who attempt to use a security policy in which Endpoint NAC is enabled.

Table 41: Endpoint NAC replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Endpoint NAC Download Portal	<i>Quarantine Hosts to User Portal (Enforce compliance)</i>	The user can download the FortiClient Endpoint Security application installer. If you modify this message, be sure to retain the %%LINK%% tag which provides the download URL for the FortiClient installer.
Endpoint NAC Recommendation Portal	<i>Notify Hosts to Install FortiClient (Warn only)</i> (Endpoint profile)	The user can either download the FortiClient Endpoint Security application installer or select the <i>Continue to</i> link to access their desired destination. If you modify this message, be sure to retain both the %%LINK%% tag which provides the download URL for the FortiClient installer and the %%DST_ADDR%% link that contains the URL that the user requested.
Endpoint NAC Block Page	N/A	This message appears when FortiClient is opened before FortiClient has time to see the HTTP message.

Table 41: Endpoint NAC replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Endpoint NAC Recommendation Block Page	<code>set recommendation-disclaimer enable (CLI)</code>	This message displays when it is recommended that the endpoint be compliant so that the user may gain access to the network. Select Continue to link to access the desired destination. If you modify this message, be sure to retain both the %%LINK%% tag which provides the download URL for the FortiClient installer and the %%DST_ADDR%% link that contains the URL that the user requested.
Endpoint NAC Feature Block Page	N/A	This message displays when endpoint security is required and the FortiClient security check failed.
Endpoint NAC Recommendation Feature Block Page	FortiClient's antivirus settings enabled	This message displays when endpoint security is required and the FortiClient security check failed. Select Continue to link to access the desired destination.

NAC quarantine replacement messages

The page that is displayed for the user depends on whether NAC quarantine blocked the user because a virus was found, a DoS sensor detected an attack, an IPS sensor detected an attack, or a DLP rule with action set to *Quarantine IP address* or *Quarantine Interface* matched a session from the user.

The default messages inform the user of why they are seeing this page and recommend they contact the system administrator. You can customize the pages as required, for example to include an email address or other contact information or if applicable a note about how long the user can expect to be blocked.

Table 42: NAC quarantine replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
Virus Message	<i>Quarantine Virus Sender</i> (antivirus profile; the FortiGate unit adds a source IP address or FortiGate interface to the banned user list)	The FortiGate unit displays this message as a web page when the blocked user attempts to connect through the FortiGate unit using HTTP on port 80 or when any user attempts to connect through a FortiGate interface added to the banned user list using HTTP on port 80.

Table 42: NAC quarantine replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
DoS Message	<code>attack</code> or <code>interface</code> (CLI) (DoS sensor) AND applied to a DoS security policy	For a DoS Sensor the CLI <code>quarantine</code> option set to <code>attacker</code> or <code>interface</code> and the DoS Sensor added to a DoS security policy adds a source IP, a destination IP, or FortiGate interface to the banned user list. The FortiGate unit displays this message as a web page when the blocked user attempts to connect through the FortiGate unit using HTTP on port 80 or when any user attempts to connect through a FortiGate interface added to the banned user list using HTTP on port 80. This replacement message is not displayed if <code>quarantine</code> is set to <code>both</code> .
IPS Message	<i>Quarantine Attackers</i> (IPS sensor)	This message displays if <i>Quarantine Attackers</i> enabled in an IPS sensor filter or override and the IPS sensor applied to a security policy adds a source IP address, a destination IP address, or a FortiGate interface to the banned user list. The FortiGate unit displays this message as a web page when the blocked user attempts to connect through the FortiGate unit using HTTP on port 80 or when any user attempts to connect through a FortiGate interface added to the banned user list using HTTP on port 80. This message is not displayed if <i>method</i> is set to <i>Attacker and Victim IP Address</i> .
DLP Message	Action set to <i>Quarantine IP address</i> OR <i>Quarantine Interface</i> (DLP sensor)	This message displays if Action set to <i>Quarantine IP address</i> or <i>Quarantine Interface</i> in a DLP sensor adds a source IP address or a FortiGate interface to the banned user list. The FortiGate unit displays this message as a web page when the blocked user attempts to connect through the FortiGate unit using HTTP on port 80 or when any user attempts to connect through a FortiGate interface added to the banned user list using HTTP on port 80.

Traffic quota control replacement messages

When user traffic is going through the FortiGate unit and it is blocked by traffic shaping quota controls, users see the *Traffic shaper block message* or the *Per IP traffic shaper block message* when they attempt to connect through the FortiGate unit using HTTP.

The traffic quota HTTP pages should contain the %%QUOTA_INFO%% tag to display information about the traffic shaping quota setting that is blocking the user.

SSL VPN replacement message

The SSL VPN login replacement message is an HTML replacement message that formats the FortiGate SSL VPN portal login page. You can customize this replacement message according to your organization's needs. The page is linked to FortiGate functionality and you must construct it according to the following guidelines to ensure that it will work.

- The login page must be an HTML page containing a form with `ACTION="%%SSL_ACT%%"` and `METHOD="%%SSL_METHOD%%"`
- The form must contain the %%SSL_LOGIN%% tag to provide the login form.
- The form must contain the %%SSL_HIDDEN%% tag.

MM1 replacement messages

MM1 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile. The replacement messages for MM1 are listed in [Table 43](#).



You must have *Remove Blocked* selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the FortiGate unit.

Table 43: MM1 replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM1 send-req virus message	<i>Virus Scan</i> (MMS profile)	This message is sent to the client when a virus is detected within multiple messages during the scan of the client's m-send.req HTTP post request.
MM1 send-req file block message	<i>Virus Scan</i> (MMS profile)	This message is sent to the client when during the scan of the client's m-send.req HTTP post request a banned file was found in multiple messages, the file is blocked and this message is sent to the client.
MM1 send-req carrier end point filter message	<i>Carrier end-point filter</i> for (MMS profile)	This message is sent to the client when, during the scan of the client's m-send.req HTTP post request, a banned user and/or recipient was being contacted.
MM1 send-req content checksum block message	<i>Content checksum</i> (MMS profile)	This message is sent to the client when, during the scan of the client's m-send.req HTTP post request, banned content was found.

Table 43: MM1 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM1 send-req banned word message	<i>Banned Word Check</i> (Email filter profile)	This message is sent to the client when, during the scan of the client's m-send.req HTTP post request, banned words were found in multiple messages.
MM1 send-req flood message	<i>Message flood threshold</i> (MMS profile)	This message is sent to the client when, during the scan of the client's m-send.req HTTP post request, multiple messages that were being sent to a mobile device were flagged as message floods.
MM1 send-req duplicate message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent to the client when, during the scan of the client's m-send.req HTTP post request, multiple messages that were being sent to a mobile device were flagged as duplicates.
MM1 send request flood alert message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This is sent when a mass MMS flood event is detected by FortiOS Carrier. This message is in an email format, with to, from and subject line included.
MM1 send request duplicate alert message	<i>MMS Bulk Email Filter AND MMS Notifications</i> (MMS profile)	This message is sent when a mass MMS duplicate message event is detected by FortiOS Carrier. This message is in an email format, with to, from and subject line included as well as the hash checksum.
MM1 send-conf virus message	<i>Virus Scan</i> (MMS profile)	This message is sent to confirm that a message that has been sent was blocked because it contains a banned file. FortiOS Carrier quarantines the banned file.
MM1 send-conf file block message	<i>File Filter</i> (MMS profile)	This message is sent to confirm that a message that was previously sent was blocked because it contains a banned file. FortiOS Carrier quarantines the banned file.

Table 43: MM1 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM1 send-conf carrier end point filter message	<i>MMS Address Translation AND Carrier End Point Block</i> (MMS profile)	This message is sent to notify the person that sent a message that it was blocked because the sender or recipient is banned. FortiOS Carrier determined that the content was not acceptable because the m-send-req messages contained blocked carrier end points.
MM1 send-conf content checksum block message	<i>MMS Content Checksum</i> (with the MMS content checksum list included in MMS profile)	This message is sent to notify the person that sent a message, that message was blocked because the actual message contains banned content. FortiOS Carrier determined that the content was not acceptable because the m-send-req messages contained content checksum blocked payloads.
MM1 send-conf flood message	<i>MMS Bulk Email Filtering Detection AND MMS Notification</i> (MMS profile)	This message is sent to notify the person that sent a message that it was blocked because the sender has been banned for sending too many messages. FortiOS Carrier determined that the content was not acceptable because the m-send-req messages were flagged as flood messages.
MM1 send-conf duplicate message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that sent a message that it was blocked because the content within the message has been sent too many times. FortiOS Carrier determined that the content was not acceptable because the m-send-req messages were flagged as duplicate messages.
MM1 send-conf banned word message	<i>Banned Word Check</i> (Email Filter profile) <i>AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that sent a message that it was blocked because it contained a banned word. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf messages contained a banned word.

Table 43: MM1 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM1 send-conf virus message	<i>Virus Scan</i> (MMS profile)	This message is sent to notify the person that sent a message that is was infected with a virus. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message for retrieval contained a virus. FortiOS Carrier also quarantines the virus.
MM1 retrieve-conf virus message	<i>File Filter</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that it contained a file which contained a virus. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message for retrieval contained a virus. FortiOS Carrier also quarantines the file
MM1 retrieve-conf file block message	<i>File Filter AND MMS Notifications</i> (MMS Profile)	This message is sent to notify the person that was retrieving the message was blocked because it contains a banned file. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message contained a banned file. FortiOS Carrier also quarantines the file.
MM1 retrieve-conf carrier endpoint filter message	<i>Carrier Endpoint Block AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that it was blocked because the sender or recipient is banned. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message contained blocked carrier end points.
MM1 retrieve-conf content checksum block message	<i>MMS Content Checksum AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that is was blocked because it contained banned content. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message contained content checksum blocked payloads.

Table 43: MM1 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM1 retrieve-conf flood message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that it was blocked because the sender was banned for sending too many messages. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message was flagged as flood.
MM1 retrieve-conf duplicate message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that it was blocked because the message content was sent too many times. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-cof message contained blocked content.
MM1 retrieve-conf banned word message	<i>Banned Word Check</i> (Email filter profile) AND <i>MMS Notifications</i> (MMS profile)	This message is sent to notify the person that was retrieving the message, that it was blocked because it contained a banned word. FortiOS Carrier determined that the content was not acceptable because the m-retrieve-conf message contained a banned word.

MM3 replacement messages

MM3 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile. The replacement messages for MM3 are listed in [Table 44](#).



You must have *Remove Blocked* selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the unit.

Table 44: MM3 replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM3 virus message	<i>Virus Scan</i> (MMS profile)	This message is sent when an MM3 message contains a virus. FortiOS Carrier quarantines the file.
MM3 file block message	<i>File Filter</i> (MMS profile)	This message is sent when an MM3 message contains a blocked file. FortiOS Carrier quarantines the file.
MM3 carrier end point filter message	<i>Carrier End Point Block</i> (MMS profile)	This message is sent when an MM3 message contains a banned sender or recipient is blocked by FortiOS Carrier.
MM3 content checksum block message	<i>MMS Content Checksum</i> (MMS profile)	This message is sent when an MM3 message contains banned contain and is blocked by FortiOS Carrier.
MM3 banned word message	<i>Banned Word Check</i> (Email Filter profile) <i>AND MMS Notifications</i> (MMS profile)	This message is sent when an MM3 message contains a banned word.
MM3 flood alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when an MM3 message is scanned and found to have mass MMS flood. This replacement message is in the form of an email, containing from, to and subject lines.
MM3 duplicate alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when an MM3 message is found to have duplicate messages. This message is in the form of an email, containing form, to, subject, and hash checksum.
MM3 virus notification message	<i>Virus Scan AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device sending messages are found to contain a virus.
MM3 file block notification message	<i>File Filter AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device sending the messages are found to contain banned files.
MM3 carrier end point filter notification message	<i>Carrier End Point Block</i> (MMS profile)	This message is sent when the mobile device has sent messages that contain a banned or recipient.

Table 44: MM3 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM3 content checksum block notification message	<i>MMS Content Checksum</i> (MMS profile)	This messages is sent when the mobile device has messages that contain banned content.
MM3 banned word notification message	<i>Banned Word Check</i> (Email Filter profile) AND <i>MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages containing banned words.
MMS flood notification message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when the mobile device has sent messages that were flagged as floods.
MM3 duplication notification message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when the mobile device has sent messages that were flagged as duplicates.

MM4 replacement messages

MM4 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile. The replacement messages for MM1 are listed in [Table 45](#).

Table 45: MM4 replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM4 virus message	<i>Virus Scan</i> (MMS profile)	This message is sent when the person has sent a message containing a file to a recipient, and that file was blocked because it contained a virus. FortiOS Carrier quarantines the file.
MM4 file block message	<i>File Filter</i> (MMS profile)	This message is sent when the person that has sent a message was blocked because it contains a banned file. FortiOS Carrier quarantines the file.
MM4 carrier end point filter message	<i>Carrier End Point Block</i> (MMS profile)	This message is sent when the message is blocked because the sender or recipient is banned.
MM4 content checksum block message	<i>MMS Content Checksum</i> (MMS profile)	This message is sent when the message is blocked because it contains banned content.
MM4 flood message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when a message is blocked because the sender is banned for sending too many messages.
MM4 duplicate message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when that message is blocked. FortiOS Carrier blocked it because that message body was sent too many times.
MM4 banned word message	<i>Banned Word Check</i> (Email Filter profile) <i>AND MMS Notifications</i> (MMS profile)	This message is sent when a message is blocked. FortiOS Carrier blocked the message because it contains a banned word.
MM4 virus notification message	<i>Virus Scan AND MMS Notifications</i> (MMS profile)	This replacement message is sent when the mobile device has sent messages containing a virus.
MM4 file block notification message	<i>File Filter AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages containing banned files.

Table 45: MM4 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM4 carrier end point notification message	<i>Carrier End Point Block AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages that contain a banned sender.
MM4 content checksum block notification message	<i>MMS Content Checksum AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent message that contain banned content.
MM4 flood notification message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages that were flagged as floods.
MM4 duplication notification message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages there were flagged as duplicates.
MM4 banned word notification message	<i>Banned Word Check (Email Filter profile) AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages that contain banned words.
MM4 flood alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when a mass MMS flood is detected. This replacement message is in an email message format, with to, from and subject lines.
MM4 duplicate alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when a mass MMS duplicates is detected. This replacement message is in an email message format, with to, from, subject, and hash checksum lines.

MM7 replacement messages

MM7 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile. The replacement messages for MM7 are listed in [Table 46](#).

Table 46: MM7 replacement messages

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM7 virus message	<i>Virus Scan</i> (MMS profile)	This message is sent when a person has sent a message and FortiOS Carrier blocked the message because it contained a virus. FortiOS Carrier quarantines the file.
MM7 file block message	<i>File Filter</i> (MMS profile)	This message is sent when a person has sent a message and FortiOS Carrier blocked the message because it contained a banned file. FortiOS Carrier quarantines the file.
MM7 carrier end point filter message	<i>Carrier End Point Block</i> (MMS profile)	This message is sent when FortiOS Carrier detected that the message sent contained a banned sender or recipient. The message was blocked by FortiOS Carrier.
MM7 content checksum block message	<i>MMS Content Checksum</i> (MMS profile)	This message is sent when FortiOS Carrier detected that the message contained banned content. The message was blocked by FortiOS Carrier.
MM7 banned word message	<i>Banned Word Check</i> (Email Filter profile)	This message is sent when a message that a person sent was blocked because it contains a banned word.
MM7 virus notification message	<i>Virus Scan AND MMS Notifications</i> (MMS profile)	This replacement message is sent when the mobile device has sent messages that contain a virus.
MM7 file block notification message	<i>File Filter AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages containing banned files.
MM7 carrier end point filter notification message	<i>Carrier End Point Block AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages that contain a banned sender or recipient.
MM7 content checksum block notification message	<i>MMS Content Checksum AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent messages that contain banned content.

Table 46: MM7 replacement messages (Continued)

Replacement Message name	Enabled options/settings that triggers the replacement message	Reason for sending or displaying the replacement message
MM7 banned word notification message	<i>Banned Word Check</i> (Email Filter profile) <i>AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device sent message containing banned words.
MM7 flood notification message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device sent messages that were flagged as flood.
MM7 duplicate notification message	<i>MMS Bulk Email Filtering Detection AND MMS Notifications</i> (MMS profile)	This message is sent when the mobile device has sent message that were flagged as duplicates.
MM7 flood alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when a mass MMS flood is detected. This replacement message is in an email message format, with to, from and subject lines.
MM7 duplicate alert message	<i>MMS Bulk Email Filtering Detection</i> (MMS profile)	This message is sent when a mass MMS duplicates is detected. This message is in an email message format, with to, from, subject, and hash checksum lines.

MMS replacement messages

The MMS replacement message is sent when a section of an MMS message has been replaced because it contains a blocked file. This replacement message is in HTML format.

The message text is:

```
<HTML><BODY>This section of the message has been replaced because
it contained a blocked file</BODY></HTML>
```

Replacement message groups

You configure the default replacement message group from *System > Config > Replacement Message Group* in FortiOS Carrier. All new replacement message groups that you add inherit from the default group. Modifying messages in the default group automatically changes any messages that are unmodified in the other groups.

If you enable virtual domains (VDOMs) on the FortiGate unit, replacement message groups are configured separately for each virtual domain. Each virtual domain has its own default replacement message group, configured from *System > Config > Replacement Message Group*.

When you modify a message in a replacement message group, a Reset icon appears beside the message in the group. You can select this Reset icon to reset the message in the replacement message group to the default version.

All MM1/4/7 notification messages (and MM1 retrieve-conf messages) can contain a SMIL layer and all MM4 notification messages can contain an HTML layer in the message. These layers can be used to brand messages by using logos uploaded to the FortiGate unit via the 'Manage Images' link found on the replacement message group configuration page.

Disk

To view the status and storage information of the local disk on your FortiGate unit, go to *System > Config > Advanced*. The *Disk* menu appears only on FortiGate units with an internal hard or flash disk.

Formatting the disk

The internal disk of the FortiGate unit (if available) can be formatted by going to *System > Config > Disk* and selecting *Format*.

Formatting the disk will erase all data on it, including databases for antivirus and IPS; logs, quarantine files, and WAN optimization caches. The FortiGate unit requires a reboot once the disk has been formatted.

Setting space quotas

If the FortiGate unit has an internal hard or flash disk, you can allocate the space on the disk for specific logging and archiving, and WAN optimization. By default, the space is used on an as required basis. As such, a disk can fill up with basic disk logging, leaving less potential space for quarantine.

By going to *System > Config > Disk*, you can select the *Edit* icon for *Logging and Archiving* and *WAN Optimization & Web Cache* and define the amount of space each log, archive and WAN optimization has on the disk.

CLI Scripts

To upload bulk CLI commands and scripts, go to *System > Config > Advanced*.

Scripts are text files containing CLI command sequences. Scripts can be used to deploy identical configurations to many devices. For example, if all of your devices use identical security policies, you can enter the commands required to create the security policies in a script, and then deploy the script to all the devices which should use those same settings.

Use a text editor such as Notepad or other application that creates simple text files. Enter the commands in sequence, with each line as one command, similar to examples throughout the FortiOS documentation set.

If you are using a FortiGate unit that is not remotely managed by a FortiManager unit or the FortiGuard Analysis and Management Service, the scripts you upload are executed and discarded. If you want to execute a script more than once, you must keep a copy on your management PC.

If your FortiGate unit is configured to use a FortiManager unit, you can upload your scripts to the FortiManager unit, and run them from any FortiGate unit configured to use the FortiManager unit. If you upload a script directly to a FortiGate unit, it is executed and discarded.

If your FortiGate unit is configured to use FortiGuard Analysis and Management Service, scripts you upload are executed and stored. You can run uploaded scripts from any FortiGate unit configured with your FortiGuard Analysis and Management Service account. The uploaded script files appear on the FortiGuard Analysis and Management Service portal web site.

Uploading script files

After you have created a script file, you can then upload it through *System > Config > Advanced*. When a script is uploaded, it is automatically executed.



Commands that require the FortiGate unit to reboot when entered in the command line will also force a reboot if included in a script.

To execute a script

- 1 Go to *System > Config > Advanced*.
- 2 Verify that *Upload Bulk CLI Command File* is selected.
- 3 Select *Browse* to locate the script file.
- 4 Select *Apply*.

If the FortiGate unit is not configured for remote management, or if it is configured to use a FortiManager unit, uploaded scripts are discarded after execution. Save script files to your management PC if you want to execute them again later.

If the FortiGate unit is configured to use the FortiGuard Analysis and Management Service, the script file is saved to the remote server for later reuse. You can view the script or run it from the FortiGuard Analysis and Management Service portal web site.

Rejecting PING requests

The factory default configuration of your FortiGate unit allows the default external interface to respond to ping requests. Depending on the model of your FortiGate unit the actual name of this interface will vary. For the most secure operation, you should change the configuration of the external interface so that it does not respond to ping requests. Not responding to ping requests makes it more difficult for a potential attacker to detect your FortiGate unit from the Internet. One such potential threat are Denial of Service (DoS) attacks.

A FortiGate unit responds to ping requests if ping administrative access is enabled for that interface.

To disable ping administrative access - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Choose the external interface and select *Edit*.
- 3 Clear the *Ping Administrative Access* check box.
- 4 Select *OK*.

In the CLI, when setting the `allowaccess` settings, by selecting the access types and not including the PING option, that option is then not selected. In this example, only HTTPS is selected.

To disable ping administrative access - CLI

```
config system interface
  edit external
    set allowaccess https
  end
```

Opening TCP 113

Although seemingly contrary to conventional wisdom of closing ports from hackers, this port, which is used for ident requests, should be opened.

Port 113 initially was used as an authentication port, and later defined as an identification port (see RFC 1413). Some servers may still use this port to help in identifying users or other servers and establish a connection. Because port 113 receives a lot of unsolicited traffic, many routers, including on the FortiGate unit, close this port.

The issue arises in that unsolicited requests are stopped by the FortiGate unit, which will send a response saying that the port is closed. In doing so, it also lets the requesting server know there is a device at the given address, and thus announcing its presence. By enabling traffic on port 113, requests will travel to this port, and will most likely, be ignored and never responded to.

By default, the ident port is closed. To open it, use the following CLI commands:

```
config system interface
  edit <port_name>
    set inden_accept enable
  end
```

You could also further use port forwarding to send the traffic to a non-existent IP address and thus never have a response packet sent.

Obfuscate HTTP headers

The FortiGate unit can obfuscate the HTTP header information being sent to external web servers to better cloak the source. By default this option is not enabled. To obfuscate HTTP headers, use the following CLI command:

```
config system global
  set http-obfuscate {none | header-only | modified | no-error}
end
```

Where:

none — do not hide the FortiGate web server identity.

header-only — hides the HTTP server banner.

modified — provides modified error responses.

no-error — suppresses error responses.



Chapter 4 Logging and Reporting

This FortiOS Handbook chapter contains the following sections:

[Logging overview](#) provides general information about logging. We recommend that you begin with this chapter as it contains information for both beginners and advanced users as well.

[Log devices](#) provides information about how to configure your chosen log device. Configuring multiple FortiAnalyzer units or Syslog servers is also included.

[Logging FortiGate activity](#) provides information about the different log types and subtypes, and how to enable logging of FortiGate features.

[The SQLite log database](#) provides information about SQLite statements as well as examples that you can use to base your own custom datasets on.

[Log message usage](#) provides general information about log messages, such as what is a log header. Detailed examples of each log type are discussed as well. For additional information about all log messages recorded by a FortiGate unit running FortiOS 4.0 and higher, see the [FortiGate Log Message Reference](#).

[Reports](#) provides information about how to configure reports if you have logged to a the FortiGate unit's hard disk SQL database.



Logging overview

This section explains what logging is in relation to your FortiGate unit, what a log message is, and log management practices. These practices can help you to improve and grow your logging requirements.

This section also includes information concerning log management practices that help you to improve and grow your logging requirements.

The following topics are included in this section:

- [What is logging?](#)
- [Log messages](#)
- [Log files](#)
- [Best Practices: Log management](#)

What is logging?

Logging records the traffic passing through the FortiGate unit to your network and what action the FortiGate unit took during its scanning process of the traffic. This recorded information is called a log message.

After a log message is recorded, it is stored within a log file which is then stored on a log device. A log device is a central storage location for log messages. The FortiGate unit supports several log devices, such as a FortiGuard Analysis and Management Service, and the FortiAnalyzer unit. A FortiGate unit's system memory and local disk can also be configured to store logs, and because of this, are also considered log devices.



You must subscribe to FortiGuard Analysis and Management Service so that you can configure the FortiGate unit to send logs to a FortiGuard Analysis server.

How the FortiGate unit records log messages

The FortiGate unit records log messages in a specific order, storing them on a log device. The order of how the FortiGate unit records log messages is as follows:

- 1 Incoming traffic is scanned
- 2 During the scanning process, the FortiGate unit performs necessary actions, and simultaneously are recorded
- 3 Log messages are sent to the log device

Example: How the FortiGate unit records a DLP event

- 1 The FortiGate unit receives incoming traffic and scans for any matches associated within its firewall policies containing a DLP sensor.
- 2 A match is found; the DLP sensor, `dlp_sensor`, had a rule within it called All-HTTP with the action Exempt applied to the rule. The sensor also has Enable Logging selected, which indicates to the FortiGate unit that the activity should be recorded and placed in the DLP log file.

- 3 The FortiGate unit exempts the match, and places the recorded activity (the log message) within the DLP log file.
- 4 According to the log settings that were configured, logs are stored on the FortiGate unit's local hard drive. The FortiGate unit places the DLP log file on the local hard drive.

Log messages

Log messages are recorded information containing specific details about what is occurring on your network. Within each log message there are fields. A field is two pieces of information that explain a specific part of the log message. For example, the action field contains login (action=login).

The fields within the log message are arranged into two groups; one group, which is first, is called the log header, and the second group is the log body, which contains all other fields. A log header from the FortiGate unit appears as follows when viewed in the Raw format:

```
2011-01-08 12:55:06 log_id=24577 type=dlp subtype=dlp pri=notice
vd=root
```

The log body appears as follows when viewed in the Raw format:

```
policyid=1 identidx=0 serial=73855 src="10.10.10.1" sport=1190
src_port=1190 srcint=internal dst="192.168.1.122" dport=80
dst_port=80 dst_int="wan1" service="https" status="detected"
hostname="example.com" url="/image/trees_pine_forest/" msg="data
leak detected(Data Leak Prevention Rule matched)" rulename="All-
HTTP" action="log-only" severity=1
```

Logs from other devices, such as the FortiAnalyzer unit and Syslog server, contain a slightly different log header. For example, when viewing FortiGate log messages on the FortiAnalyzer unit, the log header contains the following log fields when viewed in the Raw format:

```
itime=1302788921 date=20110401 time=09:04:23
devname=FG50BH3G09601792 device_id=FG50BH3G09601792
log_id=0100022901 type=event subtype=system pri=notice vd=root
```

Within the log header, there is a type field, and this field indicates the type of log file the log message is put into after it is recorded. The log header also contains the log_id field. The log_id field contains the unique identification number that is associated with that particular log message. For example, 32001. All log messages have a unique number that helps to identify them within their log file.

The log header also contains information about the log severity level and is indicated in the pri field. This information is important because the severity level indicates various severities that are occurring. For example, if the pri field contains alert, you need to take immediate action with regards to what occurred. There are six log severity levels.

The log severity level is the level at which the FortiGate unit records logs at. The log severity level is defined when configuring the logging location. The FortiGate unit logs all messages at and above the logging severity level you select. For example, if you select Error, the unit logs Error, Critical, Alert, and Emergency level messages.

Table 47: Log severity levels

Levels	Description
0 - Emergency	The system has become unstable.

Table 47: Log severity levels

Levels	Description
1 - Alert	Immediate action is required.
2 - Critical	Functionality is affected.
3 - Error	An error condition exists and functionality could be affected.
4 - Warning	Functionality could be affected.
5 - Notification	Information about normal events.
6 - Information	General information about system operations.

The Debug severity level, not shown in [Table 47](#), is rarely used. It is the lowest log severity level and usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly. Debug log messages are only generated if the log severity level is set to Debug. Debug log messages are generated by all types of FortiGate activities.

The log body contains the rest of the information of the log message, and this information is unique to the log message itself. There are no two log message bodies that are alike, however, there may be the same fields in most log message bodies, such as the srcintf log field or identidix log field.

For detailed information on all log messages, see the [FortiGate Log Message Reference](#).

Explanation of a log message

The following is a detailed explanation of the fields within a log message. It explains the log header fields, as well as the log body fields of a DLP log message.

Log header:

```
2011-01-04 12:55:06 log_id=24577 type=dlp subtype=dlp pri=notice
vd=root
```

date=(2010-08-03)	The year, month and day of when the event occurred in yyyy-mm-dd format.
time=(12:55:06)	The hour, minute and second of when the event occurred in the format hh:mm:ss.
log_id=(24577)	A five-digit unique identification number. The number represents that log message and is unique to that log message. This five-digit number helps to identify the log message.
type=(dlp)	The section of system where the event occurred.
subtype=(dlp)	The subtype category of the log message. See Table 47 on page 626 .
pri=(notice)	The severity level of the event. See Table 47 on page 626 .
vd=(root)	The name of the virtual domain where the action/event occurred in. If no virtual domains exist, this field always contains root.

Log body:

```

policyid=1 identidx=0 serial=73855 src="10.10.10.1" sport=1190
src_port=1190 srcint=internal dst="192.168.1.122" dport=80
dst_port=80 dst_int="wan1" service="https" status="detected"
hostname="example.com" url="/image/trees_pine_forest/" msg="data
leak detected(Data Leak Prevention Rule matched)" rulename="All-
HTTP" action="log-only" severity=1

```

policyid=(1)	The ID number of the firewall policy that applies to the session or packet. Any policy that is automatically added by the FortiGate will have an index number of zero.
identidx=(0)	The identity-based policy identification number. This field displays zero if the firewall policy does not use an identity-based policy; otherwise, it displays the number of the identity-based policy entry that the traffic matched. This number is not globally unique, it is only locally unique within a given firewall policy.
serial=(73855)	The serial number of the firewall session of which the event happened.
src=(10.10.10.1)	The source IP address.
sport=(1190)	The source port number.
src_port=(1190)	The source port number.
srcint=(internal)	The source interface name.
dst=(192.168.1.122)	The destination IP address.
dport=(80)	The destination port number.
dst_port=(80)	The destination port number.
dst_int=(wan1)	The destination interface name.
service=(https)	The IP network service that applies to the session or packet. The services displayed correspond to the services configured in the firewall policy.
status=(detected)	The action the FortiGate unit took.
hostname=(example.com)	The home page of the web site.
url=(/image/trees_pine_forest/)	The URL address of the web page that the user was viewing.
msg=(data leak detected(Data Leak Prevention Rule matched))	Explains the FortiGate activity that was recorded. In this example, the data leak that was detected matched the rule, All-HTTP, in the DLP sensor.
rulename=(All-HTTP)	The name of the DLP rule within the DLP sensor.
action=(log-only)	The action that was specified within the rule. In some rules within sensors, you can specify content archiving. If no action type is specified, this field display log-only.
severity=(1)	The level of severity for that specific rule.

Explanation of a debug log message

Debug log messages are only generated if the log severity level is set to Debug. The Debug severity level is the lowest log severity level and is rarely used. This severity level usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly. Debug log messages are generated by all types of FortiGate features.

The following is an example of a debug log message:

```
2010-01-25 17:25:54 log_id=93000 type=webfilter subtype=urlfilter
pri=debug msg="found in cache"
```

Table 48: Explanation of an example of a Debug log message

date=(2010-01-25)	The year, month and day of when the event occurred in the format yyyy-mm-dd.
time=(17:25:54)	The hour, minute and second of when the event occurred in the format hh:mm:ss.
log_id=(93000)	A five-digit unique identification number. The number represents that log message and is unique to that log message. This five-digit number helps to identify the log message.
type=(webfilter)	The section of system where the event occurred. There are eleven log types in FortiOS 4.0.
subtype=(urlfilter)	The subtype of the log message. This represents a policy applied to the FortiGate feature in the firewall policy.
pri=(debug)	The severity level of the event. There are six severity levels to specify.
msg=("found in cache")	Explains the activity or event that the FortiGate unit recorded.

Viewing log messages

There are two viewing options, Format and Raw. When you download the log messages from within the Log Access menu, you are downloading log messages that will be viewed in the Raw format. The Raw format displays logs as they appear within the log file. You can view logs messages in the Raw format using a text editor, such as Notepad. Format is in a more readable format, and you can easily filter information when viewing log messages this way. The Format view is what you see when viewing logs in the web-based manager, in *Log&Report > Log Access*.



You can also view log messages in the Log Viewer table. The table displays the fields that are within the log message, similar to Raw, but not all fields display. If you want to view the details of a log message, you should use the Raw format. For more information, see [“Viewing log messages and archives” on page 664](#).

The web-based manager is not the only place to view log messages. You can also view log messages from the CLI. For more information about viewing log messages, see [“Viewing log messages and archives” on page 664](#).

Log files

Each log message that is recorded by the FortiGate unit is put into a log file. The log file contains the log messages that belong to that log type, for example, traffic log messages are put in the traffic log file.

When downloading the log file from *Log&Report > Log Access*, the file name indicates the log type and the device on which it is stored on. This name is in the format <logtype>log<logdevice_logtype>.log. For example, tlog0100.log. The log device and log type part are in numerical format. In the example, tlog0100.log, 01 indicates that the traffic log file was stored on the unit's local hard drive and 00 indicates that it is a traffic log file.

The log devices that are indicated in a log file's name are as follows:

- 00 – indicates that the logs are stored on memory
- 01 – indicates that the logs are stored on the unit's local hard drive
- 02 – indicates that the logs are stored on a FortiAnalyzer unit
- 04 – indicates that the logs are stored on the FortiGuard Analytics server

The log type number that comes after the log device number in the log file's name is as follows:

00 – traffic log	01 – event log
02 – antivirus	03 – web filter
04 – IDS or attack log	05 – spam log or email filter log
06 – content	09 – DLP log
10 – Application control log	

In [Table 49](#), each of the five log files are explained.

Table 49: Log Types based on network traffic

Log Type	File name	Description
Traffic	tlog.log	The traffic log records all traffic to and through the FortiGate interface.
Event	elog.log	The event log records management and activity events. For example, when an administrator logs in or logs out of the web-based manager.

Table 49: Log Types based on network traffic

Log Type	File name	Description
UTM	ulog.log	<p>The UTM log file contains all UTM-related log messages, such as antivirus and attack. The following explains the log messages that are included in the UTM log file:</p> <p>Antivirus – records virus incidents in Web, FTP, and email traffic.</p> <p>Web – records HTTP FortiGate log rating errors including web content blocking actions that the FortiGate unit performs</p> <p>Attack – records attacks that are detected and prevented by the FortiGate unit</p> <p>Email Filter – records blocking of email address patterns and content in SMTP, IMAP, and POP3 traffic</p> <p>Application Control – records data detected by the FortiGate unit and the action taken against the network traffic depending on the application that is generating the traffic, for example, instant messaging software, such as MSN Messenger</p> <p>Data Leak Prevention – records log data that is considered sensitive and that should not be made public. This log also records data that a company does not want entering their network</p>
DLP archive	clog.log	<p>The DLP archive log, or clog.log, records the following log messages:</p> <p>email messages (using protocols such as SMTP or POP3)</p> <p>FTP events</p> <p>quarantine</p> <p>IPS packet</p> <p>instant messaging</p> <p>VoIP</p>
Netscan	nlog.log	The Network Vulnerability Scan log records vulnerabilities during the scanning of the network.

Best Practices: Log management

When the FortiGate unit records FortiGate activity, valuable information is collected that provides insight into how to better protect network traffic against attacks, including misuse and abuse. There is a lot to consider before enabling logging on a FortiGate unit, such as what FortiGate activities to enable and which log device is best suited for your network's logging needs. A plan can help you in deciding the FortiGate activities to log, a log device, as well as a backup solution in the event the log device fails.

This plan should provide you with an outline, similar to the following:

- what FortiGate activities you want and/or need logged (for example, DLP archives)
- the logging device best suited for your network
- if you want or require archival of log files

- ensuring logs are not lost in the event a failure occurs.

After the plan is implemented, you need to manage the logs and be prepared to expand on your log setup when the current logging requirements are outgrown. Good log management practices help you with these tasks.

Log management practices help you to improve and manage logging requirements. Logging is an ever-expanding tool that can seem to be a daunting task to manage. The following management practices will help you when issues arise, or your logging setup needs to be expanded.

- 1** Revisit your plan on a yearly basis to verify that your logging needs are being met by your current log setup. For example, your company or organization may require archival logging, but not at the very beginning. Archival logs are stored on either a FortiGate unit's local hard drive, a FortiAnalyzer unit, or a FortiGuard Analysis server.
- 2** Configure an alert message that will notify you of activities that are important to be aware about. For example, a branch office does not have a FortiGate administrator so you need to know, at all times, that the IPSec VPN tunnel is up and running. An alert email notification message can be configured for sending only IPSec tunnel errors.
- 3** Ensure that your backup solution is up-to-date. If you have recently expanded your log setup, you should also review your backup solution. The backup solution provides a way to ensure that all logs are not lost in the event that the log device fails or issues arise with the log device itself.



The SQLite log database

The SQLite log database is used to store logs generated by most FortiGate units with hard drives. SQLite is an embedded Relational Database Management System (RDBMS). SQLite supports most of the SQL-92 standard for Structured Query Language (SQL). SQLite is designed for managing data in RDBMS.

Reports include datasets and these datasets use SQL statements to retrieve the log information needed for reports. A large set of datasets are pre-defined on FortiGate units, however, you can create custom datasets but SQL knowledge is required.

This section explains how to create the statements that you can then use in datasets. This section also includes examples of SQL statements that you can use to base your own custom datasets on.

If you require more information, see the technical note [SQL Log Database Query](#), which includes information about SQL on the FortiAnalyzer unit.

The following topics are included in this section:

- [SQL overview](#)
- [SQLite database tables](#)
- [SQLite statement examples](#)
- [Troubleshooting SQL issues](#)



SQL logging is enabled by default on models that support the local SQLite database, and are running FortiOS 4.0 MR3 or higher.



If you have disabled SQL logging and have factory defaults on the FortiGate unit, and then upgrade, the upgrade will not automatically enable SQL logging.



If you are formatting a disk that contains more than just logs, all information on the disk will be lost.

SQL overview

The syntax for SQL queries is based on the SQLite3 syntax (see <http://www.sqlite.org/lang.html> for more information).

There is an additional convenience macro, `F_TIMESTAMP`, that allows you to easily specify a time interval for the query. It takes this form:

`F_TIMESTAMP(base_timestring, unit, relative value)`. For example, `F_TIMESTAMP('now', 'hour', '-23')` means “last 24 hours” or that the hour in the timestamp is 23 less than now. The FortiGate unit will automatically translate the macro into SQLite3 syntax.

You can use the following CLI commands to write SQL statements to query the SQLite database.

```
config report dataset
  edit <dataset_name>
    set query <sql_statement>
  next
end
```

For more information about specific examples that are used in creating custom datasets, see the [“SQLite statement examples”](#) on page 634.

SQLite database tables

The FortiGate unit creates a database table for each log type, when log data is recorded. If the FortiGate unit is not recording log data, it does not create log tables for that device.

The command syntax, `get report database schema`, allows you to view all the tables, column names and types that are available to use when creating SQL statements for datasets.

If you want to view the size of the database, as well as the log database table entries, use the `get log sql status` command. This command displays the amount of free space that is available as well as the first log database entry time and date and the last log database entry time and date.

The output of the `get log sql status` command contains information similar to the following:

```
Database size: 104856576
Free size in database: 670004416
Entry number:
  Event: 1263
  Traffic: 254039
  Attack: 4
  Antivirus: 8
  WebFilter: 5291
  DLP: 76544
  Application Control: 68103
  Netscan: 75
Total: 405331
First entry time: 2011-03-21 08:25:55
Last entry time: 2011-04-21 09:10:55
```

SQLite statement examples

The following examples help to explain SQL statements and how they are configured within a dataset. Datasets contain SQL statements which are used to query the SQL database.

Distribution of Applications by Type in the last 24 hours

This dataset is created to show the distribution of the type of applications that were used, and will use the application control logs to get this information.

CLI commands

```
config report dataset
```



```

edit "appctrl.Dist.Type.last24h"
  set query "select app_type, count(*) as totalnum from
    app_control_log where timestamp >=
    F_TIMESTAMP('now','hour','-23') and (app_type is not null
    and app_type!='N/A') group by app_type order by totalnum
    desc"
next

```

Explanation about the parts of the statement

- `edit "appctrl.Dist.Type.last24h"` - creates a new dataset with descriptive title.
- `F_TIMESTAMP('now','hour','-23')` - the `F_TIMESTAMP` macro covers the last 24 hours (from now until 23 hours ago).
- The application control module classifies each firewall session in `app_type`. One firewall session may be classified to multiple `app_types`. For example, an HTTP session can be classified to: HTTP and Facebook, as well as others.
- Some `app/app_types` may not be able to detected, then the 'app_type' field may be null or 'N/A'. These will be ignored by this query.
- The result is ordered by the total session number of the same `app_type`. The most frequent `app_types` will appear first.

Top 10 Application Bandwidth Usage Per Hour Summary

This dataset is created to show the bandwidth used by the top ten applications. Application control logs are used to gather this information.

CLI commands

```

config report dataset
edit "appctrl.Count.Bandwidth.Top10.Apps.last24h"
  set query "select (timestamp-timestamp%3600) as hourstamp,
    (CASE WHEN app!='N/A\' and app!='\'\' then app ELSE service
    END) as appname, sum(sent+rcvd) as bandwidth from
    traffic_log where timestamp >=
    F_TIMESTAMP('\now\',\'hour\',\'-23\') and (appname in
    (select (CASE WHEN app!='N/A\' and app!='\'\' then app ELSE
    service END) as appname from traffic_log where timestamp >=
    F_TIMESTAMP('\now\',\'hour\',\'-23\') group by appname
    order by sum(sent+rcvd) desc limit 10)) group by hourstamp,
    appname order by hourstamp desc"
next

```

Explanation about the parts of the statement

- `(timestamp-timestamp%3600) as hourstamp` - this calculates an "hourstamp" to indicate bandwidth per hour.
- `(CASE WHEN app!='N/A\' and app!='\'\' then app ELSE service END) as appname` - use the app as 'appname', or if it's undefined, use the service instead.
- `appname in (select (CASE WHEN app!='N/A\' and app!='\'\' then app ELSE service END) as appname from traffic_log where timestamp >= F_TIMESTAMP('\now\',\'hour\',\'-23\') group by appname order by sum(sent+rcvd) desc limit 10)` - selects the top 10 apps using most bandwidth

- `order by hourstamp desc` - this orders the results by descending hourstamp
- `LIMIT 10` - this lists only the top 10 applications.

Example of how to create a dataset containing attack name instead of attack ID

If you want to create a dataset that contains the attack name instead of the attack ID, use the following as a basis.

```
config report dataset
edit top_attacks_1hr
set log-type attack
set time-period last-n-hours
set period-last-n 1
set query "SELECT attack_id, COUNT( * ) AS totalnum FROM
$log WHERE $filter and attack_id IS NOT NULL GROUP BY
attack_id ORDER BY totalnum DESC LIMIT 10"
end
```

Troubleshooting SQL issues

If the query is unsuccessful, an error message appears in the results window indicating the cause of the problem. The following are issues that may arise when creating statements for datasets that query the SQL database.

SQL statement syntax errors

Here are some example error messages and possible causes:

You have an error in your SQL syntax (remote/MySQL) or ERROR: syntax error at or near... (local/PostgreSQL)

- Verify that the SQL keywords are spelled correctly, and that the query is well-formed.
- Table and column names are demarked by grave accent (`) characters. Single (`) and double (` `) quotation marks will cause an error.

No data is covered.

- The query is correctly formed, but no data has been logged for the log type. Verify that you have configured the FortiGate unit to save that log type. On the Log Settings page, make sure that the log type is checked.

Connection problems

If well formed queries do not produce results, and logging is turned on for the log type, there may be a database configuration problem with the remote database.

Ensure that:

- MySQL is running and using the default port 3306.

- You have created an empty database and a user who has read/write permissions for the database.


Here is an example of creating a new MySQL database named fazlogs, and adding a user for the database:

```
#Mysql -u root -p
mysql> Create database fazlogs;
mysql> Grant all privileges on fazlogs.* to 'fazlogger'@'*'
identified by 'fazpassword';
mysql> Grant all privileges on fazlogs.* to
'fazlogger'@'localhost' identified by 'fazpassword';
```

SQL database error

You may encounter the following error message (Figure 63) after upgrading or downgrading the FortiGate unit's firmware image.

Figure 63: Example of an SQL database error message that appears after logging in to the web-based manager



The error message indicates that the SQL database is corrupted and cannot be updated with the SQL schemas any more. When you see this error message, you can do one of the following:

- select *Cancel* and back up all log files; then select *Rebuild* to rebuild the database
- select *Rebuild* only after verifying that all log files are backed up to a safe location.

When you select *Cancel*, no logging is recorded by the FortiGate unit regardless of the log settings that are configured on the unit. When you select *Rebuild*, all logs are lost because the SQL database is erased and then rebuilt again. Logging resumes after the SQL database is rebuilt.

If you want to view the database's errors, use the `diag debug sqlldb-error-read` command in the CLI. This command indicates exactly what errors occurred, and what tables contain those errors.

Log files are backed up using the `execute log backup {alllogs | logs}` command in the CLI. You must use the text variable when backing up log files because the text variable allows you to view the log files outside the FortiGate unit. When you back up log files, you are really just copying the log files from the database to a specified location, such as a TFTP server.



Log devices

The FortiGate unit supports a variety of logging devices, including the FortiGuard Analysis and Management Service. This provides great flexibility when choosing a log device for the first time, as well as when logging requirements change.

This section explains how to configure your chosen log device, as well as how to configure multiple FortiAnalyzer units or Syslog servers. This section also includes how to log to a FortiGuard Analysis server, which is available if you subscribed to the FortiGuard Analysis and Management Service.

The following topics are included in this section:

- [Choosing a log device](#)
- [Example: Setting up a log device and backup solution](#)
- [Configuring the FortiGate unit to store logs on a log device](#)
- [Troubleshooting issues](#)
- [Testing FortiAnalyzer and FortiGuard Analysis server connections](#)
- [Connecting to a FortiAnalyzer unit using Automatic Discovery](#)
- [Uploading logs to a FortiAnalyzer or a FortiGuard Analysis server](#)



You may need to reschedule uploading or rolloing of log files because the size of log files is reduced in FortiOS 4.0 MR1 and higher. Reduction in size provides more storage room for larger amounts of log files on log devices.

Choosing a log device

When you have developed a plan that meets your logging needs and requirements, you need to select the log device that is appropriate for that plan. A log device must be able to store all the logs you need, and if you require archiving those logs, you must consider what log devices support this option.

During this process of deciding what log device meets your needs and requirements, you must also figure out how to provide a backup solution in the event the log device that the FortiGate unit is sending logs to has become unavailable. A backup solution should be an important part of your log setup because it helps you to maintain all logs and prevents lost logs, or logs that are not sent to the log device. For example, a daily backup of log files to the FortiAnalyzer unit occurs at 5 pm.

Log devices provide a central location for storing logs recorded by the FortiGate unit. The following are log devices that the FortiGate unit supports:

- FortiGate system memory
- Hard disk or AMC
- SQL database (for FortiGate units that have a hard disk)
- FortiAnalyzer unit
- FortiGuard Analysis server (part of FortiGuard Analysis and Management Service)
- Syslog server

- NetIQ WebTrends server

These logs devices, except for the FortiGate system memory and local hard disk, can also be used as a backup solution. For example, you configure logging to the FortiGate unit's local disk, but also configure logging to a FortiGuard Analysis server and archive logs to both the FortiGuard Analysis server and a FortiAnalyzer unit.



If you are formatting a disk that contains more than just logs, all information on the disk will be lost.

Example: Setting up a log device and backup solution

The following is an example of how to set up a log device and backup solution when you are integrating them into your network. This example does not include how to enable FortiGate activities for logging.

Your company has received a FortiAnalyzer unit and three new Syslog server software licenses. You install the FortiAnalyzer unit in as stated in its install guide, and you also install the Syslog server software on three servers.

To setup a log device and backup solution

- 1 On the FortiGate unit, log in to the CLI.
- 2 Enter the following to configure the FortiGate unit to send logs to the FortiAnalyzer unit.

```
config log fortianalyzer setting
  set status enable
  set ips-archive enable
  set server 172.16.120.154
  set address-mode static
  set upload-option realtime
end
```

- 3 Enter the following commands for the first Syslog server:

```
config log syslogd setting
  set status enable
  set server 192.168.16.121
  set reliable enable
  set csv enable
  set facility local1
  set source-ip
end
```

- 4 Enter the following for the second Syslog server:


```
config log syslogd2 setting
  set status enable
  set server 192.168.17.125
  set reliable enable
  set csv enable
  set facility local2
  set source-ip
end
```

- 5 Enter the following for the third Syslog server:

```
config log syslogd setting
  set status enable
```

```
set server 192.168.16.121
set reliable enable
set csv enable
set facility local3
set source-ip
end
```

Figure 64: Example of an integrated FortiAnalyzer unit and Syslog servers in a network



Configuring the FortiGate unit to store logs on a log device

After setting up and integrating your log device into the network, you can now configure the log device's settings for logging FortiGate activities. If you have multiple FortiAnalyzer units or Syslog servers, you must configure them in the CLI. For more information, see [“Logging to multiple FortiAnalyzer units or Syslog servers”](#) on page 38.

This topic includes the following:

- [Logging to the FortiGate unit's system memory](#)
- [Logging to the FortiGate unit's hard disk](#)
- [Logging to a FortiAnalyzer unit](#)
- [Logging to a FortiGuard Analysis server](#)
- [Logging to a Syslog server](#)
- [Logging to a WebTrends server](#)
- [Logging to multiple FortiAnalyzer units or Syslog servers](#)



If you experience issues, see [“Troubleshooting issues”](#) on page 650. This topic may not contain all the information you may need when troubleshooting a logging issue; if the topic cannot help you, see the Troubleshooting chapter in the [FortiOS Handbook](#).

Logging to the FortiGate unit's system memory

The system memory displays recent log entries and stores all log types, which includes archives and traffic logs. When the system memory is full, the FortiGate unit overwrites the oldest messages. All log entries stored in system memory are cleared when the FortiGate unit restarts.



By default, logging to memory is enabled.

When a hard disk is not present on a FortiGate unit, real-time logging is enabled by default. Real-time logging is recording activity as it happens.

To send logs to the unit's system memory- web-based manager

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Logging and Archiving* section, select the check box beside *Memory*.
- 3 Select a log level from the *Minimum log level* drop-down list.
- 4 Select *Apply*.

The FortiGate unit logs all messages at and above the logging severity level you select. If you want to archive IPS logs, use the CLI to enable this log.

To send logs to the unit's system memory - CLI

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log memory setting
  set diskfull <overwrite>
  set ips-archive {enable | disable}
  set status {enable | disable}
end
config log memory global-setting
  set max-size <byte_size>
  set full-final-warning-threshold <integer>
  set full-first-warning-threshold <integer>
  set full-second-warning-threshold <integer>
end
```

Use the `config log memory filter` command to disable the FortiGate features you do not want to log. By default, most FortiGate features are enabled in the `config log memory filter` command.

Logging to the FortiGate unit's hard disk

If your FortiGate unit contains a hard disk, you can configure the FortiGate unit to store logs on the disk. When configuring logging to a hard disk, you can also configure uploading of those logs to a FortiAnalyzer unit or to a FortiGuard Analysis server. If you want to upload the logs to a FortiAnalyzer unit or to a FortiGuard Analysis server, see [“Uploading logs to a FortiAnalyzer or a FortiGuard Analysis server” on page 652](#).

When the FortiGate unit does not have a hard disk, real-time logging is enabled by default. Real-time logging is recording activity as it happens. Real-time logging is enabled only in the CLI and you must format the disk (if not already formatted) before first use.

When logging to the unit's hard disk, you must also enable SQL logging, which is enabled only in the CLI. For more information, see [“Enabling SQL logging” on page 662](#).



You must include a storage location, or logs will not be recorded. If your FortiGate unit has an SQLite log database, you must enable SQL logging as well. For more information about how to configure SQL logging, see [“Enabling SQL logging” on page 662](#).



If you have disabled SQL logging and have factory defaults on the FortiGate unit, and then upgrade, the upgrade will not automatically enable SQL logging.

To log to the hard disk on a FortiGate unit - CLI

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log disk setting
  set status {enable | disable}
  set ips-archive {enable | disable}
  set dlp-archive-quota <integer>
  set max-log-file-size <size>
  set storage <storage_location>
  set diskfull {nolog | overwrite}
  set report-quota <integer>
  set log-quota <quota_size>
  set dlp-archive-quota <quota_size>
  set report-quota <quota_size>
  set upload {enable | disable}
  set upload-format {compact | text}
  set sql-max-size <maximum_size>
  set sql-max-size-action {overwrite | nolog}
  set sql-oldest-entry <integer>
  set drive-standby-time
  set full-first-warning-threshold
  set full-second-warning-threshold
  set full-final-warning-threshold
  set ms-per-transaction <integer>
  set rows-per-transaction <integer>
end
```

Use the `config log disk filter` command to disable the FortiGate features you do not want to log. By default, most FortiGate features are enabled in the `config log disk filter` command.

Logging to a FortiAnalyzer unit

A FortiAnalyzer unit can log all FortiGate activity that is available for logging, including archiving. You can also configure the FortiGate unit to upload logs on a regular basis to the FortiAnalyzer unit.

When logging to a FortiAnalyzer unit, you do not need a hard drive to configure logging to a FortiAnalyzer unit. Encryption is supported by default and logs are sent using IPsec or SSL VPN. When the FortiAnalyzer unit and FortiGate unit have SSL encryption, both must choose a setting for the `enc-algorithm` command for encryption to take place. By default, this is enabled and the default setting is SSL communication with high and medium encryption algorithms. The setting that you choose must be the same for both.

If you are using the FortiGate and FortiAnalyzer-VM images, and these are evaluation software, you will only be able to use low encryption.

The following procedure assumes that you have only one FortiAnalyzer unit to configure. If you are configuring more than one FortiAnalyzer unit, you must configure the other units in the CLI. Use the procedures in [“Logging to multiple FortiAnalyzer units or Syslog servers” on page 38](#) to configure multiple FortiAnalyzer units.

If you want to connect to a FortiAnalyzer unit using automatic discovery, see [“Connecting to a FortiAnalyzer unit using Automatic Discovery” on page 41](#). If you are going to connect to the FortiAnalyzer using the automatic discovery feature, see [“Connecting to a FortiAnalyzer unit using Automatic Discovery” on page 652](#).

To send logs to a FortiAnalyzer unit

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log fortianalyzer setting
  set status {enable | disable}
  set server <ip_address>
  set address-mode {auto-discovery | static}
  set use-hdd {enable | disable}
  set enc-algorithm {enable | disable}
  set localid <id_number>
  set conn-timeout <seconds>
  set roll-schedul {daily | weekly}
  set roll-time <hh:mm>
  set diskfull {nolog | overwrite}
  set monitor-keepalive-period <seconds>
  set monitor-failure-retry-period <seconds>
  set upload
end
```

- 3 To send logs remotely to the FortiAnalyzer unit, use the following command syntax:

```
config log remote setting
  set status enable
  set destination FAZ
end
```

The `remote` command appears only after configuring logging to a FortiAnalyzer unit. It is also available if you configure logging to a FortiGuard Analysis server.

Logging to a FortiGuard Analysis server

You can configure logging to a FortiGuard Analysis server after registering for the FortiGuard Analysis and Management Service. The following procedure assumes that you have already configured the service account ID in *System > Maintenance > FortiGuard*.

Logging to a FortiGuard Analysis server is configured in the CLI. Uploading logs to the FortiGuard Analysis server is configured in both the web-based manager and the CLI. For more information about uploading logs to a FortiGuard Analysis server, see [“Uploading logs to a FortiAnalyzer or a FortiGuard Analysis server” on page 652](#).

The following assumes that you have already configured the FortiGuard Analysis and Management Service account ID information in *System > Maintenance > FortiGuard*.

To send logs to a FortiGuard Analysis server - CLI

- 1 Log in to the CLI.
- 2 Enter the following command syntax:


```
config log fortiguard setting
  set status {enable | disable}
  set quotafull {nolog | overwrite}
end
```
- 3 To send logs remotely to the FortiGuard Analysis server, use the following command syntax:


```
config log remote setting
  set status enable
  set destination FAZ
end
```

The remote command appears only after configuring logging to a FortiGuard Analysis server. This command is also available if you configure logging to a FortiAnalyzer unit.

Logging to a Syslog server

The Syslog server is a remote computer running syslog software. Syslog is a standard for forwarding log messages in an IP network. Syslog servers capture log information provided by network devices.

The following procedure configures one Syslog server. You can configure up to three Syslog servers. Use the procedure in [“Configuring multiple Syslog servers” on page 39](#) to configure multiple Syslog servers.

To send logs to a syslog server - web-based manager

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Logging and Archiving* section, select the check box beside *Syslog*.
After you select the check box, the *Syslog* options appear.
- 3 Enter the appropriate information for the following:

IP/FDQN	Enter the domain name or IP address of the syslog server.
Port	Enter the port number for communication with the syslog server, usually port 514.
Minimum log level	Select a log level the FortiGate unit will log all messages at and above that logging severity level.
Facility	Facility indicates to the syslog server the source of a log message. By default, the FortiGate reports facility as local7. You can change the Facility if you want to distinguish log messages from different FortiGate units.
Enable CSV Format	Select to have logs formatted in CSV format. When you enable CSV format, the FortiGate unit produces the log in Comma Separated Value (CSV) format. If you do not enable CSV format, the FortiGate unit produces plain text files.
- 4 Select *Apply*.

To log to a Syslog server - CLI

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log syslogd setting
  set status {enable | disable}
  set server <address_ipv4>
  set port <port_integer>
  set csv {enable | disable}
  set facility {alert | audit | auth | authpriv | clock | cron
    | daemon | ftp | kernel | local0 | local1 | local2 |
    local3 | local4 | local5 | local6 | local7 | lpr | mail |
    news | ntp | syslog | user | uucp}
end
```

Enabling reliable syslog

The reliable syslog feature is based on RFC 3195. Reliable syslog logging uses TCP, which ensures that connections are set up, including packets transmitted.

There are several profiles available for reliable syslog, but only the RAW profile is currently supported on the FortiGate unit. The RAW profile is designed to provide a high-performance, low-impact footprint using essentially the same format as the existing UDP-based syslog service. The reliable syslog feature is available on FortiGate units running FortiOS 4.0 MR1 or higher.

When you enable reliable syslog in the CLI, TCP is used. The default setting, `disable`, uses UDP. TCP ensures that packets are transmitted easily, as well as connections are set up.

The reliable Syslog feature automatically resets the port number to TCP 601. This is based on the recommendation in RFC 3195.

The following procedure assumes that you have already configured the Syslog server.

To enable reliable syslog

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log syslogd setting
  set status enable
  set reliable enable
end
```

Logging to a WebTrends server

A WebTrends server is a remote computer, similar to a Syslog server, running NetIQ WebTrends firewall reporting server. FortiGate log formats comply with WebTrends Enhanced Log Format (WELF) and are compatible with NetIQ WebTrends Security Reporting Center and Firewall Suite 4.1.

To send logs to a WebTrends server, log in to the CLI and enter the following commands:

```
config log webtrends setting
  set server <address_ip4>
  set status {disable | enable}
end
```

Example

This example shows how to enable logging to and set an IP address for a remote NetIQ WebTrends server.

```
config log webtrends settings
  set status enable
  set server 172.25.82.145
end
```

Logging to multiple FortiAnalyzer units or Syslog servers

FortiOS 4.0 allows you to configure multiple FortiAnalyzer units or multiple Syslog servers, ensuring that all logs are not lost in the event one of them fails.

You can configure multiple FortiAnalyzer units or Syslog servers within the CLI.

This topic includes the following:

- [Configuring multiple FortiAnalyzer units](#)
- [Configuring multiple Syslog servers](#)
- [Example of configuring multiple FortiAnalyzer units](#)

Configuring multiple FortiAnalyzer units

Fortinet recommends that you contact a FortiAnalyzer administrator first, to verify that the IP addresses of the FortiAnalyzer units you want to send logs to are correct and that all FortiAnalyzer units are currently installed with FortiAnalyzer 4.0 firmware.

If VDOMs are enabled, you can configure multiple FortiAnalyzer units or Syslog servers for each VDOM.

The following procedure does not contain how to enable logging of FortiGate features within the CLI. Most FortiGate features are, by default, enabled within the `log filter` command in the CLI. You should disable the FortiGate features that you do not want logged within the `log filter` command.

If you want to test the FortiAnalyzer unit's connection to the FortiGate unit, see [“Testing the FortiAnalyzer configuration” on page 651](#).

To enable logging to multiple FortiAnalyzer units

- 1 Log in to the CLI.
- 2 Enter the following commands, using `encrypt` and `psksecret` if you want to encrypt the connection:

```
config log fortianalyzer setting
  set status enable
  set server <faz_ip address>
  set encrypt [disable | enable]
  set psksecret <password>
  set localid <identification_ipsectunnel>
  set conn-timeout <value_seconds>
end
```

- 3 Enter the following commands for the second FortiAnalyzer unit, using `encrypt` and `psksecret` if you want to encrypt the connection:

```
config log fortianalyzer2 setting
  set status {disable | enable}
  set server <fortianalyzer_ipv4>
  set encrypt [disable | enable]
```

```

    set psksecret <password>
    set localid <identification_ipsectunnel>
    set ver-1 {disable | enable}
    set conn-timeout <value_seconds>
end

```

- 4 Enter the following commands for the last FortiAnalyzer unit, using `encrypt` and `psksecret` if you want to encrypt the connection:

```

config log fortianalyzer3 setting
    set status enable
    set server <faz_ip address>
    set encrypt [disable | enable]
    set psksecret <password>
    set localid <identification_ipsectunnel>
    set conn-timeout <value_seconds>
end

```

Configuring multiple Syslog servers

When configuring multiple Syslog servers (or one Syslog server), you can configure reliable delivery of log messages from the Syslog server. Configuring of reliable delivery is available only in the CLI.

If VDOMs are enabled, you can configure multiple FortiAnalyzer units or Syslog servers for each VDOM.

The following procedure does not contain how to enable logging of FortiGate features within the CLI. Most FortiGate features are, by default, enabled within the `log filter` command in the CLI. You should disable the FortiGate features that you do not want logged within the `log filter` command.

To enable logging to multiple Syslog servers

- 1 Log in to the CLI.

- 2 Enter the following commands:

```

config log syslogd setting
    set csv {disable | enable}
    set facility <facility_name>
    set port <port_integer>
    set reliable {disable | enable}
    set server <ip_address>
    set status {disable | enable}
end

```

- 3 Enter the following commands to configure the second third Syslog server:

```

config log syslogd2 setting
    set csv {disable | enable}
    set facility <facility_name>
    set port <port_integer>
    set reliable {disable | enable}
    set server <ip_address>
    set status {disable | enable}
end

```

- 4 Enter the following commands to configure the third Syslog server:

```

config log syslogd3 setting
    set csv {disable | enable}

```

```
set facility <facility_name>
set port <port_integer>
set reliable {disable | enable}
set server <ip_address>
set status {disable | enable}
end
```

Example of configuring multiple FortiAnalyzer units

The IT department at your organization's headquarters discovers that their one and only FortiAnalyzer unit is not working properly; however, it is soon up and running again. When it is working properly again, your department's managers realize that two day's worth of logs are lost. Your IT manager asks you to install and configure two FortiAnalyzer units onto the network so that logs are not lost again.

In this example, you have already installed and configured the two FortiAnalyzer units and are now ready to configure sending logs from the FortiGate unit to the two FortiAnalyzer units.

The three units are named as follows: FortiAnalyzer_Main (main log storage unit), FortiAnalyzer_1 (first back up), and FortiAnalyzer_2 (second back up).

To configure logging to multiple FortiAnalyzer units

- 1 Verify that FortiAnalyzer_Main configuration is correct and communication between the FortiGate unit and FortiAnalyzer_Main is working properly.

You can do this by using Test Connectivity from the FortiGate unit. See [“Testing the FortiAnalyzer configuration” on page 41](#).

- 2 Go to *System > Dashboard > Status* and locate the CLI console widget.
- 3 Log in to the CLI using the CLI Console widget on the Status page.
- 4 Enter the following command syntax for FortiAnalyzer_1:

```
config log fortianalyzer2 setting
set status enable
set server 10.10.20.125
set encrypt enable
set psksecret it123456
set localid ipsec_vpn1
set conn-timeout 1200
set max-buffer-size 800
end
```

- 5 Enter the following command syntax for FortiAnalyzer_2:

```
config log fortianalyzer3 setting
set status enable
set server 10.10.22.120
set encrypt enable
set psksecret it123456
set localid ipsec_vpn1
set conn-timeout 1200
set max-buffer-size 800
end
```

- 6 Enter the following command syntax to log the FortiGate features:

```
config log fortianalyzer2 filter
```

```
get log fortianalyzer filter
```

The `get log fortianalyzer filter` displays the FortiGate features that are enabled by default.

- 7 Enter the following variables to disable the four FortiGate features that you do not want:

```
set wanopt-traffic disable
set netscan disable
set vulnerability disable
end
```

- 8 Repeat steps 6 and 7 for FortiAnalyzer_2.

Logs are now configured to go to multiple FortiAnalyzer units.

Troubleshooting issues

From time to time, issues may arise due to connectivity or logging has stopped altogether. The following provides information on troubleshooting these issues.

Unable to connect to a supported log device

After configuring logging to a supported log device, and you have tested the connection, you find you cannot connect. Use the following to help you find out what happened and resolve the issue:

- verify that the information you entered is correct; it could be a simple mistake within the IP address or you did not select *Apply* on the Log Settings page which would not apply the settings.
- use `execute ping` to see if you can ping to the log device; if unable to ping to the log device check to see if the log device itself working and that it is on the network.
- if the unit has stopped logging, see [“FortiGate unit has stopped logging” on page 650](#).

FortiGate unit has stopped logging

If your FortiGate unit contains an SQL database, and when you log into the web-based manager and you see an SQL database error message, it is because the SQL database has become corrupted. You must make sure to back up your logs at that point and then rebuild the database.

If the FortiGate unit stopped logging to a device, test the connection between both the FortiGate unit and device using the `execute ping` command. The log device may have been turned off, is upgrading to a new firmware version, or just not working properly. After determining that there are no logs being sent, see [“Unable to connect to a supported log device” on page 650](#).

Testing FortiAnalyzer and FortiGuard Analysis server connections

When you want to verify that the connection between a FortiAnalyzer unit and a FortiGate unit is successful, or a FortiGate unit and a FortiGuard Analysis server, you can use the *Test Connectivity* button.

This topic contains the following:

- [Testing the FortiAnalyzer configuration](#)
- [Testing the FortiGuard Analysis server configuration](#)

Testing the FortiAnalyzer configuration

After configuring FortiAnalyzer settings, you can test the connection between the FortiGate unit and the FortiAnalyzer unit to ensure the connection is working properly. This enables you to view the connection settings between the FortiGate unit and the FortiAnalyzer unit.

To test the connection

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Logging and Archiving* section, under *Upload logs remotely*, select *Test Connectivity* beside the FortiAnalyzer unit's IP address field.

When you select *Test Connectivity*, a window appears with general information about the FortiAnalyzer unit, disk space available and used on the FortiAnalyzer unit, as well as privileges that the FortiGate unit while connected to the FortiAnalyzer unit.

- 3 Select *Close* when you are done viewing the connection status and general information.

Testing the FortiGuard Analysis server configuration

After configuring logging to a FortiGuard Analysis server, you can verify that they connection is working properly by selecting *Test Connectivity* which is located beside the *Account ID* field.

Similar to *Test Connectivity* for FortiAnalyzer, this information includes your contract's expiry date as well as daily volume and disk quota amounts.

To test the connection

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Logging and Archiving* section, under *Upload logs remotely*, select *Test Connectivity* beside the *Account ID* field.
- 3 Select *Close* when you are done viewing the connection status and general information.

Using diag sys logdisk usage

The `diag sys logdisk usage` command allows you to view detailed information about how much space is currently being used for logs. This is useful when you see a high percentage, such as 92 percent for the disk's capacity. The FortiGate unit uses only 75 percent of the available disk capacity to avoid a high storage amount so when there is a high percentage, it refers to the percentage of the 75 percent that is available. For example, 92 percent of the 75 percent is available.

The following is an example of what you may see when you use `diag sys logdisk usage` command on a unit with no VDOMs configured:

```
diag sys logdisk usage
```

The following appears:

```
Total HD usage: 176MB/3011 MB
Total HD logging space: 22583MB
Total HD logging space for each vdom: 22583MB
HD logging space usage for vdom "root": 30MB/22583MB
```

Connecting to a FortiAnalyzer unit using Automatic Discovery

Automatic Discovery is a method of establishing a connection to a FortiAnalyzer unit by using the FortiGate unit to find a FortiAnalyzer unit on the network. The Fortinet Discovery Protocol (FDP) is used to locate the FortiAnalyzer unit. Both units must be on the same subnet to use FDP, and they must also be able to connect using UDP.

When you select Automatic Discovery, the FortiGate unit uses HELLO packets to locate any FortiAnalyzer units that are available on the network within the same subnet. When the FortiGate unit discovers the FortiAnalyzer unit, the FortiGate unit automatically enables logging to the FortiAnalyzer unit and begins sending log data.

To connect using automatic discovery

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log fortianalyzer setting
  set status {enable | disable}
  set server <ip_address>
  set gui-display {enable | disable}
  set address-mode auto-discovery
end
```

If your FortiGate unit is in Transparent mode, the interface using the automatic discovery feature will not carry traffic. For more information about how to enable the interface to also carry traffic when using the automatic discovery feature, see the Fortinet Knowledge Base article, [Fortinet Discovery Protocol in Transparent mode](#).



The FortiGate unit searches within the same subnet for a response from any available FortiAnalyzer units

Uploading logs to a FortiAnalyzer or a FortiGuard Analysis server

You can upload logs that are stored on the FortiGate unit's local disk to a FortiAnalyzer unit or to a FortiGuard Analysis server. This provides a way to back up your logs to another storage location. You can upload all types of log files.

You need to have access to both the CLI and the web-based manager when configuring uploading of logs. The upload time and interval settings can only be done in the CLI; however, for uploading of logs to a FortiAnalyzer unit, you can modify the upload time and interval settings from the web-based manager after you have configured it in the CLI.



Uploading logs to an AMC disk is no longer available.

To upload logs to a FortiGuard Analysis server

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Logging and Archiving* section, select the check box beside *Upload logs remotely*.
- 3 Select *FortiGuard Analysis and Management Service (Daily at 00:00)*.
- 4 Enter the account ID in the *Account ID* field.
- 5 Select *Apply*.

- 6 To configure the daily upload time, log in to the CLI.
- 7 Enter the following to configure when the upload occurs, and the time when the unit uploads the logs:

```
config log fortiguard setting
    set upload-interval {daily | weekly | monthly}
    set upload-time <hh:mm>
end
```

To upload logs to a FortiAnalyzer unit

- 1 Go to *Log&Report > Log Config > Log Settings*.
- 2 In the *Logging and Archiving* section, select the check box beside *Upload logs remotely*.
- 3 Select *FortiAnalyzer (Daily at 00:60)*.
- 4 Enter the FortiAnalyzer unit's IP address in the *IP Address* field.
- 5 To configure the daily upload time, log in to the CLI.
- 6 Enter the following to configure when the upload occurs, and the time when the unit uploads the logs:

```
config log fortianalyzer setting
    set upload-interval {daily | weekly | monthly}
    set upload-time <hh:mm>
end
```

- 7 To change the upload time, in the web-based manager, select *Change* beside the upload time period, and then make the changes in the Upload Schedule window. Select *OK*.



Logging FortiGate activity

The FortiGate unit, depending on its configuration, can have a lot of network activity flowing through it. The FortiGate unit provides settings to record the activity, referred to in this document as log settings.

This section explains the FortiGate activity that can be logged, as well as how to log FortiGate activity. This section includes how to configure an alert notification email, filtering and customizing the display of logs messages on the web-based manager, as well as viewing quarantined files.

The following topics are included in this section:

- [Logs](#)
- [Configuring logging of FortiGate activity on your FortiGate unit](#)
- [Viewing log messages and archives](#)
- [Customizing the display of log messages](#)
- [Alert email messages](#)

Logs

Logs record FortiGate activity, providing detailed information about what is happening on your network. This recorded activity is found in log files, which are stored on a log device. However, logging FortiGate activity requires configuring certain settings so that the FortiGate unit can record the activity. These settings are often referred to as log settings, and are found in most UTM features, such as profiles, and include the event log settings, found on the Event Log page.

Log settings provide the information that the FortiGate unit needs so that it knows what activities to record. This topic explains what activity each log file records, as well as additional information about the log file, which will help you determine what FortiGate activity the FortiGate unit should record.

This topic includes the following:

- [Traffic](#)
- [Event](#)
- [Data Leak Prevention](#)
- [Application control](#)
- [Antivirus](#)
- [Web Filter](#)
- [IPS \(attack\)](#)
- [Packet logs](#)
- [Email filter](#)
- [Archives \(DLP\)](#)
- [Network scan](#)

Traffic

Traffic logs record the traffic that is flowing through your FortiGate unit. Since traffic needs firewall policies to properly flow through the unit, this type of logging is also referred to as firewall policy logging. Firewall policies control all traffic that attempts to pass through the FortiGate unit, between FortiGate interfaces, zones and VLAN sub-interfaces.

Logging traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic or Log Violation Traffic)
- packet comes into an inbound interface
- a possible log packet is sent regarding a match in the firewall policy, such as URL filter
- traffic log packet is sent, per firewall policy
- packet passes and is sent out an interface

Traffic log messages are stored in the traffic log file. Traffic logs can be stored any log device, even system memory.

If you have enabled and configured WAN Optimization, you can enable logging of this activity in the CLI using the `config wanopt setting` command. These logs contain information about WAN Optimization activity and are found in the traffic log file. When configuring logging of this activity, you must also enable logging within the security policy itself, so that the activity is properly recorded.

Other Traffic

The traffic log also records interface traffic logging, which is referred to as other traffic. Other traffic is enabled only in the CLI. When enabled, the FortiGate unit records traffic activity on interfaces as well as firewall policies. Logging other traffic puts a significant system load on the FortiGate unit and should be used only when necessary.

Logging other traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic or Log Violation Traffic) and other-traffic
- packet comes into an interface
- interface log packet is sent to the traffic log that is enabled on that particular interface
- possible log packet is sent regarding a match in the firewall policy, such as URL filter
- interface log packet is sent to the traffic log if enabled on that particular interface
- packet passes and is sent out an interface
- interface log packet is sent to traffic (if enabled) on that particular interface

Event

The event log records administration management as well as FortiGate system activity, such as when a configuration has changed, admin login, or high availability (HA) events occur. Event logs are an important log file to record because they record FortiGate system activity, which provides valuable information about how your FortiGate unit is performing.

Event logs help you in the following ways:

- keep track of configuration setting changes
- IPsec negotiation, SSL VPN and tunnel activity
- quarantine events, such as banned users

- system performance
- HA events and alerts
- firewall authentication events
- wireless events on models with WiFi capabilities
- activities concerning modem and internet protocols L2TP, PPP and PPPoE
- VIP activities
- AMC disk's bypass mode
- VoIP activities that include SIP and SCCP protocols.

The FortiGate unit records event logs only when events are enabled.

Data Leak Prevention

Data Leak Prevention logs, or DLP logs, provide valuable information about the sensitive data trying to get through to your network as well as any unwanted data trying to get into your network. The DLP rules within a DLP sensor can log the following traffic types:

- email (SMTP, POP3 or IMAP; if SSL content SMTPS, POP3S, and IMAPS)
- HTTP
- HTTPS
- FTP
- NNTP
- IM

A DLP sensor must have log settings enabled for each DLP rule and compound rule, as well as applied to a firewall policy so that the FortiGate unit records this type of activity. A DLP sensor can also contain archiving options, which these logs are then archived to the log device.

NAC Quarantine

Within the DLP sensor, there is an option for enabling NAC Quarantine. The NAC Quarantine option allows the FortiGate unit to record details of DLP operation that involve the ban and quarantine actions, and sends these to the event log file. The NAC Quarantine option must also be enabled within the Event Log settings. When enabling NAC quarantine within a DLP Sensor, you must enable this in the CLI because it is a CLI-only command.

Application control

Application control logs provide detailed information about the traffic that an application is generating, such as Skype. The application control feature controls the flow of traffic from a specific application, and the FortiGate unit examines this traffic for signatures that the application generates.

The log messages that are recorded provide information such as the type of application being used (for example P2P software), and what type of action the FortiGate unit took. These log messages can also help you to determine the top ten applications that are being used on your network. This feature is called application control monitoring and you can view the information from a widget on the Executive Summary page.

The application control list that is used must have *Enabled Logging* selected within the list, as well as logging enabled within each application entry. Each application entry can also have packet logging enabled. Packet logging for application control records the packet when an application type is identified, similar to IPS packet logging.

Logging of application control activity can only be recorded when an application control list is applied to a firewall policy, regardless of whether or not logging is enabled within the application control list.

Antivirus

Antivirus logs are recorded when, during the antivirus scanning process, the FortiGate unit finds a match within the antivirus profile, which includes the presence of a virus or grayware signature. Antivirus logs provide a way to understand what viruses are trying to get in, as well as additional information about the virus itself, without having to go to the FortiGuard Center and do a search for the detected virus. The link is provided within the log message itself.

These logs provide valuable information about:

- name of the detected virus
- name of the oversized file or infected file
- action the FortiGate unit took, for example, a file was blocked
- URL link to the FortiGuard Center which gives detailed information about the virus itself

The antivirus profile must have log settings enabled within it so that the FortiGate unit can record this activity, as well as having the antivirus profile applied to a firewall policy.

Web Filter

Web filter logs record HTTP traffic activity. These log messages provide valuable and detailed information about this particular traffic activity on your network. Web filtering activity is important to log because it can inform you about:

- what types of web sites are employees accessing
- if users try to access a banned web site and how often this occurs
- network congestion due to employees accessing the Internet at the same time
- alerts you to web-based threats when users surf non-business-related web sites

Web Filter logs are an effective tool to help you determine if you need to update your web filtering settings within a web filter profile due to unforeseen web-based threats, or network congestion. These logs also inform you about web filtering quotas that were configured for filtering HTTP traffic as well.

You must configure log settings within the web filter profile as well as apply it to a firewall policy so that the FortiGate unit can record web filter logs.

IPS (attack)

IPS logs, also referred to as attack logs, record attacks that occurred against your network. Attack logs contain detailed information about whether the FortiGate unit protected the network using anomaly-based defense settings or signature-based defense settings, as well as what the attack was.

The IPS or attack log file is especially useful because the log messages that are recorded contain a link to the FortiGuard Center, where you can find more information about the attack. This is similar to antivirus logs, where a link to the FortiGuard Center is provided as well and informs you of the virus that was detected by the FortiGate unit.

An IPS sensor with log settings enabled must be applied to a firewall policy so that the FortiGate unit can record the activity.

Packet logs

When you enable packet logging within an IPS signature override or filter, the FortiGate unit examines network packets, and if a match is found, saves them to the attack log. Packet logging is designed to be used as a diagnostic tool that can focus on a narrow scope of diagnostics, rather than a log that informs you of what is occurring on your network.

You should use caution when enabling packet logging, especially within IPS filters. Filter configuration that contains thousands of signatures could potentially cause a flood of saved packets, which would take up a lot of storage space on the log device. This would also take a great deal of time to sort through all the log messages, as well as consume considerable system resources to process.

You can archive packets, however, you must enable this option on the Log Setting page. If your log configuration includes multiple FortiAnalyzer units, packet logs are only sent to the primary, or first FortiAnalyzer unit. Sending packet logs to the other FortiAnalyzer units is not supported.

Email filter

Email filter logs, also referred to as spam filter logs, records information regarding the content within email messages. For example, within an email filter profile, a match is found that finds the email message to be considered spam.

Email filter logs are recorded when the FortiGate unit finds a match within the email filter profile and logging settings are enabled within the profile.

Archives (DLP)

Recording DLP logs for network use is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Archived logs are usually saved for historical use and can be accessed at any time. IPS packet logs can also be archived, within the Log Settings page.

You can use the two default DLP sensors that were configured specifically for archiving log data, Content_Archive and Content_Summary. They are available in *UTM > Data Leak Prevention > Sensor*. Content_Archive provides full content archiving, while Content_Summary provides summary archiving.

You must enable the archiving to record log archives. Logs are not archived unless enabled, regardless of whether or not the DLP sensor for archiving is applied to the firewall policy.

Network scan

Network scan logs are recorded when a scheduled scan of the network occurs. These log messages provide detailed information about the network's vulnerabilities regarding software as well as the discovery of any vulnerabilities.

A scheduled scan must be configured, as well as logging enabled within the Event Log settings, for the FortiGate unit to record these log messages.

Configuring logging of FortiGate activity on your FortiGate unit

There are many FortiGate activities that you can enable logging for on your FortiGate unit. Most are configured in the UTM menu; however, there are some activities that require configuration from other locations. For example, enabling events is done from *Log&Report > Log Config > Log Setting*.

When you are ready to enable logging of the FortiGate activities you have chosen, you must apply them to a firewall policy. When applied to a firewall policy, the FortiGate unit determines whether to record the activity or event based on what is applied on the firewall policy and configured within the Log Config page.

The following explains how to enable and configure logging on your FortiGate unit.

This topic includes the following:

- [Enabling logging within a firewall policy](#)
- [Enabling logging of events](#)
- [Enabling SQL logging](#)
- [Configuring IPS packet logging](#)
- [Configuring NAC quarantine logging](#)

Enabling logging within a firewall policy

You must enable logging within a firewall policy, as well as the options that are applied to a firewall policy, such as UTM features. Event logs are enabled within the Event Log page.

To enable logging within an existing firewall policy

- 1 Go to *Firewall > Policy > Policy*.
- 2 Expand to reveal the policy list of a policy.
- 3 Select the firewall policy you want to enable logging on and then select *Edit*.
- 4 To log all general firewall policy traffic, select the check box beside *Log Allowed Traffic*.
- 5 On the firewall policy's page, select the check box beside *UTM*.
- 6 Select the protocol option list from the drop-down list beside *Protocol Options*.
By default, the *Protocol Options* check box is selected. You must choose a list from the drop-down list.
- 7 For each row under *UTM*, select the check box beside each of the profiles and/or sensors that you want applied to the policy; then select the profile or sensor from the drop-down list as well.
- 8 To apply other options, such as application lists, repeat step 7 and choose the option from the drop-down list.
- 9 Select *OK*.



You need to set the logging severity level to *Notification* when configuring a logging location to record traffic log messages.

Enabling logging of events

The event log records management and activity events, such as when a configuration has changed, admin login, or high availability (HA) events occur.

When you are logged in to VDOMs, certain options may not be available, such as VIP ssl event or CPU and memory usage events. You can enable event logs only when you are logged in to a VDOM; you cannot enable event logs in the root VDOM.

To enable the event logs

- 1 Go to *Log&Report > Log Config > Log Setting*.

2 In the *Event Log Settings* section, select the check box beside *Enable*.

3 Select one or more of the following:

System activity event	All system-related events, such as ping server failure and gateway status.
IPSec negotiation event	All IPSec negotiation events, such as process and error reports.
DHCP service event	All DHCP-events, such as the request and response log.
L2TP/PPTP/PPPoE service event	All protocol-related events, such as manager and socket create processes.
Admin event	All administration events, such as user logins, resets, and configuration updates.
HA activity event	All high availability events, such as link, member, and stat information.
Firewall authentication event	All firewall-related events, such as user authentication.
Pattern update	All pattern update events, such as antivirus and IPS pattern updates and update failures.
Configuration event	All configuration changes
GTP service event (FortiOS Carrier only)	All GTP service events.
Notification event	All notification events.
MMS statistics event (FortiOS Carrier only)	All MMS statistics events.
Explicit web proxy event	All explicit web proxy events.
SSL VPN user authentication event	All administrator events related to SSL VPN, such as SSL configuration and CA certificate loading and removal.
SSL VPN administration event	All administration events related to SSL VPN, such as SSL configuration and CA certificate loading and removal.
SSL VPN session event	All session activity such as application launches and blocks, timeouts, verifications and so on.
VIP ssl event	All server-load balancing events that are happening during SSL session, especially details about handshaking.
VIP server health monitor event	All related VIP server health monitor events that occur when the VIP health monitor is configured, such as an interface failure.
WiFi activity event	All WiFi activities, such as Rogue AP.

CPU & memory usage (every 5 min)	Real-time CPU and memory events only, at 5-minute intervals.
VoIP event	All VoIP activity, such as SIP and SCCP violations.
AMC interface enters bypass mode event	All AMC interface bypass mode events that occur.
NAC Quarantine event	All endpoint activity that have quarantined hosts when Endpoint NAC is checking hosts.
DNS lookup	All lookups that are DNS-related.

4 Select *Apply*.

5 Select *OK*.

Enabling SQL logging

SQL logging is enabled by default on the FortiGate unit; however, if you have disabled SQL logging and have factory defaults on the FortiGate unit, and then you upgrade the FortiGate unit, the upgrade will not automatically enable SQL logging. When this occurs, use the following procedure to enable SQL logging.

The following procedure enables all logs except for DLP archives.

To enable SQL logging

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config log disk setting
config sql-logging
set app-crtl enable
set attack enable
set dlp enable
set event enable
set netscan enable
set spam enable
set traffic enable
set virus enable
set webfilter enable
end
```

Configuring IPS packet logging

IPS packet log messages are recorded in the attack log file. You must enable packet logging within the IPS sensor's filter itself because the FortiGate unit will not log the IPS packets with only logging enabled within the IPS sensor or the filter itself.

To configure IPS packet logging

- 1 Go to *UTM > Intrusion Protection > IPS Sensor*.
- 2 Select the IPS sensor that you want to enable IPS packet logging on, and then select *Edit*.
- 3 Within the sensor, select the filter to enable IPS packet logging for and then select *Edit*.
- 4 In *Signature Settings*, select *Enable all for Packet Logging*.

- 5 Select *OK*.
- 6 Repeat steps 3 to 5 for each filter that you want IPS packet logging enabled for.

If you want to configure the packet quota, number of packets that are recorded before alerts and after attacks, use the following procedure.

To configure additional settings for IPS packet logging

- 1 Log in to the CLI.
- 2 Enter the following to start configuring additional settings:

```
config ips settings
    set ips-packet quota <integer>
    set packet-log-history <integer>
    set packet-log-post-attack <integer>
end
```

Configuring NAC quarantine logging

NAC quarantine logging is enabled within the Event Log settings page. However, you must have a NAC profile applied to a firewall policy for the FortiGate unit to record these log messages.

To configure NAC quarantine logging

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 If not enabled, select the check box beside *Enable* which is located in the Event Log Settings section of the page.
- 3 If not already enabled, select *NAC Quarantine event*.
- 4 Select *Apply*.
- 5 Go to *Firewall > Policy > Policy*.
- 6 Select the firewall policy that you want to apply the NAC profile to, and then select *Edit*.
- 7 Within the UTM section, select the check box beside *Enable Endpoint NAC* and then select the endpoint profile from the drop-down list.
- 8 To add a disclaimer, select the check box beside *Enable Disclaimer*, and then enter the comment in the *Comments (maximum 63 characters)* field.
- 9 Select *OK*.
- 10 Log in to the CLI.
- 11 Enter the following to enable NAC quarantine in the DLP sensor:

```
config dlp sensor
    edit <dlp_sensor_name>
        set nac-quar-log enable
    end
```

Viewing log messages and archives

Depending on the log device, you may be able to view logs within the web-based manager or CLI on the FortiGate unit. If you have configured a FortiAnalyzer unit, local hard disk, or system memory, you can view log messages from within the web-based manager or CLI. If you have configured either a Syslog or WebTrends server, you will not be able to view log messages from the unit's web-based manager or CLI. The unit also does not currently support viewing logs from FortiGuard Analysis and Management Service from within the web-based manager or CLI.

Log messages and log archives can be viewed from the Log & Archive Access menu. Archived logs are stored on FortiAnalyzer units, a FortiGate unit's local disk or system memory, and a FortiGuard Analysis server. The Log & Archive Access menu displays the archived logs only when archiving is enabled and logs are being archived by the unit. The submenus are based on the log file, for example the UTM log file, which contains log messages that contain information regarding UTM activity, such as virus activity and application control activity. Viewing log archives is the same as viewing log messages, for example, *Log & Report > Log & Archive Access > E-mail Archive*.

From the Log & Archive Access menu, you can view detailed information about the log message in a log view table, located (by default) at the bottom of the page. Each page contains this log message viewer table, as well as the option of downloading the raw log file. Raw log format is how the log message displays in the log file, and contains sometimes contains additional log fields that are not present in the log message viewer table on the page. When viewing log messages on the web-based manager, you are viewing them in Formatted format.

Viewing logs in Raw format allows you to view all log fields, as well as making a log file available regardless of whether you are archiving logs or not. You download the log file by selecting *Download Raw Log*. The log file is named in the following format: `<log_type><log_number><log_subtype_number>.log`. For example, `clog0112.log`, which is the archive log downloaded. The time period is the day and month of when the log was downloaded, not the time period of the log messages within the file itself.

You do not have to view log messages from only the web-based manager. You can view log messages from the CLI as well using the `execute log display` command. This command allows you to see specific log messages that you already configured within the `execute log filter` command. The `execute log filter` command configures what log messages you will see, how many log messages you can view at one time (a maximum of 1000 lines of log messages), and the type of log messages you can view. For example, log messages from the DLP log file.

This topic contains the following:

- [Viewing log messages from the web-based manager and CLI](#)
- [Viewing log messages using the log table](#)

Viewing log messages from the web-based manager and CLI

If you have configured logging to a FortiAnalyzer unit, FortiGate unit's local disk or system memory, you can view log messages from either the web-based manager or CLI. The following procedures explain how to view log messages from the Log Access menu in the web-based manager, and how to view log messages from within the CLI.



If the FortiGate unit is running FortiOS 4.0 MR2 or lower, when viewing log messages in the Raw format in *Memory*, the ten-digit log ID number is used; however, when viewing the same log messages, in Raw format, in *Disk*, the five-digit log ID number is used (except for traffic logs which have only one-digit log IDs). This five-digit log identification number is used because of log size reduction that occurred in FortiOS 4.0 MR1.

To view log messages from the web-based manager

- 1 Go to *Log&Report > Log & Archive Access*.
- 2 Select the log menu that you want to view log messages in.
For example, the attack log messages in *Log&Report > Log & Archive Access > Traffic Log*.
- 3 Within the page, use any one of the following to view each log message:
 - *Download Raw Log* – downloads the log file to your PC. See “[Downloading log messages and viewing them from your computer](#)” on page 667.
 - *Column Settings* – customize what columns display on the page. See “[Customizing the display of log messages](#)” on page 668
 - *Filter Settings* – filter the information within the page. See “[Customizing the display of log messages](#)” on page 668
 - *Detailed Information* – display the log table on the right side of the page, at the bottom (default), or hide the log table. See
- 4 To display current log messages on the page, select *Refresh*.

To view log messages from the CLI

- 1 Enter the following to configure how the log messages will be displayed, as well as what log messages you want to display:


```
execute log filter category <category_number>
execute log filter start-line <line_number>
execute log filter view-lines <lines_per_view>
```
- 2 Enter the following to display the logs messages within the CLI:


```
execute log display
```
- 3 Log messages appear and stop when the maximum number of view-lines is reached.

Example: Viewing DLP log messages from the CLI

The following is an example of viewing DLP log messages from the CLI.

To view log messages from the CLI - example

- 1 Enter the following to configure the display of the DLP log messages:


```
execute log filter category 9
execute log filter start-line 1
execute log filter view-lines 20
```
- 2 Enter the following to view DLP log messages:

```
execute log display
600 logs found
20 logs returned
```

The first 20 log messages from the DLP log file display.

Quarantine

Within the Log & Archive Access menu, you can view detailed information about each quarantined file. The information can either be sorted or filtered, depending on what you want to view.

You must enable quarantine settings within an antivirus profile and the destination must be configured in the CLI using the `config antivirus quarantine` command. The destination can be either a FortiAnalyzer unit or local disk.

Sort the files by file name, date, service, status, duplicate count (DC), or time to live (TTL). Filter the list to view only quarantined files with a specific status or from a specific service.

On *Log&Report > Log & Archive Access > Quarantine*, the file quarantine list displays the following information about each quarantined file.

Quarantine page

Lists all files that are considered quarantined by the unit. On this page you can filter information so that only specific files are displayed on the page.

Source	Either <i>FortiAnalyzer</i> or <i>Local disk</i> , depending where you configure to quarantined files to be stored.
Sort by	Sort the list. Choose from: <i>Status</i> , <i>Service</i> , <i>File Name</i> , <i>Date</i> , <i>TTL</i> , or <i>Duplicate Count</i> . Select <i>Apply</i> to complete the sort. Filter the list. Choose either <i>Status</i> (infected, blocked, or heuristics) or <i>Service</i> (IMAP, POP3, SMTP, FTP, HTTP, IM, or NNTP). Select <i>Apply</i> to complete the filtering. Heuristics mode is configurable through the CLI only.
Filter	If your unit supports SSL content scanning and inspection Service can also be IMAPS, POP3S, SMTPS, or HTTPS. For more information, see the UTM chapter of the FortiOS Handbook.
Apply	Select to apply the sorting and filtering selections to the list of quarantined files.
Delete	Select to delete the selected files.
Page Controls	Use the controls to page through the list.
Remove All Entries	Removes all quarantined files from the local hard disk. This icon only appears when the files are quarantined to the hard disk.
File Name	The file name of the quarantined file.
Date	The date and time the file was quarantined, in the format dd/mm/yyyy hh:mm. This value indicates the time that the first file was quarantined if duplicates are quarantined.
Service	The service from which the file was quarantined (HTTP, FTP, IMAP, POP3, SMTP, IM, NNTP, IMAPS, POP3S, SMTPS, or HTTPS).
Status	The reason the file was quarantined: <i>infected</i> , <i>heuristics</i> , or <i>blocked</i> .

Status	Specific information related to the status, for example, "File is infected with "W32/Klez.h"" or "File was stopped by file block pattern."
Description	
DC	Duplicate count. A count of how many duplicates of the same file were quarantined. A rapidly increasing number can indicate a virus outbreak.
TTL	Time to live in the format hh:mm. When the TTL elapses, the FortiGate unit labels the file as EXP under the TTL heading. In the case of duplicate files, each duplicate found refreshes the TTL. The TTL information is not available if the files are quarantined on a FortiAnalyzer unit.
Upload status	Y indicates the file has been uploaded to Fortinet for analysis, N indicates the file has not been uploaded. This option is available only if the FortiGate unit has a local hard disk.
Download	Select to download the corresponding file in its original format. This option is available only if the FortiGate unit has a local hard disk.
Submit	Select to upload a suspicious file to Fortinet for analysis. This option is available only if the FortiGate unit has a local hard disk.

Downloading log messages and viewing them from your computer

You can view log messages in Raw format by downloading them from the web-based manager to your computer. The log file that is downloaded contains all the log messages that were recorded by the unit. When you are ready to view log messages in their Raw format, use a text editor, such as Notepad.

The following procedures can be used when you want to view log archives in Raw format.

To download log messages

- 1 Go to *Log&Report > Log & Archive Access* and then select the submenu that you want to download log messages from.

For example, *Log&Report > Log & Archive Access > UTM Log*

- 2 Within the page, select *Download Raw Log*.
- 3 Save the log file to your computer.

To view downloaded log messages

- 1 On your computer, locate the log file.
- 2 Open the log file with a text editor.
- 3 Scroll up and down to view each log message within the file.

Viewing log messages using the log table

You can view log messages using the log table, located at the bottom of the page within the Log & Archive Access menu. The log table can be positioned on the page either at the right or at the bottom. You can also hide the log table as well.

To view log messages using the log table

- 1 Go to *Log&Report > Log & Archive Access*.
- 2 Select the submenu that you want to access logs from.

For example, DLP log messages in *Log&Report > Log & Archive Access > UTM Log*.

- 3 From within the page, select inside the row of the log message that you want to view.
The log message row is highlighted and the log message appears in the log view table, located at the bottom of the page.
- 4 To close the table, select the arrow beside *Detailed Information* and then select *Hidden*.
- 5 To view the next log message from the table, use your keyboard's up and down arrows, or select the next log message row.
- 6 To view the log table at the right side of the page, select the arrow beside *Detailed Information* and then select *On Right*.

Monitoring the recording activity of logs on the FortiGate unit

When viewing the logs messages, you may want to know the activity of what is being recorded by the FortiGate unit. You can view this information in a user-friendly way by going to *Log&Report > Monitor > Logging Monitor*. The information displays in a bar chart on the page.

The Logging Monitor allows you to view the log file entries that have been recorded by the unit on a daily basis. For example, on Monday the Logging Monitor shows there are 20 entries. The information on the page is not saved when the week is up; it resets the information every week.

You can view detailed information about the log entries by selecting a bar in the chart; the Log Activity for <day> appears to show a break down of the log entries for that day. You can select Return to go back to the Logging Monitor page at any time.

Customizing the display of log messages

After the configuration is completed, and logging of FortiGate activities begins, you can customize and filter the log messages you see in the Log Access menu. Filtering and customizing the display provides a way to view specific log information without scrolling through pages of log messages to find the information.

Filtering information is similar to customizing, however, filtering allows you to enter specific information that indicates what should appear on the page. For example, including only log messages that appeared on April 4 between the hours of 8:00 and 8:30 am.

Customizing log messages allows you to remove or add columns to the page, allowing you to view certain information. The most columns represent the fields from within a log message, for example, the user column represents the user field, as well as additional information. If you want to reset the customized columns on the page to their defaults, you can within *Column Settings* by selecting *Default*.

The following is an example of how to filter and customize the display of application control log messages. Use the following example to help you to filter and customize the display of log messages in the web-based manager.

Filtering and customizing application control log messages example

The following example displays all application control log messages that do not contain the source IP address 10.10.10.1, as well as not displaying the Message, Date and Time columns on the page.

To filter and customize log messages

- 1 Go to *Log&Report > Log & Archive Access > UTM Log*.

- 2 On the page, select *Column Settings*.
- 3 In the *Show these fields in this order*, remove each of the following by selecting the column name and then using the <- arrow:
 - *Message*
 - *Date*
 - *Time*
- 4 Select *OK*.
- 5 Select *Filter Settings*.
- 6 In *Filters*, select *Add new filter*.
- 7 In the *Field* drop-down list, select *UTM (type)*, and in *UTM Type*, select *Application Control* and use the -> arrow to move it to the other column.
- 8 Select *Add new filter*.
- 9 In the *Field* drop-down list, select *Src*, enter 10.10.10.1 in the field and then select the check box beside *NOT*.
- 10 Select *OK*.

Alert email messages

Alert email messages provide notification about activities or events logged. These email messages also provide notification about log severities that are recorded, such as a critical or emergency.

You can send alert email messages to up to three email addresses. Alert messages are also logged and can be viewed from the Log Access menu. Alert messages are recorded in the event-system log file.

This topic contains the following:

- [Configuring an alert email message](#)
- [Configuring an alert email for notification of FortiGuard license expiry](#)

Configuring an alert email message

You can use the alert email feature to monitor logs for log messages, and to send email notification about a specific activity or event logged. For example, you require notification about administrators logging in and out. You can also base alert email messages on the severity levels of the logs. The FortiGate unit does not currently support SSL/TLS connections for SMTP servers, so you must choose an SMTP server that does not need SSL/TLS when configuring the SMTP server settings.

Before configuring alert email, you must configure at least one DNS server if you are configuring with an Fully Qualified Domain Server (FQDN). The FortiGate unit uses the SMTP server name to connect to the mail server, and must look up this name on your DNS server. You can also specify an IP address.

To configure an alert email message

- 1 Go to *Log&Report > Log Config > Alert E-mail*.

- 2 On the Alert E-mail page, enter the information for the SMTP server.
If the SMTP user requires authentication, enter the information after selecting the check box beside *Enable*.



The FortiGate unit currently does not support SSL/TLS connections with email servers, for example, Gmail. You must use an SMTP server that does not need an SSL/TLS connection.

- 3 Select *Apply* to apply the SMTP server information.
- 4 To verify these settings are correct, select *Test Connectivity*.
- 5 To send an alert email message based on a log's severity, select *Send to alert email for logs based on severity*, and then select a severity from the *Minimum log level* drop-down list.
- 6 To send an alert email message based on different activities, such as an administrator logging in and out, select from the following options:

Interval Time (1-9999 minutes)	Enter the minimum time interval between consecutive alert emails. Use this to rate-limit the volume of alert emails.
Intrusion detected	Select if you require an alert email message based on attempted intrusion detection.
Virus detected	Select if you require an alert email message based on virus detection.
Web access blocked	Select if you require an alert email message based on blocked web sites that were accessed.
HA status changes	Select if you require an alert email message based on HA status changes.
Violation traffic detected	Select if you require an alert email message based on violated traffic that is detected by the FortiGate unit.
Firewall authentication failure	Select if you require an alert email message based on firewall authentication failures.
SSL VPN login failure	Select if you require an alert email message based on any SSL VPN logins that failed.
Administrator login/logout	Select if you require an alert email message based on whether administrators log in or out.
IPSec tunnel errors	Select if you require an alert email message based on whether there is an error in the IPSec tunnel configuration.
L2TP/PPTP/PPPoE errors	Select if you require an alert email message based on errors that occurred in L2TP, PPTP, or PPPoE.
Configuration changes	Select if you require an alert email message based on any changes made to the FortiGate configuration.
FortiGuard license expiry time (1-100 days)	Enter the number of days before the FortiGuard license expiry time notification is sent.

Disk usage (1-99%) Enter a number for the disk usage threshold, in percent.

FortiGuard log quota usage Select if you require an alert email message based on the FortiGuard Analysis server log disk quota getting full.

7 Select *Apply*.



The default minimum log severity level is Alert. If the FortiGate unit collects more than one log message before an interval is reached, the FortiGate unit combines the messages and sends out one alert email.

Configuring an alert email for notification of FortiGuard license expiry

You can configure an alert email to notify you prior to when the FortiGuard license will actually expire. By sending this type of alert email, users are reminded sooner about their license requiring renewal soon, rather than later.

To configure an alert email for notification of FortiGuard license expiry - web-based manager

- 1 Go to *Log&Report > Log Config > Alert Email*.
- 2 Configure the SMTP server, email from and to fields, and if applicable, authentication.
- 3 Verify that *Select Send Alert email* is selected.
- 4 Select the check box beside *FortiGuard license expiry time: n (1-100 days)*.
The default for the expiry time is 15 days. This means that 15 days before the actual expiry date, an alert message is sent informing you that your license will expire in 15 days.
- 5 Enter a number for the number of days prior to the expiry date in the field provided.
For example, you want to be notified five days before the expiry date (December 31), an email is sent to the specified email address on December 27, five days before December 31.
- 6 Select *Apply*.

If you have configured FortiGate system memory as your log device, logging alert email notifications for FortiGuard license expiry requires you to enable event and admin in the log memory filter command. Use the following procedure when you want to log this event and if your log device is system memory. If you have enabled system memory on a FortiGate unit that has a local disk, you do not have to use the following procedure.

All other log devices, including the FortiGate unit's local disk, log alert messages by default. You can find the alert email logs within the event-system log file.

To configure logging of an alert email notification of FortiGuard license expiry (memory only)- CLI

- 1 Log in to the CLI.
- 2 To enable logging of an alert email notification using system memory, enter the following command syntax:

```
config log memory setting
    set status enable
end
config log memory filter
    set event enable
```

```
set admin enable  
end
```



Log message usage

FortiGate log messages present detailed accounts of an event or activity that happened on your network recorded by the FortiGate unit. These log messages provide valuable information about your network that inform you about attacks, misuse and abuse, and traffic activity that may provide useful when troubleshooting an issue, testing a configuration setting, or verifying that a setting is working properly.

If you require information about specific log messages, such as log message 32001, see the [FortiGate Log Message Reference](#).

This section explains how to use logs when troubleshooting, as well as verify settings and testing purposes.

The following topics are included in this section:

- [Using log messages to help when issues arise](#)
- [How to use log messages to help verify settings and for testing purposes](#)

Using log messages to help when issues arise

Log messages are very useful when an issue arises, such as when an SSL VPN tunnel does down. Log messages can help you when locating the problem, and can provide a starting point to go towards solving the problem. They can also indicate that an issue has arisen. For example, when an SSL VPN connection is not working properly, an easy way to find out what happened is to go directly to the log messages and view the ones for SSL VPN, in the event log.

HA log messages indicate lost neighbor information

When the FortiGate unit is in HA mode, you may notice the following log messages within the event log:

```
2011-04-16 11:05: 34 device_id=FG50B11111111 log_id=35001
type=event subtype=ha pri=critical vd=root msg= "HA slave
heartbeat interface internal lost neighbor information"
```

OR

```
2011-04-16 11:16 11:05: 40 device_id=FG50B11111111 log_id=35001
type=event subtype=ha pri=critical vd=root msg= "Virtual cluster 1
of group 0 detected new joined HA member"
```

OR

```
2011-04-16 11:16 11:05: 40 device_id=FG50B11111111 log_id=35001
type=event subtype=ha pri=critical vd=root msg= "HA master
heartbeat interface internal get peer information"
```

The log messages occur within a given time, and indicate that the units within the cluster are not aware of each other anymore. These log messages provide the information you need to fix the problem.

Alert email test configuration issues example

You have just configured the settings for an alert email message that will be sent to your inbox whenever web access is blocked, violation traffic is detected, and IPSec tunnel errors. You select *Test Connectivity* to test the configuration; however, there is no test email in your inbox.

You immediately go to the log messages to see what is going on.

1 *Log&Report > Log Access > Event.*

2 The following message appears in the *Message* field: Failed to send alert email from mail.example.com.

Raw format of the log message in the email:

```
2011-04-05 13:34:31 log_id=0100200003 type=event subtype=system
vd=root pri=notice user=system ui=system action=alert-email
status=failure count=5 msg="Failed to send alert email from
mail.example.com to myemailaddress@example.com"
```

The above log message indicates that the alert email message could not be sent to your inbox. You must verify the SMTP server, and if incorrect, enter the correct SMTP server and try again.

3 If you get the following log message, it means that the address you entered in the SMTP server field was not correct. You need to verify that it is the correct address.

```
2011-04-05 13:34:31 log_id=200003 type=event subtype=system
vd=root pri=notice user="system" ui="system" action="unknown"
status="failure" msg="Can't resolve the IP address of
mail.example.com" vd=root pri=notice
```

An alert email log message is not recorded if the alert email configuration is correct. If it is correct, you should see a test alert email message in your inbox and no log messages concerning alert email configuration settings, similar to those in this example.

How to use log messages to help verify settings and for testing purposes

You can use log message to verify settings, as well as for testing connections. The following are examples that explain various situations where you can use log message to verify settings within a feature, and also for testing purposes.

Verifying to see if a network scan was performed example

When you have configured an asset for a scan and a schedule for when to scan, you may want to verify that the asset is working properly. The log file that records network scanning is the netscan log.

You have just configured an asset to scan. You verify the scan by selecting *Discover Assets* on the Asset page (*Endpoint > Network Vulnerability Scan > Asset*). After the scan is complete, you view the netscan logs. The following four log messages appear:

4097

This log message indicates that the scanned was performed.

```
2011-04-05 13:34:31 log_id=4097 type=netscan subtype=discovery
pri=notice vd=root action=scan start=1275363661 end=1275363719
engine=1.053 plugin=1.098
```

4100

Both these log messages indicate that the scan discovered two separate service-detection events.

```
2011-04-05 13:34:31 log_id=4100 type=netscan subtype=discovery
pri=notice vd=root action=service-detection ip=10.10.20.3
service=microsoft-ds proto=tcp prot=445
```

```
2011-04-05 13:34:31 log_id=4100 type=netscan subtype=discovery
pri=notice vd=root action=service-detection ip=10.10.20.3
service=netbios-ssn proto=tcp port=139
```

4099

This log message indicates that the scan discovered a host, and explains the host's operating system, family that the OS belongs to, the type of generation, and the vendor or company that created the OS.

```
2010-04-05 13:34:31 log_id=1600004099 type=netscan
subtype=discovery pri=notice vd=root action=host-detection
ip=10.10.20.3 os="Windows XP" os_family="Windows"
os_gen="NT/2K/XP" os_vendor="Microsoft"
```

Testing for the FortiGuard license expiry log message example

You recently configured an alert message that would notify when the FortiGuard Analysis service license expires. You know that the expiry date is 2010-08-30, but want to verify that the alert message will notify you both by email and by log message, so you entered 100 days. You made sure that the log configuration includes logging alert messages as well.

After a while, you receive an alert email message in your inbox stating the following:

```
Message meets Alert condition
date=2011-06-08 time=11:55:52 devname=FG50BH3G09601792
device_id=FG50BH3G09601792 log_id=0104032014 type=event
subtype=admin pri=warning vd=root msg="FortiGuard analysis
service license will expire in 82 day(s) "
```

You go to *Log&Report > Log Access > Event* to view the license expiry log message using a filter on the Message column:

- 1 In *Log&Report > Log Access > Event*, select the filter icon beside the *Message* column's name.
- 2 In the Edit Filters window, select the check box beside *Enable*.
- 3 In the *Text* field, enter the sentence found in the alert email message: FortiGuard analysis service license will expire in 82 day(s).
- 4 Select *OK*.

The following log message appears in the first line on the first page of the Event page (Raw format):

```
2011-06-03 13:55:22 log_id=32014 type=event subtype=admin
pri=warning vd=root msg="FortiGuard analysis service license
will expire in 82 day(s) "
```

Using diag log test to verify logs are sent to a log device

After setting up a log device, you may want to verify that everything is working properly, including sending of logs to the log device. The diag log test command is used to create various test logs.


```
diag log test
```

The command can also be used when a Syslog server is not receiving certain logs. When the command is entered in the CLI, an output similar to the following appears below the line:

```
generating a system event message with level - warning
generating an infected virus message with level - warning
generating a blocked virus message with level - warning
generating a URL block message with level - warning
generating a DLP message with level - warning
generating an attack detection message with level - warning
generating an application control IM message with level -
information
generating an antispam message with level - notification
generating an allowed traffic message with level - notice
generating a wanopt traffic log message with level - notification
generating a HA event message with level - warning
generating netscan log messages with level - notice
generating a VOIP event message with level - information
generating authentication event messages
```

These log messages can be viewed from the Log Access menu. The following is a test log message that may be generated and recorded by the FortiGate unit. It is shown as it would be displayed in Raw format in system memory.

```
2011-08-10 09:34:22 log_id=0508020480 type=emailfilter
subtype=smtp pri=notice policyid=12345 identidx=67890 serial=312
user="user" group="group" vd="root" src=1.1.1.1 sport=2560
src_port=2560 src_int="lo" dst=2.2.2.2 dport=5120 dst_port=5120
dst_int="eth0" service=mm1 carrier_ep="carrier endpoint"
profile="N/A" profilegroup="N/A" profiletype="N/A" status=detected
from="from@xxx.com" to="to@xxx.com" tracker="Tracker"
msg="SpamEmail"
```



Reports

Reports provide a way to analyze log data without manually going through a large amount of logs to get to the information you need. This section explains how to configure a FortiOS report and how to modify the existing default FortiOS UTM report. The FortiOS default UTM report is a report that gathers UTM activity information and compiles it into a report. This section also explains how to view these reports.

The following topics are included in this section:

- [FortiOS reports](#)
- [Configuring a FortiOS report](#)
- [Viewing reports](#)



You can only configure reports if the FortiGate unit has a hard disk and SQL logging is enabled.



Configuring reports from other log devices, such as a Syslog server, are not supported. If you want to configure a report from logs stored on a FortiAnalyzer unit, you must go directly to the FortiAnalyzer unit itself; starting in FortiOS 4.0 MR3, support for configuring a FortiAnalyzer schedule is no longer available. From FortiOS 4.0 MR3 and onward, executive summary reports are no longer supported and existing summary reports are not carried forward.

FortiOS reports

Reports provide a clear, concise overview of what is happening on your network based on log data, without manually going through large amounts of logs. Reports can be configured on a FortiGate unit or a FortiAnalyzer unit. However, in this document only FortiOS reports are explained. FortiOS reports are the reports that are generated on the FortiGate unit. FortiAnalyzer reports are configured on a FortiAnalyzer unit and for information about those reports, see the [FortiAnalyzer Administration Guide](#).

FortiOS reports are configured from logs stored on the FortiGate unit's hard drive. These reports, generated by the FortiGate unit itself, provide a central location for both configuring and generating reports. A default FortiOS report, called the FortiGate UTM Daily Activity Report, is available for you to modify to your requirements. The default report provides a way to quickly and easily set up your own report from within the web-based manager. The default FortiOS report is a report that compiles UTM activity from various UTM-related logs, such as virus and attack logs.

FortiOS reports consist of multiple parts, regardless of whether its the default FortiOS report or a report that you have configured from scratch, and these parts are configured separately and added to the layout. These parts of a FortiOS report are:

- charts (including datasets within the charts themselves)
- themes (including styles which are within the themes themselves)
- images

- layout

Charts are used to display the log information in a clear and concise way using graphs and tables. Charts contain datasets, which are SQLite queries and help the FortiGate unit to add specific log information into the chart using the log information that is stored in the SQLite database on the local hard disk. If you want to configure a chart, you must configure the dataset first. Datasets are required for each chart, and if there is no dataset included in a chart, the chart will not be saved.

Themes provide a one-step style application for report layouts. Themes contain various styles, including styles for the table of contents, headings, headers and footers, as well as the margins of the report's pages. Themes are applied to layouts. The styles that are applied to themes are configured separately.

You can easily upload your company or organization's logo to use within a report. By uploading your company or organization's logo and applying it to a report, you provide a personalized report that is recognizable as your company or organization's report. The image must be in JPEG, JPG or PNG format.

Layouts provide a way to incorporate the charts, images, and themes that are configured to create a formatted report. A layout is used as a template by the FortiGate unit to compile and then generate the report.

You can reset the reports you have configured, as well as the default FortiOS report you modified, to default settings. When you reset reports to default settings, any configured reports that you created from scratch are lost. The `execute report-config reset` command resets the reports to default settings. If you are going to reset the reports to their default settings, you should back up the current configuration file before doing so, in the event you want to revert back to the reports you previously created and/or modified.

Configuring a FortiOS report

The following explains how to configure a FortiOS report as well as how to modify the existing default FortiOS report.

This topic contains the following:

- [Modifying the default FortiOS report](#)
- [Configuring charts, datasets, themes and styles for a report](#)
- [Importing images for the report](#)



You can only configure reports if the FortiGate unit has a hard disk and SQL logging is enabled.

Modifying the default FortiOS report

The default FortiOS report can be modified so that it meets your requirements for a report. This default report is located in *Log&Report > Report Access*.

The default report is broken up into the following submenus:

- Cover Page
- Bandwidth and Application Usage
- Web Usage
- Emails
- Threats

- VPN Usage

Each submenu is a page of the default FortiOS report. For example, the Bandwidth and Application Usage is the page in the report that contains information regarding bandwidth usage for WAN Opt and web cache, as well as application usage. These pages can be removed or modified, or new pages added by selecting *Options*. This allows you to configure when the report will be generated, whether it will include a table of contents, and whether to add or removed a page from the layout. When you add a new page to the layout, it becomes a submenu in *Log&Report > Report Access*.

Each page can be modified to suite your requirements for a default report, or removed. You can also add pages to the report. Adding or removing pages is done by selecting *Options* in Edit mode. The Edit mode allows you to modify the default report. The View mode is the mode when you first access the pages, and when you are viewing the saved changes to the default report.

The FortiOS default report contains information about the FortiGate unit in two text boxes, one on the cover page and one as an appendix that comes at the end of the generated report. The appendix is located in the VPN Usage page.

After generating a report, you can view it from the Historical Reports page by selecting *Historical Reports*. You must be in View mode to view generated reports.

Example for creating a new default report from the existing default report

The following is an example of how to create a new default from the existing FortiOS default report. The new default report will be generated on a daily basis, and include only email activity information and application usage.

To modify the cover page of the layout

- 1 Go to *Log&Report > Report Access > Cover Page*.
- 2 Select *Edit*.
The Editing Section: Cover Page page appears.
- 3 Within the *FortiGate UTM* text box, select inside the box so that the cursor appears; delete *FortiGate UTM* and enter *CompanyX*.
- 4 Within the *Daily Activity Report* text box, delete *Daily Activity Report* and enter *Email, Traffic and Application Usage Activity Report*.
- 5 Select *Save* to save the current changes.

The following assumes that you are still in Editing mode.

To remove and add pages within the report

- 1 In the Cover Page page, select *Options*.
When you select *Options*, the Report Options window appears. This window provides the settings for adding or removing pages, scheduling when the report generates, and whether to include a table of contents or not.
- 2 In the Report Options window, under *Sections*, select the check boxes beside *VPN Usage*, *Threats*, and *Web Usage*.
- 3 Select *Delete*.
This removes these three submenus from the Report Access menu.
- 4 Rename the Bandwidth and Application Usage to Application Usage by selecting *Bandwidth and Application Usage*, and then removing *Bandwidth* from the name.

- 5 Press Enter on your keyboard to apply the new name.
- 6 Select *Create New*.
- 7 Enter `Traffic Usage`.
This will add the new Traffic Usage submenu to the list.
- 8 Select *OK*.
- 9 Select *Save*.

To configure email addresses and enable sending the report attached in an email

- 1 In the Report Options window, select the check box beside *Email Generated Reports*.
The *Email Recipients* table appears.
- 2 Select within the table below the title *Add Email Recipient*; the blinking cursor appears, allowing you to enter the first email address.
- 3 Enter the first email address.
- 4 On your keyboard, press Enter to add another row so that you can add another email address.
- 5 Repeat steps 2 to 4 until all email addresses are included.
- 6 Select *OK* to save the email addresses.

The following adds the charts and text to the Emails page in *Log&Report > Report Access > Emails*.

To modify the information in the Email, Traffic and Application Usage page

- 1 Go to *Log&Report > Report Access > Application Usage*.
- 2 On the page, select *Edit*.
- 3 Remove the bandwidth usage information from the page.
This is done by moving your mouse over the top right-hand corner of a text box and then selecting the red x that appears.
- 4 Go to *Log&Report > Report Access > Traffic Usage*.
- 5 In *Text*, (located in the right-hand pane of the page), drag *H1* icon onto the page.
A text box appears.
- 6 Enter the chart heading's name in the text box, `Top Users By Bandwidth`.
The heading style automatically configures for you as you type.
- 7 In *Text*, drag *T* icon onto the page.
The *T* icon allows you to enter sentences or phrases in a normal font style, with a smaller font size than is available with *H1* or *H2*.
- 8 Enter the following in the text box: The top users by bandwidth for today.
- 9 In *Chart* (located in the right-hand pane of the page), drag the bar chart onto the page.
The Chart Chooser window appears.
- 10 Select *Traffic* in the right-hand column, and then select *traffic.bandwidth.user*; select *OK* to put the chart on the page.

11 Repeat steps 5 to 10 to add each of the following charts with a description of each chart, to the page:

- *traffic.sessions.app_cats*
- *traffic.sessions.users*
- *traffic.bandwidth.users*
- *traffic.bandwidth.app_cats*

12 Select *Save* to save the current changes.

The following assumes that you are still editing the report. You will be configuring when the FortiGate unit generates the default report as well as including a table of contents.

To configure the report schedule and include a table of contents

- 1** On the page you're in, select *Options*.
- 2** In the Report Options window, select *Daily* for *Schedule Type*.
- 3** Enter *08:30* in *Time*.
- 4** Verify that the check box is selected beside *Include Table of Contents*; if the check box is not selected, select it.
- 5** Select *OK*.
- 6** Select *Save* to save the current changes.
- 7** To generate the report immediately, select *Run*.

The report is generated and appears in the Historical Reports list on the Historical Reports page.

The FortiGate unit generates the report at the scheduled time.

If you want to have an on-demand report, you must select *Demand* from the drop-down list in *Schedule Type*.

Configuring charts, datasets, themes and styles for a report

If you want to create a report from scratch, you must configure it within the CLI. The CLI contains default summary and themes, as well as default charts and datasets. All of these can be modified or created from scratch. The following explains how to configure these parts for a report.

This topic contains the following:

- [Configuring datasets](#)
- [Configuring themes](#)
- [Configuring styles](#)
- [Adding charts to the layout](#)

Configuring datasets

You must configure datasets because they are required when configuring a chart. You can use the default datasets that are available when configuring a chart. Datasets require knowledge of SQL because the logs are stored in an SQLite database. You can view the SQLite schema using the `get report database schema` CLI command syntax.



If you are creating a chart from scratch, you must create a dataset for that chart. The chart cannot be configured without a dataset.

Use the following to configure a dataset that will be applied to a chart.

```
config report dataset
edit <report_dataset>
set query <SQL_statement>
config field
edit <field_id>
set displayname <string>
set type {text | integer | date | ip}
next
end
end
```

If you need more information about queries required for datasets, see [“The SQLite log database” on page 633](#).

Configuring themes

When you are configuring a layout for a report, you can also add a theme. A theme is a group of settings that create the general style of a report. For example, the styles that are applied to the table of contents section of the report. Themes are configured only in the CLI.

You may want to configure your own styles for a theme, such as the type of alignment for the text. Styles are configured within the CLI, and you can also customize the default styles as well.

Use the following procedure to configure a theme for a report, which can then be applied to a report's layout.

To configure a theme for a report

- 1 Log in to the CLI.
- 2 Enter the following command syntax:

```
config report theme
edit <theme_name>
set column-count [ 1 | 2 | 3]
set default-html-style <string>
set default-pdf-style <string>
set graph-chart-style <string>
set heading1-style <string>
set heading2-style <string>
set heading3-style <string>
set heading4-style <string>
set hline-style <string>
set image-style <string>
set normal-text-style <string>
set page-footer-style <string>
set page-header-style <string>
set page-orient {landscape | portrait}
set page-style <string>
set report-subtitle-style <string>
set report-title-style <string>
set table-chart-caption-style <string>
set table-chart-even-row-style <string>
set table-chart-head-style <string>
set table-chart-odd-row-style <string>
```



```

    set table-chart-style <string>
    set toc-heading1-style <string>
    set toc-heading2-style <string>
    set toc-heading3-style <string>
    set toc-heading4-style <string>
    set toc-title-style <string>
end

```

- 3 To choose a style for any one of the above commands, except for `column-count` and `page-orient`, enter `?` to view the available choices.
- 4 To change the style of any one of the above commands, except for `column-count` and `page-orient`, go to [“Configuring styles” on page 683](#).

Configuring styles

You can customize the default styles or create your own styles for reports. There are default styles and summary styles to choose from. Default styles use a default style scheme, and the summary styles are for summary reports that contain one or two pages with a small graph or table.

To customize an existing style

- 1 Log in to the CLI.
- 2 Enter the following command syntax:


```

config report style
    edit style <style_name>

```

 For example `default.graph`.
- 3 To view a list of available styles, enter `?` after entering `edit`.

To create a new style

- 1 Log in to the CLI.
- 2 Enter the following command syntax:


```

config report style
    edit <new_style_name>
        set options {align | border | color | column | font | margin
            | padding | size | text }
        set align {center | justify | left | right}
        set bg-color {colour_name1 | color_name2 | color_name3 | ...}
        set border-bottom <border width_pixels> <border_style_{solid
            | dotted | dashed}> <border_color>
        set border-left <border width_pixels> <border_style_{solid |
            dotted | dashed}> <border_color>
        set border-right <border width_pixels> <border_style_{solid
            | dotted | dashed}> <border_color>
        set border-top <border width_pixels> <border_style_{solid |
            dotted | dashed}> <border_color>
        set column-gap <pixels>
        set column-span {all | none}
        set fg-color <color>
        set font-family {Arial | Courier | Helvetica | Times |
            Verdana}
        set font-size {xx-small | x-small | small | medium | large |
            x-large | xx-large | <pixels>}

```

```

    set font-style {italic | normal}
    set font-weight {bold | normal}
    set height <pixels or percentage>
    set line-height <pixels or percentage>
    set margin-bottom <pixels>
    set margin-left <pixels>
    set margin-right <pixels>
    set margin-top <pixels>
    set padding-bottom <pixels>
    set padding-left <pixels>
    set padding-right <pixels>
    set padding-top <pixels>
    set width <pixels or percentage>
end

```

Example of a style for a theme

The following example shows how to configure a specific style that is then applied to a theme.

```

config report style
  edit style_1
    set options align color font margin text
    set align center
    set bg-color navy
    set fg-color white
    set font-family Verdana
    set font-size medium
    set font-weight bold
    set line-height 100
    set margin-bottom 20
    set margin-left 20
    set margin-right 30
    set margin-top 50
  end
config report theme
  edit theme_1
    set column-count 2
    set page-style style_1
  end

```

Configuring a report layout

After you have configured a theme and style, as well as your custom datasets and charts, you can then put them all together in a layout. Similar to the default FortiOS report's layout, you must specify the page orientation, footers and headers, as well as the style of these footers and headers.

A layout includes a scheduled time period so that the report will be generated at a specific time. For example, a weekly report is generated every Thursday at 4:00 p.m.

The report layout can also include email addresses of people that you want the generated report to be sent to. The generated report is immediately sent to the

To configure a report's layout

- 1 Log in to the CLI.
- 2 Enter the following to configure the layout:

```

config report layout
  edit <layout_name>
    set title <text>
    set cutoff-option {run-time | custom}
    set cutoff-time {time_str}
    set cache-time-out <seconds>
    set email-recipients <email_addr1, email_addr2, ...>
    set email-send {enable | disable}
    set description <text>
    set format {html | pdf}
    set schedule-type {demand | once | daily | weekly}
    set time <hh:mm>
    set day {sunday | monday | tuesday | wednesday | thursday |
      friday | saturday}
    set style-theme <theme_name>
    set options {include-table-of-contents | auto-numbering-
      heading | view-chart-as-heading | show-html-navbar-
      before-heading1}
  config page
    set paper {A4 | letter}
    set column-break-before {heading1 | heading2 | heading3}
    set options {header-on-first-page | footer-on-first-page}
    set style <style_name>
  config body-item
    edit <item_id>
    set type {text | image | chart | misc}
    set description <text>
    set style <style_name>
    set text-component {heading1 | heading2 | heading3 | normal
      | text}
    set content <text>
    set img-src <text>
    set chart <chart_name>
    set chart-options {hide-title | include-no-data | show-
      caption}
    set parameter1 <value_str>
  end
end

```

Charts

Charts are used to display the log information in a clear and concise way using a graph. The information for charts is gathered from the log tables in the SQLite database.

When you need to find information about a default chart or a chart that was created from scratch you can use the `get` command in the following way.

```

config report chart
  edit <chart_name>
  get

```

The information that displays is similar to the following:

```

name: web.allowed-request.sites.user
policy: 0
type: graph
period: last7d

```

```

drill-down-chart: (null)
comments: (null)
dataset: web.allowed-request.sites.user
category: webfilter
favorite: no
graphy-type: bar
style: auto
dimension: 3D
  x-series
    caption: (null)
    databind: field(1)
    is-category: yes
    label-angle: 45-degree
    unit: (null)
  y-series
    caption: Requests
    databind: field(2)
    extra-y: disable
    group: (null)
    label-angle: horizontal
    unit: (null)
title: Top Allowed Web Sites for User by Request
legend: enable

```

This information provides what the time period of the chart is, the title that will display for that chart (in the example, title is *Top Allowed Web Sites for User by Request*), as well as the dataset that is included with the chart. The chart may also include a drill-down chart to be used in a HTML report, where you can view additional detailed information. This is the same drill-down feature that is available in most monitoring pages, such as in *Log&Report > Monitor > Logging Monitor*, where you select a bar on a specific day and view the activity of each log type that was recorded.

When you need to view information about a dataset, as well as what log table is used to gather that information, you can use the `get` command in the following way:

```

config report dataset
  edit <dataset_name>
  get

```

The information that displays is similar to the following:

```

name: appctrl.Count.Bandwidth.Top10.Apps
policy: 0
query: select (timestamp-timestamp%3600) as hourstamp, (CASE
  WHEN app!='N/A' and app!=' ' then app ELSE service END) as
  appname, sum(sent+rcvd) as bandwidth from traffic_log where
  timestamp between ###start_time### and ###end_time### and
  (appname in (select (CASE WHEN app!='N/A' and app!='N/A' and
  app!=' ' then app END ELSE service END) as appname from
  traffic_log where timestamp between ###start_time### and
  ###end_time### group by appname order by sum(sent+rcvd) desc
  limit 10)) group by hourstamp, appname order by hourstamp
  desc field

```

In the example output, you can see that the log table used to gather the information is from the traffic log table indicated by the `from traffic_log` in the query statement.

Charts that are used in reports do not have “last24h” in the name, however, these charts are carried forward and are not usually used in FortiOS 4.0 MR3 and higher.

Adding charts to the layout

Charts are added to the report's layout using the `config body-item` command.

Before configuring a chart for a report, you may want to see the list of available default charts in the CLI, using the `config report chart` command. You can create your own chart; however, you must use either a default dataset or create your own dataset. SQL knowledge is required when configuring a dataset.

You should verify that the datasets you want to use are configured and available before configuring a chart because a chart cannot be configured without a dataset. For information about configure datasets, see [“Configuring datasets” on page 681](#).

The following explains how to add a chart to a report layout.

To add a chart to a report layout

1 Log in to the CLI.

2 Enter the following to access the report layout:

```
config report layout
edit <layout_name>
```

3 Enter the following to add the chart to the layout:

```
config body-item
edit 1
set type chart
set chart <chart_name>
set chart-options {hide-title | include-no-data | show-
caption}
end
```

Configuring a chart

The following explains how to configure a chart from scratch.

To create a new chart

1 Log in to the CLI.

2 Enter the following to create a new chart:

```
config report chart
edit <chart_name>
set type {graph | table}
set period {last24h | last7d}
set drill-down-chart <name_drill_down_chart>
set comments <text>
set dataset <dataset_name>
set category {app-crt1 | attack | dlp | event | misc | spam
| traffic | virus | vulnerability | webfilter}
set favorite {yes | no}
set graph-type {pie | flow | bar | line | none}
set style {auto | manual}
set title <title_chart>
set legend {enable | disable}
config x-series
set caption <text>
set databind <value_expression>
set is-category {yes | no}
set label-angle {45-degree | horizontal | vertical}
```

```

set scale-format {HH-MM | MM-DD | YYYY | YYYY-MM | YYYY-
MM-DD| YYYY-MM-DD HH | YYYY-MM-DD-HH-MM}
set scale-number-of-step <scale_totalnum>
set scale-origin {max | min}
set scale-start <time&date>
set scale-step <integer>
set scale-type <datetime>
set scale-unit {day | hour | minute | month | year}
set unit <integer>
end
config y-series
set caption <text>
set databind <value_expression>
set extra-y {enable | disable}
set group <y-series_group>
set label-angle {45-degree | horizontal | vertical}
set unit <integer>
end
end

```

Importing images for the report

Images can be imported for use in reports. The supported images are JPEG, JPG and PNG.

To import images for the report

- 1 Go to *Log&Report > Report Access > Report*.
- 2 Select *Edit* on the Viewing default layout page.
- 3 In *Image*, drag the image icon to the page.

The Graphic Chooser window appears.



All pages, except for the Table of Contents page, can import an image.

- 4 In the Graphic Chooser window, select *Upload*.
- 5 Locate the image file and then select *OK*.

Viewing reports

Generated reports are viewed from the Historical Reports page in *Log&Report > Report Access > Report*.

Historical Reports page

Displays each generated report. Reports are removed using the *Delete* icon. You can view either HTML reports or PDF reports directly from this page. A HTML report opens up in a separate window, while a PDF report opens within the Disk page.

Return to Layout Select to return to the report layout page, Viewing default layout.

Delete	Removes one, multiple or all reports from the list. If you select the check box in the check box column, you can remove all reports from the list at one time.
Report File	The name of the report file, which includes the date and time. Note: To view a HTML report, select the name in this column. The HTML report appears in a separate window.
Started	The time when the report began generating. This format includes the date and is displayed in this type of format, yyyy-mm-dd hh:mm:ss. The hour is in the 24 hour format.
Finished	The time when the report stopped generating. This format includes the date and is displayed in this type of format, yyyy-mm-dd hh:mm:ss. The hour is in the 24 hour format.
Size (bytes)	The size of the report file, in bytes.
Other Formats	Displays PDF formatted generated reports. Select the format in this column to view the report in PDF.

Report example

The following is an example of a report created from within the CLI. The type of report is what is known as “a report created from scratch” because you are configuring the style, theme, layout and if applicable, the datasets for charts. In the following report, no datasets are configured.

Report for analyzing web activity on the FortiGate unit

Your manager wants to view the web activity on the network. The following example configures a report from scratch, and uses the existing web filter charts. This example also includes how to add email addresses and enable sending the report by email.

Configuring the style

The style of the web activity report must include the company’s font and image. The following procedures will create a style for:

- footers and headers
- the pages within the report
- cover page
- charts
- table of contents headings and title
- page headings

To configure the footers and headers

- 1 Log in to the CLI.
- 2 Enter the following to configure a style:

```
config report style
edit web-footerheader
set options align font
set align center
set font-size xx-small
```

```
set font-family Arial
set font-weight normal
set font-style normal
next
```

To configure the pages

After entering next, enter the following for the margins and columns:

```
edit web-pages
set options align margin column
set align justify
set margin-top 5
set margin-bottom 5
set margin-right 6
set margin-left 7
set column-gap 3
set column-span all
next
```

To configure the table of contents headings and title

After entering next in the above procedure, enter the following to configure the table of contents headings and title

```
edit web-toc-title
set options align font
set font-family Arial
set font-weight bold
set font-style normal
set font-size x-large
set align center
next
edit web-toc-heading1
set options align font
set font-family Arial
set font-weight bold
set font-size large
set font-style normal
set align left
next
edit web-toc-heading2
set options align font
set font-family Arial
set font-size medium
set font-style normal
set font-weight bold
set align left
next
edit web-toc-heading3
set options align font
set font-family Arial
set font-size medium
set font-style italic
set font-weight bold
set align left
next
```


To configure the page headings

After entering next in the above procedure, enter the following to configure the page headings:

```
edit web-page-heading1
  set options align font
  set font-family Arial
  set font-size large
  set font-style normal
  set font-weight bold
  set align left
next
edit web-page-heading2
  set options align font
  set font-family Arial
  set font-size medium
  set font-style normal
  set font-weight bold
next
edit web-page-heading3
  set options align font
  set font-family Arial
  set font-size medium
  set font-style italic
  set font-weight bold
next
```

To configure the chart style

After entering next in the above procedure, enter the following to configure the style of the font and alignment of the chart:

```
edit web-chart
  set options align font
  set font-family Arial
  set font-size small
  set font-style normal
  set font-weight bold
  set align left
next
```

To configure the cover page

After entering next in the above procedure, enter the following to configure the style of the cover page of the report:

```
edit web-cover
  set options align font
  set font-family Arial
  set font-size xx-large
  set font-style normal
  set font-weight bold
  set align center
end
end
```

Configuring the theme

The theme applies the styles that were configured in the procedures in “[Configuring styles](#)” on page 683. The theme will then be applied to the report’s layout.

To configure a theme

Enter the following to configure a theme:

```
config report theme
edit web-theme
    set column-count 2
    set default-pdf-style web-cover
    set graph-chart-style web-chart
    set page-orient landscape
    set page-style web-pages
    set table-chart-style web-chart
    set toc-title-style web-toc-title
    set toc-heading1-style web-toc-heading1
    set toc-heading2-style web-toc-heading2
    set toc-heading3-style web-toc-heading3
    set heading1-style web-heading1
    set heading2-style web-heading2
    set heading3-style web-heading3
    set page-footer-style web-footerheader
    set page-header-style web-footerheader
    set report-title-style web-cover
end
```

Uploading the company graphic for the report

You need to upload the company graphic for the report. You must use the web-based manager to upload the image to the FortiGate unit so that it can be used in the web activity report.

To upload the company graphic

- 1 Log in to the web-based manager.
- 2 Go to *Log&Report > Report Access > Cover Page*.
- 3 Select *Edit*.
- 4 In the Edit Section: Cover Page page, select and drag the image icon onto the page. The Graphic Chooser window appears.
- 5 In the Graphic Chooser window, select *Upload*.
- 6 Browse and locate the company’s image file. The image file immediately uploads to the FortiGate unit.
- 7 Select the image and then select *OK* to save the image.
- 8 In the Edit Section: Cover Page, select *Save* to save the uploaded image.

Modifying the time period for the charts

The charts that will be used for the report must be modified to gather information over the past seven days. The charts that will be modified are:

- web-allowed-request.sites.user
- web.blocked-request.sites.user

- web.blocked-request.web_cats
- web.bandwidth.sites.user
- web.bandwidth.stream-sites.user
- web.bandwidth.stream-sites

To modify the time period for the charts

- 1 In the CLI, enter the following:

```
config report chart
```

- 2 Enter the following to modify the time period of each chart:

```
edit web.allowed-request.sites.user
set period last7d
next
edit web.blocked-request.sites.user
set period last7d
next
edit web.blocked-requests.web_cats
set period last7d
next
edit web.bandwidth.sites.
set period last7d
next
edit web.bandwidth.stream-sites.user
set period last7d
next
edit web.bandwidth.stream-sites
set period last7d
end
end
end
```

Configuring the layout

The layout of the report includes who will be receiving the report by email, as well as style and when the report will be generated by the FortiGate unit.

To configure layout

- 1 Log into the CLI.
- 2 Enter the following to configure the layout:

```
config report layout
edit web-activity-layout
set cutoff-option custom
set cutoff-time 04:50
set day monday
set description "Web activity report for the week of March
7"
set format pdf
set options include-table-of-contents
set style-theme web-theme
set subtitle "Web Activity Report"
set time 08:30
set title "Web Activity Report for March 7"
set email-send enable
```

```
    set email-recipients user1@example.com, user2@example.com
next
config body-item
  edit 1
    set type image
    set img-src image_company.jpg
  next
  edit 2
    set content "Web Activity: Blocked and Allowed Requested
    Web Sites"
    set style web-heading1
  next
  edit 3
    set content "The blocked and allowed web sites that were
    requested by users over the last seven days"
  next
  edit 4
    set type chart
    set chart web.allowed-request.sites.user
    set description "Allowed Web Sites Requested by Users"
    set style web-heading2
  next
  edit 5
    set type chart
    set chart web.blocked-request.sites.user
    set description "Blocked Web Sites Requested by User"
    set style web-heading2
  next
  edit 6
    set type chart
    set chart web.blocked-requests.web_cats
    set description "Blocked Web Sites Requested, per-Web
    Filter Category"
    set style web-heading2
  next
  edit 7
    set description "Web Bandwidth Activity"
    set style web-heading1
  next
  edit 8
    set content "The bandwidth that was used for web activity
    for the last seven days"
  next
  edit 9
    set type chart
    set chart web.bandwidth.sites.user
    set description "Bandwidth Used per User"
    set style web-heading2
  next
  edit 10
    set type chart
    set chart web.bandwidth.stream-sites.user
    set description "Bandwidth Used for Web Streaming per
    User"
```

```
        set style web-heading2
    next
    edit 11
        set type chart
        set chart web.bandwidth.stream-sites
        set description "Bandwidth Used for Web Streaming"
        set style web-heading2
    end
    config paage
        set options footer-on-first-page
        set paper letter
        set column-break-before heading1
    end
end
```




Chapter 5 Troubleshooting

-

This handbook chapter describes concepts of troubleshooting and solving issues that may occur with FortiGate units.

This FortiOS Handbook chapter contains the following chapters:

[Life of a Packet](#) explains the different layers and modules a packet goes through in FortiOS, including the order of operations.

[Troubleshooting process](#) walks you through best practice concepts of FortiOS troubleshooting.

[Troubleshooting tools](#) describes some of the basic commands and parts of FortiOS that can help you with troubleshooting.

[Technical Support Organization Overview](#) describes how Fortinet Support operates, what they will need from you if you contact them, and what you can expect in general.

[Troubleshooting common issues](#) walks you through how to troubleshoot some common issues you may have, and some more in-depth coverage of other issues.

[Troubleshooting advanced](#) walks through some more advanced features including traffic shaping, user logons, and VPN.

[Troubleshooting 'get' commands](#) provides a list of diagnostic commands that can help you troubleshoot your FortiOS unit.

[Troubleshooting bootup and FSSO](#) addresses potential problems your unit may have when booting up. Also covered is troubleshooting an FSSO installation. The format is an easy to follow step by step question and answer format.



Life of a Packet

Directed by security policies, a FortiGate unit screens network traffic from the IP layer up through the application layer of the TCP/IP stack. This chapter provides a general, high-level description of what happens to a packet as it travels through a FortiGate security system.

The FortiGate unit performs three types of security inspection:

- stateful inspection, that provides individual packet-based security within a basic session state
- flow-based inspection, that buffers packets and uses pattern matching to identify security threats
- proxy-based inspection, that reconstructs content passing through the FortiGate unit and inspects the content for security threats.

Each inspection component plays a role in the processing of a packet as it traverses the FortiGate unit en route to its destination. To understand these inspections is the first step to understanding the flow of the packet.

This section contains the following topics:

- [Stateful inspection](#)
- [Flow inspection](#)
- [Proxy inspection](#)
- [Comparison of inspection layers](#)
- [FortiOS functions and security layers](#)
- [Packet flow](#)
- [Example 1: client/server connection](#)
- [Example 2: Routing table update](#)
- [Example 3: Dialup IPsec VPN with application control](#)

Stateful inspection

With stateful inspection, the FortiGate unit looks at the first packet of a session to make a security decision. Common fields inspected include TCP SYN and FIN flags to identify the start and end of a session, the source/destination IP, source/destination port and protocol. Other checks are also performed on the packed payload and sequence numbers to verify it as a valid communication and that the data is not corrupted or poorly formed.

What makes it stateful is that one or both ends must save information about the session history in order to communicate. In stateless communication, only independent requests and responses are used, that do not depend on previous data. For example, UDP is stateless by nature because it has no provision for reliability, ordering, or data integrity.

The FortiGate unit makes the decision to drop, pass or log a session based on what is found in the first packet of the session. If the FortiGate unit decides to drop or block the first packet of a session, then all subsequent packets in the same session are also dropped or blocked without being inspected. If the FortiGate unit accepts the first packet of a session, then all subsequent packets in the same session are also accepted without being inspected.

Connections over connectionless

A connection is established when two end points use a protocol to establish connection through use of various methods such as segment numbering to ensure data delivery, and handshaking to establish the initial connection. Connections can be stateful because they record information about the state of the connection. Persistent connections reduce request latency because the end points do not need to re-negotiate the connection multiple times, but instead just send the information without the extra overhead. By contrast, connectionless communication does not keep any information about the data being sent or the state. It is based on an autonomous response/reply that is independent of other responses/replies that may have gone before. One example of connectionless communication is IP.

Benefits of connections over connectionless include being able to split data up over multiple packets, the data allows for a best-effort approach, and once the connection is established subsequent packets are not required to contain the full addressing information which saves on bandwidth. Connections are often reliable network services since acknowledgements can be sent when data is received.

What is a session?

A session is established on an existing connection, for a defined period of time, using a determined type of communication or protocol. Sessions can have specific bandwidth, and time to live (TTL) parameters.

You can compare a session to a conversation. A session is established when one end point initiates a request by establishing a TCP connection on a particular port, the receiving end is listening on that port, and replies. You could telnet to port 80 even though telnet normally uses port 23, because at this level, the application being used cannot be determined.


However, the strong points of sessions and stateful protocols can also be their weak points. Denial of service (DoS) attacks involve creating so many sessions that the connection state information tables are full and the unit will not accept additional sessions.

Differences between connections and sessions

In almost all cases, established sessions are stateful and all involve connections. However, some types of connections, such as UDP, are stateless, and are not sessions.


This means that not all traffic can be inspected by stateful inspection, because some of it is stateless. For example IP packets are stateless. Communications using HTTP are stateless, but HTTP often uses cookies to store persistent data in a way that approaches stateful.

Stateful inspection of sessions has the benefit of being able to apply the initial connection information to the packets that follow — the end points of the session will remain the same as will the protocol for example. That information can be examined for the first packet of the session and if it is malicious or not appropriate, the whole session can be dropped without committing significant resources.

Figure 65: Stateful inspection of packets through the FortiGate unit

Flow inspection

With flow inspection, the FortiGate unit samples multiple packets in a session and multiple sessions, and uses a pattern matching engine to determine the kind of activity that the session is performing and to identify possible attacks or viruses. For example, if application control is operating, flow inspection can sample network traffic and identify the application that is generating the activity. Flow-based antivirus can sample network traffic and determine if the content of the traffic contains a virus, IPS can sample network traffic and determine if the traffic constitutes an attack. The security inspection occurs as the data is passing from its source to its destination. Flow inspection identifies and blocks security threats in real time as they are identified.

Figure 66: Flow inspection of packets through the FortiGate unit

Flow-based inspection typically requires less processing than proxy-based inspection, and therefore flow-based antivirus performance can be better than proxy-based antivirus performance. However, some threats can only be detected when a complete copy of the payload is obtained so, proxy-based inspection tends to be more accurate and complete than flow-based inspection.


Proxy inspection

With flow inspection, the FortiGate unit will pass all the packets between the source and destination, and keeps a copy of the packets in its memory. It then uses a reconstruction engine to build the content of the original traffic. The security inspection occurs after the data has passed from its source to its destination.

Proxy inspection examines the content contained a content protocol session for security threats. Content protocols include the HTTP, FTP, and email protocols. Security threats can be found in files and other content downloaded using these protocols. With proxy inspection, the FortiGate unit downloads the entire payload of a content protocol session and re-constructs it. For example, proxy inspection can reconstruct an email message and its attachments. After a satisfactory inspection the FortiGate unit passes the content on to the client. If the proxy inspection detects a security threat in the content, the content is removed from the communication stream before the it reaches its destination. For example, if proxy inspection detects a virus in an email attachment, the attachment is removed from the email message before its sent to the client. Proxy inspection is the most thorough inspection of all, although it requires more processing power, and this may result in lower performance.

If you enable ICAP in a security policy, HTTP traffic intercepted by the policy is transferred to the ICAP servers in the ICAP profile added to the policy. The FortiGate unit is the surrogate, or “middle-man”, and carries the ICAP responses from the ICAP server to the ICAP client; the ICAP client then responds back, and the FortiGate unit determines the action that should be taken with these ICAP responses and requests.

Figure 67: Proxy inspection of packets through the FortiGate unit



Comparison of inspection layers

The three inspection methods each have their own strengths and weaknesses. The following table looks at all three methods side-by-side.

Table 50: Inspection methods comparison

Feature	Stateful	Flow	Proxy
Inspection unit per session	first packet	selected packets	complete content

Feature	Stateful	Flow	Proxy
Memory, CPU required	low	medium	high
Level of threat protection	good	better	best
Authentication	yes		
IPsec and SSL VPN	yes		
Antivirus protection		on some models	yes
Application control		yes	some
Delay in traffic			small
Reconstruct entire content			yes

FortiOS functions and security layers

Within these security inspection types, FortiOS functions map to different inspections. The table below outlines when actions are taken as a packet progresses through its life within a FortiGate unit.

Table 51: FortiOS security functions and security layers

Security Function	Stateful	Flow	Proxy
Firewall	ü		
IPsec VPN	ü		
Traffic Shaping	ü		
User Authentication	ü		
Management Traffic	ü		
SSL VPN	ü		
Intrusion Prevention		ü	
Flow-based Antivirus		ü	
Application Control		ü	
VoIP inspection			ü
Proxy Antivirus			ü
Email Filtering			ü
Web Filtering (Antispam)			ü
Data Leak Prevention			ü

Packet flow

After the FortiGate unit's external interface receives a packet, the packet proceeds through a number of steps on its way to the internal interface, traversing each of the inspection types, depending on the security policy and UTM profile configuration. The diagram in [Figure 68 on page 704](#) is a high level view of the packet's journey.

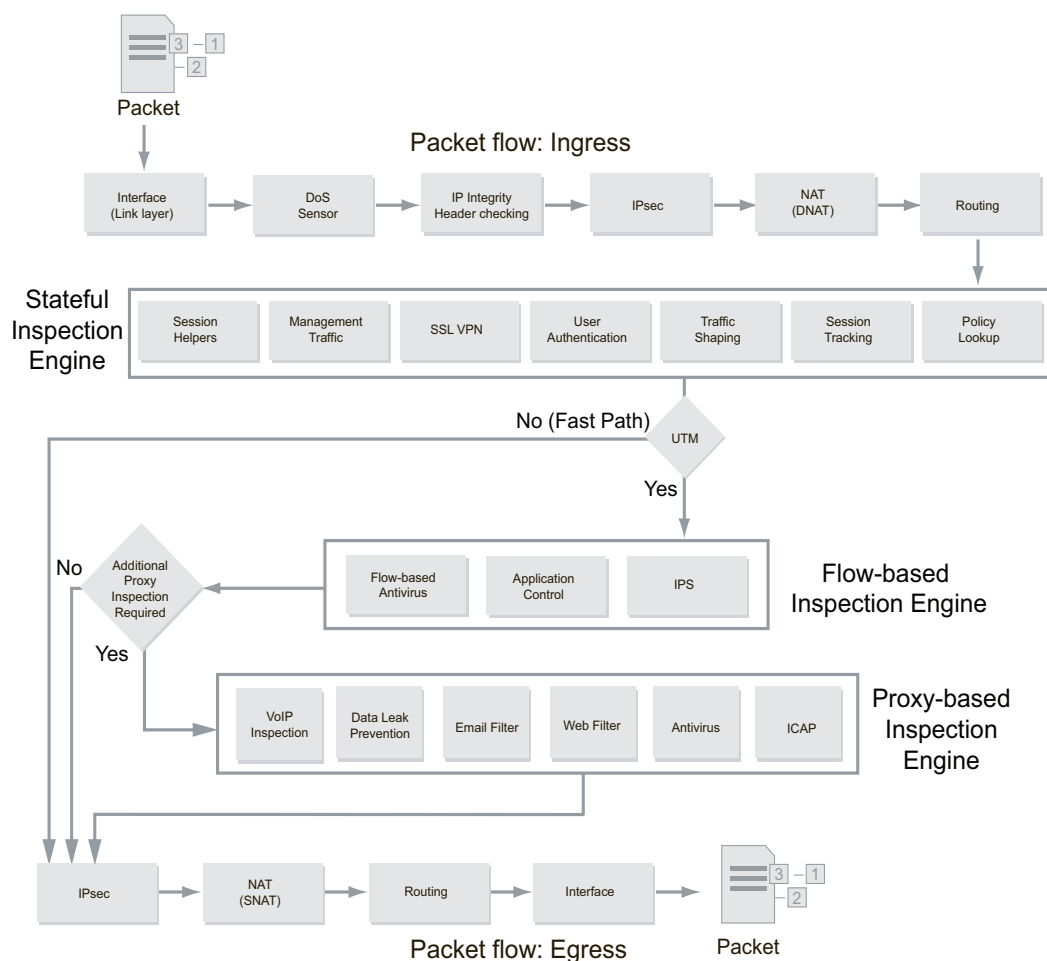
The description following is a high-level description of these steps as a packet enters the FortiGate unit towards its destination on the internal network. Similar steps occur for outbound traffic.

Packet inspection (Ingress)

In the diagram in [Figure 68 on page 704](#), in the first set of steps (ingress), a number of header checks take place to ensure the packet is valid and contains the necessary information to reach its destination. This includes:

- Packet verification - during the IP integrity stage, verification is performed to ensure that the layer 4 protocol header is the correct length. If not, the packet is dropped.
- Session creation - the FortiGate unit attempts to create a session for the incoming data
- IP stack validation for routing - the firewall performs IP header length, version and checksum verifications in preparation for routing the packet.
- Verifications of IP options - the FortiGate unit validates the routing information

Figure 68: Packet flow



Interface

Ingress packets are received by a FortiGate interface. The packet enters the system, and the interface network device driver passes the packet to the Denial of Service (DoS) sensors, if enabled, to determine whether this is a valid information request or not.

DoS sensor

DoS scans are handled very early in the life of the packet to determine whether the traffic is valid or is part of a DoS attack. Unlike signature-based IPS which inspects all the packets within a certain traffic flow, the DoS module inspects all traffic flows but only tracks packets that can be used for DoS attacks (for example TCP SYN packets), to ensure they are within the permitted parameters. Suspected DoS attacks are blocked, other packets are allowed.

IP integrity header checking

The FortiGate unit reads the packet headers to verify if the packet is a valid TCP, UDP, ICMP, SCTP or GRE packet. The only verification that is done at this step to ensure that the protocol header is the correct length. If it is, the packet is allowed to carry on to the next step. If not, the packet is dropped.

IPsec

If the packet is an IPsec packet, the IPsec engine attempts to decrypt it. The IPsec engine applies the correct encryption keys to the IPsec packet and sends the unencrypted packet to the next step. IPsec is bypassed when for non-IPsec traffic and for IPsec traffic that cannot be decrypted by the FortiGate unit.

Destination NAT (DNAT)

The FortiGate unit checks the NAT table and determines the destination IP address for the traffic. This step determines whether a route to the destination address actually exists.

For example, if a user's browser on the internal network at IP address 192.168.1.1 visited the web site www.example.com using NAT, after passing through the FortiGate unit the source IP address becomes NATed to the FortiGate unit external interface IP address. The destination address of the reply back from www.example.com is the IP address of the FortiGate unit internal interface. For this reply packet to be returned to the user, the destination IP address must be destination NATed to 192.168.1.1.

DNAT must take place before routing so that the FortiGate unit can route packets to the correct destination.

Routing

The routing step determines the outgoing interface to be used by the packet as it leaves the FortiGate unit. In the previous step, the FortiGate unit determined the real destination address, so it can now refer to its routing table and decide where the packet must go next.

Routing also distinguishes between local traffic and forwarded traffic and selects the source and destination interfaces used by the security policy engine to accept or deny the packet.

Policy lookup

The policy look up is where the FortiGate unit reviews the list of security policies which govern the flow of network traffic, from the first entry to the last, to find a match for the source and destination IP addresses and port numbers. The decision to accept or deny a packet, after being verified as a valid request within the stateful inspection, occurs here. A denied packet is discarded. An accepted packet will have further actions taken. If IPS is enabled, the packet will go to [Flow-based inspection engine](#), otherwise it will go to the [Proxy-based inspection engine](#).

If no other UTM options are enabled, then the session was only subject to stateful inspection. If the action is accept, the packet will go to Source NAT to be ready to leave the FortiGate unit.

Session tracking

Part of the stateful inspection engine, session tracking maintains session tables that maintain information about sessions that the stateful inspection module uses for maintaining sessions, NAT, and other session related functions.

User authentication

User authentication added to security policies is handled by the stateful inspection engine, which is why Firewall authentication is based on IP address. Authentication takes place after policy lookup selects a security policy that includes authentication. This is also known as identify-based policies. Authentication also takes place before UTM features are applied to the packet.

Management traffic

This local traffic is delivered to the FortiGate unit TCP/IP stack and includes communication with the web-based manager, the CLI, the FortiGuard network, log messages sent to FortiAnalyzer or a remote syslog server, and so on. Management traffic is processed by applications such as the web server which displays the FortiOS web-based manager, the SSH server for the CLI or the FortiGuard server to handle local FortiGuard database updates or FortiGuard Web Filtering URL lookups.

SSL VPN traffic

For local SSL VPN traffic, the internal packets are decrypted and are routed to a special interface. This interface is typically called `ssl.root` for decryption. Once decrypted, the packets go to policy lookup.

Session helpers

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and adjust the firewall to allow those protocols to send packets through the firewall.

Flow-based inspection engine

Flow-based inspection is responsible for IPS, application control, flow-based antivirus scanning and VoIP inspection. Packets are sent to flow-based inspection if the security policy that accepts the packets includes one or more of these UTM features.



Flow-based antivirus scanning is only available on some FortiGate models.

Once the packet has passed the flow-based engine, it can be sent to the proxy inspection engine or egress.

Proxy-based inspection engine

The proxy inspection engine is responsible for carrying out antivirus protection, email filtering (antispam), web filtering and data leak prevention. The proxy engine will process multiple packets to generate content before it is able to make a decision for a specific packet.

IPsec

If the packet is transmitted through an IPsec tunnel, it is at this stage the encryption and required encapsulation is performed. For non-IPsec traffic (TCP/UDP) this step is bypassed.

Source NAT (SNAT)

When preparing the packet to leave the FortiGate unit, it needs to NAT the source address of the packet to the external interface IP address of the FortiGate unit. For example, a packet from a user at 192.168.1.1 accessing www.example.com is now using a valid external IP address as its source address.

Routing

The final routing step determines the outgoing interface to be used by the packet as it leaves the FortiGate unit.

Egress

Upon completion of the scanning at the IP level, the packet exits the FortiGate unit.

Example 1: client/server connection

The following example illustrates the flow of a packet of a client/web server connection with authentication and FortiGuard URL and antivirus filtering.

This example includes the following steps:

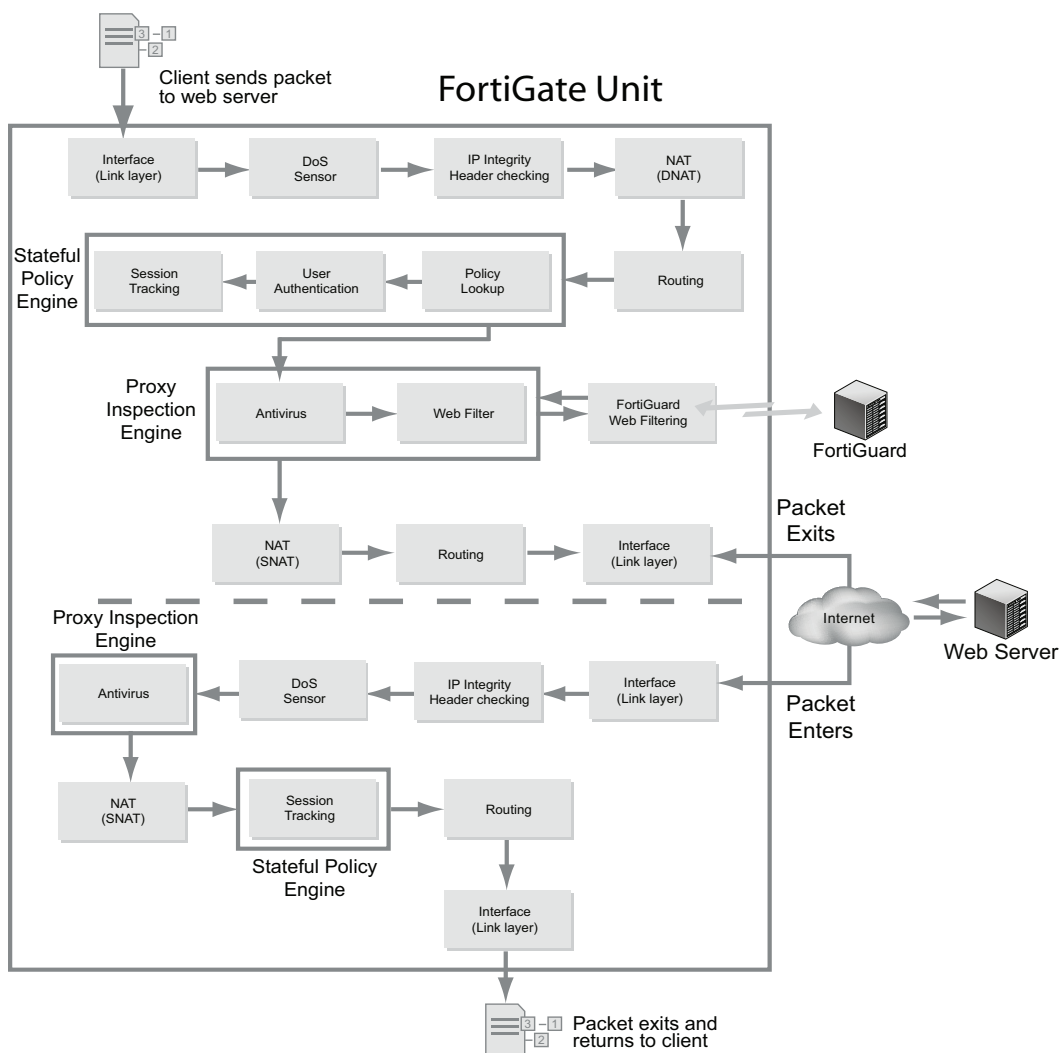
Initiating connection from client to web server

- 1 Client sends packet to web server.
- 2 Packet intercepted by FortiGate unit interface.
 - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
- 3 DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
- 4 IP integrity header checking, verifying the IP header length, version and checksums.
- 5 Next hop route
- 6 Policy lookup
- 7 User authentication
- 8 Proxy inspection
 - 8.1 Web Filtering
 - 8.2 FortiGuard Web Filtering URL lookup
 - 8.3 Antivirus scanning
- 9 Source NAT

- 10 Routing
- 11 Interface transmission to network
- 12 Packet forwarded to web server

Response from web server

- 1 Web Server sends response packet to client.
- 2 Packet intercepted by FortiGate unit interface
 - 2.1 Link level CRC and packet size checking.
- 3 IP integrity header checking.
- 4 DoS sensor.
- 5 Proxy inspection
 - 5.1 Antivirus scanning.
- 6 Source NAT.
- 7 Stateful Policy Engine
 - 7.1 Session Tracking
- 8 Next hop route
- 9 Interface transmission to network
- 10 Packet returns to client

Figure 69: Client/server connection

Example 2: Routing table update

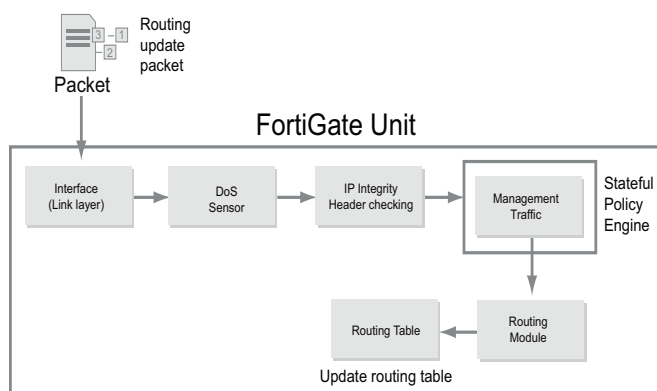
The following example illustrates the flow of a packet when there is a routing table update. As this is low level, there is no UTM involved. This example includes the following steps:

- 1 FortiGate unit receives routing update packet
- 2 Packet intercepted by FortiGate unit interface
 - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
- 3 DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
- 4 IP integrity header checking, verifying the IP header length, version and checksums.
- 5 Stateful policy engine
 - 5.1 Management traffic (local traffic)

6 Routing module

6.1 Update routing table

Figure 70: Routing table update



Example 3: Dialup IPsec VPN with application control

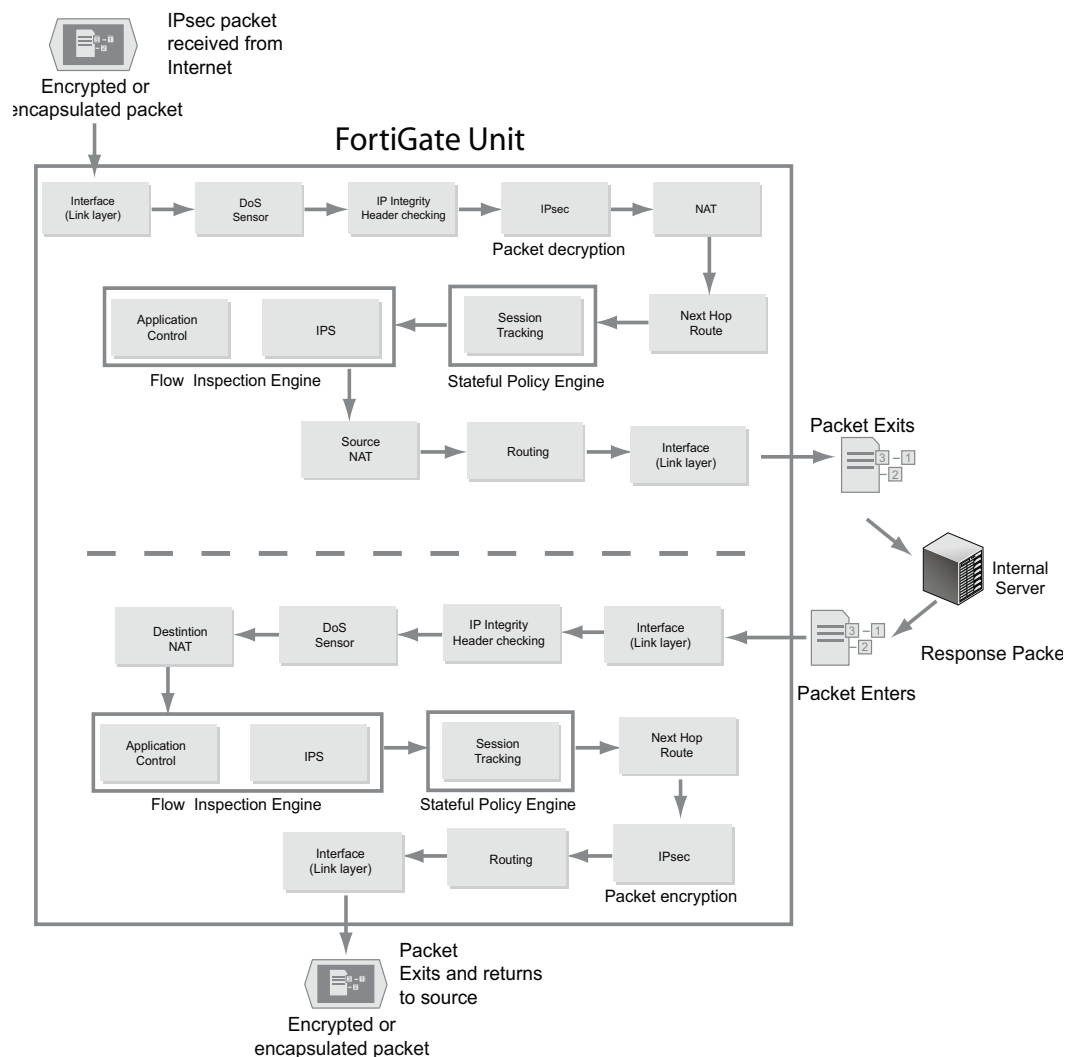
This example includes the following steps:


- 1 FortiGate unit receives IPsec packet from Internet
- 2 Packet intercepted by FortiGate unit interface
 - 2.1 Link level CRC and packet size checking. If the size is correct, the packet continues, otherwise it is dropped.
- 3 DoS sensor - checks are done to ensure the sender is valid and not attempting a denial of service attack.
- 4 IP integrity header checking, verifying the IP header length, version and checksums.
- 5 IPsec
 - 5.1 Determines that packet matched IPsec phase 1 configuration
 - 5.2 Unencrypted packet
- 6 Next hop route
- 7 Stateful policy engine
 - 7.1 Session tracking
- 8 Flow inspection engine
 - 8.1 IPS
 - 8.2 Application control
- 9 Source NAT
- 10 Routing
- 11 Interface transmission to network
- 12 Packet forwarded to internal server

Response from server

- 1 Server sends response packet
- 2 Packet intercepted by FortiGate unit interface
 - 2.1 Link level CRC and packet size checking

- 3 IP integrity header checking.
- 4 DoS sensor
- 5 Flow inspection engine
 - 5.1 IPS
 - 5.2 Application control
- 6 Stateful policy engine
 - 6.1 Session tracking
- 7 Next hop route
- 8 IPsec
 - 8.1 Encrypts packet
- 9 Routing
- 10 Interface transmission to network
- 11 Encrypted Packet returns to internet

Figure 71: Dialup IPsec with application control



Troubleshooting process

Before you begin troubleshooting anything but the most minor issues, you need to prepare. Doing so will shorten the time to solve your issue. This section helps to explain how you prepare before troubleshooting, as well as creating a troubleshooting plan and contacting support.

This section contains the following topics:

- Establish a baseline
- Search for a solution
- Create a troubleshooting plan
- Obtain any required additional equipment
- Ensure you have administrator level access to required equipment
- Contact Fortinet customer support for assistance

Establish a baseline

FortiGate units operate at all layers of the OSI model. For this reason troubleshooting problems can become complex. If you establish a normal operation parameters, or baseline, for your system before the problem occurs it will help reduce the complexity when you are troubleshooting.

Many of the guiding questions in the following sections are some form of comparing the current problem situation to normal operation on your FortiGate unit. For this reason it is a best practice that you know what your normal operating status is, and have a record of it you can refer to. This can easily be accomplished by monitoring the system performance with logs, SNMP tools, or regularly running information gathering commands and saving the output. This regular operation data will show trends, and enable you to see when changes happen and there may be a problem.



Back up your FortiOS configuration on a regular basis. This is a good practice for everyday as well as when troubleshooting. You can restore the backed up configuration when needed and save the time and effort of re-creating it from the factory default settings.

Some basic commands you can use to obtain normal operating data for your system:

<code>get system status</code>	Displays versions of firmware and FortiGuard engines, and other system information.
<code>diagnose firewall statistic show</code>	Displays the amount of network traffic broken down into categories such as email, VoIP, TCP, UDP, IM, Gaming, P2P, and Streaming.
<code>get router info routing-table all</code>	Displays all the routes in the routing table including their type, source, and other useful data.
<code>get ips session</code>	Displays memory used and max available to IPS as well and counts.
<code>get webfilter ftgd-statistics</code>	Displays list of FortiGuard related counts of status, errors, and other data.

These commands are just a sample. Feel free to include any extra information gathering commands that apply to your system. For example if you have active VPN connections, record information about them using the `get vpn *` series of commands. See [“Troubleshooting ‘get’ commands” on page 805](#).

For an extensive snapshot of your system, run the CLI command used by TAC to gather extensive information about a system — `exec tac report`. It runs many diagnose commands that are for specific configurations. This means no matter what features you are using, this command will record their current state. Then if you need to perform troubleshooting at a later date, you can run the same command again and compare the differences to quickly locate suspicious output you can investigate. See [“exec tac report” on page 806](#).

Define the problem

Before starting to troubleshooting a problem, ask the following questions:

- What is the problem?
Do not assume that the problem is being experienced is the actual problem. First determine that the problem does not lie elsewhere before starting to troubleshoot the FortiGate device.
- Has it ever worked before?
If the device never worked from the first day, you may not want to spend time troubleshooting something that could well be defective. See [“Troubleshooting bootup and FSSO” on page 871](#).
- Can the problem be reproduced at will or is it intermittent?
If the problem is intermittent, it may be dependent on system load. Also an intermittent problem can be very difficult to troubleshoot due to the difficulty reproducing the issue.
- What has changed?
Do not assume that nothing has changed in the network. Use the FortiGate event log to see if any configuration changes were made. The change could be in the operating environment, for example, a gradual increase in load as more sites are forwarded through the firewall.
If something has changed, see what the affect is if the change is rolled back.
- Determine the scope of the problem - after you have isolated the problem what applications, users, devices, and operating systems does it effect?

Before you can solve a problem, you need to understand it. Often this step can be the longest in this process.

Ask questions such as:

- What is not working? Be specific.
- Is there more than one thing not working?
- Is it partly working? If so, what parts are working?
- Is it a connectivity issue for the whole device, or is there an application that isn't reaching the Internet?

Be as specific as possible with your answers, even if it takes awhile to find the answers.

These questions will help you define the problem. Once the problem is defined, you can search for a solution and then create a plan on how to solve it.

Gathering Facts

Fact gathering is an important part of defining the problem. Record the following information as it affects your problem:

- Where did the problem occur?
- When did the problem occur and to whom?
- What components are involved?
- What is the affected application?
- Can the problem be traced using a packet sniffer?
- Can the problem be traced in the session table?
- Can log files be obtained that indicate a failure has occurred?

Answers to these questions will help you narrow down the problem, and what you have to check during your troubleshooting. The more things you can eliminate, the fewer things you need to check during troubleshooting. For this reason, be as specific and accurate as you can while gathering facts.

Search for a solution

An administrator can save time and effort during the troubleshooting process by first checking if the issue has been experienced before. Several self-help resources are available to provide valuable information about FortiOS technical issues, including:

Technical Documentation

Installation Guides, Administration Guides, Quick Start Guides, and other technical documents are available online at the following URL:

<http://docs.fortinet.com>

Release Notes

Issues that are uncovered after the technical documentation has been published will often be listed in the Release Notes that accompany the device.

Knowledge Base

The Fortinet Knowledge Base provides access to a variety of articles, white papers, and other documentation providing technical insight into a range of Fortinet products. The Knowledge Base is available online at the following URL:

<http://kb.fortinet.com>

Fortinet Technical Discussion Forums

An online technical forums allow administrators to contribute to discussions about issues related to their Fortinet products. Searching the forum can help the administrator identify if an issue has been experienced by another user. The support forums can be accessed at the following URL:

<http://support.fortinet.com/forum>

Fortinet Training Services Online Campus

The Fortinet Training Services Online Campus hosts a collection of tutorials and training materials which can be used to increase knowledge of the Fortinet products.

<http://campus.training.fortinet.com>

Create a troubleshooting plan

Once you have defined the problem, and searched for a solution you can create a plan to solve that problem. Even if your search didn't find a solution to your problem you may have found some additional things to check to further define your problem.

The plan should list all the possible causes of the problem that you can think of, and how to test for each possible cause.

Your troubleshooting plan will act as a checklist so that you know what you have tried and what is left to check. This is important to have if more than one person will be doing the troubleshooting. Without a written plan, people will become easily confused and steps will be skipped. Also if you have to hand over the problem to someone else, providing them with a detailed list of what data has been gathered and what solutions have been already tried demonstrates a good level of professionalism.

Be ready to add to your plan as needed. After you are part way through, you may discover that you forgot some tests or a test you performed discovered new information. This is normal.

Also if you contact support, they will require information about your problem as well as what you have already tried to fix the problem. This should all be part of your plan.

Providing Supporting Elements

If the Fortinet Technology Assistance Center (TAC) needs to be contacted to help you with your issue, be prepared to provide the following information:

- The firmware build version (use the `get system status` command)
- A network topology diagram
- A recent configuration file
- Optionally, a recent debug log
- Tell the support team what troubleshooting steps have already been performed and the results.



Do not provide the output from `exec tac` report unless Support requests it. The output from that command is very large and is not required in many cases.

For additional information about contacting Fortinet Customer Support, see “[Technical Support Organization Overview](#)” on page 751.

All of this is your troubleshooting plan.

Obtain any required additional equipment

You may require additional networking equipment, computers, or other equipment to test your solution.

Normally network administrators have additional networking equipment available either to loan you, or a lab where you can bring the FortiGate unit to test.

If you do not have access to equipment, check for shareware applications that can perform the same task. Often there are software solutions when hardware is too expensive.

Ensure you have administrator level access to required equipment

Before troubleshooting your FortiGate unit, you will need administrator access to the equipment. If you are a client on a FortiGate unit with virtual domains enabled, often you can troubleshoot within your own VDOM. However, you should inform your FortiGate unit's super admin that you will be doing troubleshooting.

Also, you may need access to other networking equipment such as switches, routers, and servers to help you test. If you do not normally have access to this equipment, contact your network administrator for assistance.

Contact Fortinet customer support for assistance

You have defined your problem, researched a solution, put together a plan to find the solution, and executed that plan. At this point if the problem has not been solved, its time to contact Fortinet Customer Support for assistance.

For more information, see [“Technical Support Organization Overview”](#) on page 751.



Troubleshooting tools

FortiOS provides a number of tools that help with troubleshooting both hardware and software issues. These tools include diagnostics and ports; ports are used when you need to understand the traffic coming in or going out on a specific port, for example, UDP 53, which is used by the FortiGate unit for DNS lookup and RBL lookup.

This section also contains information about troubleshooting FortiGuard issues.

This section contains the following topics:

- [FortiOS diagnostics](#)
- [FortiGate ports](#)
- [FortiAnalyzer/FortiManager ports](#)
- [FortiGuard troubleshooting](#)

FortiOS diagnostics

A collection of diagnostic commands are available in FortiOS for troubleshooting and performance monitoring. While some of these areas have web-based manager areas, all have relevant CLI commands with the main commands listed in this section.

Within the CLI commands, the two main groups of diagnostic commands are `get` and `diagnose` commands. Both commands display information about system resources, connections, and settings that enable you to locate and fix problems, or to monitor system performance.

The one exception to these two main groups is the command `exec tac report`. This is an execute command that runs an exhaustive series of diagnostic commands. It runs commands that are only needed if you are using certain features like HA, VPN tunnels, or a modem. The report takes a few minutes to complete due to the amount of output generated. If you have your CLI output logged to a file, you can run this command to familiarize yourself with the CLI commands involved. Do not include the output from this command in FortiCare tickets unless it is specifically requested. See [“exec tac report” on page 806](#).

When you call Fortinet Customer Support, you will be asked to provide information about your unit and its current state using the output from these CLI commands.

This topic includes diagnostics commands to help with:

- [Check date and time](#)
- [Resource usage](#)
- [Proxy operation](#)
- [Hardware NIC](#)
- [Conserve mode](#)
- [Traffic trace](#)
- [Session table](#)
- [Firewall session setup rate](#)
- [Finding object dependencies](#)

- [Flow trace](#)
- [Packet sniffing and packet capture](#)
- [Debug command](#)
- [Other commands](#)

Additional diagnostic commands are covered in “[Troubleshooting ‘get’ commands](#)” on [page 805](#), and commands related to specific features are covered in the chapter for that specific feature. For example in-depth diagnostics for dynamic routing are covered in the dynamic routing chapter.

Check date and time

The system date and time are important for FortiGuard services, when logging events, and when sending alerts. The wrong time will make the log entries confusing and difficult to use.

Use Network Time Protocol (NTP) to set the date and time if possible. This is an automatic method that does not require manual intervention. However, you must ensure the port is allowed through the firewalls on your network. FortiToken synchronization requires NTP in many situations.

How to check the date and time - web-based manager

1 Go to **System Information > System Time** on the dashboard.

Follow *System > Dashboard > Status* to the dashboard. Then in the System Information widget, check the date and time. Alternately, you can check the date and time using the CLI commands `execute date` and `execute time`.

2 If required, select **Change** to adjust the date and time settings.

To make changes to the date and time, select Change. Then you can set the time zone, date and time, and select NTP usage. In the CLI, use the following commands to change the date and time:

```
config system global
    set timezone (use ? to get a list of IDs and descriptions of
        their timezone)
    set
config system ntp
    config ntpserver
    edit 1
        set server "ntp1.fortinet.net"
    next
    edit 2
        set server "ntp2.fortinet.net"
    next
end
set ntpsync enable
set syncinterval 60
end
```

Resource usage

Each program running on a computer has one or more processes associated with it. For example if you open a Telnet program, it will have an associated telnet process. The same is true in FortiOS. All the processes have to share the system resources in FortiOS including memory and CPU.

Monitor the CPU/memory usage of internal processes using the following command:

```
get system performance top <delay> <max_lines>
```

The data listed includes the name of the daemon, the process ID, whether the process is sleeping or running, the CPU percentage being used, and the memory percentage being used. See [“Check CPU and memory resources” on page 772](#).

How to troubleshoot high memory usage

FortiOS has limited system resources such as memory. All the processes running, share that memory. Depending on their workload each process will use more or less as needed, usually more in high traffic situations. If some processes use all the available memory, other processes will have no memory available and not be able to function.

When high memory usage happens, you may experience services that appear to freeze up and connections are lost or new connections are refused.

If you are seeing high memory usage in the system resources widget, it could mean that the unit is dealing with high traffic volume, which may be causing the problem, or it could be when the unit is dealing with connection pool limits affecting a single proxy. If the unit is receiving large volumes of traffic on a specific proxy, it is possible that the unit will exceed the connection pool limit. If the number of free connections within a proxy connection pool reaches zero, problems may occur.

Use the following CLI command, which uses the antivirus failopen feature. Setting it to idledrop will drop connections based on the clients that have the most connections open. This helps to determine the behavior of the FortiGate antivirus system if it becomes overloaded in high traffic.

```
config system global
    set av_failopen idledrop
end
```

How to troubleshoot high CPU usage

FortiOS has many features. If many of them are used at the same time, it can quickly use up all the CPU resources. When this happens, you will experience connection related problems stemming from the FortiOS unit trying to manage its workload by refusing new connections, or even more aggressive methods.

Some examples of features that are CPU intensive are VPN high level encryption, having all traffic undergo all possible scanning, logging all traffic, and packets, and dashboard widgets that frequently update their data.

1 Determine how high the CPU usage is currently.

There are two main ways to do this. The easiest is to go to System > Dashboard and look at the resource monitor. This is a dial gauge that displays a percentage use for the CPU. If its at the red-line, you should take action. The other method is to use the Dashboard CLI widget to enter `diag sys top`.

Sample output:

```
Run Time: 13 days, 13 hours and 58 minutes
0U, 0S, 98I; 123T, 25F, 32KF
newcli    903      R      0.5      5.5
```

```
sshd      901      S      0.5      4.0
```

Where the codes displayed on the second output line mean the following:

- **U** is % of user space applications using CPU. In the example, **0U** means 0% of the user space applications are using CPU.
- **S** is % of system processes (or kernel processes) using CPU. In the example, **0S** means 0% of the system processes are using the CPU.
- **I** is % of idle CPU. In the example, **98I** means the CPU is 98% idle.
- **T** is the total FortiOS system memory in Mb. In the example, **123T** means there are 123 Mb of system memory.
- **F** is free memory in Mb. In the example, **25F** means there is 25 Mb of free memory.
- **KF** is the total shared memory pages used. In the example, **32KF** means the system is using 32 shared memory pages.

Each additional line of the command output displays information for each of the processes running on the FortiGate unit. For example, the third line of the output is:

```
newcli    903      R      0.5      5.5
```

Where:

- **newcli** is the process name. Other process names can include **ipsengine**, **sshd**, **cmdbsrv**, **httpsd**, **scanunitd**, and **miglogd**.
- **903** is the process ID. The process ID can be any number.
- **R** is the current state of the process. The process state can be:
 - **R** running
 - **S** sleep
 - **Z** zombie
 - **D** disk sleep.
- **0.5** is the amount of CPU that the process is using. CPU usage can range from 0.0 for a process that is sleeping to higher values for a process that is taking a lot of CPU time.
- **5.5** is the amount of memory that the process is using. Memory usage can range from 0.1 to 5.5 and higher.

Enter the following single-key commands when **diagnose sys top** is running:

- Press **q** to quit and return to the normal CLI prompt.
- Press **p** to sort the processes by the amount of CPU that the processes are using.
- Press **m** to sort the processes by the amount of memory that the processes are using.

2 Determine what features are using most of the CPU resources.

There is a command in the CLI to let you see the top few processes currently running that use the most CPU resources. The CLI command `get system performance top` outputs a table of information. You are interested in the second most right column — CPU usage by percentage. If the top few entries are using most of the CPU, note which processes they are and investigate those features to try and reduce their CPU load. Some examples of processes you will see are

- `ipsengine` — the IPS engine that scans traffic for intrusions
- `scanunitd` — antivirus scanner
- `httpsd` — secure HTTP
- `iked` — internet key exchange (IKE) in use with IPsec VPN tunnels
- `newcli` — active whenever you are accessing the CLI
- `sshd` — there are active secure socket connections
- `cmdbsrv` — the command database server application

Go to the features that are at the top of the list and look for evidence of them overusing the CPU. Generally the monitor for a feature is a good place to start.

3 Check for unnecessary CPU “wasters”.

These are some best practises that will reduce your CPU usage, even if you are not experiencing high CPU usage. Note that if you require a feature this section tell you to turn off, ignore it.

- Use hardware acceleration wherever possible to offload tasks from the CPU. Offloading tasks such as encryption frees up the CPU for other tasks.
- Avoid the use of GUI widgets that require computing cycles, such as the Top Sessions widget. These widgets are constantly polling the system for their information which uses CPU and other resources.
- Schedule antivirus, IPS, and firmware updates during off peak hours. Usually these don't consume CPU resources but they can disrupt normal operation.
- Check the log levels and which events are being logged. This is the severity of the messages that are recorded. Consider going up one level to reduce the amount of logging. Also if there are events you do not need to monitor, remove them from the list.
- Log to disk instead of memory. Logging to memory quickly uses up resources. Logging to local disk is fast and doesn't take much CPU.
- If the disk is almost full, transfer logs or data off the disk to free up space. When a disk is almost full it consumes a lot of resources to find the free space and organize the files.
- If you have packet logging enabled, consider disabling it. When its enabled it records every packet that comes through that policy.
- Halt all sniffers and traces.
- Ensure you are not scanning traffic twice. If traffic enters the FortiGate unit on one interface, goes out another, and then comes back in again that traffic does not need to be rescanned. Doing so is a waste of resources. However, ensure that traffic truly is being scanned once.
- Reduce the session timers to close unused sessions faster. To do this in the CLI enter the following commands and values. These values reduce the values from defaults. Note that `tcp-timewait` has 10 seconds added by the system by default.

```

config system global
    set tcp-halfclose-timer 30
    set tcp-halfopen-timer 30
    set tcp-timewait-timer 0 set udp-idle-timer 60
end

```

- Remove dns-udp firewall session helper (number 14) if not used.
- Do not enable “nice to have” features.

4 When CPU usage is under control, use SNMP to monitor CPU usage. Alternately, use logging to record CPU and memory usage every 5 minutes.

Once things are back to normal, you should set up a warning system to alert you of future CPU overusage. A common method to do this is with SNMP. SNMP monitors many values on the FortiOS and allows you to set high water marks that will generate events. You run an application on your computer to watch for and record these events. Go to *System > Config > SNMP* to enable and configure an SNMP community. If this method is too complicated, you can use logging to record CPU usage every 5 minutes. However this method will not alert you to problems; it will just record them as they happen.

Proxy operation

Monitor proxy operations using the following command:

```
diag test application <application> <option>
```

The <application> value can include the following:

```

ftpd ftp proxy
http http proxy
im im proxy
imapi map proxy
ipsengine ips sensor
ipsmonitor ips monitor
pop3 pop3 proxy
smtp smtp proxy
urlfilter urlfilter daemon

```

The <option> value for use with this command can include:

- 1 Dump Memory Usage
- 2 Drop all connections
- 4 Display connection state *
- 44 Display info per connection *
- 444 Display connections per state *
- 5 Toggle AV Bypass mode
- 6 Toggle Print Stat mode every ~40 seconds
- 7 Toggle Backlog Drop
- 8 Clear stats
- 88 Toggle statistic recording - stats cleared

9 Toggle Accounting info for display

99 Restart proxy

These commands, except for the ones identified with an “*”, should only be used under the guidance of Fortinet Support.

Hardware NIC

Monitor hardware network operations using the following command:

```
diag hardware deviceinfo nic <interface>
```

The information displayed by this command is important as errors at the interface are indicative of data link or physical layer issues which may impact the performance of the FortiGate unit.

The following is sample output when `<interface> = internal`:

```
System_Device_Name  port5
Current_HWaddr      00:09:0f:68:35:60
Permanent_HWaddr    00:09:0f:68:35:60
Link                up
Speed               100
Duplex              full
[.....]
Rx_Packets=5685708
Tx_Packets=4107073
Rx_Bytes=617908014
Tx_Bytes=1269751248
Rx_Errors=0
Tx_Errors=0
Rx_Dropped=0
Tx_Dropped=0
[... .]
```

Hardware troubleshooting

The `diag hardware deviceinfo nic` command displays a list of hardware related error names and values. The following table explains the items in the list and their meanings.

Table 52: Possible hardware errors and meanings

Field	Definition
Rx_Errors = rx error count	Bad frame was marked as error by PHY.
Rx_CRC_Errors + Rx_Length_Errors - Rx_Align_Errors	This error is only valid in 10/100M mode.
Rx_Dropped or Rx_No_Buffer_Count	Running out of buffer space.

Table 52: Possible hardware errors and meanings

Field	Definition
Rx_Missed_Errors	Equals Rx_FIFO_Errors + CEXTERR (Carrier Extension Error Count). Only valid in 1000M mode, which is marked by PHY.
Tx_Errors = Tx_Aborted_Errors	ECOL (Excessive Collisions Count). Only valid in half-duplex mode.
Tx_Window_Errors	LATECOL (Late Collisions Count). Late collisions are collisions that occur after 64-byte time into the transmission of the packet while working in 10 to 100Mb/s data rate and 512-byte time into the transmission of the packet while working in the 1000Mb/s data rate. This register only increments if transmits are enabled and the device is in half-duplex mode.
Rx_Dropped	See Rx_Errors.
Tx_Dropped	Not defined.
Collisions	Total number of collisions experienced by the transmitter. Valid in half-duplex mode.
Rx_Length_Errors	Transmission length error.
Rx_Over_Errors	Not defined.
Rx_CRC_Errors	Frame CRC error.
Rx_Frame_Errors	Same as Rx_Align_Errors. This error is only valid in 10/100M mode.
Rx_FIFO_Errors	Same as Rx_Missed_Errors - a missed packet count.
Tx_Aborted_Errors	See Tx_Errors.
Tx_Carrier_Errors	The PHY should assert the internal carrier sense signal during every transmission. Failure to do so may indicate that the link has failed or the PHY has an incorrect link configuration. This register only increments if transmits are enabled. This register is not valid in internal SerDes 1 mode (TBI mode for the 82544GC/EI) and is only valid when the Ethernet controller is operating at full duplex.
Tx_FIFO_Errors	Not defined.
Tx_Heartbeat_Errors	Not defined.
Tx_Window_Errors	See LATECOL.
Tx_Single_Collision_Frames	Counts the number of times that a successfully transmitted packet encountered a single collision. The value only increments if transmits are enabled and the Ethernet controller is in half-duplex mode.
Tx_Multiple_Collision_Frames	A Multiple Collision Count which counts the number of times that a transmit encountered more than one collision but less than 16. The value only increments if transmits are enabled and the Ethernet controller is in half-duplex mode.
Tx_Deferred	Counts defer events. A defer event occurs when the transmitter cannot immediately send a packet due to the medium being busy because another device is transmitting, the IPG timer has not expired, half-duplex deferral events are occurring, XOFF frames are being received, or the link is not up. This register only increments if transmits are enabled. This counter does not increment for streaming transmits that are deferred due to TX IPG.
Rx_Frame_Too_Longs	The Rx frame is over size.

Table 52: Possible hardware errors and meanings

Field	Definition
Rx_Frame_Too_Short	The Rx frame is too short.
Rx_Align_Errors	This error is only valid in 10/100M mode.
Symbol Error Count	Counts the number of symbol errors between reads - SYMERRS. The count increases for every bad symbol received, whether or not a packet is currently being received and whether or not the link is up. This register only increments in internal SerDes mode.



The counters displayed depend on the type of the NIC interface. Please see the following website for more information:

<http://kc.forticare.com/default.asp?id=1979&Lang=1&SID=>

Conserve mode

The FortiOS antivirus and IPS systems operate in one of two modes, depending on the available shared memory. If the shared memory utilization is below a defined upper threshold, the system is in non-conserve mode. If shared memory usage goes above this threshold, the system enters conserve mode and remains in this state until the shared memory usage drops below a second threshold, slightly lower than the original. These thresholds are non-configurable; the threshold above which the system enters conserve mode is 80 percent, the system will not go back to non-conserve mode until the shared memory usage goes below 70 percent.

When does it happen

Conserve mode occurs under high usage and traffic conditions. It is expected to be a temporary condition that is self-correcting when bursty traffic subsides. When it occurs, users will not be able to open any new sessions until the old ones are closed. This means they will not be able to open new web browser pages, download new files, or other activities over the network. It is possible that communications sessions such as VoIP calls will be cutoff.

When in conserve mode, any new sessions are ignored (no SYN-ACK from the FortiGate unit) or passed without being scanned.

The following is sample output from the event log when entering conserve mode.

```
66 2008-04-16 14:01:31 critical The system has entered system
conserve mode
```

Antivirus failopen

Antivirus failopen is a safeguard feature that determines the behavior of the FortiGate unit antivirus system if it becomes overloaded in high traffic.

When does it happen

When there is high network traffic volume, it may cause low memory situations or it may limit the connection pool for a single proxy. If a FortiGate unit is receiving large volumes of traffic on a specific proxy (such as ikev2, the IPsec vpn IKE proxy), it is possible that the unit will exceed the connection pool limit. If the number of free connections within a proxy connection pool reaches zero, problems may occur. The connection pool is a reserved number of connections that a proxy has available to use. This is intended to reserve connections so no single proxy can consume all the resources.

The feature is configurable only through the CLI.

```
config system global
  set av_failopen {off|one-shot|pass |idledrop}
end
```

`av-failopen-session` controls the behavior when the proxy connection pool is exhausted. Again in this case, the FortiGate unit does not send the SYN-ACK. Failopen is only available on FortiGate models 300A and higher. On other lower FortiGate models, the failopen action is configured to pass.

The `set av-failopen` command has the following four options:

- off

If the FortiGate unit enters conserve mode, the antivirus system will stop accepting new AV sessions but will continue to process current active sessions.

- one-shot

If the FortiGate unit enters conserve mode, all subsequent connections bypass the antivirus system but current active sessions will continue to be processed. One-shot is similar to pass but will not automatically turn off once the condition causing `av-failopen` has stopped.

- idledrop

When configured in this mode, the antivirus failopen mechanism will drop connections based on the clients that have the most open connections.

- pass

Pass becomes the default setting when the `av-failopen-session` command has been run. If the system enters conserve mode, connections bypass the antivirus system until the system enters non-conserve mode again. Current active sessions will continue to be processed.



The one-shot and pass options do not content filter traffic. Therefore, the data stream could contain malicious content.

Traffic trace

Traffic tracing allows a specific packet stream to be followed. This is useful to confirm packets are taking the route you expected on your network.

View the characteristics of a traffic session through specific security policies using:

```
diag sys session
```

Trace per-packet operations for flow tracing using:

```
diag debug flow
```

Trace per-Ethernet frame using:

```
diag sniffer packet
```

Session table

A session is a communication channel between two devices or applications across the network. Sessions enable FortiOS to inspect and act on a sequential group of packets in a session all together instead of inspecting each packet individually. Each of these sessions has an entry in the session table that includes important information about the session.

Use as a tool

Session tables are useful troubleshooting tools because they allow you to verify connections that you expect to see open. For example, if you have a web browser open to browse the Fortinet website, you would expect a session entry from your computer, on port 80, to the IP for the Fortinet website. Another troubleshooting method is if there are too many sessions for FortiOS to process, you can examine the session table for evidence why this is happening.

The FortiGate session table can be viewed from either the CLI or the web-based manager. The most useful troubleshooting data comes from the CLI. The session table in web-based manager also provides some useful summary information, particularly the current policy number that the session is using.

Web-based manager session information

In the web-based manager there are actually two places to view session information — the policy session monitor, and the dashboard Top Sessions Widget.

Top Sessions widget

In Web Config, click *Add Content* and select *Top Sessions*. In the *Top Sessions* pane, click the *Details* link. (If there are not enough entries in the session table, try browsing to a different web site and re-examine the table.) The *Policy ID* shows which security policy matches the session. The sessions that do not have a *Policy ID* entry originate from the FortiGate device

Session monitor

The session monitor is the session table. It lists the protocol used, source and destination addresses, source and destination ports, what policy ID was matched (if any), how long until the session expires, and how long it has been established.

If there is no policy ID listed in the session entry, the traffic originated from the FortiGate unit. Otherwise all sessions must match a security policy to pass through the FortiGate unit.

As there are potentially many sessions active at one time, there are different methods you can use to filter unimportant sessions out of your search. The easiest filter is to display only IPv4 or IPv6 sessions. By default both are displayed. For this option, you must have IPv6 displayed (to enable go to *System > Admin > Settings*).

Refresh Filter Settings										IPv4 IPv6 Both		
#	Protocol	Src Address	Src Port	Src NAT IP	Src NAT Port	Dst Address	Dst Port	Policy ID	Expiry (sec)	Duration (sec)		
1	tcp	172.20.120.230	49287			172.20.120.136	443		3594	6		
2	tcp	172.20.120.230	49284			172.20.120.136	443		3594	8		
3	tcp	172.20.120.230	49283			172.20.120.136	443		3593	8		
4	tcp	172.20.120.230	49286			172.20.120.136	443		3593	6		
5	tcp	172.20.120.230	49285			172.20.120.136	443		3593	6		
6	tcp	172.20.120.230	49282			172.20.120.136	443		3600	10		
7	tcp	172.20.120.136	12266			172.20.120.23	541		25	97		
8	tcp	172.20.120.136	12267			172.20.120.23	541		43	79		
9	tcp	172.20.120.136	12268			172.20.120.23	541		61	61		
10	tcp	172.20.120.136	12269			172.20.120.23	541		79	43		
11	tcp	172.20.120.136	12270			172.20.120.23	541		97	25		
12	tcp	172.20.120.136	12271			172.20.120.23	541		115	7		
13	tcp	172.20.120.136	12265			172.20.120.23	541		7	115		

How to find which security policy a specific connection is using

Every program and device on your network must have a communication channel, or session, open to pass information. The FortiGate unit manages these sessions with its many features from traffic shaping, to antivirus scanning, and even blocking known bad web sites. Each session has an entry in the session table. In the web , you can use the Session Monitor or Top Session widget to view session information.

You may want to find information for a specific session, say a secure web browser session, for troubleshooting. For example if that web browser session is not working properly, you can check the session table to ensure the session is still active, and that it is going to the proper address. It can also tell you the security policy number it matches, so you can check what is happening in that policy.

1 Know your connection information.

You need to be able to identify the session you want. For this you need the source IP address (usually your computer), the destination IP address if you have it, and the port number which is determined by the program being used. Some commons ports are:

- port 80 (HTTP for web browsing),
- port 22 (SSH used for secure login and file transfers)
- port 23 (telnet for a text connection)
- port 443 (HTTPS for secure web browsing)

2 Find your session and policy ID.

Follow *Policy > Monitor > Session Monitor* to the session table monitor. Find your session by finding your source IP address, destination IP address if you have it, and port number. The policy ID is listed after the destination information. If the list of sessions is very long, you can filter the list to make it easier to find your session.

3 When there are many sessions, use a filter to help you find your session.

If there are multiple pages of sessions it is difficult to find a single session. To help you in your search you can use a filter to block out sessions that you don't want. Select the filter icon next to Src Address. In the window that pops up, enter your source IP address and select Apply. Now only sessions that originate from your IP address will be displayed in the session table. If the list is still too long, you can do the same for the Src port. That will make it easy to find your session and the security policy ID. When you are finished remember to clear the filters.

CLI session information

The session table output from the CLI (`diag sys session list`) is very verbose. Even on a system with a small amount of traffic, displaying the session table will generate a large amount of output. For this reason, filters are used to display only the session data of interest.

You can filter a column in the web-based manager by clicking the funnel icon on the column heading or from the CLI by creating a filter.

An entry is placed in the session table for each traffic session passing through a security policy. The following command will list the information for a session in the table:

```
diag sys session list
```

For extended information about all aspects of sessions, see [“get system session-info full-stat” on page 842](#).

Sample Output:

```
FGT# diag sys session list
```



```

session info: proto=6 proto_state=05 expire=89 timeout=3600
              flags=00000000 av_idx=0 use=3
bandwidth=204800/sec    guaranteed_bandwidth=102400/sec
              traffic=332/sec prio=0  logtype=session ha_id=0 hakey=4450
tunnel=/
state=log shape may_dirty
statistic(bytes/packets/err): org=3408/38/0 reply=3888/31/0
              tuples=2
origin->sink: org pre->post, reply pre->post oif=3/5
              gwy=192.168.11.254/10.0.5.100
hook=post dir=org act=snat 10.0.5.100:1251-
              >192.168.11.254:22(192.168.11.105:1251)
hook=pre dir=reply act=dnat 192.168.11.254:22-
              >192.168.11.105:1251(10.0.5.100:1251)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 domain_info=0 auth_info=0 ftgd_info=0 ids=0x0 vd=0
              serial=00007c33 tos=ff/ff

```

Since output can be verbose, the filter option allows specific information to be displayed, for example:

```
diag sys session filter <option>
```

The <option> values available include the following:

clear	clear session filter
dport	destination port
dst	destination IP address
negate	inverse filter
policy	policy ID
proto	protocol number
sport	source port
src	source IP address
vd	index of virtual domain. -1 matches all

Even though UDP is a sessionless protocol, the FortiGate unit still keeps track of the following two different states:

- UDP reply not seen with a value of 0
- UDP reply seen with a value of 1

The following illustrates FW session states from the session table:

State	Meaning
log	Session is being logged.
local	Session is originated from or destined for local stack.
ext	Session is created by a firewall session helper.

State	Meaning
may_dirty	Session is created by a policy. For example, the session for <code>ftp control channel</code> will have this state but <code>ftp data channel</code> will not. This is also seen when NAT is enabled.
ndr	Session will be checked by IPS signature.
nds	Session will be checked by IPS anomaly.
br	Session is being bridged (TP) mode.

Firewall session setup rate

The number of sessions that can be established in a set period of time is useful information. A session is an end-to-end TCP/IP connection for communication with a limited lifespan. If you record the setup rate during normal operation, when you experience problems you have that setup rate with the current number to see if its very different. While this will not solve your problems, it can be a useful step to help you define your problem.

A reduced firewall session setup rate could be the result of a number of things from a lack of system resources on the FortiGate unit, to reaching the limit of your session count for your VDOM.

To view your session setup rate - web-based manager

- 1 Got to *System > Dashboard > Dashboard*.
- 2 Maximize the *Top Sessions* widget.
- 3 Read the *New Sessions per Second* value displayed at the bottom of the widget.

If the *Top Sessions* widget is not visible on your dashboard, go to the + *Widget* button at the top of the window. When a window pops up, select *Top Sessions* for it to be added to the dashboard.

To view your session setup rate method 1- CLI

```
FGT# get sys performance status
CPU states: 0% user 0% system 0% nice 100% idle
Memory states: 10% used
Average network usage: 0 kbps in 1 minute, 0 kbps in 10 minutes,
13 kbps in 30 minutes
Average sessions: 31 sessions in 1 minute, 30 sessions in 10
minutes, 31 sessions in 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 44 days, 18 hours, 42 minutes
```

The information you are looking for is the Average sessions section, highlighted in the above output. In this example you can see there were 31 sessions in 1 minute, or an average of 0.5 sessions per second. The values for 10 minutes and 30 minutes allow you to take a longer average for a more reliable value if your FortiGate unit is working at maximum capacity. The smallest FortiGate unit can have 1 000 sessions established per second across the unit.

Remember that session setup rate is a global command. If you have multiple VDOMs configured with many sessions in each one, the session setup rate per VDOM will be slower than if there were no VDOMs configured.

Finding object dependencies

An administrator may not be permitted to delete a configuration object if there are other configuration objects that depend on it. This command identifies other objects which depend on or *make reference to* the configuration object in question. If an error is displayed that an object is in use and cannot be deleted, this command can help identify the source of the problem.

Another use is if you have a virtual interface with objects that depend on it, you need to find and remove those dependencies before you delete that interface.

CLI method

When running multiple VDOMs, this command is run in the Global configuration only and it searches for the named object both in the Global and VDOM configuration most recently used:

```
diag sys checkused <path.object.mkey>
```

For example, to verify which objects are referred to in a security policy with an ID of 1, enter the command as follows:

```
diag sys checkused firewall.policy.policyid 1
```

To check what is referred to by interface `port1`, enter the following command:

```
diag sys checkused system.interface.name port1
```

To show all the dependencies for an interface, enter the command as follows:

```
diag sys checkused system.interface.name <interface name>
```

Sample Output:

```
entry used by table firewall.address:name '10.98.23.23_host'
entry used by table firewall.address:name 'NAS'
entry used by table firewall.address:name 'all'
entry used by table firewall.address:name 'fortinet.com'
entry used by table firewall.vip:name 'TORRENT_10.0.0.70:6883'
entry used by table firewall.policy:policyid '21'
entry used by table firewall.policy:policyid '14'
entry used by table firewall.policy:policyid '19'
```

In this example, the interface has dependent objects, including four address objects, one VIP, and three security policies.

Web-based manager method

In the web-based manager, the object dependencies for an interface can be easily checked and removed.

To remove interface object dependencies - web-based manager

- 1 Go to *System > Interface*.

The number in the *Ref.* column is the number of objects that refer to this interface.

- 2 Select the number in the *Ref.* column for the desired interface.
A Window listing the dependencies will appear.
- 3 Use these detailed entries to locate and remove object references to this interface.
The trash can icon will change from gray when all object dependencies have been removed.
- 4 Remove the interface by selecting the check box for the interface, and select *Delete*.

Flow trace

To trace the flow of packets through the FortiGate unit, use the following command:

```
diag debug flow trace start
```

Follow packet flow by setting a flow filter using this command:

```
diag debug flow filter <option>
```

Filtering options include the following:

addr	IP address
clear	clear filter
daddr	destination IP address
dport	destination port
negate	inverse filter
port	port
proto	protocol number
saddr	source IP address
sport	source port
vd	index of virtual domain, -1 matches all

Enable the output to be displayed to the CLI console using the following command:

```
diag debug flow show console
```



diag debug flow output is recorded as event log messages and are sent to a FortiAnalyzer unit if connected. **Do not let this command run longer than necessary since it generates significant amounts of data.**

Start flow monitoring with a specific number of packets using this command:

```
diag debug flow trace start <N>
```

Stop flow tracing at any time using:

```
diag debug flow trace stop
```

The following is an example of the flow trace for the device at the following IP address:
203.160.224.97

```
diag debug enable
diag debug flow filter addr 203.160.224.97
diag debug flow show console enable
diag debug flow show function-name enable
diag debug flow trace start 100
```

Flow trace output example - HTTP

Connect to the web site at the following address to observe the debug flow trace. The display may vary slightly:

```
http://www.fortinet.com
```

Comment: SYN packet received:

```
id=20085 trace_id=209 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

SYN sent and a new session is allocated:

```
id=20085 trace_id=209 func=resolve_ip_tuple line=2799
msg="allocate a new session-00000e90"
```

Lookup for next-hop gateway address:

```
id=20085 trace_id=209 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.11.254 via port6"
```

Source NAT, lookup next available port:

```
id=20085 trace_id=209 func=get_new_addr line=1219
msg="find SNAT: IP-192.168.11.59, port-31925"
direction"
```

Matched security policy. Check to see which policy this session matches:

```
id=20085 trace_id=209 func=fw_forward_handler line=317
msg="Allowed by Policy-3: SNAT"
```

Apply source NAT:

```
id=20085 trace_id=209 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

SYN ACK received:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6, 203.160.224.97:80-
>192.168.11.59:31925) from port6."
```

Found existing session ID. Identified as the reply direction:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, reply
```

```
direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=210 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Lookup for next-hop gateway address for reply traffic:

```
id=20085 trace_id=210 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.3.221 via port5"
```

ACK received:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, original
direction"
```

Apply source NAT:

```
id=20085 trace_id=211 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from client:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
original direction"
```

Apply source NAT:

```
id=20085 trace_id=212 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from server:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
203.160.224.97:80->192.168.11.59:31925) from port6."
```

Match existing session in reply direction:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=213 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Flow trace output example - IPsec (policy-based)

```
id=20085 trace_id=1 msg="vd-root received a packet(proto=1,
10.72.55.240:1->10.71.55.10:8) from internal."
id=20085 trace_id=1 msg="allocate a new session-00001cd3"
id=20085 trace_id=1 msg="find a route: gw-66.236.56.230 via
wan1"
id=20085 trace_id=1 msg="Allowed by Policy-2: encrypt"
id=20085 trace_id=1 msg="enter IPsec tunnel-RemotePhase1"
id=20085 trace_id=1 msg="encrypted, and send to 15.215.225.22
with source 66.236.56.226"
id=20085 trace_id=1 msg="send to 66.236.56.230 via intf-wan1"
id=20085 trace_id=2 msg="vd-root received a packet (proto=1,
10.72.55.240:1-1071.55.10:8) from internal."
id=20085 trace_id=2 msg="Find an existing session, id-00001cd3,
original direction"
id=20085 trace_id=2 msg="enter IPsec ="encrypted, and send to
15.215.225.22 with source 66.236.56.226" tunnel-RemotePhase1"
id=20085 trace_id=2 msgid=20085 trace_id=2 msg="send to
66.236.56.230 via intf-wan1"
```

Packet sniffing and packet capture

FortiOS devices can sniff packets using commands in the CLI or capture packets using the web-based manager. The differences between the two methods are not large.

Packet sniffing in the CLI is well suited for spot checking traffic from the CLI, but if you have complex filters to enter it can be a lot of work to enter them each time. You can also save the sniffing output; however, you must log to a file and then analyze the file later by hand.

Packet capture in the web-based manager makes it easy to set up multiple filters at once and just run one or two as you need them. You also have controls to start and stop capturing as you wish. Packet capture output is downloaded to your local computer as a *.pcap file which requires a third party application to read the file, such as Wireshark. This method is useful to send Fortinet support information to help resolve an issue.

Features	Packet sniffing	Packet capture
Command location	CLI	web-based manager
Third party software required	puTTY to log plaintext output	Wireshark to read *.pcap files
Read output in plain text file	yes	no
Read output as *.pcap file using Wireshark	no	yes
Easily configure single quick and simple filter	yes	no

Features	Packet sniffing	Packet capture
Record packet interface	yes	no
Configure complex sniffer filters on multiple interface	no	yes
sniff IPv6	hard	easy
sniff non-IP packets	no	yes
Filter packets by protocol and/or port	easy	easy
Filter packets by source and/or destination address	easy	easy

Packet sniffing

Before you start sniffing packets on the CLI, you should be prepared to capture the output to a file — there can be huge amounts of data that you will not be able to see without saving it to a file. One method is to use a terminal program like puTTY to connect to the FortiGate unit's CLI. Then once the packet sniffing count is reached you can end the session and analyze the output in the file.

Details within packets passing through particular interfaces can be displayed using the packet sniffer with the following command:

```
diag sniffer packet <interface> <filter> <verbose> <count>
```

The <interface> value is required, with the rest being optional. If not included the default values will be "none" 4 0 .

For example the simplest valid sniffer command would be:

```
diag sniffer packet any
```

The <interface> value can be any physical or virtual interface name. Use *any* to sniff packets on all interfaces.

The <filter> value limits the display of packets using filters, including Berkeley Packet Filtering (BPF) syntax. The <filter> value must be enclosed in quotes.

```
'[[src|dst] host <host_name_or_IP1> [[src|dst] host  
  <host_name_or_IP2>] [[arp|ip|ip6|gre|esp|udp|tcp] [port_no]]  
  [[arp|ip|ip6|gre|esp|udp|tcp] [port_no]] '
```

If a second host is specified in the filter, only the traffic between the two hosts will be displayed. Optionally, you can use logical OR to match only one of the hosts, or match one of multiple protocols or ports. When defining a port, there are up to two parts — protocol and port number.

For example, to display UDP 1812 traffic or TCP 8080 traffic, use the following:

```
'udp port 1812 or tcp port 8080'
```

To display all IP traffic that has a source of 192.168.1.2 and a destination of 192.168.2.3:

```
'ip src host 192.168.1.2 and dst host 192.168.2.3'
```

The <verbose> option allows different levels of information to be displayed. The verbose levels include:

- 1 Print header of packets

- 2 Print header and data from the IP header of the packets
- 3 Print header and data from the Ethernet header of the packets
- 4 Print header of packets with interface name
- 5 Print header and data from ip of packets with interface name
- 6 Print header and data from ethernet of packets with interface name

The `<count>` value indicates the number of packets to sniff before stopping. If this variable is not included, or is set to zero, the sniffer will run until you manually halt it with Ctrl-C.

Packet capture

FortiOS 4.0 MR3 Patch 2 introduced packet capture to the web-based manager. To configure packet capture filters, go to *System > Config > Advanced*.

When you add a packet capture filter, enter the following information and select *OK*.

Interface	Select the interface to sniff from the dropdown menu. You must select one interface. You cannot change the interface without deleting the filter and creating a new one, unlike the other fields.
Max Packets to Capture	Enter the number of packets to capture before the filter stops. This number cannot be zero. You can halt the capturing before this number is reached.
Source Address	Enter the source address as an IP address and netmask. This field cannot be empty and it cannot be a hostname. To capture traffic from all source addresses enter 0.0.0.0/0.0.0.0.
Source Port(s)	Enter one or more ports to capture on the source interface. Separate multiple ports with commas. Enter a range using a dash without spaces, for example 88-90.
Destination Address	Enter the destination address as an IP address and netmask. This field cannot be empty and it cannot be a hostname. To capture traffic from all destination addresses enter 0.0.0.0/0.0.0.0.
Destination Port(s)	Enter one or more ports to capture on the source interface. Separate multiple ports with commas. Enter a range using a dash without spaces, for example 88-90.
Protocol	Select a protocol to capture from the drop down list. Or select ALL from the list to capture all protocols.

Interface	Select the interface to sniff from the dropdown menu. You must select one interface. You cannot change the interface without deleting the filter and creating a new one, unlike the other fields.
Include IPv6 packets	Select this option if you are troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.
Capture Non-IP packets	The protocols available in the list are all IP based except for ICMP (ping). To capture non-IP based packets select this feature. Some examples of non-IP packets include IPsec, IGMP, ARP, and as mentioned ICMP.

If you select a filter and go back to edit it, you have the added option of starting and stopping packet capture in the edit window, or downloading the captured packets. You can also see the filter status and the number of packets captured.

You can also select the filter and select *Start* to start capturing packets. While the filter is running, you will see the number of captured packets increasing until it reaches the max packet count or you select *Stop*. While the filter is running you cannot download the output file.

When the packet capture is complete, you can select *Download* to send the packet capture filter captured packets to your local computer as a *.pcap file. To read this file format, you will need to use [Wireshark](#) or a similar third party application. Using this tool you will have extensive analytics available to you and the full contents of the packets that were captured.

FA2 and NP2 based interfaces

Many Fortinet products contain network processors. Some of these products contain FortiAccel (FA2) network processors while others contain NP2 network processors. Network processor features, and therefore offloading requirements, vary by network processor model.

When using the FA2- and NP2-based interfaces, only the initial session setup will be seen through the `diag debug flow` command. If the session is correctly programmed into the ASIC (fastpath), the debug flow command will no longer see the packets arriving at the CPU. If the NP2 functionality is disabled, the CPU will see all the packets, however, this should only be used for troubleshooting purposes.

First, obtain the NP2 and port numbers with the following command:

```
diag npu np2 list
```

Sample output:

```
ID PORTS
-- -----
0 port1
0 port2
0 port3
0 port4
ID PORTS
-- -----
1 port5
```

```
1 port6
1 port7
1 port8
ID PORTS
-- -----
2 port9
2 port10
2 port11
2 port12
ID PORTS
-- -----
3 port13
3 port14
3 port15
3 port16
```

Run the following commands:

```
diag npu np2 fastpath disable 0
```

(where 0 is the NP2 number)

Then, run this command:

```
diag npu np2 fastpath-sniffer enable port1
```

Sample output:

```
NP2 Fast Path Sniffer on port1 enabled
```

This will cause all traffic on *port1* of NP2 0 to be sent to the CPU meaning a standard sniffer trace can be taken and other diag commands should work if it was a standard CPU driven port.

These commands are only for the newer NP2 interfaces. FA2 interfaces are more limited as the sniffer will only capture the initial packets before the session is offloaded into HW (FA2). The same holds true for the `diag debug flow` command as only the session setup will be shown, however, this is usually enough for this command to be useful.

Debug command

Debug output provides continuous, real-time event information. Debugging output continues until it is explicitly stopped or until the unit is rebooted. Debugging output can affect system performance and will be continually generated even though output might not be displayed in the CLI console.

Debug information displayed in the console will scroll in the console display and may prevent CLI commands from being entered, for example, the command to disable the debug display. To turn off debugging output as the display is scrolling by, press the \uparrow key to recall the recent `diag debug` command, press backspace, and type "0", followed by *Enter*.

Debug output display is enabled using the following command:

```
diag debug enable
```

When finished examining the debug output, disable it using:

```
diag debug disable
```

Once enabled, indicate the debug information that is required using this command:

```
diag debug <option> <level>
```

Debug command options include the following:

application	method of debugging output from many FortiGate daemons
authd	configure FSSO or clear authentication daemon
cli	configure cli debug level
console	configure console settings for debugging
crashlog	get or clear the crash log info
disable	halt debug output
enable	start outputting filtered debug output
flow	trace packet flow in kernel
info	show active debug level settings
kernel	configure kernel and ha debug levels
rating	display website rating server list and information
report	report for tech support
reset	reset all debug level to default

The debug level can be set at the end of the command. Typical values are 2 and 3, for example:

```
diag debug application DHCPD 2
diag debug application spamfilter 2
```

Fortinet support will advise which debugging level to use.

Timestamps can be enabled to the debug output using the following command:

```
diag debug console timestamp enable
```

Debug output example

This example shows the IKE negotiation for a secure logging connection from a FortiGate unit to a FortiAnalyzer system.

```
diag debug reset
diag debug application ike 3 192.168.11.2
diag debug enable
```

Sample Output:

```
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2-
>192.168.10.201:500, natt_mode=0 rekey=0 phase2=FGh_FtiLog1
FGh_FtiLog1: using existing connection, dpd_fail=0
FGh_FtiLog1: found phase2 FGh_FtiLog1
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2 ->
192.168.10.201:500 negotiating
FGh_FtiLog1: overriding selector 225.30.5.8 with 192.168.11.2
FGh_FtiLog1: initiator quick-mode set pfs=1536...
FGh_FtiLog1: try to negotiate with 1800 life seconds.
```

```
FGh_FtiLog1: initiate an SA with selectors:
  192.168.11.2/0.0.0.0->192.168.10.201, ports=0/0,
  protocol=0/0
Send IKE Packet(quick_outI1):192.168.11.2:500(if0) ->
  192.168.10.201:500, len=348
Initiator: sent 192.168.10.201 quick mode message #1 (OK)
FGh_FtiLog1: set retransmit: st=168, timeout=6.
```

In this example:

192.168.11.2->192.168.10.201:500	Source and Destination gateway IP address
dpd_fail=0	Found existing Phase 1
pfs=1536...	Create new Phase 2 tunnel

Other commands

ARP table

To view the ARP cache, use the following command:

```
get sys arp
```

To view the ARP cache in the system, use this command:

```
diag ip arp list
```

Sample output:

```
index=14 ifname=internal 224.0.0.5 01:00:5e:00:00:05
  state=00000040 use=72203 confirm=78203 update=72203 ref=1
index=13 ifname=dmz 192.168.3.100 state=00000020 use=1843
  confirm=650179 update=644179 ref=2      ? VIP
index=13 ifname=dmz 192.168.3.109 02:09:0f:78:69:ff
  state=00000004 use=71743 confirm=75743 update=75743 ref=1
index=14 ifname=internal 192.168.11.56 00:1c:23:10:f8:20
  state=00000004 use=10532 confirm=10532 update=12658 ref=4
```

To remove the ARP cache, use this command:

```
execute clear system arp table
```

To remove a single ARP entry, use:

```
diag ip arp delete <interface name> <IP address>
```

To remove all entries associated with a particular interface, use this command:

```
diag ip arp flush <interface name>
```

To add static ARP entries, use the following command:

```
config system arp-table
```

Time and date settings

Check time and date settings for log message timestamp synchronization (the Fortinet support group may request this) and for certificates that have a time requirement to check for validity. Use the following commands:

```
execute time
current time is: 12:40:48
last ntp sync:Thu Mar 16 12:00:21 2006
execute date
current date is: 2006-03-16
```

To force synchronization with an NTP server, toggle the following command:

```
set ntpsync enable/disable
```

If all devices have the same time, it helps to correlate log entries from different devices.

IP address

There may be times when you want to verify the IP addresses assigned to the FortiGate unit interfaces are what you expect them to be. This is easily accomplished from the CLI using the following command.

```
diag ip address list
```

The output from this command lists the IP address and mask if available, the index of the interface (a sort of ID number) and the devname is the name of the interface. While physical interface names are set, virtual interface names can vary. Listing all the virtual interface names is a good use of this command. For vsys_ha and vsys_fgfm, the IP addresses are the local host — these are internally used virtual interfaces.

```
# diag ip address list
IP=10.31.101.100->10.31.101.100/255.255.255.0 index=3
devname=internal
IP=172.20.120.122->172.20.120.122/255.255.255.0 index=5
devname=wan1
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=8 devname=root
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=11 devname=vsys_ha
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=13 devname=vsys_fgfm
```

Other related commands include flushing the IP addresses (`diag ip address flush`), which will force a reload of the IP addresses. This can be useful if you think an IP address is wrong and don't want to reboot the unit. You can add or delete a single IP address (`diag ip address add <ipv4_addr>` or `diag ip address delete <ipv4_addr>`).

For more information on useful diagnostic commands, see [“Troubleshooting ‘get’ commands” on page 805](#) and [“Troubleshooting common issues” on page 763](#).

FortiGate ports

In the TCP and UDP stacks, there are 65 535 ports available for applications to use when communicating with each other. Many of these ports are commonly known to be associated with specific applications or protocols. These known ports can be useful when troubleshooting your network.

Use the following ports while troubleshooting the FortiGate device:

Port(s)	Functionality
UDP 53	DNS lookup, RBL lookup
UDP 53 or UDP 8888	FortiGuard Antispam or Web Filtering rating lookup
UDP 53 (default) or UDP 8888 and UDP 1027 or UDP 1031	FDN Server List - source and destination port numbers vary by originating or reply traffic. See the article "How do I troubleshoot performance issues when FortiGuard Web Filtering is enabled?" in the Knowledge Base.
UDP 123	NTP Synchronization
UDP 162	SNMP Traps
UDP 514	SYSLOG - All FortiOS versions can use syslog to send log messages to remote syslog servers. FortiOS v2.80 and v3.0 can also view logs stored remotely on a FortiAnalyzer unit.
TCP 22	Configuration backup to FortiManager unit or FortiGuard Analysis and Management Service.
TCP 25	SMTP alert email, encrypted virus sample auto-submit
TCP 389 or TCP 636	LDAP or PKI authentication
TCP 443	FortiGuard Antivirus or IPS update - When requesting updates from a FortiManager unit instead of directly from the FDN, this port must be reconfigured as TCP 8890.
TCP 443	FortiGuard Analysis and Management Service
TCP 514	FortiGuard Analysis and Management Service log transmission (OFTP)
TCP 541	SSL Management Tunnel to FortiGuard Analysis and Management Service (FortiOS v3.0 MR6 or later)
TCP 514	Quarantine, remote access to logs and reports on a FortiAnalyzer unit, device registration with FortiAnalyzer units (OFTP)
TCP 1812	RADIUS authentication
TCP 8000 and TCP 8002	FSSO
TCP 10151	FortiGuard Analysis and Management Service contract validation

Diagnostic commands

FortiAnalyzer/FortiManager ports

If you have a FortiAnalyzer unit or FortiManager unit on your network you may need to use the following ports for troubleshooting network traffic.

Functionality	Port(s)
DNS lookup	UDP 53
NTP synchronization	UDP 123
Windows share	UDP 137-138
SNMP traps	UDP 162
Syslog, log forwarding	UDP 514
Log and report upload	TCP 21 or TCP 22
SMTP alert email	TCP 25
User name LDAP queries for reports	TCP 389 or TCP 636
RVS update	TCP 443
RADIUS authentication	TCP 1812
Log aggregation client	TCP 3000



For more information about FortiAnalyzer/FortiManager ports, see the Fortinet Knowledge Base at the following address:

<http://kc.forticare.com/default.asp?SID=&Lang=1&id=773>.

FortiGuard troubleshooting

The FortiGuard service provides updates to Antivirus, IPSec, Webfiltering, and more. The FortiGuard Distribution System (FDS) involves a number of servers across the world that provide updates to your FortiGate unit. Problems can occur both with connection to FDS, and its configuration on your local FortiGate unit. Some of the more common troubleshooting methods are listed here including

- [Troubleshooting process for FortiGuard updates](#)
- [FortiGuard server settings](#)
- [FortiGuard URL rating](#)

Troubleshooting process for FortiGuard updates

The following process are the logical steps to take when troubleshooting FortiGuard update problems. This includes antivirus (AV), intrusion protection services (IPS), antispyware (AS), and web filtering (WB).

- 1 Does the device have a valid licence that includes these services?

Each device requires a valid FortiGuard license to access updates for some or all of these services. You can verify the support contract status for your devices at the Fortinet Support website — <https://support.fortinet.com/>.

- 2 If the device is part of an HA cluster, do all members of the cluster have the same level of support?

As with the previous step, you can verify the support contract status for all the devices in your HA cluster at the Fortinet Support website.

- 3 Have services been enabled on the device?

To see the FortiGuard information and status for a device, in the web-based manager go to *System > Config > FortiGuard*. On that page you can verify the status of each component, and if required enable each service. If there are problems, see the FortiGuard section of the FortiOS Handbook.

- 4 Is the device able to communicate with FortiGuard servers?

At *System > Config > FortiGuard* you can also attempt to update AV and IPS, or test the availability of WF and AS default and alternate ports. If there are problems, see the FortiGuard section of the FortiOS Handbook.

- 5 Is there proper routing to reach the FortiGuard servers?

Ensure there is a static or dynamic route that enables your FortiGate unit to reach the FortiGuard servers. Usually a generic default route to the internet is enough, but you may need to verify this if your network is complex.

- 6 Are there issues with DNS?

An easy way to test this is to attempt a traceroute from behind the FortiGate unit to an external network using the FQDN for a location. If the traceroute FQDN name does not resolve, you have general DNS problems. See [“DNS settings” on page 771](#).

- 7 Is there anything upstream that might be blocking FortiGuard traffic, either on the network or ISP side?

Many firewalls block all ports by default, and often ISPs block ports that are low. There may be a firewall between the FortiGate unit and the FortiGuard servers that is blocking the traffic. FortiGuard uses port 53 by default, so if it is being blocked you need to either open a hole for it, or change the port it is using.

- 8 Is there an issue with source ports?

It is possible that ports used to contact FortiGuard are being changed before reaching FortiGuard or on the return trip before reaching your FortiGate unit. A possible solution for this is to use a fixed-port at NATd firewalls to ensure the port remains the same. Packet sniffing can be used to find more information on what is happening with ports. See [“Perform a sniffer trace” on page 780](#).

- 9 Are there security policies that include antivirus?

If no security policies include antivirus, the antivirus database will not be updated. If antivirus is included, only the database type used will be updated.

FortiGuard server settings

Your local FortiGate unit connects to remote FortiGuard servers get updates to FortiGuard information such as new viruses that may have been found or other new threats. This section demonstrates ways to display information about FortiGuard server information on your FortiGate unit, and how to use that information and update it to fix potential problems. This includes

- [Displaying the server list](#)
- [Sorting the server list](#)
- [Calculating weight](#)

Displaying the server list

The `get webfilter status` command shows the list of FDS servers the FortiGate unit is using to send web filtering requests. Rating requests are only sent to the server on the top of the list in normal operation. Each server is probed for Round Trip Time (RTT) every two minutes.

You can optionally add a refresh rate to the end of this command and that will determine how often the server list will be refreshed.

Rating may not be enabled on your FortiGate unit.

```
get webfilter status
```

Sample Output:

```
Locale      : english
License     : Contract
Expiration  : Thu Oct  9 02:00:00 2011
Hostname    : fortiguard.example.com
-- Server List (Mon Feb 18 12:55:48 2008) --
```

IP	Weight	RTT	Flags	TZ	Packets	CurrLost	TotalLost
a.b.c.d	0	1	DI	2	1926879	0	11176
10.1.101.1	10	329		1	10263	0	633
10.2.102.2	20	169		0	16105	0	80
10.3.103.3	20	182		0	6741	0	776
10.4.104.4	20	184		0	5249	0	987
10.5.105.5	25	181		0	12072	0	178

Output Details

Hostname is the name of the FortiGuard server the FortiGate unit will attempt to contact. The Server List includes the IP addresses of alternate servers if the first entry cannot be reached. In this example the IP addresses are not public addresses

The following flags in `get webfilter status` indicate the server status:

- **D** - the server was found through the DNS lookup of the hostname. If the hostname returns more than one IP address, all of them will be flagged with D and will be used first for INIT requests before falling back to the other servers.
- **I** - the server to which the last INIT request was sent.
- **F** - the server has not responded to requests and is considered to have failed.
- **T** - the server is currently being timed.

Sorting the server list

The server list is sorted first by weight. The server with the smallest RTT is put at the top of the list, regardless of weight. When a packet is lost (there has been no response in 2 seconds), it will be resent to the next server in the list. Therefore, the top position in the list is selected based on RTT while the other list positions are based on weight.

Calculating weight

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a remote server, the weight is not allowed to dip below a base weight, calculated as the difference in hours between the FortiGate unit and the server times 10. The further away the server is, the higher its base weight and the lower in the list it will appear.



The output for the `diag debug rating` command will vary based on the state of the FortiGate device.

The following output is from a FortiGate device that has no DNS resolution for `service.fortiguard.net`.

```
COM1 - PuTTY
class-fgt (global) # diagnose debug rating
Locale      : english
License     : Unknown
Expiration  : N/A
Hostname    : service.fortiguard.net

-- Server List (Fri Sep  5 11:53:03 2008) --

IP          Weight  RTT Flags TZ   Packets  Curr Lost Total Lost
None

class-fgt (global) #
```

If only three IP addresses appear with the `D` flag, it means that DNS is good but probably the FortiGuard ports 53 and 8888 are blocked.

When the license is expired, an INIT request will be sent every 10 minutes for up to six attempts. If a license is not found after this limit is reached, the INIT requests will be sent every day.

A low source port number may appear which means that ports 1024 and 1025 could be blocked on the path to the FDS. Increase the source port on the FortiGate device with the following commands:

```
config sys global
set ip-src-port-range <start-end> (Default 1024-25000)
```

Be careful moving ports like this as it may cause some services to stop working if they can't access their original ports. If you make this change, ensure all services that use ports are checked and updated to new port numbers if needed.

FortiGuard URL rating

The following commands can be used to troubleshoot issues with FortiGuard URL ratings:

```
diag debug enable
diag debug application urlfilter -1
```

Sample output:

```
id=93000 msg="pid=57 urlfilter_main-723 in main.c received
pkt:count=91, a=/tmp/.thttp.socket/21" id=22009
msg="received a request /tmp/.thttp.socket, addr_len=21: d=
="www.goodorg.org:80, id=12853, vfid=0, type=0,
client=192.168.3.90, url="/" id=99501 user="N/A"
src=192.168.3.90 sport=1321 dst=<dest_ip> dport=80
service="http" cat=43 cat_desc="Organisation"
hostname="www.goodorg.org" url="/" status=blocked msg="URL
belongs to a denied category in policy"
```

Sample output:

```
id=22009 msg="received a request /tmp/.thttp.socket,
addr_len=21: d=pt.dnstest.google.com:80, id=300, vfid=0,
type=0, client=192.168.3.12, url=/gen_204"
id=93003 user="N/A" src=192.168.3.12 sport=21715 dst=<dest_ip>
dport=80 service="http" cat=41 cat_desc="Search Engines"
hostname="pt.dnstest.google.com" url="/gen_204"
status=passthrough msg="URL belongs to an allowed category in
the policy"
```

Table 53: Breakdown of sample output parts from URL rating command

id=93000 msg="pid=57 urlfilter_main-723 in main.c received pkt:count=91, a=/tmp/.thttp.socket/21"	The process ID (PID) is listed along with the function in the file running (main.c). Then it lists the number of packets received and the associated socket where the packets came from.
id=22009 msg="received a request /tmp/.thttp.socket, addr_len=21: d= ="www.goodorg.org:80, id=12853, vfid=0, type=0, client=192.168.3.90, url="/"	Received a request on a particular socket (/tmp/.thttp.socket). The website to be rated is "www.goodorg.org:80" and the client browser that wants the verification is 192.169.3.90.
id=99501 user="N/A" src=192.168.3.90 sport=1321 dst=<dest_ip> dport=80 service="http" cat=43 cat_desc="Organisation" hostname="www.goodorg.org" url="/" status=blocked msg="URL belongs to a denied category in policy"	No user associated with this source address (192.168.3.90) and port (1321). The destination IP is unknown and the port is the standard HTTP port 80, which is confirmed by service=http. The cat keyword gives the category of the URL being checked, which turns out to be an organization. This is confirmed by the hostname of "goodorg.org". The status is stated as blocked with the reason stated as "URL belongs to a denied category in policy".



Technical Support Organization Overview

This section explains how Fortinet's technical support works, as well as how you can easily create an account to get technical support for when issues arise that you cannot solve yourself.

This section contains the following topics:

- [Fortinet Global Customer Services Organization](#)
- [Creating an account](#)
- [Registering a device](#)
- [Reporting problems](#)
- [Assisting technical support](#)
- [Support priority levels](#)
- [Return material authorization process](#)

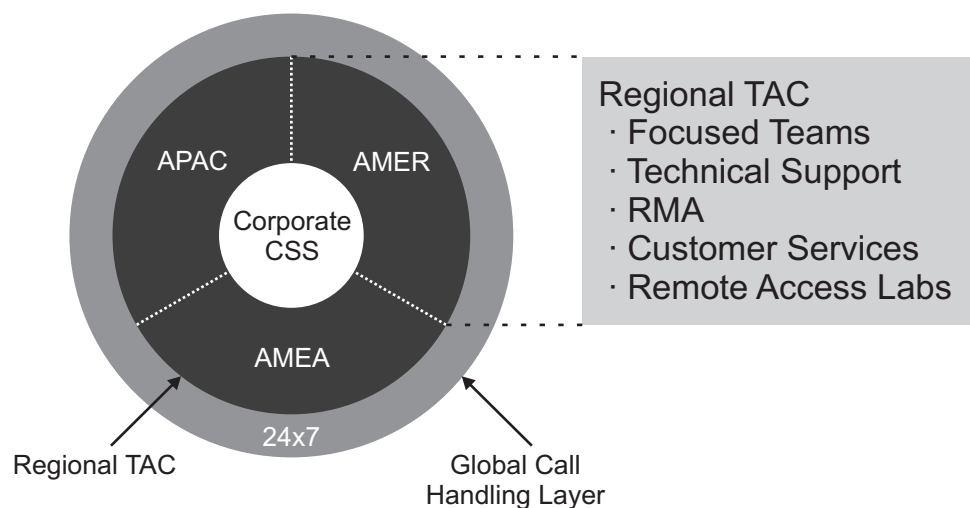
Fortinet Global Customer Services Organization

The Fortinet Global Customer Services Organization is composed of three regional *Technical Assistance Centers* (TAC):

- The Americas (AMER)
- Europe, Middle East, and Africa (EMEA)
- Asia Pacific (APAC)

The regional TACs are contacted through a global call center. Incoming service requests are then routed to the appropriate TAC.

Each regional TAC delivers technical support to the customers in its regions during its hours of operation. These TACs also combine to provide seamless, around-the-clock support for all customers.

Figure 72: Fortinet regions and TAC.

Creating an account

To receive technical support and service updates, Fortinet products in the organization must be registered. The **Product Registration Form** on the support website will allow the registration to be completed online.

Creating an account on the support website is the first step in registering products

Figure 73: Support account product registration form

Home > Support

PRODUCTS REGISTRATION FORM

Product Registration
» Registration FAQ

Knowledge Base

Support Login

Premium RMA

Contact Support

Training & Certification

Product Registration Form

Thank you for purchasing your FortiGate Antivirus Firewall and for taking the opportunity to register the product online. Completing the information below will enable you to receive updates for your threat detection and prevention databases (Antivirus, Intrusion Detection, etc.) and will also ensure your access to technical support.

Please note that we will use your email address submitted below to send your registration. If you have already received a username and password from Fortinet Support, please [login](#) here.

Contact Information

First Name * Last Name *

Company * Title

Email *

Address 1 *

Address 2

City * State/ Province *

Post Code * Country/ Region *

Contact Phone * Fax Number

Password *

Retype *

CONTACT SUPPORT

AMERICAS
Tel: 1-800-048-4030
Hours: M-F 8:00 AM - 6:00 PM PDT

EMEA
Tel: +33-4-3987-0555
Hours: M-F 9:00 AM - 6:00 PM CEST

APAC
Tel: +603-2711-7391
Hours: M-F 9:00 AM-6:00 PM MYT

Japan
Tel: +81 (0)3-8434-8535
Hours: M-F 9:00 AM-6:00 PM JST

China
Tel: +86 10-62976084
Hours: M-F 9:00 AM-6:00 PM CCT

Site Index | Legal | Privacy | Worldwide Offices | Copyright ©2009 Fortinet, Inc. All Rights Reserved.

Once the account has been created, the *Product Registration Form* will be displayed and the product details can be provided. Alternately, the product registration can be completed at a later time.

Registering a device

Complete the following steps when registering a device for support purposes:

- 1 Log in using the *Username* and *Password* defined when the account was created

Figure 74: Support account login screen.

- 2 Select *Add Registration* on the left-hand side.
- 3 Select *New Fortinet Product/License Registration*.

Figure 75: Adding a product to a support account

- 4 Select the appropriate *Product Model*.
- 5 Enter the *Serial Number*.

- 6 Enter the *Support Contract No.* provided by Fortinet when the support contract was purchased.
- 7 In the *Product Description* field, explain where this unit is physically located.
- 8 Click *Next* and accept the *End User License Agreement* (EULA) to complete the registration.

Reporting problems

Problems can be reported to a Fortinet Technical Assistance Center in the following ways:

- By logging an online ticket
- By phoning a technical support center

Logging online tickets

Problem reporting methods differ depending on the type of customer.

Fortinet partners

Fortinet Partners are entitled to priority web-based technical support. This service is designed for partners who provide initial support to their customers and who need to open a support ticket with Fortinet on their behalf. We strongly encourage submission and follow up of support tickets using this service.

The support ticket can be submitted after logging into the partner website using one of the following links using FortiPartner account details:

<http://partners.fortinet.com>

This link will redirect to the general *Partner Extranet* website. Click *Support > Online Support Ticket*.

<https://www.forticare.com:1443/customersupport/login/partnerlogin.aspx>

This link redirects to the *Partner Online Support Ticket* section also known as *FortiCare*.





The Partner Online Support Ticket section is accessed through HTTPS on port 1443. Ensure that the firewall allows external access this port. Also, a customer's Fortinet device must have a valid support contract to be able submit the support request as a Partner.

Fortinet customers

Fortinet customers should complete the following steps to report a technical problem online:

- 1 Log in to the support web site at the following address with the account credentials used when the account was created:
<https://support.fortinet.com>
- 2 Click *View Products*.
- 3 In the *Products List*, select the product that is causing the problem.
- 4 Complete the *Create Support Ticket* fields.

Figure 76: Account products support details


Support


- Home
- View Products
- Add Registration
- View Support Tickets
- Download FortiGuard Services Updates
- Firmware Images
- Beta Program
- Product Registration FAQ
- Product Life Cycle Info
- Technical Forum
- Fortinet Knowledge Base
- My Profile
- CS Reference Guide
- Registration Help
- Logout

Product Support Details

Product Info

Product Model	FortiGate 51B
Serial Number	FG50BXXXXXX
Registration Date	4/6/2009
Ship Date	4/2/2009
Warranty	Fortinet Internal Order
Description	CDN KIT - ITF00000064
Fortinet Partner	Unknown

OS Version	4.00 092
AV Engine Version	3.011
AV DB Version	10.535
NIDS Version	2.657
Last Update	6/26/2009 2:35 PM

Current Support Coverages

Note: Contract starts in the future may not include in this list.

Support Type	Hours	Activation Date	Expiration Date
Hardware Coverage	Advanced HW	4/6/2009	4/6/2010
Firmware Updates	-	4/6/2009	4/6/2010
Enhanced Support	24x7	4/6/2009	4/6/2010
Telephone Support	24x7	4/6/2009	4/6/2010
Virus Definitions Updates	-	4/6/2009	4/6/2010
Attack Definitions Updates	-	4/6/2009	4/6/2010
FortiGuard Web Filtering	-	4/6/2009	4/6/2010
FortiGuard AntiSpam	-	4/6/2009	4/6/2010

Registered Support Contract(s) Info

No registered contract

Product/Contract Maintenance

Add/Renew Contract

RMA Replacement

RMA Replacement Notes:
 1. Enter the replacement unit serial number in the above field "New Serial No." and click "RMA Replace" to complete RMA replacement. Previous support contract will be transferred from the defective unit to the replacement unit.
 2. In case the defective products were purchased/registered with a VDOM license key, a new key will automatically be generated upon RMA replacement completion. The VDOM key for the defective unit will be disabled.

Register VDOM license

Register VDOM License

Support Ticket List

No ticket was created

Create Support Ticket

Title *

Product Type *

Category *

S/W Version

Build

Case Priority

In order for Fortinet Technical Support to provide you with the optimum level of service, we recommend that the ticket be initially opened with:

1. A clear problem description
2. The problem history (Has this configuration worked in the past? Is it a new configuration for the device? Were any changes made on the device or on the network recently?)
3. A network diagram with IP schema
4. The configuration file
5. The debug log of the unit
6. A description and the results of the troubleshooting steps already performed

Attachments

Note: The maximum size of all attachment(s) in total is 4M bytes. But you can upload file that exceeds 4M bytes to temporary storage after you submit the ticket.
 * - indicates Required Fields

SUBSCRIBE TO FORTINET'S CORPORATE e-LETTER [CLICK HERE TO SUBSCRIBE](#)

[SITE MAP](#) | [LEGAL NOTICES](#)

©2009 FORTINET INC. ALL RIGHTS RESERVED

Following up on online tickets

Perform the following steps to follow up on an existing issue.

Partners should log into the following web site:

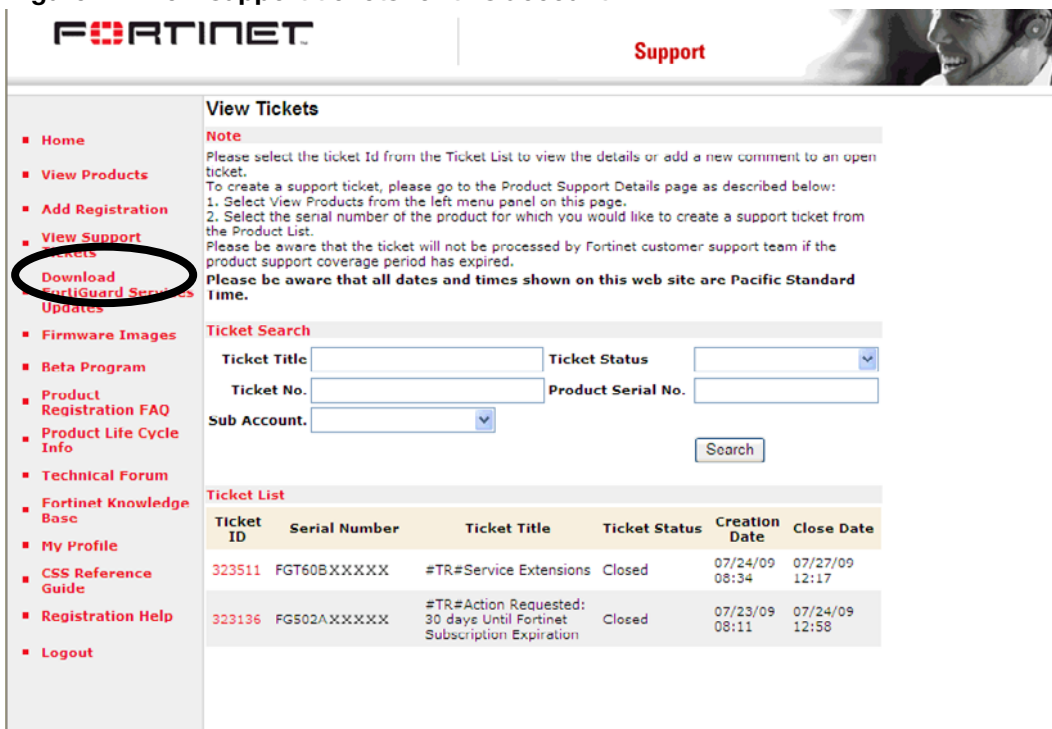
<http://partners.fortinet.com>

Customers should log into the following site:

<http://support.fortinet.com>.

- 1 Log in with the account credentials used when the account was created.
- 2 Click *View Support Tickets* on the left-hand side. Use the *Search* fields on the *View Tickets* form to locate the tickets assigned to the account.

Figure 77: View support tickets for this account



FORTINET Support

View Tickets

Note
Please select the ticket Id from the Ticket List to view the details or add a new comment to an open ticket.
To create a support ticket, please go to the Product Support Details page as described below:
1. Select View Products from the left menu panel on this page.
2. Select the serial number of the product for which you would like to create a support ticket from the Product List.
Please be aware that the ticket will not be processed by Fortinet customer support team if the product support coverage period has expired.
Please be aware that all dates and times shown on this web site are Pacific Standard Time.

Ticket Search

Ticket Title: Ticket Status:

Ticket No. Product Serial No.

Sub Account:

Ticket List

Ticket ID	Serial Number	Ticket Title	Ticket Status	Creation Date	Close Date
323511	FGT60BXXXXX	#TR#Service Extensions	Closed	07/24/09 08:34	07/27/09 12:17
323136	FGS02AXXXXX	#TR#Action Requested: 30 days Until Fortinet Subscription Expiration	Closed	07/23/09 08:11	07/24/09 12:58

- 3 Select the appropriate ticket number. Closed tickets cannot be updated. A new ticket must be submitted if it concerns the same problem.
- 4 Add a *New Comment* or *Attachment*.
- 5 Click *Submit* when complete.



Every web ticket update triggers a notification to the ticket owner, or ticket queue supervisor.

Telephoning a technical support center

The Fortinet Technical Assistance Centers can also be contacted by phone.

Americas

- Telephone: 1-866-648-4638
- Hours: Monday to Friday 6:00 AM to 6:00 PM (Pacific Daylight Time)

EMEA

- Telephone: +33-4-898-0555

If a support call is placed outside of EMEA business hours (Monday to Friday 9:00 AM to 6:00 PM (Central European Daylight Time)), priority 1 issues will be transferred to another Fortinet Technical Solutions center according to the *Follow the Sun* policy, meaning wherever it's daylight, that TAC will be taking all support calls.

- Hours: Monday to Friday 9:00 AM to 6:00 PM (Central European Daylight Time)

APAC

- Telephone: +603-2711-7391
- Hours: Monday to Friday 9:00 AM to 6:00 PM (Malaysia Time)

Assisting technical support

The more information that can be provided to Fortinet technical support, the better they can assist in resolving the issue. Every new support request should contain the following information:

- A valid contact name, phone number, and email address.
- A clear and accurate problem description.
- A detailed network diagram with complete IP address schema.
- The configuration file, software version, and build number of the Fortinet device.
- Additional log files such as *Antivirus* log, *Attack* log, *Event* log, *Debug* log or similar information to include in the ticket as an attachment. If a third-party product is involved, for example, email server, FTP server, router, or switch, please provide the information on its software revision version, configuration, and brand name.

The following *Knowledge Base* article provides an example of what type of technical information and network diagram details should be submit to receive the quickest resolution time to a problem:

<http://kb.forticare.com/default.asp?id=1068&SID=&Lang=1>

Support priority levels

Fortinet technical support assigns the following priority levels to support cases:

Priority 1

This **Critical** priority is assigned to support cases in which:

- The network or system is down causing customers to experience a total loss of service.
- There are continuous or frequent instabilities affecting traffic-handling capability on a significant portion of the network.
- There is a loss of connectivity or isolation to a significant portion of the network.
- This issue has created a hazard or an emergency.

Priority 2

This **Major** priority is assigned to support cases in which:

- The network or system event is causing intermittent impact to end customers.
- There is a loss of redundancy.
- There is a loss of routine administrative or diagnostic capability.
- There is an inability to deploy a key feature or function.
- There is a partial loss of service due to a failed hardware component.

Priority 3

This **Medium** priority is assigned to support cases in which:

- The network event is causing only limited impact to end customers.
- Issues seen in a test or pre-production environment exist that would normally cause adverse impact to a production network.
- The customer is making time sensitive information requests.
- There is a successful workaround in place for a higher priority issue.

Priority 4

This **Minor** priority is assigned to support cases in which:

- The customer is making information requests and asking standard questions about the configuration or functionality of equipment.

Customers must report Priority 1 and 2 issues by phone directly to the Fortinet EMEA Support Center.

For lower priority issues, you may submit an assistance request (ticket) via the web system.

The web ticket system also provides a global overview of all ongoing support requests.

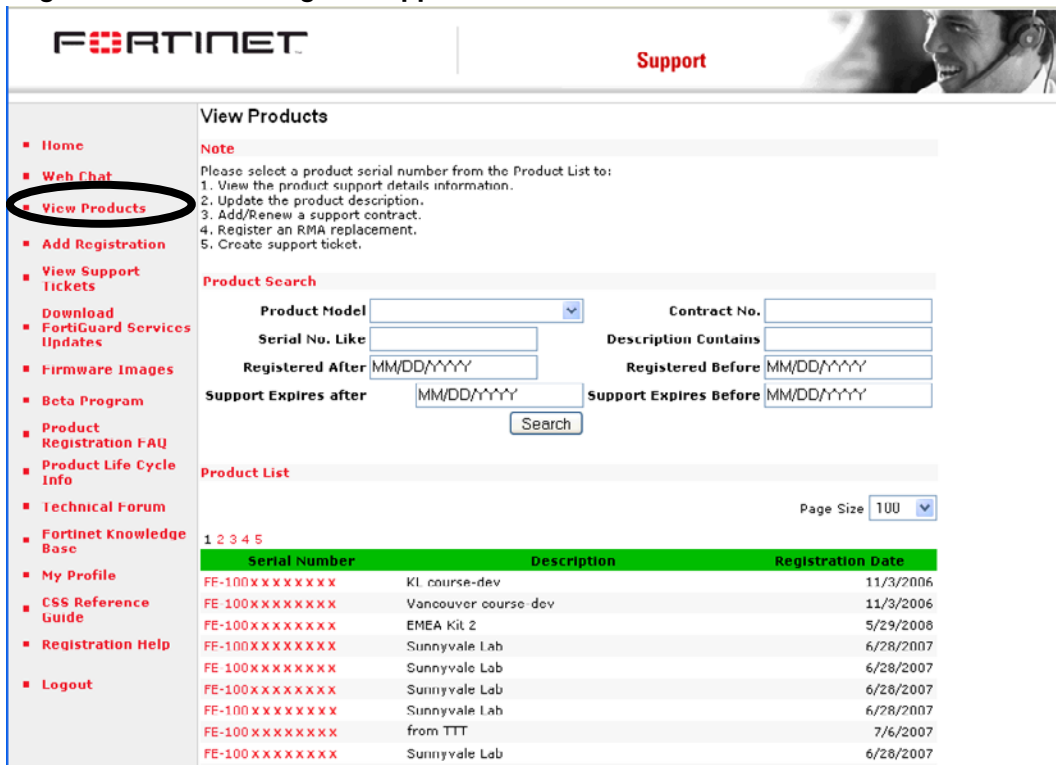
Return material authorization process

In some cases hardware issues are experienced and a replacement unit must be sent. This is referred to as a Return Material Authorization (RMA). In these cases or RMAs, the support contract must be moved to the new device. Customers can move the support contract from the failing production unit to the new device through the support web site.

To move the support contract to a new device

- 1 Log in to the support web site with the credentials indicated when the account was created.
- 2 From *View Products*, locate the serial number of the defective unit from the list of devices displayed for the account

Figure 78: From Moving the support to a new unit.



FORTINET Support

- Home
- Web Chat
- View Products**
- Add Registration
- View Support Tickets
- Download FortiGuard Services Updates
- Firmware Images
- Beta Program
- Product Registration FAQ
- Product Life Cycle Info
- Technical Forum
- Fortinet Knowledge Base
- My Profile
- CSG Reference Guide
- Registration Help
- Logout

View Products

Note
Please select a product serial number from the Product List to:
1. View the product support details information.
2. Update the product description.
3. Add/Renew a support contract.
4. Register an RMA replacement.
5. Create support ticket.

Product Search

Product Model Contract No.
 Serial No. Like Description Contains
 Registered After Registered Before
 Support Expires after Support Expires Before


Product List


Page Size

Serial Number	Description	Registration Date
FF-100XXXXXXX	KL course-dev	11/3/2006
FE-100XXXXXXX	Vancouver course dev	11/3/2006
FE-100XXXXXXX	EMEA Kit 2	5/29/2006
FF-100XXXXXXX	Sunnyvale Lab	6/28/2007
FE-100XXXXXXX	Sunnyvale Lab	6/28/2007
FE-100XXXXXXX	Sunnyvale Lab	6/28/2007
FE-100XXXXXXX	Sunnyvale Lab	6/28/2007
FE-100XXXXXXX	from TTT	7/6/2007
FE-100XXXXXXX	Sunnyvale Lab	6/28/2007

The *Product Support Details* for the selected device will be displayed.

Figure 79: Support details for a registered product





- Home
- Web Chat
- View Products
- Add Registration
- View Support Tickets
- Download FortiGuard Services Updates
- Firmware Images
- Beta Program
- Product Registration FAQ
- Product Life Cycle Info
- Technical Forum
- Fortinet Knowledge Base
- My Profile
- CSS Reference Guide
- Registration Help
- Logout

Product Support Details

Product Info

Product Model	FortiMail 100
Serial Number	FE-100XXXXXX
Registration Date	11/3/2006
Ship Date	9/21/2006
Warranty	Fortinet Internal Order
Description	Vancouver course-dev
Fortinet Partner	Unknown

Current Support Coverages

Note: Contract starts in the future may not include in this list.

Support Type	Hours	Activation Date	Expiration Date
Hardware Coverage	Advanced HW	11/3/2006	11/3/2007
Firmware Updates	-	11/3/2006	11/3/2007
Enhanced Support	24x7	11/3/2006	11/3/2007
Telephone Support	24x7	11/3/2006	11/3/2007
Virus Definitions Updates	-	11/19/2008	11/19/2009
Attack Definitions Updates	-	11/3/2006	11/3/2007
FortiGuard Web Filtering	-	11/3/2006	11/3/2007
FortiGuard AntiSpam	-	11/19/2008	11/19/2009

Registered Support Contract(s) Info

Contract Number	Part Number	Description	Registration Date
XXXXXXXXXX	FCX-15-00000-103-02-12	Virus Definitions Updates 11/19/2008 11/19/2009 FortiGuard AntiSpam 11/19/2008-11/19/2009	11/19/2008

Product/Contract Maintenance

Add/Renew Contract

RMA Replacement

New Serial No.

RMA Replacement Notes:
 1. Enter the replacement unit serial number in the above field "New Serial No." and click "RMA Replace" to complete RMA replacement. Previous support contract will be transferred from the defective unit to the replacement unit.
 2. In case the defective products were purchased/registered with a VDOM license key, a new key will automatically be generated upon RMA replacement completion. The VDOM key for the defective unit will be disabled.

Register VDOM license

Register VDOM License

VDOM License Number.

VDOM License Confirmation Number.

Support Ticket List

Ticket ID	Ticket Title	Ticket Status	Creation Date	Close Date
Z58540	PLEASE update subscription for our Fortimail	Closed	11/19/08 13:54	11/20/08 15:04

Create Support Ticket

Title *

Product Type *

S/W Version

Build

Case Priority

In order for Fortinet Technical Support to provide you with the optimum level of service, we recommend that the ticket be initially opened with:

1. A clear problem description
2. The problem history (Has this configuration worked in the past? Is it a new configuration for the device? Were any changes made on the device or on the network recently?)
3. A network diagram with IP schema
4. The configuration file
5. The debug log of the unit
6. A description and the results of the troubleshooting steps already performed

- 3 In the *RMA Replacement* section of the *Product Support Details* page, enter the serial number of the replacement device and click *RMA Replace*.

Figure 80: RMA Replacement section of Product Support Details

The screenshot shows a web interface with two main sections. The first section, titled "Product/Contract Maintenance" in red, contains a sub-section "Add/Renew Contract" with a text input field for "Contract Number" and an "Add/Renew" button. The second section, titled "RMA Replacement", contains a text input field for "New Serial No." and an "RMA Replace" button. Below these fields, there are red "RMA Replacement Notes" with two numbered instructions: 1. Enter the replacement unit serial number in the above field "New Serial No." and click "RMA Replace" to complete RMA replacement. Previous support contract will be transferred from the defective unit to the replacement unit. 2. In case the defective products were purchased/registered with a VDOM license key, a new key will automatically be generated upon RMA replacement completion. The VDOM key for the defective unit will be disabled.

This will transfer the support contract from the defective unit to the new unit with the serial number provided.



Troubleshooting common issues

Connectivity problems can be caused by a number of reasons including hardware problems such as cabling or interfaces, routing problems, firewall configuration problems, and more. Connectivity problems can vary from total lack of connection, to some applications not working or some users having problems, to streaming video or VoIP connections being too slow to be useful.

These general troubleshooting tips provide a starting point for you to determine why your network is behaving unexpectedly. This section includes general troubleshooting methods. More in depth coverage of these topics, such as transparent mode and firewall sessions, can be found in the [FortiGate Fundamentals handbook chapter](#).

If you are experiencing problems when starting your FortiGate unit, see [“FortiGate unit bootup issues” on page 871](#).

This section contains the following topics:

- [How to troubleshoot cabling](#)
- [How to troubleshoot no Internet connection](#)
- [How to troubleshoot a connection in Transparent mode](#)
- [Common issues and questions](#)

How to troubleshoot cabling

An integral part of setting up networking hardware is connecting it to the network via network cables. When you have connection problems, it may be a cabling issue.

1 Ensure devices are powered on and booted up.

If a device is powered off, it will appear as if the cabling has a problem. To this end, ensure that the device at each end of the cabling connection has powered on, and has completed booting up. It's possible during boot up that a diagnostic on the network interface may appear as an outage.

2 Look for the connection LEDs to light up.

When you plug a network cable into a FortiGate unit and the other end is connected to a functioning network device, the FortiGate unit has LEDs that light up to indicate if there is a good connection on that interface and what the configured speed is for that interface. If there are no lights, there are cable problems of some kind.

3 Ensure the cable is completely inserted into the jack.

While most network cables have the tab that clicks when the cable is fully inserted into the jack, some cables do not. These cables make it more difficult to determine if the cable is properly inserted into the jack. The best way to verify this is to watch the LED connection lights while you wiggle the cable. If they blink on and off, the cable is not properly inserted. If this is the case, you may have to try a different cable.

4 Ensure you are using the proper cable type.

There are different types of ethernet network cable with the main difference being the wiring — is it a crossover or patch cable? However there are even different accepted cross-over wiring orders, so its possible that with two cross over cables, one may not work. Also if the cable is hand-made instead of purchased, there are multiple places for problems — cable rating (Cat5, Cat5e, Cat6), bad connections in the ends, cable too long, or even breaks in the cable wiring.

5 Try using a different cable.

When all else fails, try a different cable. It is possible there are breaks in the wires, or the cable was wired differently than you expect. Changing to a new cable will quickly let you know if the cable was bad.

6 Check the interfaces are connected.

Use the CLI command `get hardware nic <interface_name>` to check the interface with the connectivity problems. For example checking wan1 would be `get hardware nic wan1`. In the output of this command look for a line that states `Link: up`. If the Link is down, it is a physical connection problem. If the Link is up, it is an internal problem.

7 Try connecting to a different network device.

It is possible there is a problem with the configuration of the device you are trying to connect to. If it has a dead interface for example, even if everything else is working you will not see the LED connection lights on the FortiGate unit. If you have tried everything else, try connecting to a different network device.

How to troubleshoot no Internet connection

The most common use of FortiGate units is to connect internal networks to the Internet. Sometimes when you think you should be able to connect to the Internet, you can't.

If you have not already done so, you may want to try [“How to troubleshoot cabling” on page 763](#) before trying these other things.

1 Run ping to an external domain name.

On your FortiGate unit, enter the CLI command `execute ping www.fortinet.com`. This will determine two things: if you have a connection to the Fortinet web site or not, and if your DNS service is working. If DNS is not working, you will see a message telling you it was unable to resolve the domain name (FQDN). If you see this error, ensure you have valid DNS entries (System > Network > DNS), and ensure the security policies allow DNS packets through the firewall. If you do not have a connection, continue to the following steps.

2 Check the modem status.

Almost all internal networks connect to the Internet through a modem. The modem converts the network traffic from your ISPs format to ethernet traffic that your network can understand. The modem can be a cable modem, DSL, ADSL, a 3G cellular modem, or even a serial port modem. In all cases, you need to verify that the modem is working properly, and that you have its configuration information properly recorded. You should always run the `diagnose sys modem detect` CLI command after connecting the modem to your FortiGate unit, and you can view the modem configuration by using the `get system modem` command. It is possible your ISP is having connection problems. If you have modem problems, contact your ISP or modem documentation for help.

3 Check the default route.

For network traffic to find its way across the network, it must have a route. Without a route, the traffic doesn't know where to go. Ensure there is a valid static route on your FortiGate unit that will direct outbound traffic to the Internet. Ensure the proper interface is used, and that the gateway used does exist. Also ensure the default route admin distance is lower than other routes you may have configured.

4 Check DHCP settings.

If you are using DHCP for any computers trying to connect to the Internet, the settings may have problems. If your computers are having IP address or DNS problems, it's possibly DHCP server related. Things to check include the range of addresses, the DHCP server is assigned to the correct interface, and any overrides are assigned to the proper IP and MAC address. It is also possible there is an interface waiting for a DHCP server that isn't there.

5 Check the interface settings.

Typically the FortiGate unit interface connected to the Internet is Wan1 or Wan2. Ensure that interface has an IP address and subnet mask, that the IP address exists on the subnet the interface is connected to, and check that both the admin and link status are up. If everything else is working but the admin status is down, the interface will not work. If your FortiGate unit interfaces are labelled port1, port2, and so on consider setting alias names for the interfaces based on their function. This will save you time when troubleshooting.

6 Check the firewall settings for the interfaces.

For traffic to flow between two FortiGate unit interfaces, there must be a security policy to allow that traffic. Check there is a policy to allow traffic between your two interfaces. Also check that the policy is not restricting traffic to a specific time of day, specific users, or such. If those restrictions exist, and you remove them for testing remember to put them back in place.

7 Check the firewall settings for basic networking protocols.

It is common practice to have a DENY security policy at the top of the list that will block unwanted protocols. If this type of policy exists on your FortiGate unit, consider disabling it to check if it is preventing access to the Internet. If it is the problem, go through the list of blocked protocols and remove any common internet protocols needed for common traffic such as FTP, HTTP, HTTPS, PING, DNS, and so on.

8 Check the logs.

If you cannot find the source of a problem, before you contact support you should always check your logs. The easiest log to start with is the event log. This log records major events that happen on the FortiGate unit such as configuration changes, failed logins, FortiGuard and firmware updates, and critical problems. These log messages may help you identify the problem, and give a place to start trying to fix it. Other logs such as UTM, traffic, and vulnerability logs may prove useful but they are more specific and complex. If you do not have logging configured, go to *Log&Report > Log Config > Log Setting*. Select *Memory* and a minimum log level of *Information*. Select *Event Logging*, and *Enable All*. This will start logging events on your FortiGate unit and give you a valuable troubleshooting tool.

9 Perform a packet capture on the traffic.

If everything else looks properly configured, try capturing packets from the traffic to get a better idea what is happening at the low level. To enable packet capturing on the web-based interface go to *System > Config > Advanced*. Select *Create New* to create a packet capture filter by selecting the interface, protocol, and source and destination address and port. When you run the capture filter it will capture packets that match, and save them to a *.pcap file on your local computer, named for the interface. To read this file, you must use a third-party application such as Wireshark. This will show you the packet source, destination, protocol, and other information that may help you solve your connection issues.

10 Contact support.

If you have gone through all these steps, and you have not determined why you have no Internet connection it is time to contact support. Remember to tell support everything you have tried up to this point to ensure they do not duplicate work you have already done.

How to troubleshoot intermittent connection problems

You have an Internet connection that works, but from time to time you are unable to access the Internet for a period of time before it starts working again. This type of problem can be annoying because it is unpredictable. Here are some things to try.

1 Check the cabling.

Often this problem is a bad cable with a bad connector or a loose wire. If you are in doubt, try a new cable. See [“How to troubleshoot cabling” on page 763](#).

2 Check the information on the Dashboard.

The dashboard (*System > Dashboard > Status*) has a number of useful widgets that display useful system information. Most useful are the System Resources Widget which shows the CPU and memory usage, the Top Sessions widget with not only lists the addresses with the most open connections but also lists the total sessions, and the Log and Archive Statistics widget that lists many details about DLP archive, and Log.

3 Check your DNS information.

If DNS can not resolve domain names, you will not be able to connect to Internet resources. The two places to check DNS entries are the two main entries for the FortiGate unit (*System > Network > DNS*) and any interface or wireless SSID where you have the choice to use the system DNS or enter a separate value. You can also check your ISP DNS information to confirm it is correct and current.

4 Check the session table and new sessions.

For a new connection to the Internet you must be able to create a new session. Intermittent connection problems may mean the session table is full, or the FortiGate unit is unable to create more sessions for some reason. A common reason is if the session table is full, and the FortiGate unit goes into conserve mode where no new sessions are created. To determine this, you can check the log messages for “conserve” and “failover” messages, or you can go to *Policy > Monitor > Session Table* to see the session setup rate and the number of current sessions. If the setup rate is very low or there are only a handful of sessions, it is not likely a failover problem.

How to troubleshoot a connection in Transparent mode

Transparent mode behaves differently from NAT/Route mode. Unless specifically stated, all FortiOS documentation refers to NAT/Route mode by default.

In Transparent mode, all interfaces are bridged to each other—traffic is sent out over all interfaces at once. The FortiGate unit essentially becomes a layer-2 switch, where in NAT/Route mode it is a layer-3 router. This difference extends to troubleshooting.

Before troubleshooting Transparent mode, you should try [“How to troubleshoot cabling” on page 763](#) and [“How to troubleshoot no Internet connection” on page 764](#).

1 Check the interfaces are connected.

Use the CLI command `get hardware nic <interface_name>` to check the interface with the connectivity problems. For example, checking wan1 would be `get hardware nic wan1`. In the output of this command look for a line that states `Link:` up. If the Link is down, it is a physical connection problem. If the Link is up, it is an internal problem.

2 Check the bridge table.

The bridge table is important when in Transparent mode. If a MAC is not in the bridge table, then packets to that MAC will be blocked. This may appear as traffic going to a MAC with no reply traffic coming back. In this situation, check the bridge table to ensure the MACs involved are in the table. Use the CLI command `diag netlink brctl name host root.b` to check the bridge table associated with the root VDOM. The root VDOM exists no matter if VDOMs are enabled or not. This command lists MAC addresses in the table. If your device's address is not listed, that means the FortiGate unit cannot find the device when it floods its interfaces. Ensure your device cabling is properly connected to the active network.

3 Check the session table.

The session table has an entry for each active session on the FortiGate unit. If your device's MAC address is in the bridging table, the next place to look is the session table. Go to *Policy > Monitor > Session Monitor*. If there are many sessions, filter them based on the source IP address of your device. If there are no sessions listed for your device, security policies may be blocking it. You can also run the CLI command `get test app proxyworker 4` to see the active sessions and related errors for the various proxies. For more information, see [“Check number of sessions used by UTM proxy” on page 783](#).

4 Three things to check if you are running HA.

Ensure Bridge Protocol Data Units (BPDUs) are forwarded between HA units. BPDUs are exchanged across bridges to detect loops in a network topology, but HA units sharing the same MAC address can cause problems with BPDUs. This is not an issue if HA units are directly connected.

Configure multiple redundant interfaces to the switch when operating in active-passive HA mode. In an HA cluster, virtual MAC addresses are used so if there is a failover, the new primary unit will have the same virtual MAC and IP address as the primary unit that failed. The problem is that switches usually see this as the same MAC on multiple interfaces on the same subnet, which is usually associated with a hacker masquerading on the network.

Form the HA cluster carefully. As the cluster is being configured, the network interfaces will momentarily lose connectivity as the FGCP assigns virtual MACs to the interfaces. If this is the case, try deleting the arp table on your management computer using the command “arp -d”. If the priorities of the subordinate units are not lower than the primary unit, the cluster may take longer to determine the primary unit and lost connectivity for a longer period as a result. Also if you power on a secondary unit first, it will become the primary, and then when you power on the intended primary there will be a switchover. This should be avoided.

5 Capture packets for in-depth information.

If you have checked the bridge and session tables without finding the problem, the next step is to capture some packets. Captured packets to and from an IP address can tell you exactly what is going on at a lower level than most other troubleshooting methods. To enable packet capturing on the FortiGate unit's web-based interface, go to *System > Config > Advanced*. Select *Create New* to create a packet capture filter by selecting the interface, protocol, and source and destination address and port. When you run the capture filter it will capture packets that match the filter, and save them to a *.pcap file on your local computer, named for the interface. To read this file, you must use a third-party application such as Wireshark. This will show you the packet source, destination, protocol, and other information that may help you solve your connection issues.

Common issues and questions

The general troubleshooting tips include, and can help answer the following questions.

1 “Check hardware connections” on page 770

Are all the cables and interfaces connected properly?

Is the LED for the interface green?

2 “Check FortiOS network settings” on page 770

If you are having problems connecting to the management interface, is your protocol enabled on the interface?

Is there an IP address on the interface?

3 “Check CPU and memory resources” on page 772

Is your CPU running at almost 100 percent usage?

Are you running low on memory?

- 4 [“Check modem status” on page 773](#)
Is the modem connected?
Are there PPP issues?
- 5
- 6 [“Run ping and traceroute” on page 773](#)
Are you experiencing complete packet loss?
- 7 [“Check the logs” on page 777](#)
- 8 [“Verify the contents of the routing table \(in NAT mode\)” on page 778](#)
Are there routes in the routing table for default and static routes?
Do all connected subnets have a route in the routing table?
Does a route wrongly have a higher priority than it should?
- 9 [“Check the bridging information in Transparent mode” on page 778](#)
Are you having problems in transparent mode?
- 10 [“Perform a sniffer trace” on page 780](#)
Is traffic entering the FortiGate unit and does it arrive on the expected interface?
Is the ARP resolution correct for the next-hop destination?
Is the traffic exiting the FortiGate unit to the destination as expected?
Is the traffic being sent back to the originator?
- 11 [“Debug the packet flow” on page 782](#)
Is the traffic entering or leaving the FortiGate unit as expected?
- 12 [“Check number of sessions used by UTM proxy” on page 783](#)
Have you reached the maximum number of sessions for a protocol?
Are new sessions not starting for a certain protocol?
- 13 [“Examine the firewall session list” on page 787](#)
Are there active firewall sessions?
- 14 [“Checking wireless information” on page 788](#)

For troubleshooting tips for specific non-connectivity areas, see [“Troubleshooting advanced” on page 789](#).

In addition to these steps, you may find other diagnose commands useful. See [“Other diagnose commands” on page 788](#).

Check hardware connections

If there is no traffic flowing from the FortiGate unit, it may be a hardware problem.

To check hardware connections

- ensure the network cables are properly plugged into the interfaces
- ensure there are connection lights for the network cables on the unit
- change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality—such as straight through or crossover, or possibly exposed wires at the connector.
- connect the FortiGate unit to different hardware
- ensure the link status is set to *Up* for the interface, see *Status > Network > Interface*. The link status is based on the physical connection and cannot be set in FortiOS

If any of these solve the problem, it was a hardware connection problem. You should still perform some basic software connectivity tests to ensure complete connectivity. It might also be that the interface is disabled, or have its Administrative Status set Down.

To enable an interface - web-based manager

- 1 Using the web-based management interface, go to *System > Network > Interface*.
- 2 Select and edit the interface to enable, such as *port1*.
- 3 Find *Administrative Status* at the bottom of the screen, and select *Up*.
- 4 Select *Apply*.

To enable an interface - CLI

```
config system interface
  edit port1
    set status enable
  next
end
```

Check FortiOS network settings

FortiOS network settings are present in both the web-based manager interface and the CLI. The following information includes troubleshooting and best practice information. The network settings include:

- [Interface settings](#)
- [DNS settings](#)
- [DHCP Server settings](#)

Interface settings

If you can access the FortiGate unit with the management cable only, the first step is to display the interface settings. To display the settings for the internal interface, use the following CLI command:

```
FGT# show system interface internal
```

or for a complete listing of all the possible interface settings, use the following CLI command:

```
config system interface
  edit internal
  get
```


end

Check the interface settings to ensure they are not preventing traffic. Specific things to check include (web-based manager names are shown, CLI names may vary slightly):

- **Link Status** — Down until a valid cable is plugged into this interface, after which it will be *Up*. The Link Status is shown physically by the connection LED for the interface — if it lights up green, it is a good connection. If Link Status is down, the interface does not work. Link Status is also displayed on the *System > Network > Interface* screen by default.
- **Addressing mode** — do not use *DHCP* if you don't have a DHCP server — you will not be able to logon to an interface in DHCP mode as it will not have an IP address.
- **IP/Netmask** — an interface needs an IP address to be able to connect to other devices. Ensure there is a valid IP address in this field. The one exception is if *DHCP* is enabled for this interface to get its IP address from an external DHCP server.
- **IPv6 address** — unless specifically stated all IP addresses are IPv4. The same protocol must be used by both ends to complete the connection. Ensure both this interface and the remote connection are both using IPv4 or both using IPv6 addressing.
- **Administrative access** — If no protocols are selected, you will have to use the local management cable to connect to the unit. If you are using IPv6, configure the IPv6 administrative access protocols.
- **Administrative status** — set to *Up* or interface will not work.

DNS settings

While this section is not complicated, many networking problems can be traced back to DNS problems. Things to check in this area include:

- Are there values for both primary and secondary entries?
- Is the local domain name correct?
- Are you using IPv6 addressing? If so, are the IPv6 DNS settings correct?
- Are you using Dynamic DNS (DDNS)? If so, it is using the correct server, credentials, and interface?
- Can you contact both DNS servers to verify the servers are operational?
- If an interface addressing mode is set to DHCP and is set to override the internal DNS, is that interface receiving a valid DNS entry from the DHCP server — is it a reasonable address and can it be contacted to verify its operational?
- Are there any DENY security policies that need to allow DNS?
- Can any internal device perform a successful traceroute to a location using the FQDN? See [“Traceroute” on page 775](#).

DHCP Server settings

DHCP Servers are common on internal and wireless networks. If the a DHCP server is not configured properly it can cause problems. Things to check in this area include:

- Is the DHCP server entry set to *Relay*? If so, verify there is another DHCP server to relay the requests to. Otherwise, it should be set to *Server*.
- Is the DHCP server enabled?

- Is the range of IP addresses this DHCP server uses valid? Are those addresses in use by other devices? If one or more devices are using IP addresses in this range, you can use the IP reservation feature to ensure the DHCP server does not use these addresses.
- Is there a gateway entry? Include a gateway entry to ensure clients of this server have a default route.
- Is the system DNS setting being used? The best practice is to whenever possible use the system DNS to avoid confusion. However, the option to specify up to three custom DNS servers is available, and all three entries should be used for redundancy.



There are some situations, such as a new wireless interface, or during the initial FortiGate unit configuration, where interfaces override the system DNS entries. When this happens, it often shows up as intermittent Internet connectivity. To fix the problem, go to the problem interface and ensure its set to use the system DNS entries.

Check CPU and memory resources

System resources are shared but a number of processes running at the same time on the FortiGate unit. If one of these processes consumes nearly all the resources.

A quick way to monitor CPU and memory usage is on the System Dashboard using the System Resources widgets. They have both a visual gauge and a percent number displayed to show you the usage.

To check the system resources on your FortiGate unit, run the CLI command

```
FGT# get system performance status
```

This command provides a quick and easy snapshot of the FortiGate.

The first line of output shows the CPU usage by category. A FortiGate that is doing nothing will look like:

```
CPU states: 0% user 0% system 0% nice 100% idle
```

However, if your network is running slow you might see something like:

```
CPU states: 1% user 98% system 0% nice 1% idle
```

This line shows that all the CPU is used up by system processes. Normally this should not happen as it shows the FortiGate is overloaded for some reason. If you see this overloading, you should investigate farther as its possible a process such as scanunitid is using all the resources to scan traffic in which case you need to reduce the amount of traffic being scanned by blocking unwanted protocols, configuring more security policies to limit scanning to certain protocols, or similar actions. It is also possible that a hacker has gained access to your network and is overloading it with malicious activity such as running a spam server or using zombie PCs to attack other networks on the Internet. You can get additional CPU related information with the CLI command `get system performance top`. This command shows you all the top processes running on the FortiGate unit (names on the left) and their CPU usage. If a process is using most of the CPU cycles, investigate it to determine if its normal activity.

The second line of output from `get system performance status` shows the memory usage. Memory usage should not exceed 90 percent. If memory is too full some processes will not be able to function properly. For example if the system is running low on memory, antivirus scanning will go into failopen mode where it will start dropping connections or bypass the antivirus system.

The other lines of output such as network usage, average session setup rate, and viruses caught, IPS attacks blocked can also help you determine why system resource usage is high. For example, if network usage is high it will result in high traffic processing on the FortiGate, or if the session setup rate is very low or zero the proxy may be overloaded and not able to do its job.

Check modem status

Sometimes the modem may not work properly, or the unit may not be detecting the modem. Use the following diagnostic commands to help you troubleshoot issues with the modem.

```
diagnose sys modem {cmd | com | detect | history | wireless-id}
```

You should always run the following diagnose command after inserting the USB modem into the unit:

```
diagnose sys modem detect
```

You can view the modem configuration by using the `get system modem` command. You can also view the modem's vendor identification as well as the custom product identification number from the information output from the `get system modem` command.

When the modem is not being detected by the unit, use the following command:

```
diagnose sys modem wireless-id
```

When there are connectivity issues, use the following to help you resolve them:

- `diag debug enable` – activates the debug on the console
- `diag debug application modemd` – dumps communication between the modem and the unit
- `diag debug application pppd` – dumps the PPP negotiating messages.
- `execute modem dial` – displays modem debug output

The modem diagnose output should not contain any error on the way to initializing. You should also verify the number that is used to dial with your ISP.

Run ping and traceroute

Ping and traceroute are useful tools in network troubleshooting. Alone, either one can determine network connectivity between two points. However, ping can be used to generate simple network traffic to view with diagnose commands on the FortiGate unit. This combination can be a very powerful one in locating network problems.

In addition to their normal uses, ping and traceroute can tell you if your computer or network device has access to a domain name server (DNS). While both tools can use IP addresses alone, they can also use domain names for devices. This is an added troubleshooting feature that can be useful in determining why particular services, such as email or web browsing, may not be working properly.



If ping does not work, you likely have it disabled on at least one of the interface settings, and security policies for that interface.

Both ping and traceroute require particular ports to be open on firewalls, or they cannot function. Since you typically use these tools to troubleshoot, you can allow them in the security policies and on interfaces only when you need them, and otherwise keep the ports disabled for added security.

Ping

The ping command sends a very small packet to the destination, and waits for a response. The response has a timer that may expire, indicating the destination is unreachable. The behavior of ping is very much like a sonar ping from a submarine, where the command gets its name.

Ping is part of Layer-3 on the OSI Networking Model. Ping sends Internet Control Message Protocol (ICMP) “echo request” packets to the destination, and listens for “echo response” packets in reply. However, many public networks block ICMP packets because ping can be used in a denial of service (DoS) attack (such as Ping of Death or a smurf attack), or by an attacker to find active locations on the network. By default, FortiGate units have ping enabled and broadcast-forward is disabled on the external interface.

What ping can tell you

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is some packet loss detected, you should investigate:

- possible ECMP, split horizon, network loops
- cabling to ensure no loose connections
- verify which security policy was used (use the packet count column on the *Policy* > *Policy* page)

If there is total packet loss, you should investigate:

- **hardware** — ensure cabling is correct, and all equipment between the two locations is accounted for
- **addresses and routes** — ensure all IP addresses and routing information along the route is configured as expected
- **firewalls** — ensure all firewalls, including FortiGate unit security policies allow PING to pass through

How to use ping

Ping syntax is the same for nearly every type of system on a network.

To ping from a FortiGate unit

- 1 Connect to the CLI either through telnet or through the CLI widget on the web-based manager dashboard.
- 2 Enter `exec ping 10.11.101.101` to send 5 ping packets to the destination IP address. There are no options for this command.

Output appears as:

```
Head_Office_620b # exec ping 10.11.101.101
PING 10.11.101.101 (10.11.101.101): 56 data bytes
64 bytes from 10.11.101.101: icmp_seq=0 ttl=255 time=0.3 ms
64 bytes from 10.11.101.101: icmp_seq=1 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=2 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=3 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=4 ttl=255 time=0.2 ms

--- 10.11.101.101 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
round-trip min/avg/max = 0.2/0.2/0.3 ms
```

To ping from an MS Windows PC

- 1 Open a command window.
 - In Windows XP, select *Start > Run*, enter *cmd*, and select *OK*.
 - In Windows 7, select the Start icon, enter *cmd* in the search box, and select *cmd.exe* from the list.
- 2 Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate unit with four packets.

Other options include:

- `-t` to send packets until you press “Control-C”
- `-a` to resolve addresses to domain names where possible
- `-n X` to send X ping packets and stop

Output appears as:

```
C:\>ping 10.11.101.101
```

```
Pinging 10.11.101.101 with 32 bytes of data:
Reply from 10.11.101.101: bytes=32 time=10ms TTL=255
Reply from 10.11.101.101: bytes=32 time<1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255
```

```
Ping statistics for 10.11.101.101:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 10ms, Average = 3ms
```

To ping from a Linux PC

- 1 Go to a shell prompt.
- 2 Enter `"/bin/etc/ping 10.11.101.101"`.

Traceroute

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, traceroute can be used to locate exactly where the problem is.

What is traceroute

Traceroute works by sending ICMP packets to test each hop along the route. It will send out three packets, and then increase the time to live (TTL) setting by one each time. This effectively allows the packets to go one hop farther along the route. This is the reason why most traceroute commands display their maximum hop count before they start tracing the route — that is the maximum number of steps it will take before declaring the destination unreachable. Also the TTL setting may result in steps along the route timing out due to slow responses. There are many possible reasons for this to occur.

Traceroute by default uses UDP datagrams with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP echo request (type 8) instead, as used by the Windows `tracert` utility. If you have a firewall and if you want traceroute to work from both machines (Unix-like systems and Windows) you will need to allow both protocols inbound through your FortiGate security policies (UDP with ports from 33434 to 33534 and ICMP type 8).

You can also use the packet count column of the *Policy > Policy > Policy* page to track traceroute packets. This allows you to verify the connection, but also confirm which security policy the traceroute packets are using.

What traceroute can tell you

ping and traceroute have similar functions — to verify connectivity between two points. The big difference is that traceroute shows you each step of the way, where ping doesn't. Also, ping and traceroute use different protocols and ports, so one may succeed where the other fails.

You can verify your DNS connection using traceroute. If you enter an FQDN instead of an IP address for the traceroute, DNS will try to resolve that domain name. If the name does not get resolved, you know you have DNS issues.

How to use traceroute

The traceroute command varies slightly between operating systems. Note that in MS Windows the command name is shortened to “`tracert`”. Also, your output will list different domain names and IP addresses along your route.

To use traceroute on an MS Windows PC

- 1 Open a command window.
 - In Windows XP, select *Start > Run*, enter `cmd`, and select *OK*.
 - In Windows 7, select the Start icon, enter `cmd` in the search box, and select `cmd.exe` from the list.
- 2 Enter “`tracert fortinet.com`” to trace the route from the PC to the Fortinet web site.

Output will appear as:

```
C:\>tracert fortinet.com
```

```
Tracing route to fortinet.com [208.70.202.225]
over a maximum of 30 hops:
  1  <1 ms    <1 ms    <1 ms    172.20.120.2
  2  66 ms     24 ms     31 ms    209-87-254-xxx.storm.ca
    [209.87.254.221]
  3  52 ms     22 ms     18 ms    core-2-g0-0-1104.storm.ca
    [209.87.239.129]
  4  43 ms     36 ms     27 ms    core-3-g0-0-1185.storm.ca
    [209.87.239.222]
  5  46 ms     21 ms     16 ms    te3-x.1156.mpd01.cogentco.com
    [38.104.158.69]
  6  25 ms     45 ms     53 ms    te8-7.mpd01.cogentco.com
    [154.54.27.249]
  7  89 ms     70 ms     36 ms    te3-x.mpd01.cogentco.com
    [154.54.6.206]
  8  55 ms     77 ms     58 ms    sl-st30-chi-.sprintlink.net
    [144.232.9.69]
```

```

9      53 ms      58 ms      46 ms  sl-0-3-3-x.sprintlink.net
[144.232.19.181]
10     82 ms      90 ms      75 ms  sl-x-12-0-1.sprintlink.net
[144.232.20.61]
11    122 ms     123 ms     132 ms  sl-0-x-0-3.sprintlink.net
[144.232.18.150]
12    129 ms     119 ms     139 ms  144.232.20.7
13    172 ms     164 ms     243 ms  sl-321313-0.sprintlink.net
[144.223.243.58]
14     99 ms      94 ms      93 ms  203.78.181.18
15    108 ms     102 ms      89 ms  203.78.176.2
16     98 ms      95 ms      97 ms  208.70.202.225

```

Trace complete.

The first, or the left column, is the hop count, which cannot go over 30 hops. When that number is reached, the traceroute ends.

The second, third, and fourth columns display how much time each of the three packets takes to reach this stage of the route. These values are in milliseconds and normally vary quite a bit. Typically a value of <1ms indicates a local connection.

The fifth, or the column farthest to the right, is the domain name of that device and its IP address or possibly just the IP address.

To perform a traceroute on a Linux PC

- 1 Go to a command line prompt.
- 2 Enter `"/bin/etc/traceroute fortinet.com"`.

The Linux traceroute output is very similar to the MS Windows tracert output.

Check the logs

This step in troubleshooting can be forgotten, but its an important one. Logging records the traffic passing through the FortiGate unit to your network and what action the FortiGate unit took during its scanning process of the traffic. This recorded information is called a log message.

When you configure FortiOS initially, log as much information as you can. If needed, logging of unused features can be turned off or scaled back if too the logs generate are too large.

As with most troubleshooting steps, before you can determine if the logs indicate a problem, you need to know what logs result from normal operation. Without a baseline it is difficult to properly troubleshoot.

When troubleshooting with log files:

- compare current logs to a recorded baseline of normal operation
- if needed increase the level of logging (such as from Warning to Information) to obtain more information

When increasing logging levels, ensure that alert email is configured and both disk usage and log quota are selected. This ensures you will be notified if the increased logging causes problems. You can also use Logging Monitor (located in *Log&Report > Monitor > Logging Monitor*) to determine the activities that generate the most log entries.

- check all logs to ensure important information is not overlooked

- filter or order log entries based on different fields (such as level, service, or IP address) to look for patterns that may indicate a specific problem (such as frequent blocked connections on a specific port for all IP addresses)
- use log reporting to help you visualize large amounts of log data — go to *Log&Report > Report Access > Cover Page*, select *Option*, and configure the date, time, and contents of the log report. You also have the option to email generated log reports to an administrator.

Logs will help identify and locate any problems, but they will not solve the problems. The job of logs is to speed up your problem solving and save you time and effort.

For more information on Logging and Log Reports, see the [Logging and Reporting handbook chapter](#).

Verify the contents of the routing table (in NAT mode)

When you have some connectivity, or possibly none at all a good place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route is not used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. If your FortiGate unit is in Transparent mode, you are unable to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the web-based manager, use the Routing Monitor by going to *System > Routing > Monitor*.

In the CLI, use the command `get router info routing-table all`. Sample output:

```
FGT# get router info routing-table all
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
        O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type
2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-
IS inter area
        * - candidate default

S*      0.0.0.0/0 [10/0] via 172.20.120.2, wan1
C       10.31.101.0/24 is directly connected, internal
C       172.20.120.0/24 is directly connected, wan1
```

Check the bridging information in Transparent mode

When FortiOS is in Transparent mode, the unit acts like a bridge sending all incoming traffic out on the other interfaces. The bridge is between interfaces on the FortiGate unit.

What checking the bridging information can tell you

Each bridge listed is a link between interfaces. Where traffic is flowing between interfaces, you expect to find bridges listed. If you are having connectivity issues, and there are no bridges listed that is a likely cause. Check for the MAC of the interface or device in question.

How to check the bridging information

To list the existing bridge instances on the FortiGate unit, use the following command:

```
diagnose netlink brctl list
```

Sample Output:

```
#diagnose netlink brctl list
list bridge information
1. root.b fdb: size=256  used=6      num=7      depth=2      simple=no
Total 1 bridges
```

How to display forwarding domain information

Forwarding domains, or collision domains, are used in routing to limit where packets are forwarded on the network. Layer-2 broadcasts are limited to the same group. By default, all interfaces are in group 0. For example, if the FortiGate unit has 12 interfaces, only two may be in the same forwarding domain which will limit packets that are broadcast to only those two interfaces. This reduces traffic on the rest of the network.

Collision domains prevent the forwarding of ARP packets to all VLANs on an interface. Without collision domains, duplicate MAC addresses on VLANs may cause ARP packets to be duplicated. Duplicate ARP packets can cause some switches to reset.

It is important to know what interfaces are part of which forwarding domains as this determines which interfaces can communicate with each other.

To manually configure forwarding domains in transparent mode, use the FortiOS CLI command `config interface edit <intf_name> set forward-domain <int>`

To display the information for forward domains

```
diagnose netlink brctl domain <name> <id>
```

where <name> is the name of the forwarding domain to display.

Sample Output

```
diagnose netlink brctl domain ione 101
show bridge root.b ione forward domain.
id=101 dev=trunk_1 6
```

To list the existing bridge MAC table, use the following command:

```
diagnose netlink brctl name host <name>
```

Sample Output

```
show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
```

Bridge root.b host table

port no	device	devname	mac addr	ttl	attributes
2	7	wan2	02:09:0f:78:69:00	0	Local Static
5	6	vlan_1	02:09:0f:78:69:01	0	Local Static
3	8	dmz	02:09:0f:78:69:01	0	Local Static
4	9	internal	02:09:0f:78:69:02	0	Local Static
3	8	dmz	00:80:c8:39:87:5a	194	
4	9	internal	02:09:0f:78:67:68	8	
1	3	wan1	00:09:0f:78:69:fe	0	Local Static

To list the existing bridge port list, use this command:

```
diagnose netlink brctl name port <name>
```

Sample Output:

```
show bridge root.b data port.
trunk_1 peer_dev=0
internal peer_dev=0
dmz peer_dev=0
wan2 peer_dev=0
wan1 peer_dev=0
```

Perform a sniffer trace

When troubleshooting networks and routing in particular, it helps to look inside the headers of packets to determine if they are traveling along the route you expect that they are. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate unit has NP2 interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP2 interfaces, you should disable offloading on those interfaces.

What can sniffing packets tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the FortiGate unit is silently dropping packets for reasons such as Reverse Path Forwarding (RPF), also called Anti Spoofing, which prevents an IP packet from being forwarded if its Source IP does not either belong to a locally attached subnet (local interface), or be part of the routing between the FortiGate unit and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the FortiGate unit and cause many features to be turned off.



If you configure virtual IP addresses on your FortiGate unit, it will use those addresses in preference to the physical IP addresses. You will notice this when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

How do you sniff packets

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose>
<count>
```

To stop the sniffer, type CTRL+C.

<interface_name>	The name of the interface to sniff, such as “port1” or “internal”. This can also be “any” to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. “none” indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<count>	The number of packets the sniffer reads before stopping. If you do not put a number here, the sniffer will run forever until you stop it with <CTRL C>.

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next three packets on the port1 interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh
3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh
1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh
4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface travelling between a computer with the host name of “PC1” and the computer with the host name of “PC2”. With verbosity 4 and above, the sniffer trace will display the interface names where traffic enters or leaves the FortiGate unit.

Remember to stop the sniffer, type CTRL+C.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
or
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and
icmp" 4
```

The following sniffer CLI command includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for instance PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

Debug the packet flow

Traffic should come in and leave the FortiGate unit. If you have determined that network traffic is not entering and leaving the FortiGate unit as expected, debug the packet flow.

Debugging can only be performed using CLI commands. Debugging the packet flow requires a number of debug commands to be entered as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate unit has FortiASIC NP2 interface pairs that are offloading traffic, this will change the packet flow. Before performing the debug on any NP2 interfaces, you should disable offloading on those interfaces.

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate unit and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```
FGT# diag debug disable
FGT# diag debug flow filter add <PC1>
FGT# diag debug flow show console enable
FGT# diag debug flow trace start 100
FGT# diag debug enable
```

The `start 100` argument in the above list of commands will limit the output to 100 packets from the flow. This is useful for looking at the flow without flooding your log or display with too much information.

To stop all other debug activities, enter the command:

```
FGT# diag debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching security policy, and is in turn blocked by the FortiGate unit. The denied message indicates the traffic was blocked.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825
msg="vd-root received a packet(proto=6,
192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924
msg="allocate a new session-013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597
msg="find a route: gw-192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg="
Denied by forward policy check"
```

Check number of sessions used by UTM proxy

Each FortiGate model has a set limit on the maximum number of sessions the UTM proxy supports. The UTM proxy handles all the traffic for the following protocols HTTP, SMTP, POP3, IMAP, FTP, and NNTP. If the proxy for a protocol fills up its session table, the FortiGate unit will enter conserve mode where it behaves differently until entries and memory free up again.

Conserve or failopen mode

Once you reach the limit, depending on your FortiGate unit's conserve mode configuration, no new sessions are created until an old ones end. You can configure your FortiGate unit's behavior when memory is running low or the proxy connection limit has been reached. There are two related commands for this in the CLI:

```
config system global
    set av-failopen-session {enable | disable}
    set av-failopen { idledrop | off | one-shot | pass}
end
```

`av-failopen-session` must be enabled to set the behavior for these conditions. When it is enabled, and a proxy for a protocol runs out of room in its session table that protocol goes into failopen mode and behaves as defined in the `av-failopen` command.

`av-failopen` determines the behavior of the proxy until entries are free in the session table again for that proxy.

- **idledrop** — This option removes idle sessions from the session table, starting with the clients that have the most sessions currently open. This method assumes that idle sessions are not being used and it will not cause problems to close these sessions. This is usually true, but some applications may have problems with this and start either complaining about not having or being able to open a session. If this occurs, try another method to check if this is really the problem. This is a secure option as no unscanned traffic is allowed to pass.
- **off** — This option turns off accepting any new AV sessions, but will continue to process any existing AV sessions that are currently active. All the protocols listed (HTTP, SMTP, POP3, IMAP, FTP, and NNTP) are scanned by FortiGate Antivirus. If AV scanning is enabled, `av-failopen off` is selected, and the proxy session table fills up no new sessions of that type will be accepted. For example, if POP3 session table is filled and email AV scanning is enabled, no more POP3 connections will be allowed until the session table gets some free space. This is a secure option because no unscanned traffic is allowed to pass.
- **one-shot** — When memory is low, bypass the antivirus system. The name one-shot comes from the fact that once you are in one-shot `av-failopen` mode, you must set `av-failopen` to either `pass` or `off` to restart AV scanning. This is a very insecure option because it allows all traffic without AV scanning, and it never reverts to normal without manual assistance.
- **pass** — When memory is low, bypass the antivirus system much as one-shot. The difference is that when memory is freed up, the system will start AV scanning automatically again. This is an insecure option because it allows traffic to pass without AV scanning. However, it is better than one-shot because it automatically restarts AV scanning when possible.

If the proxy session table is full for one or more protocols and your FortiGate unit enters into conserve or failopen mode, it will appear as if you have lost connections, network services are intermittent or non-existent, and yet other services work normally for a while until their sessions end and they join the queue of session starved applications.

Checking sessions in use

To make troubleshooting this type of problem easier, sessions are broken down by which protocol they use. This provides you with statistics and errors specific to one of the protocols.



Due to the amount of output from this command, you should connect to the CLI with a terminal program, such as puTTY, that logs output. Otherwise, you will likely not be able to access all the output information from the command.

In the following output, only the HTTP entries are displayed. The other protocols have been removed in an attempt to shorten the output. There will be separate entries for each supported protocol (HTTP, SMTP, POP3, IMAP, FTP, and NNTP) in each section of the output.

For more information on session statistics, see [“Session table” on page 728](#).

To check sessions in use and related errors - CLI

```
FGT# # get test app proxyworker 4

Worker[0]
HTTP Common
Current Connections                8/8032
Max Concurrent Connections         76

Worker Stat
Running time (HH:MM:SS:usec)      29:06:27:369365
Time in loop scanning              2:08:000198
Error Count (accept)               0
Error Count (read)                 0
Error Count (write)                0
Error Count (poll)                 0
Error Count (alloc)                0
Last Error                        0
Acceptor Read                      6386
Acceptor Write                     19621
Acceptor Close                     0

HTTP Stat
Bytes sent                        667012 (kb)
Bytes received                     680347 (kb)
Error Count (alloc)                0
Error Count (accept)               0
Error Count (bind)                 0
Error Count (connect)              0
Error Count (socket)               0
Error Count (read)                 134
Error Count (write)                0
Error Count (retry)                40
```

```
Error Count (poll) 0
Error Count (scan reset) 2
Error Count (urlfilter wait) 3
Last Error 104
Web responses clean 17950
Web responses scan errors 23
Web responses detected 16
Web responses infected with worms 0
Web responses infected with viruses 0
Web responses infected with susp 0
Web responses file blocked 0
Web responses file exempt 0
Web responses bannedword detected 0
Web requests oversize pass 16
Web requests oversize block 0
Last Server Scan errors 102
URL requests exempt 0
URL requests blocked 0
URL requests passed 0
URL requests submit error 0
URL requests rating error 0
URL requests rating block 0
URL requests rating allow 10025
URL requests infected with worms 0
Web requests detected 0
Web requests file blocked 0
Web requests file exempt 0
POST requests clean 512
POST requests scan errors 0
POST requests infected with viruses 0
POST requests infected with susp 0
POST requests file blocked 0
POST requests bannedword detected 0
POST requests oversize pass 0
POST requests oversize block 0
Web request backlog drop 0
Web response backlog drop 0
```

```
Worker Accounting
poll=721392/649809/42 pollfail=0 cmdb=85 scan=19266
acceptor=25975
```

```
HTTP Accounting
setup_ok=8316 setup_fail=0 conn_ok=0 conn_inp=8316
urlfilter=16553/21491/20 uf_lookupf=0
scan=23786 clt=278876 srv=368557
```

```
SMTP Accounting
setup_ok=12 setup_fail=0 conn_ok=0 conn_inp=12
scan=12 suspend=0 resume=0 reject=0 spamadd=0 spamdel=0 clt=275
srv=279
```

```
POP3 Accounting
```

```

setup_ok=30 setup_fail=0 conn_ok=0 conn_inp=30
scan=3 clt=5690 srv=5836

IMAP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0

FTP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0 datalisten=0 dataclt=0 datasrv=0

NNTP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0

```

The output from this command falls into the following sections:

- **HTTP Common current connections** — There is an entry for each protocol that displays the connections currently used, and the maximum connections allowed. This maximum is for the UTM proxy which means all the protocols connections combined cannot be larger than this number. To support this, note that the maximum session count for each protocol is the same. You may also see a line titled `Max Concurrent Connections` for each protocol. This number is the maximum connections of this type allowed at one time. If VDOMs are enabled, this value is defined either on the global or per-VDOM level at `VDOM > Global Resources` or in the CLI at `config system resource-limits`.
- **Worker Stat** — This is statistics about the UTM proxy including how long it has been running, and how many errors it has found.
- **HTTP Stat** — This section includes statistics about the HTTP protocol proxy. This is a very extensive list covering errors, web responses, and any UTM positive matches. There are similar sections for each protocol, but the specific entries in each vary based on what UTM scanning is looking for in each — spam control for email, file transfer blocking for FTP, and so on.
- **Worker Accounting** — Lists accounting information about the UTM proxy such as polling statistics, how many sessions were scanned, and how many were just accepted. This information can tell you if expect AV scanning is taking place or not. Under normal operation there should be no errors or fails.
- **HTTP Accounting** — The accounting sections for each protocol provide information about successful session creation, failures, how many sessions are being scanned or filtered, and how many are client or server originated. If `setup_fail` is larger than zero, run the command again to see if it is increasing quickly. If it is, your FortiGate unit may be in conserve mode.

Related commands

To clear the UTM proxy statistics

```
# get test app proxyworker 8
```

To drop all idle connections

```
# get test app proxyworker 222
```

To display statistics per VDOM

```
# get test app proxyworker 4444
```

For additional related commands, see “[get test](#)” on page 853.

Examine the firewall session list

One further step is to examine the firewall session. The firewall session list displays all the sessions the FortiGate unit has open. You will be able to see if there are strange patterns such as no sessions apart from the internal network, or all sessions are only to one IP address.

When examining the firewall session list in the CLI, filters may be used to reduce the output. In the web-based manager, the filters are part of the interface.

To examine the firewall session list - web-based manager

- 1 Go to *System > status > Dashboard > Top Sessions*.

This widget can display information as either a bar graph or a table.

- 2 Select *Detach*, and then *Details*.
- 3 Expand the session window to full screen to display the information.
- 4 Change filters, view associated security policy, column ordering, and so on to analyze the sessions in the table.
- 5 Select the delete icon to terminate the session.

To examine the firewall session list - CLI

When examining the firewall session list, there may be too many sessions to display. In this case it will be necessary to limit or filter the sessions displayed by source or destination address, or NATed address or port. If you want to filter by more than one of these, you need to enter a separate line for each value.

The following example shows filtering the session list based on a source address of 10.11.101.112.

```
FGT# diag sys session filter src 10.11.101.112
FGT# diag sys session list
```

The following example shows filtering the session list based on a destination address of 172.20.120.222.

```
FGT# diag sys session filter dst 172.20.120.222
FGT# diag sys session list
```

To clear all sessions corresponding to a filter - CLI

```
FGT# diag sys session filter dst 172.20.120.222
FGT# diag sys session clear
```

Check source NAT information

Remember NAT when troubleshooting connections. NAT is especially important if you are troubleshooting from the remote end of the connection outside the FortiGate unit firewall. On the dashboard session list, pay attention to *Src address after NAT*, and *Src port after NAT*. These columns display the IP and port values after NAT has been applied.

The NAT values can be helpful to ensure they are the values you expect, and to ensure the remote end of the sessions can see the expected IP address and port number.

When displaying the session list in the CLI, you can match the NATed source address (*nsrc*) and port (*nport*). This can be useful if multiple internal IP addresses are NATed to a common external facing source IP address.

```
FGT# diag sys session filter nsrc 172.20.120.122
FGT# diag sys session filter nport 8888
```

```
FGT# diag sys session list
```

Checking wireless information

Wireless connections, stations, and interfaces have different issues that other physical interfaces.

Troubleshooting station connection issue

To check whether station entry is created on Access Control:

```
FG600B3909600253 # diagnose wireless-controller wlaac -d sta
* vf=0 wtp=70 rId=2 wlan=open ip=0.0.0.0 mac=00:09:0f:db:c4:03
  rssi=0 idle=148 bw=0 use=2
  vf=0 wtp=70 rId=2 wlan=open ip=172.30.32.122
  mac=00:25:9c:e0:47:88 rssi=-40 idle=0 bw=9 use=2
```

Enable diagnostic for particular station

This example uses the station MAC address to find where it is failing:

```
FG600B3909600253 # diagnose wireless-controller wlaac sta_filter
00:25:9c:e0:47:88 1
Set filter sta 00:25:9c:e0:47:88 level 1
FG600B3909600253 # 71419.245 <ih> IEEE 802.11 mgmt::disassoc <==
00:25:9c:e0:47:88 vap open rId 1 wId 0 00:09:0f:db:c4:03
71419.246 <dc> STA del 00:25:9c:e0:47:88 vap open rId 1 wId 0
71419.246 <cc> STA_CFG_REQ(34) sta 00:25:9c:e0:47:88 del ==> ws
(0-192.168.35.1:5246) rId 1 wId 0
71419.246 <cc> STA del 00:25:9c:e0:47:88 vap open ws (0-
192.168.35.1:5246) rId 1 wId 0 00:09:0f:db:c4:03 sec open reason
I2C_STA_DEL
71419.247 <cc> STA_CFG_RESP(34) 00:25:9c:e0:47:88 <== ws (0-
192.168.35.1:5246) rc 0 (Success).
```

Other diagnose commands

Diagnose commands are a series of commands available on all FortiGate units. These commands can help you troubleshoot network activity. The packet sniffer mentioned earlier is only one of many useful diagnose commands.

For additional diagnostic commands, see [“Troubleshooting ‘get’ commands” on page 805](#).



Troubleshooting advanced

There are issues other than connectivity that require troubleshooting. These are more advanced issues that may require investigating a number of different issues before the problem is solved.

The advanced troubleshooting sections include, and can help answer the following questions:

- [“Traffic shaping issues” on page 789](#)

Is an application or user consuming all your network bandwidth?

Do some applications need more bandwidth than they are getting?

Is less bandwidth allowed than expected on a connection?

Are sessions or packets being dropped?

- [“User and administrator logon issues” on page 793](#)

Are one or more of your users unable to log on to the network?

Are one or more of your users unable to initiate VPN connections?

Are you having problems logging in to administrator accounts?

- [“IPsec VPN issues” on page 797](#)

Are VPN negotiations slow?

Are VPN connections not completing the initial Phase 1 handshake?

Is the IPsec VPN tunnel down?

For connectivity troubleshooting tips, see [“Troubleshooting common issues” on page 763](#).

In addition to these steps, you may find diagnose commands useful. See [“Other diagnose commands” on page 803](#).

Traffic shaping issues

Traffic shaping, or Internet Traffic Management Practices (ITMP), is often thought of as a preventative measure, but it can also be used effectively in troubleshooting when you have some application or some user consuming all the bandwidth.

This topic includes the following:

- [Use traffic shapers to limit traffic in testing and network simulations](#)
- [Monitoring traffic](#)
- [Displaying configured traffic shaping](#)
- [Troubleshooting protocols and users using traffic shaping](#)
- [Displaying current bandwidth and dropped packets for a traffic shaper](#)

Use traffic shapers to limit traffic in testing and network simulations

Some applications require a minimum bandwidth to function properly. One example of this is VoIP — even if the packets are high priority, if too few are delivered the call quality suffers. When signs like this happen and you suspect it is a lack of bandwidth that is causing the problem, you can use traffic shapers to help confirm your problem.

One of the hardest parts of troubleshooting is reproducing the problem consistently, especially if the problem is intermittent. The understood best practice is to set up a test network to try and reproduce the behavior you have been seeing on your network so you will not be interrupting business.

Traffic shapers allow you to limit the bandwidth over a connection. This, in turn, allows you to simulate limited bandwidth to different parts of the network. You can experiment with your traffic shaper settings to discover at what point VoIP calls, for example, are unacceptably poor quality, and then you know where to set your lowest VoIP guaranteed minimum level. This can be done for many different applications, computers on the network, or users or user groups.

Monitoring traffic

Traffic shaping is best used with a traffic monitor, such as the Traffic History widget that is available on the Dashboard. Configure two — one to watch the internal interface and another for the external interface. Run these monitors for a while before trying the traffic shaping procedures to get a feel for your bandwidth usage.

To configure two traffic history monitors - web-based management

- 1 Go to *System > Dashboard > Status*.
 - 2 Select *Widget > Traffic History*.
 - 3 Select the edit icon in the widget's menu bar.
 - 4 Name the widget Internal, and select the Internal interface to monitor.
 - 5 Enable *Refresh*, and select *OK*.
 - 6 Repeat steps 2 through 5 selecting and naming this second monitor external.
- Let the monitors run for a while to display network traffic before starting troubleshooting. The times you start and stop traffic on the FortiGate unit is displayed in the monitors. It may be advisable to take screen captures of the monitors at each stage for easier comparison later.

Displaying configured traffic shaping

It can be difficult to either find the traffic shaping settings that are needed, or to see all the settings in one place. Both of these can be accomplished through CLI commands.

```
# diagnose firewall shaper traffic-shaper { list | stats | state }
```

This CLI diagnose command provides information about shared traffic shaper.

- **list** — Lists all the configured shared traffic-shapers. This is the same as going to *Firewall > Traffic Shapers > Shared* in the web-based manager. However, you can specify additional parameters to narrow down the traffic shapers that are matched. For example if you add `policy 7 ipv4`, it will only display traffic shapers that are associated with security policy number seven and then only the IPv4 traffic numbers. You must specify either `ipv4` or `ipv6` since security policies can be either IPv4 or IPv6. Apart from current-bandwidth and packets dropped, all the information is from the configuration.

Sample output:

```
# diag firewall shaper traffic-shaper list policy 4 ipv4

name traffic_test_app
maximum-bandwidth 12 KB/sec
guaranteed-bandwidth 0 KB/sec
current-bandwidth 0 B/sec
priority 3
policy 4
packets dropped 0
```

- **stats** — Display statistics for the traffic shaper. Add `list` to the command to view the stats. The stats are very brief — number of shapers, number of IPv4 shapers, number of IPv6 shapers, and number of packets dropped. Or, you can wipe the stats by adding the word `clear` on the end of the command. This will remove all store statistics and start fresh.
- **state** — Displays the total number of global traffic shapers. This is useful if VDOMs are enabled on your FortiGate unit as it will list all traffic shapers on your FortiGate unit. This command has no additional arguments.

Troubleshooting protocols and users using traffic shaping

To find a protocol or user who is using all the bandwidth, create two security policies and have them handle all the traffic on the network — one has all the users or protocols in it and the other has none to start with. The first security policy will have traffic shaping applied to all its members, so the problem will go away. However, to locate the offender, move users or protocols from the first policy with traffic shaping to the second policy that has no traffic shaping. It is easy to see who the offender is because when that protocol or user is moved to the unlimited policy, the bandwidth usage will sky rocket again. Once you have found the offender, you can either finish checking all the others, or assume that one was the only problem and put the regular security policies back in place.

One way to accomplish this test is create a VDOM, and funnel all traffic through the VDOM using an inbound and outbound inter-VDOM links. This will allow you to keep most of your security policies in place and create the new traffic shaping policies in the new VDOM. Once the test is finished, you simply delete the VDOM and its configuration.

To create a traffic shaping security policy to limit FTP bandwidth usage

- 1 Go to *Firewall > Traffic Shaper > Shared*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*.

Name	traffic_app_test
Apply Shaper	Per Policy
Maximum Bandwidth	Enable and enter 100.
Guaranteed Bandwidth	Disable
Traffic Priority	Medium

- 4 Go to *Firewall > Policy > Policy*.

The following steps will override any existing security policies. **Do not perform these steps on a live network without notifying users first.** Optionally, use the VDOM method mentioned earlier to limit disruption of users.

- 5 Select *Create New*.
- 6 Select any and all for source and destination interfaces and addresses.
- 7 Select *always* for *schedule*.
- 8 For service, select *Multiple* and enter all the FTP related entries.
- 9 Select *ACCEPT* for *Action*.
- 10 Enable NAT.
- 11 Enable *Traffic Shaping*, and select *traffic_test_app*.
- 12 Select *OK*.
- 13 Select the new security policy in the list.
- 14 Select *Move*, and insert it before Seq. No. 1.
- 15 Repeat steps 5 through 10, but for service select all services.

This works because the first policy will catch all FTP traffic first before it gets to this policy.

- 16 Select *OK*.
- 17 Select the new security policy in the list.
- 18 Select *Move*, and insert it below Seq. No. 1.

This procedure creates two security policies that will catch all the traffic — the first one catches all FTP traffic and limits its bandwidth, and the second one matches everything else. If FTP is found to not be the problem, it is easy to remove FTP from the first policy and add a different protocol or group of protocols to see if they are responsible.

Once the bandwidth consuming protocol is discovered, it can be traffic shaped using security policies (clone existing policies and change the first one to only apply to that protocol with the traffic shaping applied), users can be informed to limit their use of that protocol or application, or that protocol could be blocked completely.

Using traffic shaping to troubleshoot users

Troubleshooting users with traffic shaping is very much like troubleshooting protocols. The difference is that the list of users is changed instead of the list of services.

Displaying current bandwidth and dropped packets for a traffic shaper

If you believe the wrong amount of traffic is going through a specific traffic shaper, you may want to see how much traffic is going through.

```
FGT # get firewall shaper traffic

[myTrafficShaper]
maximum-bandwidth: 1200 KB/sec
guaranteed-bandwidth: 100 KB/sec
current-bandwidth: 227328 B/sec
priority: medium
packets dropped: 25
```

This command is per-VDOM, so if VDOMs are enabled ensure you are in the correct VDOM before running the command.

`myTrafficShaper` is a shared traffic shaper. Some of the output displayed has different meanings depending how this shaper is applied. If this shaper is applied to only one security policy, then the numbers displayed are for the traffic flowing through that one security policy. But if this shaper is applied to multiple policies, the `current-bandwidth` and `packets dropped` will be combined for all instances, and will be less useful for troubleshooting.

For this reason, if possible, either create multiple copies of the traffic shaper with different names to allow for easier troubleshooting, or if possible disable all security policies but one until you get the information you need.

Traffic shaper drops packets

If `packets dropped` is above zero, that means the maximum-bandwidth was reached and traffic was dropped to stay at or under the limit. To prevent dropped packets, increase the `maximum-bandwidth` or create a new traffic shaper for this traffic. If the traffic must never be dropped, consider increasing the `priority` to high or not using a traffic shaper at all for this traffic.

No effect for shaper priority change

If changing the `priority` does not affect which traffic is being dropped first when the `maximum-bandwidth` level is reached, it may be that the traffic is going through a virtual interfaces do not have priority queues, only physical interfaces have multiple queues for priority traffic. Examples of virtual interfaces include inter-VDOM links, VLANs, aggregated links, and redundant links.

User and administrator logon issues

Few problems are more troublesome for users or administrators than when they are unable to log onto the company network. There are a number of potential reasons for these problems. Once the reason for the issues are found, the administrator should follow the appropriate policies to resolve the problems, notifying affected users if warranted.

Administrator logon issues can be more serious. However, in most cases if remote admin access is not available the console interface is still available. Many problems such as limited IP address access, limited protocol access, and so on are not an issue when using the console interface.

This topic includes the following:

- [User logon issues](#)
- [Administrator logon issues](#)

User logon issues

This topic includes the following:

- [Use correct username and password combination for user](#)
- [Check authentication security policies](#)
- [Use proper two-factor authentication code \(FortiToken or delivered code\)](#)
- [User credentials must exist on the remote server \(remote authentication\)](#)

Use correct username and password combination for user

This may be obvious, but it should be the first thing to check. While there are valid reasons for a user to forget their information or enter the wrong information, they may be trying to use the information of another user to gain illegal access to the company network. If this is the case, you do not want to waste time on any additional troubleshooting. Also if this is the case it will generally be a single user with problems instead of a group of users.

Check authentication security policies

One or more security policies allow or deny a user group access to the network. If one or more users have authentication problems, it is possible there were firewall changes or the user may have had their group memberships recently changed resulting in changes to network access.

To check the security policies that apply to the group test

- 1 Go to *User > User Group > User Group*.
- 2 Expand *Firewall* if required.



You should pay attention to the user groups that include the affected user or users. Take special attention if multiple affected users are all part of one group.

- 3 Go to *Policy > Policy > Policy*.

Ensure the *Authentication* column is visible. If not, do the following to show the column:

- Select *Column Settings*.
 - Move *Authentication* to the list on the right.
 - Move the *Authentication* entry to your preferred place in the list using the *Move up* and *Move down* buttons.
 - Select *OK*.
- 4 Locate one or more authentication policies that authenticate the user group or groups this user belongs to as noted earlier. The user group or groups that a security policy authenticates will appear in the *Authentication* column.
 - 5 Edit these security policies to inspect the details of the security policy or policies to determine what methods of authentication are required, and other relevant information.
 - 6 Either adjust the requirements to enable user access, move user to their own group with a new security policy to troubleshoot the problems without affecting other users, or use the information in the policy to farther troubleshoot the user problems.

Use proper two-factor authentication code (FortiToken or delivered code)

When two-factor authentication is enabled for a user or admin account, the user requires more than just their username and password to logon. They require a certificate, a FortiToken device, or an authentication code sent either as email or a text message. When in place, this information is prompted for on the logon screen no matter if the user is connecting with FortiClient, VPN, or other logon screen. If there is no prompt for your token code, two-factor authentication is not enabled and not required.

It is possible with no prompt that two-factor authentication may be enabled using certificates.

To verify there is a certificate to authenticate

- 1 Go to *User > User Group > User Group*.
- 2 Expand *Firewall* if required.



You should pay attention to the user groups that include the affected user or users. Take special attention if multiple affected users are all part of one group.

- 3 Go to *Policy +> Policy > Policy*.

Ensure the Authentication column is visible. If not, do the following to show the column:

- Select *Column Settings*.
 - Move *Authentication* to the list on the right.
 - Move the *Authentication* entry to your preferred place in the list using the *Move up* and *Move down* buttons.
 - Select *OK*.
- 4 Locate one or more authentication policies that authenticate the user group or groups this user belongs to as noted earlier. The user group or groups that a security policy authenticates will appear in the Authentication column.
 - 5 Edit these security policies to look for any that use certificate authentication.



Identify which certificate is required, and ensure the user has that certificate installed on their system.

User credentials must exist on the remote server (remote authentication)

Before configuring FSSO or the FortiGate unit for remote authentication, the user credentials must exist on the remote server.

When FSSO is being configured, it must be able to contact the remote authentication server, such as Windows AD server, for user login information.

Administrator logon issues

In addition to the common user issues such as security policies, administrators have some unique problems users do not experience. Some common administrator access issues include:

- [Allow access for interface is not enabled](#)
- [Trusted hosts for admin account will not allow current IP](#)
- [FortiGate asking for password when creating a remotely authenticated administrator account](#)

Allow access for interface is not enabled

Each interface in FortiOS has a set of administrator access protocols — HTTP, HTTPS, SSH, TELNET, PING, and SNMP. These are the methods an administrator can use to connect to FortiOS, and any or all of them may be disabled on any interface.

IPv4 protocols are treated differently from IPv6 protocols and must be enabled separately. This provides added security.

When a FortiGate unit is shipped, all access on external interfaces is disabled by default except for PING for troubleshooting purposes. This means any attempt by an administrator to logon over a dmz interface will be blocked by default.

The following procedures demonstrate enabling full access on dmz interface. For security purposes, you should only enable access that is required. If you open access for troubleshooting, remember to disable it afterwards. Failure to do so will leave a gap in your security that hackers may be able to exploit.

To enable administrator access on the dmz interface - web-based manager

- 1 Log on as administrator.
- 2 Go to *System > Network > Interface*, select the dmz interface, and select *Edit*.
- 3 Under *Administrative Access*, select HTTPS, PING, HTTP, SSH, SNMP, and TELNET.
- 4 If you use IPv6, select the above protocols to be used for IPv6 administrative access.
- 5 Select *OK*.
- 6 Repeat for each interface where administrative access is required.

To enable administrator access on the dmz interface - CLI

```
config system interface
  edit dmz
    set allowaccess HTTPS PING HTTP SSH SNMP TELNET
    set allowaccess6 HTTPS PING HTTP SSH SNMP TELNET
  next
end
```

Trusted hosts for admin account will not allow current IP

A trusted host is a location that is secure where an administrator is allowed to logon. For example on a secure network an administrator would be able to logon from an internal subnet but not from the Internet.

If external administrator logon is required, a secure VPN tunnel can be established with a set IP address or range of addresses that are entered as a trusted host address.

Trusted host logon issues occur when an administrator attempts to logon from an IP address that is not included in the trusted host list.

To verify trusted host logon issues - web-based manager

- 1 Record the IP address where the administrator is attempting to logon to the FortiGate unit.
- 2 Logon to the FortiGate unit using a working admin account.
- 3 Go to *System > Admin > Administrators*.
- 4 Select the admin account in question and select *Edit*.
- 5 Compare the listed trusted hosts to the IP address the attempted logon is coming from. If there is a match, the problem is not due to trusted hosts.
- 6 If there is no match and the new address is valid (secure), add it to the list of trusted hosts.
- 7 Select *OK*.

If the problem was due to trusted hosts, the admin account will be able to logon now.

To verify trusted host logon issues - CLI

- 1 Record the IP address where the administrator is attempting to logon to the FortiGate unit.
- 2 Log on to the FortiGate unit using the console connection for your FortiGate unit.
- 3 Enter the following CLI commands:

```
config system admin
  edit <admin_account>
  get
```

This will display all the settings for this admin account, including trusted hosts

- 4 Compare the listed trusted hosts to the IP address the attempted logon is coming from. If there is a match, the problem is not due to trusted hosts.
- 5 If there is no match and the new address is valid (secure), add it to the list of trusted hosts using the following CLI commands:

```
set trusthost3 <ipv4_addr> <ipv4_netmask>
next
end
```

Use the numbered trusthost that follows the last configured trusthost. For example if trusthost, trusthost1, and trusthost2 are in use you will use trusthost3.

If the problem was due to trusted hosts, the admin account will be able to log on now.

FortiGate asking for password when creating a remotely authenticated administrator account

Local administrator accounts require you to enter a password so the administrator can be locally authenticated when logging on. A remote administrator account uses a remote authentication server to authenticate the administrator when logging on instead of using a local password.

When you are creating a remote administrator account, if you leave out the password you will not be able to create the administrator account. This password field is a required field.

For remote administrator accounts, the password field is backup password. During normal operations this account will authenticate remotely. However, if the authentication server is not reachable, authentication will occur at the local level instead to ensure administrator access is available.

IPsec VPN issues

VPN problems are more complex than many networking issues because of the layers of configuration required.

In general when troubleshooting IPsec VPN connections you should check:

- Is the tunnel up? If not, check phase 1 and phase 2 settings
- Are packets using the proper route? See [“Verify the contents of the routing table \(in NAT mode\)” on page 778](#).
- Is there a proper security policy for the tunnel? See [“Examine the firewall session list” on page 787](#).

This topic focuses on the following VPN tunnel related issues:

- [VPN negotiations appear to be slow](#)

- VPN tunnel proposal will not connect
- VPN Tunnel up but no traffic going over it
- Other useful VPN IKE related commands

VPN negotiations appear to be slow

Possible causes of slow VPN negotiations include the following:

- Many VPN negotiations take time
- Keep VPN information up to date
- Check for routing problems
- Limit number of P1 proposals

Many VPN negotiations take time

When a VPN connection is being negotiated, it takes a while to establish the connection. The upper limit on the number of VPN IKE negotiations is roughly 60 per second. This assumes there are no problems with the negotiations. Large installations may notice this time as a delay. For example, if there are 1000 connections to negotiate it would take roughly 17 seconds. This amount of delay cannot be improved upon.

Keep VPN information up to date

Many failed negotiations will result in longer delays. Possible out of date information can include changed routing information, computer for failed negotiations include

Check for routing problems

Routing problems on your network will increase delays as well. Problems such as split horizon and asymmetric routing will cause problems for all network traffic, including VPN negotiating.

For example, if you run the CLI command `diag ip route list`, the IPsec phase1 route entry should be before the system default entry. Otherwise, it is possible all traffic is wrongly using the system default route.

If you use the phase 1 CLI command `set default-go` but don't have a static route entry, traffic coming through the tunnel will use the system default route.

To solve this problem, add a static route for the appropriate VPN destination subnet before the system default in the table. This will provide a more specific match that will use the route you want.

Limit number of P1 proposals

If there are a number of P1 proposal combinations on each end, it may take a while to negotiate the connection. Every combination may be tried. For example if each end has two proposals that will result in four attempts, three results in nine, and five in 25.

VPN tunnel proposal will not connect

When attempting to establish a VPN tunnel, both ends must have at least one proposal that matches. This will be the settings they use to establish the tunnel and additional security and protocols.

Ensure you are not using a loopback for the local VPN interface. If you are, the tunnel will not be established.

The Local Interface field in Phase 1 VPN configuration should not be configured to point to a loopback interface. By design, the IPsec tunnel will not be established.

Each end of the attempted connection may have multiple protocols configured. While this gives a broader range of settings for a better chance of a match, it can result in long rambling attempts to connect that are difficult to troubleshoot.

Best practices are to set the proposal information on each end to exactly the same and it will match with only one proposal definition required. However, frequently you do not have access to both ends of the proposed tunnel and instead have to match the other end.



If you are trying to offload VPN processing to a network processing unit (NPU), remember that only SHA1 authentication is supported. For high levels of authentication such as SHA256, SHA384, and SHA512 hardware offloading is not an option — all VPN processing must be done in software.

To determine what the other end of the VPN tunnel is proposing

- 1 Start a terminal program such as puTTY and set it to log all output.
When necessary refer to the logs to locate information when output is verbose.
- 2 Log on to the FortiGate unit using a super_admin account.
- 3 Enter the following CLI commands.
- 4 Display all the possible IKE error types and the number of times they have occurred:
`diag vpn ike errors`
- 5 Check for existing debug sessions:
`diag debug info`
If a debug session is running, to halt it enter:
`diag debug disable`
- 6 Confirm your proposal settings:
`diag vpn ike config list`
- 7 If your proposal settings do not match what you expect, make a change to it and save it to force an update in memory. If that fixes the problem, stop here.
- 8 List the current vpn filter:
`diag vpn ike filter`
- 9 If all fields are set to any, there are no filters set and all VPN ike packets will be displayed in the debug output. If your system has only a few VPNs, skip setting the filter.

If your system has many VPN connections this will result in very verbose output and make it very difficult to locate the correct connection attempt.
- 10 Set the VPN filter to display only information from the destination IP address for example 10.10.10.10 :
`diag vpn ike log-filter dst-addr4 10.10.10.10`

To add more filter options, enter them one per line as above. Other filter options are displayed in [Table 54](#).

Table 54: Filter options for diag vpn ike filter

clear	Erases the current filter.
dst-addr6	The IPv6 destination address range to filter by.
dst-port	The destination port range to filter by.
interface	Interface that IKE connection is negotiated over.

Table 54: Filter options for diag vpn ike filter

list	Displays the current filter.
name	The phase1 name to filter by.
negate	Negate the specified filter parameter.
src-addr4	The IPv4 source address range to filter by.
src-addr6	The IPv6 source address range to filter by.
src-port	The source port range to filter by.
vd	Index of virtual domain. 0 matches all.

11 Start debugging:

```
diag debug app ike 255
diag debug enable
```

12 Have the remote end attempt a VPN connection.

If the remote end attempts the connection, they become the initiator. This situation makes it easier to debug VPN tunnels because then you have the remote information and all of your local information. By initiating the connection, you will not see the other end's information.

13 If possible go to the web-based manager on your FortiGate unit, go to the VPN monitor and try to bring the tunnel up.**14 Stop the debug output:**

```
diag debug disable
```

15 Go back through the output to determine what proposal information the initiator is using, and how it is different from your VPN P1 proposal settings.

Things to look for in the debug output of attempted VPN connections are shown in [Table 55](#).

Table 55: Important terms to look for in VPN debug output

initiator	Starts the VPN attempt, in the above procedure that is the remote end.
responder	Answers the initiator's request.
local ID	In aggressive mode, this is not encrypted.
error no SA proposal chosen	There was no proposal match — there was no encryption-authentication pair in common, usually occurs after a long list of proposal attempts.
R U THERE and R U THERE ack	dead peer detection (dpd), also known as dead gateway detection — after three failed attempts to contact the remote end it will be declared dead, no farther attempts will be made to contact it.
negotiation result	Lists the proposal settings that were agreed on.
SA_life_soft and SA_life_hard	Negotiating a new security association (SA) key, and the key life.

Table 55: Important terms to look for in VPN debug output

R U THERE	This is the keep alive message. The reply is R U THERE ack. If you see this, it means Phase 1 was successful.
tunnel up	The negotiation was successful, the VPN tunnel is operational.

VPN Tunnel up but no traffic going over it

You can see in the VPN monitor that your tunnel is up, but the packet or byte count isn't is still at zero even though you are sending packets.

Some possible causes include:

- **NAT issues** — for single hops, you do not want to enable NAT on an IPsec policy, but for return packets from multiple hops you may need incoming NAT turned on or the packets won't know how to return.
- **Routing** — if you have default routing, it may not be enough if you are trying to reach a computer two or more hops away over the VPN tunnel. In that case, include a specific route for that subnet you are attempting to contact.
- **Check IP address of tunnel** — confirm the IP address of the tunnel ends as its possible they were assigned addresses you were not expecting.
- **Security policy** — you need a security policy to handle VPN traffic, either route based or policy based. Ensure if you specify a range of IP addresses that the VPN tunnel is part of that range.

Other useful VPN IKE related commands

There are a series of useful diagnose commands for VPN IKE.

They all start with `diag vpn ike <arg>`. [Table 56](#) lists options for <arg>.

Table 56: List of arguments for diag vpn ike CLI command

config list	Lists all the IKE configurations.
counts	Displays list of IKE objects and their current, maximum, and total counts. For example: crypto.md5: now 0 max 1 total 24
crypto	Either sets or displays hardware or software crypto settings. crypto hardware - use hardware crypto (where possible) crypto software - use software crypto status - show number of hardware and software crypto objects
filter log filter log-filter	Sets the IKE filter. See Table 54 on page 799 .
gateway	Displays VPN IKE gateways. Optionally clear or flush (same result) the gateways in memory.

Table 56: List of arguments for diag vpn ike CLI command

<code>restart</code>	Restarts the IKE daemon. This is useful if during troubleshooting you discover that settings in memory do not match what you set. All VPN connections will be lost when restarting. If you have active VPN connections, give those connections sufficient notice to close gracefully.
<code>routes list</code>	Displays all routes in memory for IKE VPN tunnels.
<code>status</code>	Displays status of IKE objects. <code>detailed</code> - lists vdom, name, version, IKE SA, and IPsec SA <code>summary</code> - lists number of IKE SA and IPsec SA objects

Logging

The following contains information regarding troubleshooting logging issues on the unit. This topic contains the following:

- [Cannot log to a supported log device](#)
- [The alert email did not send an email to the email address](#)
- [The FortiGate unit stopped logging: what happened?](#)

Cannot log to a supported log device

If you are unable to log to a log device that is supported by the FortiGate unit, try the following:

- verify that the information you entered is correct; it could be a simple mistake within the IP address or you did not select *Apply* on the Log Settings page which would not apply the settings.
- use `execute ping` to see if you can ping to the log device; if unable to ping to the log device check to see if the log device itself working.
- if the unit has stopped logging, see [“The FortiGate unit stopped logging: what happened?” on page 802](#).

The alert email did not send an email to the email address

Verify that the alert email configuration settings are correct; an alert email will not be sent if the email address requires authentication and *Authentication* and its settings are not configured. Use *Test Connectivity* after verifying the alert email configuration settings see if a test alert email is sent to the address or addresses.

The FortiGate unit stopped logging: what happened?

If your FortiGate unit contains an SQL database, and when you log into the web-based manager you see an SQL database error message, it is because the SQL database has become corrupted. You must make sure to back up your logs at that point and then rebuild the database.

If the FortiGate unit stopped logging to a device, test the connection between both the FortiGate unit and device using the `execute ping` command. The log device may have been turned off, is upgrading to a new firmware version, or just not working properly. After determining that there are no logs being sent, see [“Cannot log to a supported log device” on page 802](#).

Other diagnose commands

Diagnose and get commands are available on all FortiGate units. They can help you troubleshoot network activity. The `diag debug` command mentioned earlier is only one of many useful diagnose commands.

For additional diagnostic commands, see [“Troubleshooting ‘get’ commands” on page 805](#).



Troubleshooting 'get' commands

This section lists CLI `get` commands that you can use to troubleshoot problems with your FortiGate unit. This is not an exhaustive list, the grammar of the commands will change without notice, and this list does not include any diagnose commands.

Each CLI command includes its scope when VDOMs are enabled on the FortiGate unit as either VDOM or global. This indicates where in the CLI you must be to run the command. For more detailed information on the areas of FortiOS these commands relate to, see the corresponding sections in the FortiOS Handbook.

This section contains the following topics:

<code>exec tac report</code>	<code>get system session-helper</code>
<code>get firewall iprope appctrl</code>	<code>get system session-info full-stat</code>
<code>get firewall iprope list</code>	<code>get system session-info list</code>
<code>get firewall proute</code>	<code>get system session-info ttl</code>
<code>get firewall shaper</code>	<code>get system startup-error-log</code>
<code>get hardware memory</code>	<code>get system status</code>
<code>get hardware memory</code>	<code>get test</code>
<code>get hardware nic</code>	<code>get test urlfilter</code>
<code>get hardware npu list</code>	<code>get vpn ipsec stats crypto</code>
<code>get hardware npu performance</code>	<code>get vpn ipsec tunnel details</code>
<code>get hardware npu status</code>	<code>get vpn ipsec tunnel summary</code>
<code>get hardware status</code>	<code>get vpn status ssl hw-acceleration-status</code>
<code>get ips session</code>	<code>get vpn status ssl list</code>
<code>get router info kernel</code>	<code>get webfilter ftgd-statistics</code>
<code>get router info routing-table all</code>	<code>get webfilter status</code>
<code>get system arp</code>	
<code>get system auto-update status</code>	
<code>get system auto-update versions</code>	
<code>get system ha status</code>	
<code>get system performance firewall</code>	
<code>get system performance status</code>	
<code>get system performance top</code>	

exec tac report

Displays in-depth debugging information including hardware and software status as well as information, configuration, and debug outputs.

This command runs a series of other CLI get, show, and diagnose commands that include:

```
get system status
get system performance status
show system interface
diagnose ip address list
show full-configuration system dns
show full-configuration system global
show full-configuration system settings
diagnose hardware lspci -v
get hardware memory
get hardware memory
diagnose hardware sysinfo shm
diagnose ip arp list
get router info kernel
diagnose ip router command show show int
diagnose ipv6 neighbor-cache list
diagnose ipv6 route list
diagnose ipv6 ipv6-tunnel list
diagnose ipv6 sit-tunnel list
diagnose ips anomaly list
diagnose ips anomaly status
diagnose ips dissector status
diagnose ips packet status
diagnose ips raw status
get ips session
diagnose sys session6 stat (similar ipv4 output from get system
    session-info full-stat)
get system auto-update status
get system auto-update versions
diagnose test update info
diagnose sys flash list
diagnose sys logdisk smart
diagnose sys logdisk status
diagnose sys ha status
diagnose sys ha showcsum
diagnose sys ha hadiff status
diagnose sys ha dump 1 - 8
get sys session-info statistics (similar output to get system
    session-info full-stat)
```

Syntax

```
exec tac report
```

Parameters

None.

Usage/Remarks

This command provides extremely detailed information about your FortiGate unit. Fortinet support may request that you run this command and send them the output when they are troubleshooting a FortiGate problem with you.

When you run this command, it will take a few minutes to collect and display all the information. Your console program (telnet, ssh, or other such program) should log the output to file. Otherwise, the output will scroll by too fast to see.

No output is listed for this command because the output is extensive, and many of the commands from this report are described elsewhere in this section.

Do not include the output from this command in FortiCare tickets unless it is specifically requested by support personnel.

Scope: Global

get firewall iprope appctrl

Displays the list of applications defined in the application control list.

Syntax

```
get firewall iprope appctrl list
get firewall iprope appctrl status
```

Parameters

None.

Usage/Remarks

Use this command to view the rules defined for peer to peer, or the status of those rules.

Scope: Vdom

Output Example

```
FGT # get firewall iprope appctrl list
app-id=17953      list-id=2004  action=Pass
app-id=17954      list-id=2004  action=Pass
app-id=17956      list-id=2004  action=Pass
app-id=17957      list-id=2004  action=Pass
app-id=107347980  list-id=2004  action=Pass
app-id=108855300  list-id=2004  action=Pass
app-id=109051910  list-id=2004  action=Pass
app-id=109051912  list-id=2004  action=Pass
```

```
FGT # get firewall iprope appctrl status
appctrl table 3 list 1 app 1083 shaper 0
```

Keyword/Variable	Description
list	Display all the application IDs, which list they are in by ID, and the action taken for each application.
status	Displays the number of ipropes for application control tables, lists, applications, and traffic shapers.

get firewall iprope list

Displays the rules defined in the selected policy group. If no parameter is entered, all the rules of every policy group are displayed. The information is extensive details about the selected firewall policy.

Syntax

```
get firewall iprope list <policy group number>
```

Parameters

<policy group >	The number of the policy group as defined in the web-based manager, in <i>Policy > Policy > Policy</i> .
------------------------------	--

Usage/Remarks

Use this command to view the view the rules defined for a policy group. This is useful to understand the behavior of the policy group per protocol.

The policy group number is easiest to find by just running this command without the number (get firewall iprope list) and finding the group that includes your policy. For example if you have a traffic shaper for your security policy called shaper_test, look for that information in one of the policy group entries to ensure it is the correct one. After that you can use the policy group number to limit the output to only the group you want.

Scope: Vdom

Output Example

```
FGT # get firewall iprope list 0020005
policy flag (10809): log redir d_rm master
flag2 (0): shapers:
shaper_test(3/12800/64000)/shaper_test(3/12800/64000) per_ip=
imflag: sockport: 0 action: accept index: 1
schedule(always) group=00200005 av=00004e20 au=00000000 host=0
split=00000000
chk_client_info=0x0 app_list=0 ips_view=0 misc=0 grp_info=0 seq=0
hash=0
tunnel=
zone(0): ->zone(0):
source(0):
dest(0):
source wildcard(0):
destination wildcard(0):
vip(2): 2 1
service(4):
[1:0x0:0/(0,65535)->(8,8)]
[6:0x0:0/(0,65535)->(443,443)]
[6:0x0:0/(0,65535)->(22,22)]
[6:0x0:0/(0,65535)->(80,80)]
nat(0):
```

mms: 0 0

policy flag (10809): log redir d_rm master	A series of options that are either on or off, including — log, redir, d_rm, master, pol_stats, auth, a_i, nlb, and nat. The number (10809) represents what flags are used. Using 10809 as an example, log redir, d_rm, and master flags are set.
flag2 (0): shapers: shaper_test(3/12800/64000)/shaper_test(3/12800/64000) per_ip=	Flags for various things including fc_chk and disclaimer. Shapers lists the configured traffic shapers by name, separated for shared and per ip shapers.
imflag: sockport: 0 action: accept index: 1	action — This is the action the firewall policy will take, defined as one of ACCEPT, DENY, IPSEC, or SSL-VPN.
tunnel=	
source(0): dest(0):	The source or destination IP address or IP address range. For example an entry of 0.0.0.0-255.255.255.255 would be an address of all.
service(4):	This is a list of the services for this policy. The number is the number of services with entries that immediately follow.
[1:0x0:0/(0,65535)->(8,8)]	The first service entry, in this case from a list of four. It covers all ports on the source that are going to port 8 on the destination. If there was a range of ports, for example port 8000-8888, the entry would be ->(8000,8888).
nat(0):	
mms: 0 0	These values are only used for MMS related firewall entries. Only FortiOS Carrier support MMS traffic.

get firewall proute

Displays configured policy routes.

Syntax

```
get firewall proute
```

Parameters

None.

Usage/Remarks

Use this command to view the policy routes configured on this FortiGate unit's current VDOM.

Scope: Vdom

Output Example

```
FGT # get firewall proute

list route policy info(vf=root):

iff=12 src=10.10.10.0/255.255.255.0 tos=0x00 tos_mask=0x00
dst=10.10.11.0/255.255.255.0 protocol=11 port=1:65535
oif=13 gwy=0.0.0.0
```

Keyword/Variable	Description
vf	The current virtual domain (VDOM). All policy routes for this VDOM are displayed.
iff	Incoming interface
src	The source IP and netmask for incoming traffic to be matched to the policy
tos	Type of service bit pattern in hexadecimal. This is matched and is good for a specific TOS.
tos_mask	Type of service bit mask in hexadecimal. Masks out unwanted bits. This is good for matching multiple TOS values.
dst	The destination IP and netmask for incoming traffic to be matched to the policy.
protocol	The protocol number to be matched.
port	The range of ports to match for incoming traffic.
oif	Outgoing interface where traffic is being directed to.
gwy	Outgoing gateway where traffic is being directed to

get firewall shaper

Displays traffic shaper bandwidths or lists configured traffic shapers.

Syntax

```
get firewall shaper [ per-ip| per-ip-shaper| traffic| traffic-shaper ]
```

Parameters

Keyword/Variable	Description
per-ip	Displays all configured per-ip traffic shapers with their configurations.
per-ip-shaper	Lists all configured per-ip traffic shapers.
traffic	Displays all configured shared traffic shapers with their configurations including current bandwidth usage.
traffic-shaper	Lists all configured shared traffic shapers.

Usage/Remarks

Use this command to view traffic shaper bandwidths with traffic levels, and list configured traffic shapers for either shared (traffic) or per-ip traffic shapers.

For per-ip and traffic commands, if there is no shaper configured none will be displayed.

Scope: Vdom

Output Example

```
FGT # get firewall shaper per-ip
```

```
[ip_test]  
maximum-bandwidth: 12 KB/sec  
maximum-concurrent-session: 100  
packets dropped: 0
```

Keyword/Variable	Description
ip_test	The name of the configured per-ip traffic shaper.
12 KB/sec	The maximum allowed bandwidth for this traffic shaper. The web-based manager units for this field are kbits/sec where this value is in kBytes/sec. Any traffic for this IP address that is over this limit will be dropped. For example entering 100 kbits/sec results in a value of 12 kBytes/sec being displayed here.
100	The maximum allowed number of concurrent sessions for this traffic shaper. If the maximum is reached, no new sessions will be started on this IP address until another session ends.
0	The number of packets that have been dropped from traffic using this shaper. Packets at rates greater than the maximum bandwidth limit are dropped.

```
FGT # get firewall shaper traffic
```

```
[guarantee-100kbps]
maximum-bandwidth: 131072 KB/sec
guaranteed-bandwidth: 12 KB/sec
current-bandwidth: 0 B/sec
priority: high
tos ff
packets dropped: 0
```

Keyword/Variable	Description
guarantee-100kbps	The name of the configured shared traffic shaper. There are a number of default configured traffic shapers that include — guarantee-100kbps, high-priority, low-priority, medium-priority, and shared-1M-pipe.
131072 KB/sec	The maximum allowed bandwidth for this traffic shaper. Any traffic for this shaper that is over this limit will be dropped. The web-based manager units for this field are kbits/sec where this value is in kBytes/sec. For example entering 1 048 576 kbits/sec results in a value of 131 072 kBytes/sec being displayed here.
12 KB/sec	The guaranteed minimum bandwidth ensures there will never be less than this amount of bandwidth available.
0 B/sec	The current amount of traffic using this traffic shaper. If this shaper is not assigned to any security policies this value will be zero. Note that the units here are B/sec and not the KB/sec of the other fields for this command.
high	The quality of service (QoS) level for traffic using this shaper. Values can be low, medium, or high. Higher priority traffic will be queued ahead of lower priority traffic to ensure its guaranteed minimum bandwidth is met when necessary.
tos ff	The type of service (TOS) used in the TCP packet headers. It is used to enact the selected priority.
0	The number of packets that have been dropped from traffic using this shaper. Packets at rates greater than the maximum bandwidth limit are dropped.

get hardware cpu

Displays the CPU information.

Syntax

```
get hardware cpu
```

Parameters

None.

Usage/Remarks

Use this command to view the specifications about CPUs on the FortiGate unit. Multiple CPUs will each have an entry. Apart from the CPU model information (vendor, cpu family, model name, MHz, and cache size) you can use the flags to determine what features are supported. For example mmx, and 3dnow indicate Single instruction, multiple data (SIMD) support.

Scope: Global

Output Example

```
FGT # get hardware cpu
processor: 0
vendor_id: GenuineIntel
cpu family: 6
model: 5
model name: Celeron (Covington)
stepping: 0
cpu MHz: 1200.078
cache size: 64 KB
fdiv_bug: no
hlt_bug: no
f00f_bug: no
coma_bug: no
fpu: yes
fpu_exception: yes
cpuid level: 2
wp: yes
flags: fpu vme de pse tsc msr pae mce cx8 sep mtrr pge mca cmov pat
clflush dts acpi mmx fxsr sse sse2 ss tm pbe
bogomips: 2392.06
```

get hardware nic

Displays information and statistics for the network interface card (NIC) specified.

Syntax

```
get hardware nic <interface>
```

Parameters

<interface>	Enter the interface name. For example, internal, wan1, wan2.
--------------------------	--

Usage/Remarks

Use this command to view the interface information and statistics. This is useful when debugging network problems on a particular interface, or providing Customer Support with hardware information about your FortiGate unit.

This command is the only way for you to see your MAC both its original value ([Permanent_HWaddr](#)) and if it has changed as its current value ([Current_HWaddr](#)) as well.

The output will be different for different NICs and may include more fields or different information fields.

Scope: Global

Output Example

```
FGT # get hardware nic wan1
Description      sundance Ethernet driver1.01+LK1.21
chip_id         6
IRQ             5
System_Device_Name wan1
Current_HWaddr   00:09:0f:78:71:32
Permanent_HWaddr 00:09:0f:78:71:32
State           up

Link            up
Speed           100
Duplex          full
FlowControl     Tx off, Rx off
MTU_Size        1500

Rx_Packets      52377
Tx_Packets      53098
Rx_Packets      47029877
Tx_Bytes        7199983
Collisions      0
Rx_Missed_Errors 0
Tx_Carrier_Errors 0
```

Keyword/Variable	Description
Description	Name of the network driver.
chip_id	Id of the chipset

Keyword/Variable	Description
IRQ	IRQ
System_Device_Name	Network device name (i.e. The physical interface name).
Current_HWaddr	Current MAC address.
Permanent_HWaddr	Permanent MAC address.
State	Administrative status (up/down).
Link	Link status (up/down).
Speed	Negotiated or configured network speed.
Duplex	Negotiated or configured duplex mode.
Rx_Packets	Packets' number received by the network device.
Tx_Packets	Packets number transmitted by the network device.
Rx_Packets	Amount of bytes received on the interface.
Tx_Bytes	Amount of bytes sent from this interface.
Collisions	Number of collisions usually due mostly to incorrect incorrect speed or duplex settings.
Rx_Missed_Errors	Equals Rx_FIFO_Errors + CEXTERR (Carrier Extension Error Count). Only valid in 1000M mode, which is marked by PHY.
Tx_Carrier_Errors	The PHY should assert the internal carrier sense signal during every transmission. Failure to do so may indicate that the link has failed, or the PHY has an incorrect link configuration. This register only increments if transmits are enabled. This register is not valid in internal SerDes1 mode (TBI mode for the 82544GC/EI), and is only valid when the Ethernet controller is operating at full duplex.

get hardware memory

Displays the status of the system memory resources.

Syntax

```
get hardware memory
```

Parameters

None.

Usage/Remarks

Use this command to view the status of the memory resources. This is useful to confirm and identify any potential memory leaks or even to simply confirm that reduced FortiOS performance is due to a shortage of free memory.

Scope: Global

Output Example

```
FGT # get hardware memory
```

```
total:      used:      free:  shared:  buffers: cached:  shm:
Mem: 261955584 97312768 164642816 0 217088 55382016 51322880
Swap:      0      0      0
MemTotal:      255816 kB
MemFree:      160784 kB
MemShared:      0 kB
Buffers:      212 kB
Cached:      54084 kB
SwapCached:      0 kB
Active:      22832 kB
Inactive:      31524 kB
HighTotal:      0 kB
HighFree:      0 kB
LowTotal:      255816 kB
LowFree:      160784 kB
SwapTotal:      0 kB
SwapFree:      0 kB
```

Keyword/Variable	Description
Mem	Memory size.
Swap	Amount of memory in the swap.
MemTotal	HighTotal + LowTotal = amount of memory available on the unit.
MemFree	Free memory on the unit.
SwapCached	Amount of memory out of the Swap, but remaining in the cache.
Active	Amount of memory recently used.
Inactive	Amount of memory which has not been used for a while.

Keyword/Variable	Description
HighTotal	Amount of memory which belongs to the zone ZONE_HIGHMEM.
LowFree	Amount of Free memory available in the zone ZONE_NORMAL.

get hardware npu list

Displays the network processing unit (NPU) devices and paired interface names.

This is important to know because the paired interfaces can offload their traffic from the main CPU to the NPU. This can only be accomplished with paired interfaces.

Syntax

```
get hardware npu { legacy | np1 | np2 | np4 } list
```

Parameters

{ legacy np1 np2 np4 }	Specify which level of NPU interfaces are to be displayed.
---------------------------------	--

Usage/Remarks

Use this command to view the NPU devices and port numbers. This is useful because some FortiOS products contain network processors. Network processor features, and offloading requirements vary by network processor model.

The output is grouped by the ID number. Each ID number has four interfaces associated with it. These are the interfaces that are available to create paired interfaces. This is because each such group is assigned to one NPU. Mixing interfaces from different NPUs is not allowed. For example, from the output example, you would not be able to pair port1 with port14. Use this information when configuring your paired interfaces for accelerated offloading.

If the specified version of NPU is not present on this FortiGate unit, a message indicating that will be displayed.

Scope: Global

Output Example

```
FGT # get hardware npu np1 list
ID   Interface
0    port9 port10
```

```
FGT # get hardware npu np2 list
ID PORTS
-- ----
0 port1
0 port2
0 port3
0 port4
```

```
ID PORTS
-- ----
1 port5
1 port6
1 port7
1 port8
```

```
ID PORTS
-- ----
```

```
2 port9
2 port10
2 port11
2 port12
```

```
ID PORTS
-- -----
3 port13
3 port14
3 port15
3 port16
```

Keyword/Variable	Description
ID	The ID of the NPU assigned to these ports.
Interface	The names of the interfaces in the NPU's paired group. An NPU can handle either two or four ports. Any incoming traffic on this group of ports that leaves the FortiGate unit on the same group of ports can be handled by the NPU. Otherwise, the FortiGate unit's CPU will be involved.
PORTS	The names of the interfaces in the NPU's group.

get hardware npu performance

Displays the NPU performance for ISCP2, messages, and NAT.

Syntax

```
get hardware npu { legacy | np1 | np2 | np4 }performance <dev_id>
```

Parameters

{ legacy np1 np2 np4 }	Specify which level of NPU interfaces are to be displayed.
<dev_id>	NPU ID number. If you are unsure you can enter a "?" to get a list of available valid NPU ID numbers.

Usage/Remarks

Use this command to view the performance numbers for NPU devices. This can be useful when you are checking NPU traffic—either for specific problems or for network optimization.

Network processor features, and therefore offloading requirements, vary by network processor model.

The values displayed are divided into two sections. The first is the ISCP2 performance numbers. In this section, the values are mostly counts for the stated values. For example BADCSUM is the number of bad checksums encountered.

The last section starts with IRQ. This is a list of 18 different registers.

Scope: Global

Output Example

```
FGT # get hardware npu np2 performance 1
ISCP2 Performance:
Nr_int      : 0x00000005    INTwoInd   : 0x00000000    RXwoDone   :
0x00000000
PKTTwoEnd   : 0x00000000    PKTCSErr  : 0x00000000
PKTidErr    : 0x00000000    PHYInt    : 0x0/0x0/0x0/0x0
CSUMOFF     : 0x00000000    BADCSUM   : 0x00000000    MSGINT     :
0x00000000
IPSEC       : 0x00000000    IPSVLAN   : 0x00000000    SESMISS    :
0x00000000
TOTUP       : 0x00000000    TCPPLTOT  : 0x00000000    TCPPLBADCT:
0x00000000
TCPPLBADL   : 0x00000000    TCPPLSESI : 0x00000000    TCPPLSESR  :
0x00000000
TCPPLBD     : 0x00000000    TXInt     : 0x0/0x0/0x0/0x0 RxI : 0x0
BDEmpty     : 0x0/0x0/0x0/0x0 Congest: 0x0/0x0/0x0/0x0
TMM_Busy    : 0x0
Poll List: 0 0 0 0
MSG Performance:
TOTMSG      : 0x00000000    BADMSG     : 0x00000000    TOUTMSG    :
0x00000000    MSGLostEvent : 0x00000000
QUERY       : 0x00000000    TAE        : 0x00000000
```

```

SAEXP-SN : 0x00000000    SAEXP-TRF : 0x00000000    OUTUPD    :
0x00000000 INUPD        : 0x00000000
NULLTK: 0x00000000
NAT Performance: BYPASS (Enable) BLOCK (Disable)

IRQ  :00000005  QFTL :00000000  DELF :00000000  FFTL :00000000
OVTH :00000005  QRYF :00000000  INSF :00000000  INVC :00000000
ALLO :00000005  FREE :00000005  ALLOF :00000000 BPENTR:00000000  B
KENTR:00000000
PBPENTR:00000000  PBKENTR:00000000  NOOP :00000000  THROT :00000000
0 (0x002625a0)
SWITOT:00000000  SWDTOT:00000000  ITDB:00000000  OTDB:00000000
SPISES:00000000  FLUSH:00000021

```

Keyword/Variable	Description
ISCP2 Performance	This section displays information about the ISCP2 performance that includes packet, TCP, and IPsec counts and errors.
MSG Performance	This section displays information about message traffic performance.
NAT Performance	This section displays information about NAT traffic performance. If BYPASS is enabled, many of the counters will be zero.

get hardware npu status

Displays the specific network processor unit (NPU) status. NPUs are used to accelerate network traffic passing through the FortiGate unit by offloading work from the CPU. More advanced NPUs have higher numbers.

Syntax

```
get hardware npu { legacy | np1 | np2 | np4 }status <dev_id>
```

Parameters

{ legacy np1 np2 np4 }	Specify which level of NPU interfaces are to be displayed.
<dev_id>	NPU ID number. If you are unsure you can enter a "?" to get a list of available valid NPU ID numbers.

Usage/Remarks

Use this command to display the status of the interfaces attached to the NPU. This can be useful for advanced users to debug traffic on the NPU interfaces.

If you enter an NPU or device ID that doesn't exist, an error message stating that it was an invalid NPU ID will be displayed.

Scope: Global

Output Example

```
FGT # get hardware npu np2 status 1
NP2 Status

ISCP2 c2160000 (Neighbor f7940000) 1a29:0703 256MB Base f8a5a000
DBG 0x00000000
RX SW Done 0 MTP 0x0
desc_alloc = c2152000
desc_size = 0x2000 count = 0x100
nxt_to_u = 0x0 nxt_to_f = 0x0
Total Interfaces: 4 Total Ports: 4
Number of Interface In-Use: 4
Interface c2160100 netdev c22bb000 0 Name port5
PHY: Attached
LB Mode 0 LB IDX 0/1 LB Ports: c2160644, 00000000, 00000000,
00000000
Port c2160644 Id 0 Status Down ictr 4
desc = c2232000
desc_size = 0x00001000 count = 0x00000100
nxt_to_u = 0x00000000 nxt_to_f = 0x00000000
Intf c2160100
Interface c2160250 netdev c2168c00 1 Name port6
PHY: Attached
LB Mode 0 LB IDX 0/1 LB Ports: c21606f8, 00000000, 00000000,
00000000
Port c21606f8 Id 1 Status Down ictr 0
desc = c2126000
desc_size = 0x00001000 count = 0x00000100
```

```
nxt_to_u    = 0x00000000 nxt_to_f = 0x00000000
Intf c2160250
Interface c21603a0 netdev c2168800 2 Name port7
PHY: Attached
LB Mode 0 LB IDX 0/1 LB Ports: c21607ac, 00000000, 00000000,
00000000
Port c21607ac Id 2 Status Down ictr 0
desc = c2123000
desc_size   = 0x00001000 count   = 0x00000100
nxt_to_u    = 0x00000000 nxt_to_f = 0x00000000
Intf c21603a0
Interface c21604f0 netdev c2168400 3 Name port8
PHY: Attached
LB Mode 0 LB IDX 0/1 LB Ports: c2160860, 00000000, 00000000,
00000000
Port c2160860 Id 3 Status Down ictr 0
desc = c2122000
desc_size   = 0x00001000 count   = 0x00000100
nxt_to_u    = 0x00000000 nxt_to_f = 0x00000000
Intf c21604f0
NAT Information:
cmdq_qw     = 0x2000 cmdq       = c2140000
head        = 0x5 tail         = 0x5
APS (Enabled) information:
Session Install when TMM TSE OOE: Disable
Session Install when TMM TAE OOE: Disable
IPS anomaly check policy: Follow config
MSG Base = c2130000 QL = 0x1000 H = 0x0
```

get hardware status

Displays basic hardware information about the FortiGate unit.

Syntax

```
get hardware status
```

Parameters

None.

Usage/Remarks

Use this command to get basic information about your FortiGate unit's hardware.

This is a short list that can be used to help guide troubleshooting efforts, especially by customer service. For example, if there is supposed to be compact flash memory available but this command shows 0MB available, then you know there is a hardware problem with the memory. Also, if there are known issues with particular chipsets, this command will display chipset information for hardware such as network cards.

The CPU information can vary from simply the number of cores, to the more in-depth information listed here.

Scope: Global

Output Example

```
FGT # get hardware status
Model name: Fortigate-3810A
ASIC version: CP6
ASIC SRAM: 64M
CPU: Intel(R) Core(TM)2 Duo CPU      E4300  @ 1.80GHz
RAM: 3532 MB
Compact Flash: 122 MB /dev/hda
USB Flash: not available
Network Card chipset: Broadcom 570x Tigon3 Ethernet Adapter
(rev.0x8003)
Network Card chipset: Intel(R) PRO/1000 Network Connection
(rev.0006)Related Commands
```

get ips session

Displays the IPS session status. An Intrusion prevention system (IPS) is a network security device that monitors network and/or system activities for malicious or unwanted behavior and can react, in real-time, to block or prevent those activities.

Syntax

```
get ips session
```

Parameters

None.

Usage/Remarks

Viewing the IPS session status allows you to see if traffic is checked by the IPS engine and if the IPS engine is working as expected. You can see the type of traffic and how many sessions are being processed by the IPS engine. You can also see how much memory the IPS engine is using. The displayed information is repeated with short delays to provide a time-lapse view of the IPS sessions status.

Scope: Vdom

Output Example

```
FGT # get ips session
```

```
SYSTEM:
memory capacity      73400320
memory used          3982780
recent pps\bps       0\0K
session in-use       0
TCP: in-use\active\total 0\0\0
UDP: in-use\active\total 0\0\0
ICMP: in-use\active\total 0\0\0
```

Keyword/Variable	Description
memory capacity	Total memory capacity in system available to IPS.
memory used	Memory used by IPS sessions. If this value is high, close to the memory capacity, it will likely result in dropped packets and other issues.
recent pps\bps	Recent IPS traffic in packets per second (pps) and bits per second (bps). Use these numbers to determine the level of IPS traffic. If they are very high, there will likely be problems with the FortiGate unit's performance.
session in-use	Current number of IPS sessions in use.
TCP: in-use\active\total	The number of in use, active, and total Transmission Control Protocol (TCP) packets sent over this session.
UDP: in-use\active\total	The number of in use, active, and total User Datagram Protocol (UDP) packets sent over this session.
ICMP: in-use\active\total	The number of in use, active and total Internet Control Message Protocol (ICMP) packets sent over this session.

get router info kernel

Displays the entries in the kernel routing table.

Syntax

```
get router info kernel
```

Parameters

None.

Usage/Remarks

Use this command to view the kernel routing table.

Computers and network devices connected to the Internet use routing tables to compute the next hop destination for a packet. The routing kernel table is an table of routes to particular network destinations that is stored in a router or a networked computer. This information contains the topology of the network immediately around it. Building the routing table is the primary goal of routing protocols and static routes.

In the routing table, the information that displays allows you to see the type of route, the protocol used, and the priority of each entry. You can use [“get router info routing-table all” on page 828](#) for more information on each routing entry.

Scope: Vdom

Output Example

```
FGT # get router info kernel  
tab=254 vf=2 scope=0 type=1 proto=11 prio=0 0.0.0.0/0.0.0.0/0-  
>192.168.0.0/24 pref=0.0.0.0 gwy=9.1.1.1 dev=5(port1)
```

Keyword/Variable	Description
192.168.0.0/24	The destination network or destination host.
gwy=9.1.1.1	The gateway address for the next hop.
dev=5(port1)	Interface to which packets for this route is sent.

get router info routing-table all

Displays the routing table.

Syntax

```
get router info routing-table all
```

Parameters

None.

Usage/Remarks

Use this command to view the routing table. This routing table may be larger and have more entries than the kernel routing table when using the `get router info kernel` command.

Computers and network devices connected to Internet use routing tables to compute the next hop destination for a packet. It is an table of routes to particular network destinations that is stored in a router or a networked computer. This information contains the topology of the network immediately around it.

Scope: Vdom

Output Example

```
FGT# get router info routing-table all

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type
       2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-
IS inter area
       * - candidate default

S*      0.0.0.0/0 [10/0] via 172.20.120.2, wan1
C       10.31.101.0/24 is directly connected, internal
C       172.20.120.0/24 is directly connected, wan1
```

The output of this command is fairly self-explanatory; however, if you select a specific routing protocol (static, rip, etc) when entering the command, only that part of this larger list of entries will be displayed.

The default route is indicated by the asterisk (*).

How a route is connected to the FortiGate unit is important in determining priority. The two C entries state "is directly connected", which indicates the lowest possible hop count to get there.

get system arp

Displays the IPv4 ARP table.

Syntax

```
get system arp
```

Parameters

None.

Usage/Remarks

This command is useful to view or alter the contents of the kernel's ARP tables. For example, when you suspect a duplicate Internet address is the cause for some intermittent network problem.

This command is not available in multiple VDOM mode.

Scope: Vdom

Output Example

```
FGT # get system arp
Address    Age(min) Hardware Addr    Interface
172.20.120.16  0    00:0d:87:5c:ab:65  internal
172.20.120.138 0    00:08:9b:09:bb:01  internal
```

Keyword/Variable	Description
Address	The IP address that is linked to the MAC address. The default is 0.0.0.0
Age	Current duration of the ARP entry in minutes. The default is 0.
Hardware Addr	The hardware, or MAC address, to link with this IP address. The default is 00:00:00:00:00:00:
Interface	The physical interface of the FortiGate unit where the address is connected.

get system auto-update status

Use this command to display information about the status FortiGuard updates on the FortiGate unit.

Syntax

```
get system auto-update status
```

Parameters

None.

Usage/Remarks

Use this command when you need to view the current antivirus and IPS status from the FortiGate unit.

The status command is used to test the current connectivity status, and can retrieve configuration information, such as push updates, update schedules, or other parameters related to the FortiGuard service operation.

The versions command provides extended information about each FortiGuard component and its version, build number, contract expiry date, last update attempts and results.



Result: `Connectivity failure` indicates the connections to the FortiGuard servers is not possible. This may be a serious problem that needs to be addressed. Until it is solved, you may not receive any FortiGuard updates leaving your network vulnerable.

These commands also allow the user to check whether the FortiGate unit is running the latest AV and IPS packages.

Scope: Global

Output Example

```
FGT # get system auto-update status
FDN availability:  available at Mon May 26 20:16:43 2008

Push update:  disable
Scheduled update:  enable
    Update every:  1 hours at 16 minutes after the hour
Virus definitions update:  enable
IPS definitions updates:  enable
Server override:  disable
Push address override:  disable
```

Web proxy tunneling: disable

Keyword/Variable	Description
FDN availability	Specify availability status and last access time (access time corresponds to the scheduled update settings). Possible values are: available/unavailable.
Push update	Specify whether push update method is enabled or disabled. Possible values are: enable/disable
Scheduled update	Specify whether scheduled update is enabled or disabled. Possible values are: enable/disable.
Update every	If scheduled update is enabled, specify the time defined to launch the update.
IPS definitions updates	Specify whether the IPS definitions update is enabled or disabled. Possible values are: enable/disable.
Server override	Specify whether the use of another FDS server is enabled or disabled. Possible values are: enable/disable. If enabled a new line is displayed showing the FDS IP address defined in the configuration. For example: <pre>Server override: enable Server: 10.0.0.1</pre>
Push address override	If push update is enabled, specify whether the Fortigate override address feature is enabled or disabled. Possible values are: enable/disable. If enabled, a new line is displayed showing the FDS IP address and the TCP port (a.b.c.d:port) defined in the configuration. Example: <pre>Push address override: enable Address: 10.0.0.2:9443</pre>
Web proxy tunneling	Specify whether FortiGate device is using a proxy to retrieve AV and IPS definitions updates. Possible values are: enable/disable. If enabled, additional lines are displayed showing the proxy settings. Example: <pre>Web proxy tunneling: enable Proxy address: 10.0.0.3 Proxy port: 8890 Username: foo Password: foo</pre>

get system auto-update versions

Displays antivirus and IPS engines and definitions, and license information.

Syntax

```
get system auto-update versions
```

Parameters

None.

Usage/Remarks

Use this command to view the antivirus and IPS engines and definitions and license information. This is useful in determining when the antivirus engine, virus definitions, attack definitions, IPS attack definitions, AS Rule set, AS engine, and FDS address were last updated, as well as when their contracts expire, which version they are using, and the result of the last update.

Scope: Global

Output Example

```
FGT # get system auto-update versions

AV Engine
-----
Version: 3.003
Contract Expiry Date: n/a
Last Updated using manual update on Wed Jan  9 18:26:00 2008
Last Update Attempt: n/a
Result: Updates Installed

Virus Definitions
-----
Version: 8.631
Contract Expiry Date: n/a
Last Updated using manual update on Tue Jan 15 14:27:00 2008
Last Update Attempt: n/a
Result: Updates Installed

Attack Definitions
-----
Version: 2.461
Contract Expiry Date: n/a
Last Updated using manual update on Fri Jan 18 11:23:00 2008
Last Update Attempt: n/a
Result: Updates Installed

IPS Attack Engine
-----
Version: 1.091
Contract Expiry Date: n/a
Last Updated using manual update on Wed Jan  9 18:22:00 2008
Last Update Attempt: n/a
```

Result: Updates Installed

FDS Address

Keyword/Variable	Description
Version	Version number of the engine or the definitions.
Contract Expiry Date	The date the contract expires.
Last Updated using manual update on	Date of the last manual update.
Last Update Attempt	The date when the last update was attempted.
Result	The status of the last update.

get system ha status

Use this command to display information about an HA cluster. The command displays general HA configuration settings. The command also displays information about how the cluster unit that you have logged into is operating in the cluster.

Syntax

```
get system ha status
```

Parameters

Scope: Global.

Usage/Remarks

Usually you would log in to the primary unit CLI using SSH or telnet. In this case the `diagnose system ha status` command displays information about the primary unit first, and also displays the HA state of the primary unit (the primary unit operates in the work state).

For a virtual cluster configuration, the `diagnose system ha status` command displays information about how the cluster unit that you have logged into is operating in virtual cluster 1 and virtual cluster 2. For example, if you connect to the cluster unit that is the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2, the output of the `diagnose system ha status` command shows virtual cluster 1 in the work state and virtual cluster 2 in the standby state. The `diagnose system ha status` command also displays additional information about virtual cluster 1 and virtual cluster 2.

See the [FortiGate CLI Reference Guide](#) or the [High Availability handbook chapter](#) more information.

Scope: Global

Output Example

```
FGT # get system ha status
Model: 5000
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4 FG50012204400045 1
Slave :100 5001_Slot_3 FG50012205400050 0
number of vcluster: 1
vcluster 1: work 10.0.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
```

Keyword/Variable	Description
Model	The FortiGate model number.
Mode	The HA mode of the cluster: a-a or a-p. If the unit is not in HA mode, standalone will be displayed.

Keyword/Variable	Description
Group	The group ID of the cluster.
Debug	The debug status of the cluster.
ses_pickup	<p>The status of the session pick-up feature, also known as session failover.</p> <p>Session failover means that a cluster maintains active network TCP and IPsec VPN sessions after a device or link failover. Session failover does not failover UDP, multicast, ICMP, or SSL VPN sessions. In some cases UDP sessions may be maintained after a failover.</p>
load_balance	The status of the <code>load-balance-all</code> keyword: enable or disable. Relevant to active-active clusters only.
schedule	The active-active load balancing schedule. Relevant to active-active clusters only.
Master Slave	<p>Master displays the device priority, host name, serial number, and cluster index of the primary (or master) unit.</p> <p>Slave displays the device priority, host name, serial number, and cluster index of the subordinate (or slave, or backup) unit or units.</p> <p>The list of cluster units changes depending on how you log into the CLI. Usually you would use <code>ssh</code> or <code>telnet</code> to log in to the primary unit CLI. In this case, the primary unit would be at the top the list followed by the other cluster units.</p> <p>If you use the <code>execute ha manage</code> command or a console connection to log in to a subordinate unit's CLI, and then enter <code>diagnose system ha status</code>, the subordinate unit that you have logged into appears at the top of the list of cluster units.</p>
number of vcluster	The number of virtual clusters. If virtual domains are not enabled, the cluster has one virtual cluster. If virtual domains are enabled, the cluster has two virtual clusters.

get system performance firewall

Displays the status of important software and hardware systems on your FortiGate unit.

Syntax

```
get system performance firewall packet-distribution
get system performance firewall statistics
```

Parameters

None

Usage/Remarks

Use this command to quickly see information about traffic through the firewall.

The `packet-distribution` command divides network traffic into ten different packet sizes and lists the number of packets received in each category. This can help you spot hacking attempts, or optimize your network.

The `statistics` command provides packet and byte counts for different major protocols and packet types through the firewall. Packet types include TCP, UDP, ICMP, and IP.

Scope: All (Global and Vdom)

Output Example

```
FGT# get sys per firewall packet-distribution
getting packet distribution statistics...
0 bytes - 63 bytes: 3624747 packets
64 bytes - 127 bytes: 802612 packets
128 bytes - 255 bytes: 371774 packets
256 bytes - 383 bytes: 712180 packets
384 bytes - 511 bytes: 30691 packets
512 bytes - 767 bytes: 57220 packets
768 bytes - 1023 bytes: 23460 packets
1024 bytes - 1279 bytes: 3055 packets
1280 bytes - 1500 bytes: 64 packets
> 1500 bytes: 0 packets

FGT# get sys per firewall statistics
getting traffic statistics...
Browsing: 708624 packets, 408018318 bytes
DNS: 411789906583486464 packets, 0 bytes
E-Mail: 0 packets, 0 bytes
FTP: 0 packets, 0 bytes
Gaming: 0 packets, 0 bytes
IM: 0 packets, 0 bytes
Newsgroups: 0 packets, 0 bytes
P2P: 0 packets, 0 bytes
Streaming: 0 packets, 0 bytes
TFTP: 201239863725391872 packets, 56530359549952 bytes
VoIP: 3543070 packets, 25 bytes
Generic TCP: 55834574848 packets, 1786706395136 bytes
Generic UDP: 0 packets, 0 bytes
```

```
Generic ICMP: 0 packets, 0 bytes  
Generic IP: 0 packets, 0 bytes
```

get system performance status

Displays the status of important software and hardware systems on your FortiGate unit.

Syntax

```
get system performance status
```

Parameters

None.

Usage/Remarks

Use this command to quickly see important information about your FortiGate unit's state. Information includes CPU usage, memory usage, network usage, number of sessions, viruses caught, IPS attacks blocked, and FortiGate unit uptime. These numbers provide a quick look at how the FortiGate unit is doing. If any one number needs attention, you can use other commands to get more information on that area.

For example, if the CPU states show 98 percent system, you know that your FortiGate unit is running at full capacity and you need to check the processes to see if any one process is using all the resources and if there are valid reasons for it.

Another example is that if memory usage is near 100 percent, AV failover may occur which, if enabled, may pass traffic without being scanned or refuse new connections.

Scope: All (Global and Vdom)

Output Example

```
FGT# get sys performance status
CPU states: 0% user 0% system 0% nice 100% idle
Memory states: 10% used
Average network usage: 0 kbps in 1 minute, 0 kbps in 10 minutes,
13 kbps in 30 minutes
Average sessions: 31 sessions in 1 minute, 30 sessions in 10
minutes, 31 sessions in 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 44 days, 18 hours, 42 minutes
```

get system performance top

Displays the processes running on your FortiGate unit.

Syntax

```
get system performance top <delay> <max_lines>
```

Parameters

delay	The amount of time, in seconds, in which the process information is polled. The default is 5 seconds.
max_lines	The maximum number of processes displayed in the output. The default is 20 lines.

Usage/Remarks

Use this command when you need to view the processes running and the information about each process. It displays as a static set of columns where the information changes in place. To exit this display, press <Ctrl-C>.

Scope: All (Global and Vdom)

Output Example

```
FGT # get system performance top 5 20
Run Time: 0 days, 1 hours and 53 minutes
1U, 1S, 97I; 248T, 99F, 73KF
newcli      113      S      1.0      2.1
  sshd      107      S      0.7      1.9
  newcli    114      R <      0.3      2.1
  thttp     48       S      0.0      5.4
ipsengine   55      S <      0.0      5.2
ipsengine   50      S <      0.0      5.2
cmdbsvr     18       S      0.0      3.7
httpsd      71      S      0.0      3.3
httpsd      92      S      0.0      3.3
httpsd      37      S      0.0      2.8
scanunitd   90      S <      0.0      2.2
scanunitd   91      S <      0.0      2.1
merged_daemons 45      S      0.0      2.1
  newcli    108      S      0.0      2.1
  updated   59      S      0.0      2.0
  newcli    112      S <      0.0      2.0
miglogd     35      S      0.0      1.9
  nsm       28      S      0.0      1.9
  imd       53      S      0.0      1.8
  authd     52      S      0.0      1.7
```

Keyword/Variable	Description
Run Time	Displays how long the FortiOS has been running as a string
U	User CPU usage (%)
S	System CPU usage (%)
I	Idle CPU usage (%)

Keyword/Variable	Description
T	Total memory
F	Free memory
KF	Kernel-free memory
Column 1	Process name
Column 2	Process identification (PID)
Column 3	One letter process status. S: sleeping process R: running process <: high priority
Column 4	CPU usage (%)
Column 5	Memory usage (%)

get system session-helper

Displays the session helper table IDs.

FortiGate units use session helpers to process sessions that have special requirements. Session helpers function like proxies by getting information from the session and performing support functions required by the session.

Syntax

```
get system session-helper
```

Parameters

None.

Usage/Remarks

Use this command to display if any of the pre-defined session-helpers are in use.

Scope: Global

Output Example

```
FGT # get system session-helper
== 1 ==
== 2 ==
== 3 ==
== 4 ==
== 5 ==
== 6 ==
== 7 ==
== 8 ==
== 9 ==
== 10 ==
== 11 ==
== 12 ==
== 13 ==
== 14 ==
== 15 ==
== 16 ==
== 17 ==
== 18 ==
== 19 ==
== 20 ==
```

Keyword/Variable	Description
1 . . . 20	If one of the pre-defined session helpers is in use, it will be listed here.

get system session-info full-stat

Displays the system's full session state.

Syntax

```
get system session-info full-stat
```

Parameters

None.

Usage/Remarks

Use this command to display in-depth information session info about the system's state. This includes session table size, expected session table size, session count, firewall error details, and more.

This data can provide an important picture of what is happening on your FortiGate unit, for example:

- if the ephemeral buffer is full, you have a very busy device which may indicate a DoS attack is under way. See [ephemeral=2/32752](#).
- if there are many sessions in a SYN_SENT state to the point of causing other problems, you may be experiencing a SYN flood type DoS attack. See [10 in SYN_SENT state](#).
- if the HTTP proxy is too busy to handle new connections, then the ACCEPT queue fills up, and resets if it completely fills up, resetting the user's connection. See [acceptqf=0](#).

Scope: Global

Output Example

```
FGT # get system session-info full-stat
session table: table_size=131072 max_depth=1 used=50
expect session table: table_size=2048 max_depth=1 used=2
misc info: session_count=167 setup_rate=0 exp_count=0 clash=0
memory_tension_drop=0 ephemeral=2/32752 removeable=134
ha_scan=0 delete=0 flush=0 dev_down=0/0
TCP sessions
 6 in ESTABLISHED state
10 in SYN_SENT state
25 in TIME_WAIT state
60 in CLOSE state
 6 in CLOSE_WAIT state
firewall error stat:
error1=00000000
error2=00000000
error3=00000000
error4=00000000
tt=00000000
cont=0000630f
ids_recv=0000dd8e
url_recv=00000000
av_recv=00000000
fqdn_count=00000000
```



```
tcp reset stat:  
syncqf=0 acceptqf=0 no-listener=0 data=0 ses=0 ips=0  
global: ses_limit=0 ses6_limit=0 rt_limit=0 rt6_limit=0
```

The session setup rate is the second item on the first line of output. The information provided by this command includes:

Keyword/Variable	Description
session table	
table_size=131072	The maximum number of entries possible in the session table.
max_depth=1	
used=50	The number of session entries in the session table
expect session table	
table_size=2048	The maximum number of entries possible in the expect session table.
max_depth=1	
used=2	The number of expect session entries in the expect session table.

Keyword/Variable	Description
misc info:	
session_count=167	The number of sessions in the kernel.
setup_rate=0	How fast sessions have been created.
exp_count=0	The number of expect sessions in the kernel.
clash=0	Count of the collisions that occurred during the creation of a new session.
memory_tension_dr op=0	Number of dropped sessions due to the system running out of memory.
ephemeral=2/32752	The ephemeral buffer is used to protect against DoS attacks. It is a type of input buffer so the real session table does not get overloaded if a DoS attack does happen. The first number is how many sessions are in use, and the second number is the maximum number allowed. If the two numbers are close, it is a good chance there is a DoS attack underway, such as a DDoS using UDP packets.

Keyword/Variable	Description
TCP sessions	During the lifetime of a TCP session, it changes state to reflect what stage it is at — starting a connection, established, or closing.
6 in ESTABLISHED state	The connection has been established and is ready for data transfer.

10 in SYN_SENT state	<p>The FortiGate unit is waiting for the remote end of the session to acknowledge the SYN that was sent.</p> <p>If there are many sessions in this state, a SYN flood DoS attack may be in progress.</p> <p>To reduce the number of sessions in the SYN_SENT state, or half open state, you can reduce the wait period using the CLI command:</p> <pre>config system global set tcp-halfopen-timer XX end</pre> <p>where xx is the number of seconds to wait for the peer to respond before closing the connection. The valid range is from 1 to 86 400, and the default is 60 seconds. Reducing this value will free up half-open sessions before the session table completely fills up.</p>
25 in TIME_WAIT state	<p>The state when a connection termination request is sent to wait for enough time to ensure the remote end received acknowledgement of the termination request.</p> <p>If there are too many sessions in TIME_WAIT state, you can reduce the wait period using the CLI command:</p> <pre>config system global set tcp-timewait-timer XX end</pre> <p>where xx is the number of seconds to wait before closing the session. The valid range is 0 to 300 seconds, with the default being 120 seconds. A value of zero indicates no wait.</p>
60 in CLOSE state	<p>The session is closed.</p>
6 in CLOSE_WAIT state	<p>The session is being closed by the FortiGate unit— there will be no more data from the sender. This waits for the last ACK before going to the CLOSE state.</p> <p>If there are a lot of sessions in this state, also called half closed state, you can reduce the timer using the CLI command:</p> <pre>config system global set tcp-halfclose timer XX end</pre> <p>where xx is the number of seconds to wait after one peer has sent a FIN packet for this session. The valid range is 0 to 86 400 seconds, with the default being 120 seconds.</p>

Keyword/Variable	Description
firewall error stat:	
error1=00000000	This indicates there was a mismatch between the encryption state of the packet received and the session the packet belongs to. Either the packet was plain text when the session says it should have been decrypted (such as during a spoof), or the packet was decrypted properly but the session has no associated IPsec tunnel.

Keyword/Variable	Description
tcp reset stat:	
syncqf=0	<p>SYN queue full — Reset happens when the syn queue is full. FortiGate sends a reset (RST) when any of the following occur:</p> <ul style="list-style-type: none"> • syn queue full • accept queue full • FortiGate unit goes into conserve mode • CPU use is too high (queues aren't served correctly, and backlog increases) <p>When this happens, the remote end will TCP timeout and close the TCP session, while on the FortiGate side the TCP session will go to a CLOSE_WAIT state.</p>
acceptqf=0	<p>ACCEPT queue full — When a new packets comes in for a new session before being transmitted to the application layer of the device they are stored in a kernel queue.</p> <p>When this queue gets full, it means that the application layer was not able to process the packet. This is a symptom of an overloaded FortiGate unit. This happens when the new session rate of traffic being handled by proxy is too high.</p>
no-listener=0	Packet was received, but there is no service to handle it.

get system session-info list

Displays detailed session configuration information for sessions.

Syntax

```
get system session-info list
```

Parameters

None.

Usage/Remarks

Use this command to tell you the default time to live setting for sessions. All sessions created will have this length of time before they expire and need to create a new session table entry.

Scope: Global

Output Example

```
FGT # get system session-info list
session info: proto=17 proto_state=00 duration=4685 expire=174
             timeout=0 flags=00000000 sockflag=00000000 sockport=0
             av_idx=0 use=4
             origin-shaper=
             reply-shaper=
             per_ip_shaper=
             ha_id=0 hakey=30944
             policy_dir=0 tunnel=/
             state=local may_dirty rem
             statistic(bytes/packets/allow_err): org=39032/119/1 reply=0/0/0
             tuples=2
             orgin->sink: org pre->in, reply out->post dev=2->9/9->2
             gwy=255.255.255.255/0.0.0.0
             hook=pre dir=org act=noop 172.20.120.225:68-
             >255.255.255.255:67(0.0.0.0:0)
             hook=post dir=reply act=noop 255.255.255.255:67-
             >172.20.120.225:68(0.0.0.0:0)
             misc=0 policy_id=0 id_policy_id=0 auth_info=0 chk_client_info=0
             vd=0
             serial=0002dcfe tos=ff/ff ips_view=0 app_list=0 app=0 dd_type=0
             dd_rule_id=0
             per_ip_bandwidth meter: addr=172.20.120.225, bps=112
```

Keyword/Variable	Description
session info: proto=17	The protocol this packet contains. Common protocol numbers include 17 (UDP), 6 (TCP), 47(GRE), 89(OSPF), and 1 (ICMP). See Assigned Internet Protocol Numbers .

Keyword/Variable	Description																										
proto_state=00	<p>proto_state has two digits to keep track of the original direction and the reply direction. There are different states for each of TCP, SCTP, and UDP. ICMP traffic has no state and will always show as proto_state=00.</p> <p>For example if a TCP session has proto_state=67, the originator has closed the session, and the responder is waiting for the last ACK or halfclose timer to expire.</p> <table> <tr> <td>TCP packet states</td><td>SCTP packet states</td></tr> <tr> <td>0 - none</td><td>0 - SCTP_S_NONE</td></tr> <tr> <td>1 - ESTABLISHED</td><td>1 - SCTP_S_ESTABLISHED</td></tr> <tr> <td>2 - SYN_SENT</td><td>2 - SCTP_S_CLOSED</td></tr> <tr> <td>3 - SYN & SYN/ACK</td><td>3 - SCTP_S_COOKIE_WAIT</td></tr> <tr> <td>4 - FIN_WAIT</td><td>4 -</td></tr> <tr> <td>5 - TIME_WAIT</td><td>SCTP_S_COOKIE_ECHOED</td></tr> <tr> <td>6 - CLOSE</td><td>5 -</td></tr> <tr> <td>7 - CLOSE_WAIT</td><td>SCTP_S_SHUTDOWN_SENT</td></tr> <tr> <td>8 - LAST_ACK</td><td>6 -</td></tr> <tr> <td>9 - LISTEN</td><td>SCTP_S_SHUTDOWN_REC'D</td></tr> <tr> <td></td><td>7 - SCTP_S_ACK_SENT</td></tr> <tr> <td></td><td>8 - SCTP_S_MAX</td></tr> </table> <p>UDP packet states</p> <p>0 - Reply Not Seen</p> <p>1 - Reply Seen</p>	TCP packet states	SCTP packet states	0 - none	0 - SCTP_S_NONE	1 - ESTABLISHED	1 - SCTP_S_ESTABLISHED	2 - SYN_SENT	2 - SCTP_S_CLOSED	3 - SYN & SYN/ACK	3 - SCTP_S_COOKIE_WAIT	4 - FIN_WAIT	4 -	5 - TIME_WAIT	SCTP_S_COOKIE_ECHOED	6 - CLOSE	5 -	7 - CLOSE_WAIT	SCTP_S_SHUTDOWN_SENT	8 - LAST_ACK	6 -	9 - LISTEN	SCTP_S_SHUTDOWN_REC'D		7 - SCTP_S_ACK_SENT		8 - SCTP_S_MAX
TCP packet states	SCTP packet states																										
0 - none	0 - SCTP_S_NONE																										
1 - ESTABLISHED	1 - SCTP_S_ESTABLISHED																										
2 - SYN_SENT	2 - SCTP_S_CLOSED																										
3 - SYN & SYN/ACK	3 - SCTP_S_COOKIE_WAIT																										
4 - FIN_WAIT	4 -																										
5 - TIME_WAIT	SCTP_S_COOKIE_ECHOED																										
6 - CLOSE	5 -																										
7 - CLOSE_WAIT	SCTP_S_SHUTDOWN_SENT																										
8 - LAST_ACK	6 -																										
9 - LISTEN	SCTP_S_SHUTDOWN_REC'D																										
	7 - SCTP_S_ACK_SENT																										
	8 - SCTP_S_MAX																										
duration=4685	The number of seconds since this session was created.																										
expire=174	<p>The number of seconds until this session expires.</p> <p>GRE sessions have a very large expire value. This is to prevent them from timing out, and dropping the GRE session.</p>																										
timeout=0	If there is a timer that will expire, this is the number of seconds until that timer expires. One example is the standard session ttl timer is 3600 seconds.																										

Keyword/Variable	Description
origin-shaper= reply-shaper= per_ip_shaper=	<p>The traffic shapers applied to traffic from the origin, from the responder (reply), or any per IP traffic shapers.</p> <p>The value is the name of the shaper that has been applied.</p>
ha_id=0	The cluster ID. in an HA cluster.
tunnel=/ 	If this session is over a tunnel, this displays information about the tunnel.

Keyword/Variable	Description
state=local may_dirty rem	<p>The status of the session. Possible states include:</p> <p>br Session is being bridged, that is, in transparent mode.</p> <p>ext Session is created by a firewall session helper.</p> <p>log Session is being logged</p> <p>local Session is originating from, or destined for, a local stack.</p> <p>may_dirty Session is created by a policy. For example, the session for FTP channel control will have this state but the FTP data channel will not.</p> <p>ndr Session will be checked by an IPS signature.</p> <p>nds Session will be checked by an IPS anomaly.</p> <p>npu Session will possibly be offloaded to NPU.</p> <p>wccp Session is handled by WCCP.</p>

Keyword/Variable	Description
misc=0 vd=0	<p>misc displays the security policy's ID. Each security policy has an identifier, such as 20004. If this session has a UTM profile applied to it, misc indicates which one it is.</p> <p>vd is the VDOM this security policy applies to.</p>
policy_id=0	The ID number of the security policy that matches this session.
auth_info=0	If an authentication policy is applied, this displays the ID of that policy. For example an IBP would be listed here.

statistic(bytes/packets/allow_err):	
org=39032/119/1	The bytes, packets, and errors sent from the original direction.
reply=0/0/0	The bytes, packets, and errors sent from the reply direction.

serial=0002dcfe	A counter that rolls over. This value is the same for parent and child sessions.
tos=ff/ff	Type of Service (TOS) from the TOS field in the packet header. The values apply to origin and reply directions accordingly.
dd_type=0 dd_rule_id=0	<p>Data re-duplication (dd) values. These are associated with WAN Opt.</p> <p>Rule ID matches the WAN Opt rule number.</p>

per_ip_bandwidth meter:	
addr=172.20.120.225,	The IP address this meter applies to.
bps=112	The bytes per second bandwidth for this IP address.

get system session-info ttl

Displays the system session time to live (ttl) configuration.

Syntax

```
get system session-info ttl
```

Parameters

None.

Usage/Remarks

Use this command to tell you the default time to live setting for sessions. All sessions created will have this length of time before they expire and need to create a new session table entry.

A similar command is `get system session-ttl`. It provides similar output.

Scope: Vdom

Output Example

```
FGT # get system session-info ttl  
default          : 3600  
port:
```

Keyword/Variable	Description
default	The default is one hour or 3600 seconds.
port	The session port.

get system startup-error-log

Displays any start-up configuration errors on the console.

Syntax

```
get system startup-error-log
```

Parameters

None.

Usage/Remarks

Use this command to view the start-up configuration errors on the console. If there are no errors, this command displays the http-err webproxy replacement message settings.

Scope: Global and Vdom

Output Example

```
FGT # get system startup-error-log
>>> config system replacemsg webproxy "http-err"
>>>     set buffer "<html><head><title>%%HTTP_ERR_CODE%%
%%HTTP_ERR_DESC%%</title></head><body><font size=2><table
width=\"100%\"><tr><td bgcolor=#3300cc align=\"center\"
colspan=2><font color=#ffffff><b>%%HTTP_ERR_CODE%%
%%HTTP_ERR_DESC%%</b></font></td></tr></table><br><br>The
webserver for %%PROTOCOL%%URL%% reported that an error occurred
while trying to access the website. Please click <u><a
href=\"javascript:history.back()\">here</a></u> to return to the
previous page.<br><br><hr></font></body></html>"
>>>     set header http
>>>     set format html
```

get system status

Displays basic information about a FortiGate unit.

Syntax

```
get system status
```

Parameters

None.

Usage/Remarks

Use this command to display basic information about your FortiGate unit such as firmware version, serial number, hostname, number of VDOMs, HA mode, and system time.

Scope: Global and Vdom

Output Example

```
FGT # get system status
Version: Fortigate-3810A v4.0,build0418,110209 (Interim)
Virus-DB: 11.00782(2010-05-07 00:42)
Extended DB: 1.00001(2010-05-21 13:37)
IPS-DB: 2.00910(2010-12-02 17:49)
FortiClient application signature package: 1.393(2011-06-06
17:13)
Serial-Number: FG3K8A3407600241
BIOS version: 04000009
Log hard disk: Not available
Hostname: FG3K8A3407600241
Operation Mode: NAT
Current virtual domain: root
Max number of virtual domains: 10
Virtual domains status: 1 in NAT mode, 0 in TP mode
Virtual domain configuration: disable
FIPS-CC mode: disable
Current HA mode: standalone
Distribution: International
Branch point: 418
Release Version Information: Interim
System time: Tue Jun 7 14:40:16 2011
```

get test

Displays information statistics, control proxies, and status.

Syntax

```
get test <application> <level>
```

Parameters

application	<p>acd — Aggregate Controller</p> <p>ddnscd — ddnscd daemon</p> <p>dhcprelay — dhcprelay</p> <p>dnsproxy — dns proxy</p> <p>ftpd — ftp proxy</p> <p>http — http proxy</p> <p>im — im proxy</p> <p>imap — imap proxy</p> <p>ipldbd — ipldbd daemon</p> <p>ipsengine — ips sensor</p> <p>ipsmonitor — ips monitor</p> <p>l2tpcd — l2tpcd</p> <p>nntp — nntp proxy</p> <p>pop3 — pop3 proxy</p> <p>proxyacceptor — proxy acceptor</p> <p>proxyworker — proxy worker</p> <p>scanunit — scanning unit</p> <p>sflowd — sflowd</p> <p>smtp — smtp proxy</p> <p>snmpd — snmpd daemon</p> <p>urlfilter — urlfilter daemon</p> <p>vs — virtual-server</p> <p>wad — wan optimization proxy</p> <p>wccpd -wccp daemon</p>
<level>	<p>1: Dump Memory Usage</p> <p>2: Drop all connections</p> <p>22: Drop max idle connections</p> <p>222: Drop all idle connections</p> <p>4: Display connection stat</p> <p>44: Display info per connection</p>

	444: Display connections per state 4444: Display per-VDOM statistics 44444: Display information about idle connections
	5: Toggle AV Bypass mode. Toggle AV bypass mode. You can use this level to diagnose AV scanning. When bypass mode is activated, no AV scanning is done on traffic handled by the proxy. Note: Antivirus scanning is disabled. 6: Toggle Print Stat mode every ~40 seconds 7: Toggle Backlog Drop 8: Clear stats 88: Toggle statistic recording — stats cleared 9: Toggle Accounting info for display 99: Restart proxy. When you suspect abnormal behavior of the proxy, you can use this level value to restore it to its normal state. Note: You will have a disruption in services. 11: Display the SSL session ID cache statistics 12: Clear the SSL session ID cache statistics 13: Display the SSL session ID cache 14: Clear the SSL session ID cache Note: Not all level numbers may be applicable to all applications. Use the command <code>diagnose test application <application> 0</code> to see a list of valid commands.

Usage/Remarks

Use this command to display information about advanced FortiOS applications such as SSL, and VoIP.

Scope: Global

Output Examples

```
FGT # get test ftpd 1
```

```
HTTP Proxy Test Usage
```

```
Fortigate-1000#
malloc memory usage
=====
pool buffer size: 2048 count: 2143 available: 2139
total memory used by buffer pools: 4286 (kB)
total allocated (malloc/realloc) memory: 0 (kB)
shm memory usage
=====
buffer pool not initialized
```

total allocated memory 0

FGT # **get test http 44**

diagnose test application http 44
id=17 clt=914(r=1, w=0) srv=915(r=1, w=0) c:192.168.50.2:3608 ->
s:192.168.135.21:80 c2s/s2c=0/0
state=RESPONSE_PASS_STATE duration=0 expire=3590
Current connections = 1/2368

Keyword/Variable	Description
44	The first number is the counter of current sessions in the connection pool of the proxy. The second number is the maximum size of the proxy connection pool table.

FG # **get test imap 4**

Current connections = 5/2670
Fortigate-1000# Running time (HH:MM:SS:usec) = 22:45:40:124613
Bytes sent = 65 (kb)
Bytes received = 2895 (kb)
Error Count (alloc) = 0
Error Count (accept) = 0
Error Count (bind) = 0
Error Count (connect) = 0
Error Count (read) = 0
Error Count (write) = 0
Error Count (poll) = 0
Last Error = 0
Scan Backlog drop = 0
Emails clean = 3
Emails detected = 2
Emails with scan errors = 0
Worms = 0
Blocked = 0
Virus = 2
Suspicious = 0
Fragmented emails = 0
Spam Detected = 0
Content Filtered emails = 0
Oversize Email Pass = 0
Oversize Email Blocked = 0
AV Bypass is off
Print is off
Drop on backlog is off
Account is on
setup_ok=7 setup_fail=0 poll_ok=2576/2576/1 sel_fail=0 conn_ok=0
conn_inp=7
step1=0 step2=0
scan=5 listen=7 cmdb=2 clt=2304 srv=258

FG # **get test ipsengine 7**

FG # PACKET STATISTICS:

```
total packets 3487023
tcp packets 87
udp packets 0
icmp packets 2476747
discard packets 0
alert packets 45
log packets 0
pass packets 0
fragment packets 0
frag_trackers 0
rebuilt_frags 0
frag_incomplete 0
frag_timeout 0
rebuild_element 0
frag_mem_faults 0
tcp_stream_pkts 87
rebuilt_tcp 5
tcp_streams 4
rebuilt_segs 0
str_mem_faults 0
```

```
FG # get test ipsmonitor 2
FG # enable ipseengine? no
FG # diagnose test ipsmonitor 2
FG # enable ipseengine? Yes
```

```
FG # get test pop3 444
FG # [OVERSIZE_STATE ] 1/1
```

```
FG # get test smtp 44
FG # id=0 clt=10(r=1, w=0) srv=11(r=1, w=0) c:192.168.200.2:60811
-> s:192.168.50.2:25 c2s/s2c=0/0
state=CONNECTED_STATE duration=0 expire=3581
Current connections = 1/2669
```

get test urlfilter

Displays statistics, clears the cached, increases the timeout values, and many other functions for the URL filter engine.

Syntax

```
get test application urlfilter <number>
```

Parameters

<number>	<ul style="list-style-type: none"> 1 — This menu 2 — Clear cache 3 — Display WF cache contents 4 — Display WF cache TTL list 5 — Display WF cache LRU list 6 — Display WF cache in tree format 7 — Toggle switch for dumping unrated packet 8 — Increase timeout for polling 9 — Decrease timeout for polling 10 — Print debug values 11 — Clear Spam Filter cache 12 — Clear AV Query cache 13 — Toggle switch for dumping expired license packets 14 — Show running timers (except request timers) 144 — Show running timers (including request timers) 15 — Send INIT requests 16 — Display WF cache contents of prefix type 99 — Restart the urlfilter daemon
----------	---

Usage/Remarks

Use this command to troubleshoot and work with the URL filter engine. If you do not include one of the numbers listed above, the list of numbers will be displayed for you.

Scope: Global

Output Example

```
FGT # get test urlfilter 3
utree_stat: keylen=223 nodes=15 leaf=11
com.cisco.www cate = 52 len=13 ch=1
dhtml_pulldown/dropdownlib-100.js cate = 52 len=33 ch=0
niffer/sniffplib-100.js cate = 52 len=22 ch=0
potlight/spotlightlib-120.js cate = 52 len=28 ch=0
offer/sp/cookie.js cate = 52 len=18 ch=0
cisco_detect.js cate = 52 len=15 ch=0
flyouts.js cate = 52 len=10 ch=0
global.js cate = 52 len=9 ch=0
hbx.js cate = 52 len=6 ch=0
```

```
sitewide_tools.js cate = 52 len=17 ch=0  
windowutil.vb cate = 52 len=13 ch=0
```


get vpn ipsec stats crypto

Displays information about IPsec VPN cryptographic statistics.

Syntax

```
get vpn ipsec stats crypto
```

Parameters

None.

Usage/Remarks

Use this command to display which crypto devices are used with VPN connections such as 3DES, SHA1, and so on. A zero indicates that type of crypto is not used with any VPN connections. This command can be useful to help enforce the use of more secure crypto.

In the output, you should be aware that there is hardware (CP6) and software crypto categories. There are also separate columns for encrypted (outbound) and decrypted (inbound), or generated (outbound) and validated (inbound).

Scope: Vdom

Output Example

```
FG10CH3G09603750 # get vpn ipsec stats crypto
IPsec crypto devices in use:
```

```
CP6 (encrypted/decrypted):
```

null:	0	0
des:	0	0
3des:	0	0
aes:	0	0

```
CP6 (generated/validated):
```

null:	0	0
md5:	0	0
sha1:	0	0
sha256:	0	0

```
SOFTWARE (encrypted/decrypted):
```

null:	0	0
des:	0	0
3des:	0	0
aes:	0	0

```
SOFTWARE (generated/validated):
```

null:	0	0
md5:	0	0
sha1:	0	0
sha256:	0	0

get vpn ipsec stats tunnel

Displays information about IPsec VPN tunnels.

Syntax

```
get vpn ipsec stats tunnel
```

Parameters

None.

Usage/Remarks

Use this command to display information about IPsec VPN tunnels including how many of each type (static, dynamic, or manual addressing), how many errors there have been, and how many selectors there are.

Scope: Vdom

Output Example

```
FG10CH3G09603750 # get vpn ipsec stats tunnel
tunnels
  total: 1
    static/ddns: 1
    dynamic: 0
    manual: 0
  errors: 0
selectors
  total: 1
  up: 1
```

get vpn ipsec tunnel details

Displays detailed information about IPsec VPN tunnel.

Syntax

```
get vpn ipsec tunnel details
```

Parameters

None.

Usage/Remarks

Use this command to display in depth information about IPsec VPN tunnels. This command is useful when troubleshooting VPN tunnels due to the amount of information displayed.

Scope: Vdom

Output Example

```
FG10CH3G09603750 # get vpn ipsec tunnel details

gateway
  name: 'ph1'
  type: policy-based
  local-gateway: 172.16.68.34:0 (dynamic)
  remote-gateway: 172.16.68.35:0 (static)
  mode: ike-v1
  interface: 'port4' (12)
  rx packets: 0 bytes: 0 errors: 0
  tx packets: 0 bytes: 0 errors: 0
  dpd: enabled/negotiated idle: 5000ms retry: 3 count: 0
  selectors
    name: 'ph2'
    auto-negotiate: enable
    mode: tunnel
    src: 0:0.0.0.0/0.0.0.0:0
    dst: 0:0.0.0.0/0.0.0.0:0
  SA
    lifetime/rekey: 180/123
    mtu: 1436
    tx-esp-seq: 1
    replay: enabled
    inbound
      spi: 585e89b2
      enc: 3des
      40a50330e01e36fbb58cbf26f2d77e4e57fc5c1fc41d20ec
      auth: sha1 1f1292da090bde71f955f4e8cb23d4e0c7768a8e
    outbound
      spi: efd4667a
      enc: 3des
      66bfeaf3fcd635d1e309be18a8eea41d53630371590ff8e0
      auth: sha1 c7ec4ac4632a22785b713644a4d1581269b29788
  SA
```

```
lifetime/rekey: 180/120
mtu: 1436
tx-esp-seq: 1
replay: enabled
inbound
  spi: 585e89b1
  enc: 3des
69e72f1bf0a0cd76aded631c0010f17e5dlacf558dad0329
  auth: sha1 ed96712db7e633bfa9720b2bf7a276d198b7661a
outbound
  spi: efd4667b
  enc: 3des
28d49a7e392487e6078907bac49560e00b21428a5bfcf8be
  auth: sha1 84eb5525df4511a36c3ac291ff5526bbf0be9ef2
```

get vpn ipsec tunnel summary

Displays brief information about IPsec VPN tunnel.

Syntax

```
get vpn ipsec tunnel summary
```

Parameters

None.

Usage/Remarks

Use this command to display brief information about IPsec VPN tunnels including their name, IP address and port, how many selectors that tunnel has, and packets received, sent, and in error. This information is displayed for each tunnel defined.

Scope: Vdom

Output Example

```
FG10CH3G09603750 # get vpn ipsec tunnel summary
'ph1' 172.16.68.35:0 selectors(total,up): 1/1 rx(pkt,err): 0/0
tx(pkt,err): 0/0
```

get vpn status ssl hw-acceleration-status

Displays the status of SSL hardware accelerated VPN connections.

Hardware acceleration can be accomplished in a number of ways. Some FortiGate models incorporate network processors (NPUs) in the main unit, others support the addition of AMC (Advanced Mezzanine Card) modules. The FortiGate-5000 series supports rear transition modules (RTMs) that incorporate network processors.

Syntax

```
get vpn status ssl hw-acceleration-status
```

Parameters

None.

Usage/Remarks

Use this command to display the status of SSL hardware accelerated VPN connections. This is useful if you need to troubleshoot a VPN connection that should be accelerated but is not.

Scope: Vdom

Output Example

```
FGT # get vpn status ssl hw-acceleration-status  
Acceleration hardware detected: kxp=on cipher=on
```

get vpn status ssl list

Displays the status of list of SSL VPN connections.

Syntax

```
get vpn status ssl list
```

Parameters

None.

Usage/Remarks

Use this command to display the status of SSL list of VPN connections.

Scope: Vdom

Output Example

```
FGT # get vpn status ssl list  
SSL VPN is disabled.
```

get webfilter ftgd-statistics

Displays the FortiGuard rating statistics.

Syntax

```
get webfilter ftgd-statistics
```

Parameters

None.

Usage/Remarks

Use this command to display many details about your FortiGuard statistics. These numbers can be useful in determining the source of a problem. For example DNS failures can reveal the source of problems that may be widespread and hard to track back to the source.

Scope: Global

Output Example

```
FGT # get webfilter ftgd-statistics
Rating Statistics:
=====
DNS failures                :          3
DNS lookups                 :         16
Data send failures         :        876
Data read failures         :          0
Wrong package type         :          0
Hash table miss            :          1
Unknown server             :          0
Incorrect CRC               :          0
Proxy request failures     :          0
Request timeout            :        292
Total requests              :          0
Requests to FortiGuard servers :          0
Server errored responses   :         11
Relayed rating              :          0
Invalid profile            :          0

Allowed                    :          0
Blocked                    :          0
Logged                     :          0
Errors                     :          0

Cache Statistics:
=====
Maximum memory             :          0
Memory usage                :          0

Nodes                      :          0
  Leaves                   :          0
  Prefix nodes             :          0
  Exact nodes              :          0
```



```

Requests                :          0
Misses                  :          0
Hits                    :          0
Prefix hits              :          0
Exact hits               :          0

No cache directives     :          0
Add after prefix        :          0
Invalid DB put          :          0
DB updates              :          0

Percent full            :          0%
  Branches              :          0%
  Leaves                :          0%
    Prefix nodes        :          0%
    Exact nodes         :          0%

Miss rate                :          0%
Hit rate                 :          0%
  Prefix hits           :          0%
  Exact hits            :          0%

```

Keyword/Variable	Description
DNS lookups	Number of DNS look-ups for the domain name of the requested URL.
Data send failures	Number of non-responsive servers.
Request timeout	Number of seconds for the request time-out. A FortiGate unit sends the URL rating request every 2 seconds.
Total requests	Total number of URL rating requests to cache and FortiGuard servers.
Requests to FortiGuard servers	Total number of URL rating requests to FortiGuard servers.
Relayed rating	The number of times the master communicates with FortiGuard servers and relays all URL rating requests from the slaves in a HA cluster.
Maximum memory	Amount of memory assigned to FortiGuard cache. The default is 2 percent.
Memory usage	The amount of memory used to store the URL tree.
Prefix nodes	The number of prefixes used for a URL to increase the cache hit rate.
Exact nodes	The number of exact matches used for a URL.

get webfilter status

Displays webfilter rating and server information from FortiGuard.

Syntax

```
get webfilter status <refresh rate>
```

Parameters

<refresh rate>	How often to refresh the server list(s).
-----------------------------	--

Usage/Remarks

Use this command to display webfilter statistics and server information if the service is enabled.

If the service is not enabled, this command displays the language of the locale and states the service is not enabled.

This command was previously called `diag debug rating`.

Scope: Global and Vdom

Output Example

```
FGT # get webfilter status 4
Locale : english
License : Contract
Expiration : Wed Feb 11 02:00:00 2009
Hostname : service.fortiguard.net

--- Server List (Mon May 26 22:36:34 2008) ---
```

IP	Weight	RTT	Flags	TZ	Packets	Curr Lost	Total Lost
212.95.252.121	10	77		0	83	0	8
62.209.40.73	0	86		1	83	0	8
62.209.40.72	0	92	DI	1	85	0	9
82.71.226.65	10	98	D	0	84	0	8
212.95.252.120	10	95		0	76	0	2
69.20.236.179	60	198		-5	84	0	9
66.117.56.42	60	189		-5	83	0	8
66.117.56.37	60	192		-5	83	0	8
69.20.236.180	60	213		-5	83	0	8
209.52.128.90	60	268		-5	84	0	9
121.111.236.179	80	383		9	83	0	8
121.111.236.180	80	404		9	83	0	8
72.52.72.243	90	289		-8	83	0	8
218.106.244.81	90	455		-8	83	0	8

69.90.198.55 90 297 D -8 85 0 9

Keyword/Variable	Description
Locale	Local environment language.
License	The license status: <ul style="list-style-type: none"> • Contract • Expired • Trial
Expiration	The date and time the license expires.
Hostname	The FortiGuard server the FortiGate unit connects to obtain the service. The FortiGuard server will return the information to the FortiGate unit. The default is service.fortiguard.net
IP	The IP address of other FortiGuard servers.
Weight	The priority value which the FortiGate unit uses to send the URL rating request. The lower weight value takes higher preference. The weight is calculated by time zone, packet round-trip time, and success rate.
RTT Flags	The round-trip time between the URL rating request and the response time from the FortiGuard server.
TZ	Time zone of the FortiGuard server (Greenwich Mean Time +/- the number).
Packets	Total packets sent to the FortiGuard server.
Curr Lost	The number of times the request is retried in a timeout period. The default is 15 seconds.
Total Lost	Total number of unresponsive requests.



Troubleshooting bootup and FSSO

When powering on your FortiGate unit, you may experience problems. This section addresses some problems you may experience in this area in rare cases. If you continue to have problems, please contact customer support for assistance.

After you have setup FSSO, you may experience problems. This section addresses many of the problems you may experience in this area.

This section contains the following topics:

- [FortiGate unit bootup issues](#)
- [FSSO issues](#)

FortiGate unit bootup issues

Bootup issues, while rare, can be very difficult to troubleshoot due to the lack of information about your issue. When the FortiGate unit is not running, you do not have access to your typical tools such as the CLI `diagnose` commands. This section walks you through some possible issues to give you direction in these situations.

To troubleshoot a bootup problem with your FortiGate unit, go to the section that lists your problem. If you have multiple problems, go the problem closest to the top of the list first, and work your way down the list.



It is rare that units experience any of the symptoms listed here. Fortinet hardware is reliable with a long expected operation life.

Basic bootup troubleshooting

If your FortiGate unit is hanging during the boot process, try the basic method here first before going through all the possible advanced troubleshooting.

If your FortiGate unit is hanging during boot, and after the “Press any key” prompt, try the following:

- 1 Reboot the FortiGate unit.
- 2 When you see the “Press any key. .” prompt, press a key.
You will see a menu of options.
- 3 If you suspect bad configuration is your problem, select [I] to view basic hardware and networking information, or to change basic networking settings.
- 4 It may be your firmware image is corrupted somehow. If you have an alternate firmware partition on the unit, select [B] to boot from the alternate partition.
- 5 If these do not solve your problem, continue on to advanced FortiGate bootup issues.

Advanced bootup troubleshooting

The issues covered in this section all refer to various potential bootup issues including:

- [A. You have text on the screen, but you have problems](#)
- [B. You do not see the boot options menu](#)

- [C. You have problems with the console text](#)
- [D. You have visible power problems](#)
- [E. You have a suspected defective FortiGate unit](#)

A. You have text on the screen, but you have problems

Solution:

- 1 If the text on the screen is garbled, ensure your Console Communication parameters are correct. Check your Quick Start Guide for settings specific to your model.
- 2 If that fixes your problem, you are done.
- 3 If not, go to [B. You do not see the boot options menu](#)

B. You do not see the boot options menu

Solution:

- 1 Ensure your serial communication parameters are set to `no flow control`, and the proper baud rate and reboot the FortiGate unit by powering off and on.



FortiGate units ship with a baud rate of 9600 by default. If you have access, verify this with the CLI command `config system console get`, or parse an archived configuration file for the term `baudrate`.

- 2 If that fixes your problem, you are done.
- 3 If it does not fix your problem, go to [E. You have a suspected defective FortiGate unit.](#)

C. You have problems with the console text

- 1 Do you have any console message?
 - If Yes, go to [D. You have visible power problems](#)
 - If No, continue.
- 2 Is there garbage text onscreen?
 - If Yes, ensure Console Communication parameters are okay.
 - If that fixes the problem, you are done.
- 3 If no, does the FortiGate unit stop before the `Press Any Key to Download Boot Image` prompt ?
 - If Yes, go to [E. You have a suspected defective FortiGate unit.](#)
 - If No, go to Step 4.
- 4 Console Message - `Press any key to Download Boot Image`
- 5 When pressing a key, do you see the following messages?


```
[G] Get Firmware image from TFTP server
[F] Format boot device
[B] Boot with backup firmware and act as default
[I] Hardware and configuration information
[Q] Quit menu and continue to boot with default firmware
[H] Display this list of options
```

 - If Yes, go to [E. You have a suspected defective FortiGate unit.](#)

- 6 If No, ensure you serial communication parameters are set to `no flow control`, and the proper baud rate and reboot the FortiGate unit by powering off and on.



FortiGate units ship with a baud rate of 9600 by default. If you have access, parse an archived configuration file for the term `baudrate` or verify this setting with the CLI command:

```
config system console
get
```

- 7 Did the reboot fix the problem?
 - If that fixes your problem, you are done.
 - If that does not fix your problem, go to [E. You have a suspected defective FortiGate unit](#)

D. You have visible power problems

- 1 Is there any LED on?
 - If No, ensure power is on. If that fixes the problem you are done. If not, continue.
 - If Yes, continue.
- 2 Do you have an external power adapter?
 - If No, go to [E. You have a suspected defective FortiGate unit](#).
 - If Yes, try replacing the power adapter.
- 3 Is the power supply defective or you can't determine one way or the other?
 - If No, go to [E. You have a suspected defective FortiGate unit](#).
 - If Yes, go to [A. You have text on the screen, but you have problems](#)

E. You have a suspected defective FortiGate unit

If you have followed these steps and determined there is a good chance your unit is defective, follow these steps.

- 1 Open a support ticket through FortiCare at <https://support.fortinet.com>.
- 2 In the ticket, document the problem or problems, and these steps that you have taken.
- 3 Provide all console messages and output.
- 4 Indicate if you have a suspected hard disk issue, and provide your evidence.

Fortinet Customer Support will contact you to help you with your ticket and issue.

FSSO issues

The following is a flow chart for troubleshooting FSSO issues.

- [A. Initial information gathering](#)
- [B. The CA is not running and not connected](#)
- [C. The CA is running but not connected](#)
- [D. The CA is connected](#)
- [E. There are at least some users logged on](#)
- [F. Test user does not appear on the FSSO list](#)

A. Initial information gathering

- 1 Is the Collector Agent (CA) connected? The best way to verify this is with the CLI commands

```
diagnose debug enable
diagnose debug authd fsso server-status
```

 - If it is connected, go to [D. The CA is connected](#).
 - If not, continue on.
- 2 Is the CA running? Check by going to *Administrative Tools > Services > Check the 'Fortinet Server Authentication Extension' Service*.
 - If the CA is running but not connected, go to step 1.
 - If the CA is not running, go to [B. The CA is not running and not connected](#).

B. The CA is not running and not connected

- 1 Is the CA running as admin? Go to *Administrative Tools > Services > Check the 'Fortinet Server Authentication Extension' Service* and see if it is using a domain administrator account. If not change the account.
- 2 Is the CA able to bind to socket? FSSO uses ports 8000 and 8002. If the ports are not available the service will not start. Check the CA logs to verify this, and stop the other application.
- 3 In general, checking the CA logs will show any other errors that are preventing the collector agent from starting.
- 4 Once CA is running if problems continue, go to [A. Initial information gathering](#) step 2.

C. The CA is running but not connected

- 1 Is the password correct? If not, reset the password
- 2 Is there a device filtering traffic? If so, change it to allow TCP port 8000.
- 3 If problems continue use the following CLI diag command to find more information.

```
diag debug enable
diag debug application authd 8256
```
- 4 Once CA is running and connected if problems continue, go to [A. Initial information gathering](#) step 2.

D. The CA is connected

- 1 Are you seeing groups on the FortiGate?
If not, check the group filter on the CA.
- 2 Are the FortiGate and CA groups using the same mode? If not, change the modes to match.
- 3 Are you seeing logon events on the FortiGate unit? You can check this with the following CLI commands.

```
diagnose debug enable
diagnose debug authd fsso list
```

 - If there are any users logged in, go to step .
 - Otherwise, continue on.
- 4 Are DC agents installed on all Domain Controllers? If not install the DC Agents.

- 5 Are you using an LDAP server on the FSSO connector?
 - To check go to *User > Directory Service > Edit FSSO connector > LDAP*.
 - If an LDAP server is configured, disable it and go to step 3.
 - If there is no LDAP server configured, contact support and open a support ticket.

E. There are at least some users logged on

- 1 Focus on a single 'test' user for farther troubleshooting. The information to collect about the 'test' user is:
 - Account username of the user currently logged in
 - IP address of the test host — you can run `ipconfig` to get the IP of the host
 - Host DNS name — you can run `hostname` to get the host name.
 - Logon server name, the domain controller the host used to authenticate — to get this, run `echo %logonserver%`.
- 2 Once you have the information about the test host, run the following CLI commands on the FortiGate unit.


```
diagnose debug enable
diagnose debug authd fssolist
```

 - If the user is not in the list, got to [F. Test user does not appear on the FSSO list](#).
 - If the user appears in the list, continue on.
- 3 Does the user have the correct IP address? If not, check the DNS settings on the DNS server. If a computer has two network interfaces (multi-homed) traffic may get mixed up and go out the wrong interface.
- 4 Does the user have the correct groups? If not, disable group caching on the CA.
- 5 If all else appears okay, check the order of the security policies. Only the first authenticated group is allowed through, which may be the wrong policy.
- 6 If problems continue, contact support and open a support ticket.

F. Test user does not appear on the FSSO list

If the user did not appear in the FSSO list, run through the following checklist.

Is the user IP showing up with a service account?	If yes, add the user to the ignore list.
Has the user been moved to a new group recently?	If yes, disable group caching.
If the user is in the CA log:	
Were there any DNS errors in the CA for the host name?	These include the collector agent unable to resolve the host name at all or resolving to an incorrect IP. If yes, check the DNS server.
Did the user time out?	If the CA logs show the user timed out, the collector agent was not able to connect to the host on port 139 & 445 to verify the user.
If the user is not in the CA log:	

Check the logon server.	Check which domain controller authenticated the host (run <code>echo %logonserver%</code> on the host) and troubleshoot that domain controller.
Does the logon server have the DC agent installed?	If not, install the DC agent. If it is installed, enable logging on the DC agent on the logon server. Use the logs produced for farther troubleshooting.



Chapter 6 UTM Guide

This FortiOS Handbook chapter contains the following sections:

[UTM overview](#) describes UTM components and their relation to firewall policies, as well as SSL content scanning and inspection. We recommend starting with this section to become familiar with the different features in your FortiGate unit.

[Network defense](#) explains basic denial of service (DoS) and distributed denial of service (DDoS) concepts and provides an overview of the best practices to use with all the UTM features to defend your network against infection and attack.

[AntiVirus](#) explains how the FortiGate unit scans files for viruses and describes how to configure the antivirus options.

[Email filter](#) explains how the FortiGate unit filters email, describes how to configure the filtering options and the action to take with email detected as spam.

[Intrusion protection](#) explains basic Intrusion Protection System (IPS) concepts and how to configure IPS options; includes guidance and a detailed table for creating custom signatures as well as several examples.

[Web filter](#) and [FortiGuard Web Filter](#) The first of these sections describes basic web filtering concepts, the order in which the FortiGate unit performs web filtering, and configuration. The second section describes enhanced features of the subscription-based FortiGuard Web Filtering service and explains how to configure them. We recommend reading both sections if you are using FortiGuard Web Filtering because settings you configure in one feature may affect the other.

[Data leak prevention](#) describes the DLP features that allow you to prevent sensitive data from leaving your network and explains how to configure the DLP rules, compound rules, and sensors.

[Application control](#) describes how your FortiGate unit can detect and take action against network traffic based on the application generating the traffic.

[DoS policy](#) describes how to use DoS policies to protect your network from DoS attacks.

[Sniffer policy](#) describes how to use your FortiGate unit as a one-armed intrusion detection system (IDS) to report on attacks.



UTM overview

Ranging from the FortiGate®-30 series for small businesses to the FortiGate-5000 series for large enterprises, service providers and carriers, the FortiGate line combines a number of security features to protect your network from threats. As a whole, these features, when included in a single Fortinet security appliance, are referred to as Unified Threat Management (UTM). The UTM features your FortiGate model includes are:

- AntiVirus
- Intrusion Prevention System (IPS)
- Anomaly protection (DoS policies)
- One-armed IPS (Sniffer policies)
- Web filtering
- E-mail filtering, including protection against spam and grayware
- Data Leak Prevention (DLP)
- Application Control (for example, IM and P2P).

Firewall policies limit access, and while this and similar features are a vital part of securing your network, they are not covered in this document.

The following topics are included in this section:

- [UTM components](#)
- [UTM profiles/lists/sensors](#)

UTM components

AntiVirus

Your FortiGate unit stores a virus signature database that can identify more than 15,000 individual viruses. FortiGate models that support additional virus databases are able to identify hundreds of thousands of viruses. With a FortiGuard AntiVirus subscription, the signature databases are updated whenever a new threat is discovered.

AntiVirus also includes file filtering. When you specify files by type or by file name, the FortiGate unit will stop the matching files from reaching your users.

FortiGate units with a hard drive or configured to use a FortiAnalyzer unit can store infected and blocked files for that you can examine later.

Intrusion Protection System (IPS)

The FortiGate Intrusion Protection System (IPS) protects your network against hacking and other attempts to exploit vulnerabilities of your systems. More than 3,000 signatures are able to detect exploits against various operating systems, host types, protocols, and applications. These exploits can be stopped before they reach your internal network.

You can also write custom signatures, tailored to your network.

Anomaly protection (DoS policies)

A complement to the signature-based IPS, anomaly protection detects unusual network traffic that can be used to attack your network. When you set thresholds for various types of network operations, the FortiGate unit will block any attempt to exceed the thresholds you have defined.

One-armed IDS (sniffer policies)

You can use sniffer policies on the FortiGate unit as a one-arm intrusion detection system (IDS). The unit examines traffic for matches to the configured IPS sensor and application control list. Matches are logged and then all received traffic is dropped. In this way, you can configure a unit to sniff network traffic for attacks without actually processing the packets.

The FortiGate unit can log all detected IPS signatures and anomalies in a traffic stream.

Web filtering

Web filtering includes a number of features you can use to protect or limit your users' activity on the web.

FortiGuard Web Filtering is a subscription service that allows you to limit access to web sites. More than 60 million web sites and two billion web pages are rated by category. You can choose to allow or block each of the 77 categories.

URL filtering can block your network users from access to URLs that you specify.

Web content filtering can restrict access to web pages based on words and phrases appearing on the web page itself. You can build lists of words and phrases, each with a score. When a web content list is selected in a web filter profile, you can specify a threshold. If a user attempts to load a web page and the score of the words on the page exceeds the threshold, the web page is blocked.

Email filtering

FortiGuard AntiSpam is a subscription service that includes an IP address black list, a URL black list, and an email checksum database. These resources are updated whenever new spam messages are received, so you do not need to maintain any lists or databases to ensure accurate spam detection.

You can use your own IP address lists and email address lists to allow or deny addresses, based on your own needs and circumstances.

Data Leak Prevention (DLP)

Data leak prevention allows you to define the format of sensitive data. The FortiGate unit can then monitor network traffic and stop sensitive information from leaving your network. Rules for U.S. social security numbers, Canadian social insurance numbers, as well as Visa, Mastercard, and American Express card numbers are included.

Application Control (for example, IM and P2P)

Although you can block the use of some applications by blocking the ports they use for communications, many applications do not use standard ports to communicate. Application control can detect the network traffic of more than 1000 applications, improving your control over application communication.

UTM profiles/lists/sensors

A profile is a group of settings that you can apply to one or more firewall policies. Each UTM feature is enabled and configured in a profile, list, or sensor. These are then selected in a security policy and the settings apply to all traffic matching the policy. For example, if you create an antivirus profile that enables antivirus scanning of HTTP traffic, and select the antivirus profile in the security policy that allows your users to access the World Wide Web, all of their web browsing traffic will be scanned for viruses.

Because you can use profiles in more than one security policy, you can configure one profile for the traffic types handled by a set of firewall policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate sets of profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

The UTM profiles include:

- antivirus profile
- IPS sensor
- Web filter profile
- Email filter profile
- Data Leak Prevention profile
- Application Control list
- VoIP profile

Although they're called profiles, sensors, and lists, they're functionally equivalent. Each is used to configure how the feature works.



Network defense

This section describes in general terms the means by which attackers can attempt to compromise your network and steps you can take to protect it. The goal of an attack can be as complex as gaining access to your network and the privileged information it contains, or as simple as preventing customers from accessing your web server. Even allowing a virus onto your network can cause damage, so you need to protect against viruses and malware even if they are not specifically targeted at your network.

The following topics are included in this section:

- [Monitoring](#)
- [Blocking external probes](#)
- [Defending against DoS attacks](#)
- [Traffic inspection](#)
- [Content inspection and filtering](#)

Monitoring

Monitoring, in the form of logging, alert email, and SNMP, does not directly protect your network. But monitoring allows you to review the progress of an attack, whether afterwards or while in progress. How the attack unfolds may reveal weaknesses in your preparations. The packet archive and sniffer policy logs can reveal more details about the attack. Depending on the detail in your logs, you may be able to determine the attackers location and identity.

While log information is valuable, you must balance the log information with the resources required to collect and store it.

Blocking external probes

Protection against attacks is important, but attackers often use vulnerabilities and network tools to gather information about your network to plan an attack. It is often easier to prevent an attacker from learning important details about your network than to defend against an attack designed to exploit your particular network.

Attacks are often tailored to the hardware or operating system of the target, so reconnaissance is often the first step. The IP addresses of the hosts, the open ports, and the operating systems the hosts are running is invaluable information to an attacker. Probing your network can be as simple as an attacker performing an address sweep or port scan to a more involved operation like sending TCP packets with invalid combinations of flags to see how your firewall reacts.

Address sweeps

An address sweep is a basic network scanning technique to determine which addresses in an address range have active hosts. A typical address sweep involves sending an ICMP ECHO request (a ping) to each address in an address range to attempt to get a response. A response signifies that there is a host at this address that responded to the ping. It then becomes a target for more detailed and potentially invasive attacks.

Address sweeps do not always reveal all the hosts in an address range because some systems may be configured to ignore ECHO requests and not respond, and some firewalls and gateways may be configured to prevent ECHO requests from being transmitted to the destination network. Despite this shortcoming, Address sweeps are still used because they are simple to perform with software tools that automate the process.

Use the `icmp_sweep` anomaly in a DoS sensor to protect against address sweeps.

There are a number of IPS signatures to detect the use of ICMP probes that can gather information about your network. These signatures include `AddressMask`, `Traceroute`, `ICMP.Invalid.Packet.Size`, and `ICMP.Oversized.Packet`. Include ICMP protocol signatures in your IPS sensors to protect against these probes/attacks.

Port scans

Potential attackers may run a port scan on one or more of your hosts. This involves trying to establish a communication session to each port on a host. If the connection is successful, a service may be available that the attacker can exploit.

Use the DoS sensor anomaly `tcp_port_scan` to limit the number of sessions (complete and incomplete) from a single source IP address to the configured threshold. If the number of sessions exceed the threshold, the configured action is taken.

Use the DoS sensor anomaly `udp_scan` to limit UDP sessions in the same way.

Probes using IP traffic options

Every TCP packet has space reserved for eight flags or control bits. They are used for communicating various control messages. Although space in the packet is reserved for all eight, there are various combinations of flags that should never happen in normal network operation. For example, the SYN flag, used to initiate a session, and the FIN flag, used to end a session, should never be set in the same packet.

Attackers may create packets with these invalid combinations to test how a host will react. Various operating systems and hardware react in different ways, giving a potential attackers clues about the components of your network.

The IPS signature `TCP.Bad.Flags` detects these invalid combinations. The default action is pass though you can override the default and set it to *Block* in your IPS sensor.

Configure packet reply and TCP sequence checking

The anti-reply CLI command allows you to set the level of checking for packet replay and TCP sequence checking (or TCP Sequence (SYN) number checking). All TCP packets contain a Sequence Number (SYN) and an Acknowledgement Number (ACK). The TCP protocol uses these numbers for error free end-to-end communications. TCP sequence checking can also be used to validate individual packets.

FortiGate units use TCP sequence checking to make sure that a packet is part of a TCP session. By default, if a packet is received with sequence numbers that fall out of the expected range, the FortiGate unit drops the packet. This is normally a desired behavior, since it means that the packet is invalid. But in some cases you may want to configure different levels of anti-replay checking if some of your network equipment uses non-RFC methods when sending packets.

Configure the anti-reply CLI command:

```
config system global
    anti-reply {disable | loose | strict}
end
```

You can set anti-replay protection to the following settings:

- **disable** — No anti-replay protection.
- **loose** — Perform packet sequence checking and ICMP anti-replay checking with the following criteria:
 - The SYN, FIN, and RST bit can not appear in the same packet.
 - The FortiGate unit does not allow more than one ICMP error packet through before it receives a normal TCP or UDP packet.
 - If the FortiGate unit receives an RST packet, and check-reset-range is set to strict, the FortiGate unit checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- **strict** — Performs all of the loose checking but for each new session also checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value for each new session. Strict anti-replay checking can also help prevent SYN flooding.

If any packet fails a check it is dropped.

Configure ICMP error message verification

```
check-reset-range {disable | strict}
```

Enable ICMP error message verification to ensure an attacker can not send an invalid ICMP error message.

```
config system global
    check-reset-range {disable | strict}
end
```

- **disable** — the FortiGate unit does not validate ICMP error messages.
- **strict** — enable ICMP error message checking.

If the FortiGate unit receives an ICMP error packet that contains an embedded IP(A,B) | TCP(C,D) header, then if FortiOS can locate the A:C->B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped. Strict checking also affects how the anti-replay option checks packets.

Protocol header checking

Select the level of checking performed on protocol headers.

```
config system global
    check-protocol-header {loose | strict}
end
```

- `loose` — the FortiGate unit performs basic header checking to verify that a packet is part of a session and should be processed. Basic header checking includes verifying that the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options are correct, etc.
- `strict` — the FortiGate unit does the same checking as above plus it verifies that ESP packets have the correct sequence number, SPI, and data length.

If the packet fails header checking it is dropped by the FortiGate unit.

Evasion techniques

Attackers employ a wide range of tactics to try to disguise their techniques. If an attacker disguises a known attack in such a way that it is not recognized, the attack will evade your security and possibly succeed. FortiGate security recognizes a wide variety of evasion techniques and normalizes data traffic before inspecting it.

Packet fragmentation

Information sent across local networks and the Internet is encapsulated in packets. There is a maximum allowable size for packets and this maximum size varies depending on network configuration and equipment limitations. If a packet arrives at a switch or gateway and it is too large, the data it carries is divided among two or more smaller packets before being forwarded. This is called fragmentation.

When fragmented packets arrive at their destination, they are reassembled and read. If the fragments do not arrive together, they must be held until all of the fragments arrive. Reassembly of a packet requires all of the fragments.

The FortiGate unit automatically reassembles fragmented packets before processing them because fragmented packets can evade security measures. Both IP packets and TCP packets are reassembled by the IPS engine before examination.

For example, you have configured the FortiGate unit to block access to the example.org web site. Any checks for example.com will fail if a fragmented packet arrives and one fragment contains `http://www.exa` while the other contains `mple.com/`. Viruses and malware can be fragmented and avoid detection in the same way. The FortiGate unit will reassemble fragmented packets before examining network data to ensure that inadvertent or deliberate packet fragmentation does not hide threats in network traffic.

Non-standard ports

Most traffic is sent on a standard port based on the traffic type. The FortiGate unit recognizes most traffic by packet content rather than the TCP/UDP port and uses the proper IPS signatures to examine it. Protocols recognized regardless of port include DHCP, DNP3, FTP, HTTP, IMAP, MS RPC, NNTP, POP3, RSTP, SIP, SMTP, and SSL, as well as the supported IM/P2P application protocols.

In this way, the FortiGate unit will recognize HTTP traffic being sent on port 25 as HTTP rather than SMTP, for example. Because the protocol is correctly identified, the FortiGate unit will examine the traffic for any enabled HTTP signatures.

Negotiation codes

Telnet and FTP servers and clients support the use of negotiation information to allow the server to report what features it supports. This information has been used to exploit vulnerable servers. To avoid this problem, the FortiGate unit removes negotiation codes before IPS inspection.

HTTP URL obfuscation

Attackers encode HTML links using various formats to evade detection and bypass security measures. For example, the URL `www.example.com/cgi.bin` could be encoded in a number of ways to avoid detection but still work properly, and be interpreted the same, in a web browser.

The FortiGate prevents the obfuscation by converting the URL to ASCII before inspection.

Table 57: HTTP URL obfuscation types

Encoding type	Example
No encoding	<code>http://www.example.com/cgi.bin/</code>
Decimal encoding	<code>http://www.example.com/&#99;&#103;&#105;&#46;&#98;&#105;&#110;&#47;</code>
URL encoding	<code>http://www.example.com/%43%47%49%2E%42%49%4E%2F</code>
ANSI encoding	<code>http://www.example.com/%u0063%u0067%u0069%u002E%u0062%u0069%u006E/</code>
Directory traversal	<code>http://www.example.com/cgi.bin/test/..</code>

HTTP header obfuscation

The headers of HTTP requests or responses can be modified to make the discovery of patterns and attacks more difficult. To prevent this, the FortiGate unit will:

- remove junk header lines
- reassemble an HTTP header that's been folded onto multiple lines
- move request parameters to HTTP POST body from the URL

The message is scanned for any enabled HTTP IPS signatures once these problems are corrected.

HTTP body obfuscation

The body content of HTTP traffic can be hidden in an attempt to circumvent security scanning. HTTP content can be GZipped or deflated to prevent security inspection. The FortiGate unit will uncompress the traffic before inspecting it.

Another way to hide the contents of HTTP traffic is to send the HTTP body in small pieces, splitting signature matches across two separate pieces of the HTTP body. The FortiGate unit reassembles these 'chunked bodies' before inspection.

Microsoft RPC evasion

Because of its complexity, the Microsoft Remote Procedure Call protocol suite is subject to a number of known evasion techniques, including:

- SMB-level fragmentation
- DCERPC-level fragmentation
- DCERPC multi-part fragmentation
- DCERPC UDP fragmentation
- Multiple DCERPC fragments in one packet

The FortiGate unit reassembles the fragments into their original form before inspection.

Defending against DoS attacks

A denial of service is the result of an attacker sending an abnormally large amount of network traffic to a target system. Having to deal with the traffic flood slows down or disables the target system so that legitimate users can not use it for the duration of the attack.

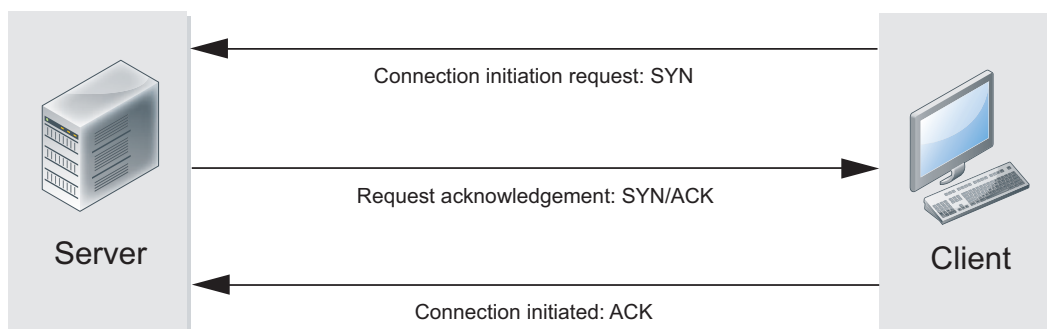
Any network traffic the target system receives has to be examined, and then accepted or rejected. TCP, UDP, and ICMP traffic is most commonly used, but a particular type of TCP traffic is the most effective. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are started between systems.

The “three-way handshake”

Communication sessions between systems start with establishing a TCP/IP connection. This is a simple three step process, sometimes called a “three-way handshake,” initiated by the client attempting to open the connection.

- 1 The client sends a TCP packet with the SYN flag set. With the SYN packet, the client informs the server of its intention to establish a connection.
- 2 If the server is able to accept the connection to the client, it sends a packet with the SYN and the ACK flags set. This simultaneously acknowledges the SYN packet the server has received, and informs the client that the server intends to establish a connection.
- 3 To acknowledge receipt of the packet and establish the connection, the client sends an ACK packet.

Figure 81: Establishing a TCP/IP connection




The three-way handshake is a simple way for the server and client to each agree to establish a connection and acknowledge the other party expressing its intent. Unfortunately, the three-way handshake can be used to interfere with communication rather than facilitate it.

SYN flood

When a client sends a SYN packet to a server, the server creates an entry in its session table to keep track of the connection. The server then sends a SYN+ACK packet expecting an ACK reply and the establishment of a connection.

An attacker intending to disrupt a server with a denial of service (DoS) attack can send a flood of SYN packets and not respond to the SYN+ACK packets the server sends in response. Networks can be slow and packets can get lost so the server will continue to send SYN+ACK packets until it gives up, and removes the failed session from the session table. If an attacker sends enough SYN packets to the server, the session table will fill completely, and further connection attempts will be denied until the incomplete sessions time out. Until this happens, the server is unavailable to service legitimate connection requests.

Figure 82: A single client launches a SYN flood attack




SYN floods are seldom launched from a single address so limiting the number of connection attempts from a single IP address is not usually effective.

SYN spoofing

With a flood of SYN packets coming from a single attacker, you can limit the number of connection attempts from the source IP address or block the attacker entirely. To prevent this simple defense from working, or to disguise the source of the attack, the attacker may spoof the source address and use a number of IP addresses to give the appearance of a distributed denial of service (DDoS) attack. When the server receives the spoofed SYN packets, the SYN+ACK replies will go to the spoofed source IP addresses which will either be invalid, or the system receiving the reply will not know what to do with it.

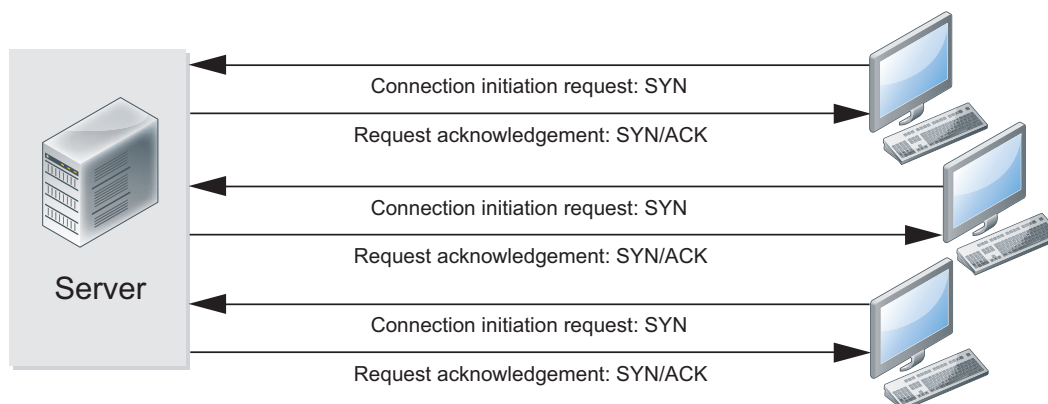
Figure 83: A client launches a SYN spoof attack



DDoS SYN flood

The most severe form of SYN attack is the distributed SYN flood, one variety of distributed denial of service attack (DDoS). Like the SYN flood, the target receives a flood of SYN packets and the ACK+SYN replies are never answered. The attack is distributed across multiple sources sending SYN packets in a coordinated attack.

Figure 84: Multiple attackers launch a distributed SYN flood



The distributed SYN flood is more difficult to defend against because multiple clients are capable of creating a larger volume of SYN packets than a single client. Even if the server can cope, the volume of traffic may overwhelm a point in the network upstream of the targeted server. The only defence against this is more bandwidth to prevent any choke-points.

Configuring the SYN threshold to prevent SYN floods

The preferred primary defence against any type of SYN flood is the DoS sensor `tcp_syn_flood` threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to *Pass*, the FortiGate unit will allow the SYN packets that exceed the threshold. If the action is set to *Block*, the FortiGate unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet when a SYN+ACK response is not received from the server. These tools will not “retry.” Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to *Block*.

For more information, see [“Creating and configuring a DoS sensor” on page 1082](#). For recommendations on how to configure DoS policies, see [“DoS policy recommendations” on page 892](#).

SYN proxy

FortiGate units with network acceleration hardware, whether built-in or installed in the form of an add-on module, offer a third action for the `tcp_syn_flood` threshold. Instead of *Block* and *Pass*, you can choose to *Proxy* the incomplete connections that exceed the threshold value.

When the `tcp_syn_flood` threshold action is set to *Proxy*, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the FortiGate unit receives an ACK response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

Other flood types

UDP and ICMP packets can also be used for DoS attacks, though they are less common. TCP SYN packets are so effective because the target receives them and maintains a session table entry for each until they time out. Attacks using UDP or ICMP packets do not require the same level of attention from a target, rendering them less effective. The target will usually drop the offending packets immediately, closing the session.

Use the `udp_flood` and `icmp_flood` thresholds to defend against these DoS attacks.

Traffic inspection

When the FortiGate unit examines network traffic one packet at a time for IPS signatures, it is performing traffic analysis. This is unlike content analysis where the traffic is buffered until files, email messages, web pages, and other files are assembled and examined as a whole.

DoS policies use traffic analysis by keeping track of the type and quantity of packets, as well as their source and destination addresses.

Application control uses traffic analysis to determine which application generated the packet.

Although traffic inspection doesn't involve taking packets and assembling files they are carrying, the packets themselves can be split into fragments as they pass from network to network. These fragments are reassembled by the FortiGate unit before examination.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats.

IPS signatures

IPS signatures can detect malicious network traffic. For example, the Code Red worm attacked a vulnerability in the Microsoft IIS web server. Your FortiGate's IPS system can detect traffic attempting to exploit this vulnerability. IPS may also detect when infected systems communicate with servers to receive instructions.

IPS recommendations

- Enable IPS scanning at the network edge for all services.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to FortiGuard IPS Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new IPS signatures as soon as they are available.

- Your FortiGate unit includes IPS signatures written to protect specific software titles from DoS attacks. Enable the signatures for the software you have installed and set the signature action to *Block*.

You can view these signatures by going to *UTM Profiles > Intrusion Protection > Predefined* and sorting by, or applying a filter to, the *Group* column.

- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to *Block*.

Suspicious traffic attributes

Network traffic itself can be used as an attack vector or a means to probe a network before an attack. For example, SYN and FIN flags should never appear together in the same TCP packet. The SYN flag is used to initiate a TCP session while the FIN flag indicates the end of data transmission at the end of a TCP session.

The FortiGate unit has IPS signatures that recognize abnormal and suspicious traffic attributes. The SYN/FIN combination is one of the suspicious flag combinations detected in TCP traffic by the `TCP.BAD.FLAGS` signature.

The signatures that are created specifically to examine traffic options and settings, begin with the name of the traffic type they are associated with. For example, signatures created to examine TCP traffic have signature names starting with TCP.

DoS policies

DDoS attacks vary in nature and intensity. Attacks aimed at saturating the available bandwidth upstream of your service can only be countered by adding more bandwidth. DoS policies can help protect against DDoS attacks that aim to overwhelm your server resources.

DoS policy recommendations

- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received.
- It is important to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, although your network may require very different values.

One way to find the correct values for your environment is to set the action to *Pass* and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

Application control

While applications can often be blocked by the ports they use, application control allows convenient management of all supported applications, including those that do not use set ports.

Application control recommendations

- Some applications behave in an unusual manner in regards to application control. For more information, see [“Application considerations” on page 1070](#).

- By default, application control allows the applications not specified in the application control list. For high security networks, you may want to change this behavior so that only the explicitly allowed applications are permitted.

Content inspection and filtering

When the FortiGate unit buffers the packets containing files, email messages, web pages, and other similar files for reassembly before examining them, it is performing content inspection. Traffic inspection, on the other hand, is accomplished by the FortiGate unit examining individual packets of network traffic as they are received.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats. Be sure to understand the effects of the changes before using the suggestions.

AntiVirus

The FortiGate antivirus scanner can detect viruses and other malicious payloads used to infect machines. The FortiGate unit performs deep content inspection. To prevent attempts to disguise viruses, the antivirus scanner will reassemble fragmented files and uncompress content that has been compressed. Patented Compact Pattern Recognition Language (CPRL) allows further inspection for common patterns, increasing detection rates of virus variations in the future.

AntiVirus recommendations

- Enable antivirus scanning at the network edge for all services.
- Use FortiClient endpoint antivirus scanning for protection against threats that get into your network.
- Subscribe to FortiGuard AntiVirus Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new antivirus signatures as soon as they are available.
- Enable the Extended Virus Database if your FortiGate unit supports it.
- Examine antivirus logs periodically. Take particular notice of repeated detections. For example, repeated virus detection in SMTP traffic could indicate a system on your network is infected and is attempting to contact other systems to spread the infection using a mass mailer.
- The *builtin-patterns* file filter list contains nearly 20 file patterns. Many of the represented files can be executed or opened with a double-click. If any of these file patterns are not received as a part of your normal traffic, blocking them may help protect your network. This also saves resources since files blocked in this way do not need to be scanned for viruses.
- To conserve system resources, avoid scanning email messages twice. Scan messages as they enter and leave your network or when clients send and retrieve them, rather than both.

FortiGuard Web Filtering

The web is the most popular part of the Internet and, as a consequence, virtually every computer connected to the Internet is able to communicate using port 80, HTTP. Botnet communications take advantage of this open port and use it to communicate with infected computers. FortiGuard Web Filtering can help stop infections from malware sites and help prevent communication if an infection occurs.

FortiGuard Web Filtering recommendations

- Enable FortiGuard Web Filtering at the network edge.
- Install the FortiClient application and use FortiGuard Web Filtering on any systems that bypass your FortiGate unit.
- Block categories such as Pornography, Malware, Spyware, and Phishing. These categories are more likely to be dangerous.
- In the email filter profile, enable *IP Address Check* in *FortiGuard Email Filtering*. Many IP addresses used in spam messages lead to malicious sites; checking them will protect your users and your network.

Email filter

Spam is a common means by which attacks are delivered. Users often open email attachments they should not, and infect their own machine. The FortiGate email filter can detect harmful spam and mark it, alerting the user to the potential danger.

Email filter recommendations

- Enable email filtering at the network edge for all types of email traffic.
- Use FortiClient endpoint scanning for protection against threats that get into your network.
- Subscribe to the FortiGuard AntiSpam Service.

DLP

Most security features on the FortiGate unit are designed to keep unwanted traffic out of your network while DLP can help you keep sensitive information from leaving your network. For example, credit card numbers and social security numbers can be detected by DLP sensors.

DLP recommendations

- Rules related to HTTP posts can be created, but if the requirement is to block all HTTP posts, a better solution is to use application control or the *HTTP POST Action* option in the web filter profile.
- While DLP can detect sensitive data, it is more efficient to block unnecessary communication channels than to use DLP to examine it. If you don't use instant messaging or peer-to-peer communication in your organization, for example, use application control to block them entirely.



AntiVirus

This section describes how to configure the antivirus options. From an antivirus profile you can configure the FortiGate unit to apply antivirus protection to HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP sessions. If your FortiGate unit supports SSL content scanning and inspection, you can also configure antivirus protection for HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.

The following topics are included in this section:

- [Antivirus concepts](#)
- [Enable antivirus scanning](#)
- [Enable the file quarantine](#)
- [Enable grayware scanning](#)
- [Testing your antivirus configuration](#)
- [Antivirus examples](#)

Antivirus concepts

The word “antivirus” refers to a group of features that are designed to prevent unwanted and potentially malicious files from entering your network. These features all work in different ways, which include checking for a file size, name, or type, or for the presence of a virus or grayware signature.

The antivirus scanning routines your FortiGate unit uses are designed to share access to the network traffic. This way, each individual feature does not have to examine the network traffic as a separate operation, and the overhead is reduced significantly. For example, if you enable file filtering and virus scanning, the resources used to complete these tasks are only slightly greater than enabling virus scanning alone. Two features do not require twice the resources.

How antivirus scanning works

Antivirus scanning examines files for viruses, worms, trojans, and malware. The antivirus scan engine has a database of virus signatures it uses to identify infections. If the scanner finds a signature in a file, it determines that the file is infected and takes the appropriate action.

The most thorough scan requires that the FortiGate unit have the whole file for the scanning procedure. To achieve this, the antivirus proxy buffers the file as it arrives. Once the transmission is complete, the virus scanner examines the file. If no infection is present, it is sent to the destination. If an infection is present, a replacement message is set to the destination.

During the buffering and scanning procedure, the client must wait. With a default configuration, the file is released to the client only after it is scanned. You can enable client comforting in the protocol options profile to feed the client a trickle of data to prevent them from thinking the transfer is stalled, and possibly cancelling the download.

Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Archives can also be expanded and the contents scanned, even if archives are nested.

Since the FortiGate unit has a limited amount of memory, files larger than a certain size do not fit within the memory buffer. The default buffer size is 10 MB. You can use the `uncompsizelimit` CLI command to adjust the size of this memory buffer.

Files larger than the buffer are passed to the destination without scanning. You can use the *Oversize File/Email* setting to block files larger than the antivirus buffer if allowing files that are too large to be scanned is an unacceptable security risk.

Flow-based antivirus scanning

If your FortiGate unit supports flow-based antivirus scanning, you can choose to select it instead of proxy-based antivirus scanning. Flow-based antivirus scanning uses the FortiGate IPS engine to examine network traffic for viruses, worms, trojans, and malware, without the need to buffer the file being checked.

The advantages of flow-based scanning include faster scanning and no maximum file size. Flow-based scanning doesn't require the file be buffered so it is scanned as it passes through the FortiGate unit, packet-by-packet. This eliminates the maximum file size limit and the client begins receiving the file data immediately.

The trade-off for these advantages is that flow-based scans detect a smaller number of infections. Viruses in documents, packed files, and some archives are less likely to be detected because the scanner can only examine a small portion of the file at any moment. Also, the file archive formats flow-based scanning will examine are limited to ZIP and GZIP.

Antivirus scanning order

The antivirus scanning function includes various modules and engines that perform separate tasks.

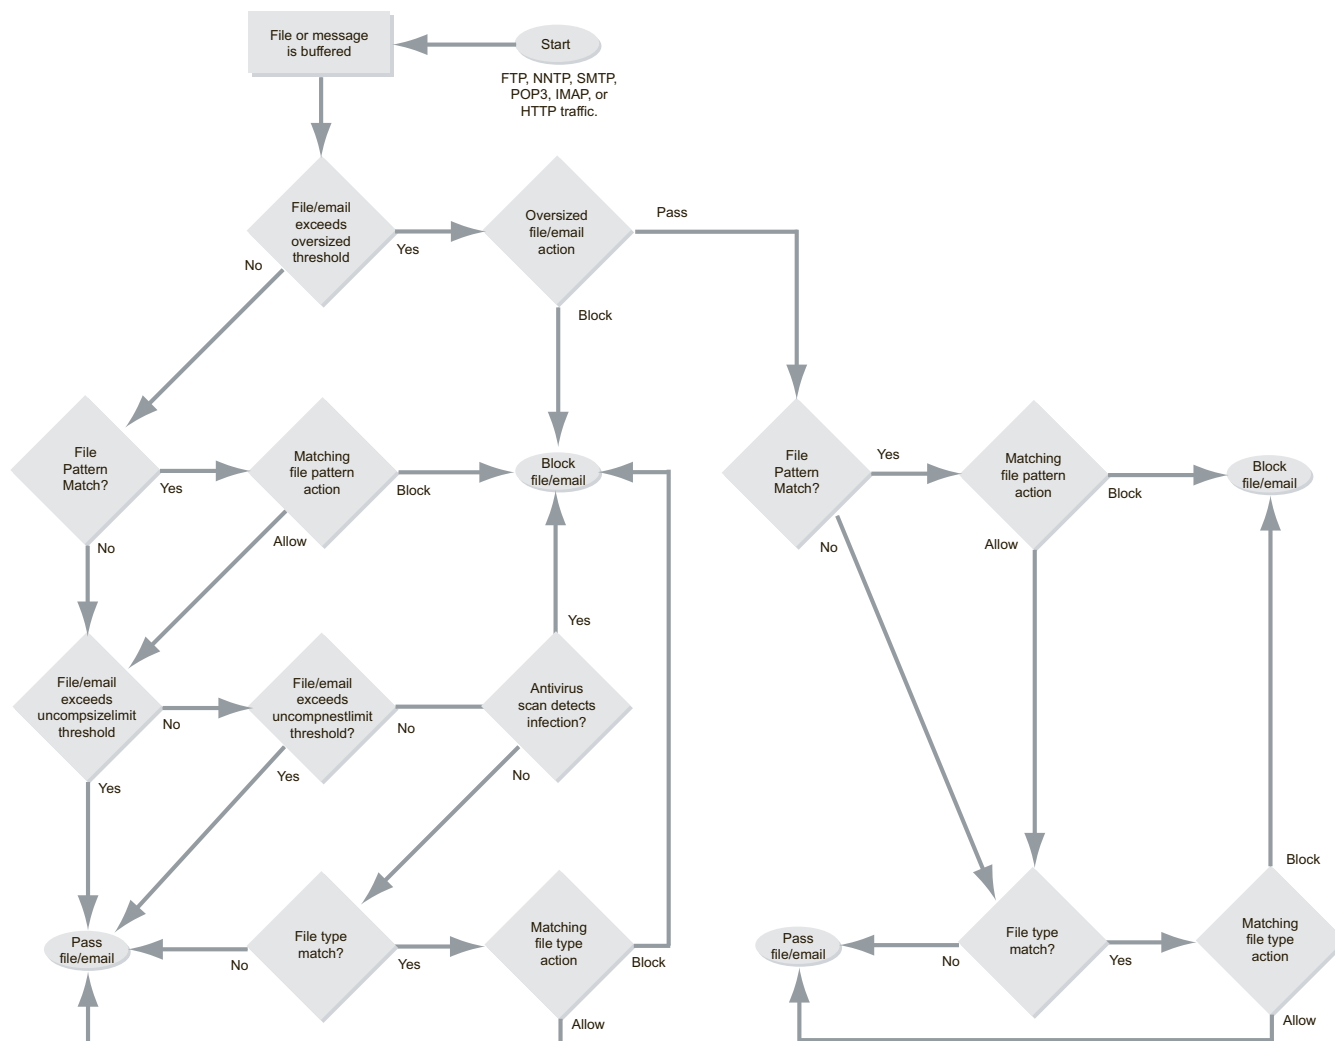
Proxy-based antivirus scanning order

Figure 85 on page 897 illustrates the antivirus scanning order when using proxy-based scanning (i.e. the normal, extended, or extreme databases). The first check for oversized files/email is to determine whether the file exceeds the configured size threshold. The `uncompsizelimit` check is to determine if the file can be buffered for file type and antivirus scanning. If the file is too large for the buffer, it is allowed to pass without being scanned. For more information, see the `config antivirus service` command in the [FortiGate CLI Reference](#). The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.



File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

Figure 85: Antivirus scanning order when using the normal, extended, or extreme database



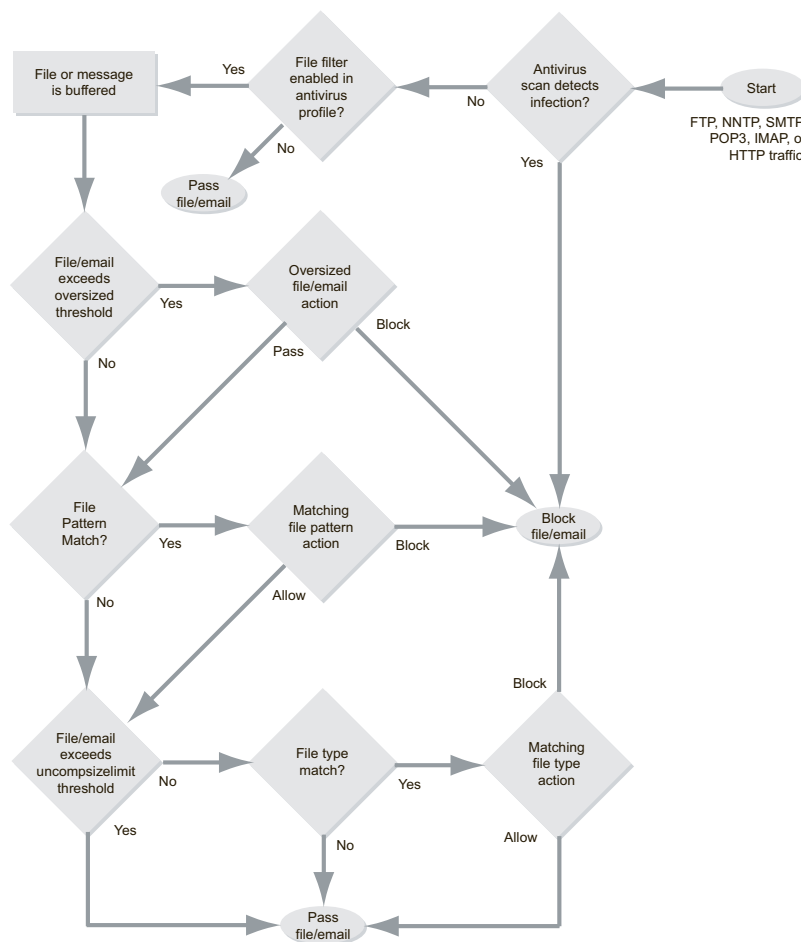
If a file fails any of the tasks of the antivirus scan, no further scans are performed. For example, if the file `fakefile.EXE` is recognized as a blocked file pattern, the FortiGate unit will send the end user a replacement message, and delete or quarantine the file. The unit will not perform virus scan, grayware, heuristics, and file type scans because the previous checks have already determined that the file is a threat and have dealt with it.

Flow-based antivirus scanning order

Figure 86 on page 898 illustrates the antivirus scanning order when using flow-based scanning (i.e. the flow-based database). The antivirus scan takes place before any other antivirus-related scan. If file filter is not enabled, the file is not buffered. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.



File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

Figure 86: Antivirus scanning order when using the flow-based database

Antivirus databases

The antivirus scanning engine relies on a database to detail the unique attributes of each infection. The antivirus scan searches for these signatures, and when one is discovered, the FortiGate unit determines the file is infected and takes action.

All FortiGate units have the normal antivirus signature database but some models have additional databases you can select for use. Which you choose depends on your network and security needs.

Normal	Includes viruses currently spreading as determined by the FortiGuard Global Security Research Team. These viruses are the greatest threat. The Normal database is the default selection and it is available on every FortiGate unit.
Extended	Includes the normal database in addition to recent viruses that are no-longer active. These viruses may have been spreading within the last year but have since nearly or completely disappeared.

Extreme	Includes the extended database in addition to a large collection of 'zoo' viruses. These are viruses that have not spread in a long time and are largely dormant today. Some zoo viruses may rely on operating systems and hardware that are no longer widely used.
Flow	The flow-based database is a subset of the extreme database. Flow-based scans can not detect polymorphic and packed-file viruses so those signatures are not included in the flow-based database. Note that flow-based scanning is not just another database, but a different type of scanning. For more information, see "How antivirus scanning works" on page 895 .

Antivirus techniques

The antivirus features work in sequence to efficiently scan incoming files and offer your network optimum antivirus protection. The first four features have specific functions, the fifth, heuristics, protects against any new, previously unknown virus threats. To ensure that your system is providing the most protection available, all virus definitions and signatures are updated regularly through the FortiGuard antivirus services. The features are discussed in the order that they are applied, followed by FortiGuard antivirus.

Virus scan

If the file passes the file pattern scan, the FortiGate unit applies a virus scan to it. The virus definitions are kept up-to-date through the FortiGuard Distribution Network (FDN). For more information, see ["FortiGuard Antivirus" on page 899](#).

Grayware

If the file passes the virus scan, it will be checked for grayware. Grayware configurations can be turned on and off as required and are kept up to date in the same manner as the antivirus definitions. For more information, see ["Enable grayware scanning" on page 906](#).

Heuristics

After an incoming file has passed the grayware scan, it is subjected to the heuristics scan. The FortiGate heuristic antivirus engine, if enabled, performs tests on the file to detect virus-like behavior or known virus indicators. In this way, heuristic scanning may detect new viruses, but may also produce some false positive results.



You can configure heuristics only through the CLI. See the [FortiGate CLI Reference](#).

FortiGuard Antivirus

FortiGuard Antivirus services are an excellent resource which includes automatic updates of virus and IPS (attack) engines and definitions, as well as the local spam DNS black list (DNSBL), through the FDN. The [FortiGuard Center](#) web site also provides the FortiGuard Antivirus virus and attack encyclopedia.

The connection between the FortiGate unit and FortiGuard Center is configured in *System > Maintenance > FortiGuard*.

Enable antivirus scanning

Antivirus scanning is enabled in the antivirus profile. Once the antivirus profile is enabled and selected in one or more firewall policies, all the traffic controlled by those firewall policies will be scanned according to your settings.

To enable antivirus scanning — web-based manager

- 1 Go to *UTM Profiles > AntiVirus > Profile*.
- 2 Select *Create New* to create a new antivirus profile, or select an existing antivirus profile and choose *Edit*.
- 3 In the row labeled *Virus Scan and Removal*, select the check boxes associated with the traffic you want scanned for viruses.
- 4 Select *OK*.

To enable antivirus scanning — CLI

You need to configure the scan option for each type of traffic you want scanned. In this example, antivirus scanning of HTTP traffic is enabled in the profile.

```
config antivirus profile
  edit my_av_profile
    config http
      set options scan
    end
  end
```

Viewing antivirus database information

The FortiGate antivirus scanner relies on up-to-date virus signatures to detect the newest threats. To view the information about the FortiGate unit virus signatures, check the status page or the *Virus Database* information page:

- **Status page:** Go to *System > Dashboard > Dashboard*. In the *License Information* section under *FortiGuard Services*, the *AV Definitions* field shows the regular antivirus database version as well as when it was last updated.

If your FortiGate unit supports extended and extreme virus database definitions, the database versions and date they were last updated is displayed in the *Extended set* and *Extreme DB* fields.

The flow-based virus database is distributed as part of the IPS signature database. Its database version and date it was last updated is displayed in the *IPS Definitions* field.

- **Virus Database:** Go to *UTM Profiles > AntiVirus > Virus Database*. This page shows the version number, number of included signatures, and a description of the regular virus database.

If your FortiGate unit supports extended, extreme, or flow-based virus database definitions, the version numbers, number of included signatures, and descriptions of those databases are also displayed.

Changing the default antivirus database

If your FortiGate unit supports extended, extreme, or flow-based virus database definitions, you can select the virus database most suited to your needs.

In most circumstances, the regular virus database provides sufficient protection. Viruses known to be active are included in the regular virus database. The extended database includes signatures of the viruses that have become rare within the last year in addition to those in the normal database. The extreme database includes legacy viruses that have not been seen in the wild in a long time in addition to those in the extended database.

The flow-based database contains a subset of the virus signatures in the extreme database. Unlike the other databases, selecting the flow-based database also changes the way the FortiGate unit scans your network traffic for viruses. Instead of the standard proxy-based scan, network traffic is scanned as it streams through the FortiGate unit. For more information on the differences between flow-based and proxy-based antivirus scanning, see [“How antivirus scanning works” on page 895](#).

If you require the most comprehensive antivirus protection, enable the extended virus database. The additional coverage comes at a cost, however, because the extra processing requires additional resources.

To change the antivirus database — web-based manager

- 1 Go to *UTM Profiles > AntiVirus > Virus Database*.
- 2 Select the antivirus database the FortiGate unit will use as the default database to perform antivirus scanning of your network traffic.
- 3 Select *Apply*.

To change the antivirus database — CLI

```
config antivirus settings
    set default-db extended
end
```

Overriding the default antivirus database

The default antivirus database is used for all antivirus scanning. If you have a particular policy or traffic type that requires scanning using a different antivirus database, you can override the default database. Antivirus database overrides are applied to individual traffic types in an antivirus profile. The override will affect only the traffic types to which the override is applied for the traffic handled by the security policy the antivirus profile is applied to. Antivirus database overrides can be set using only the CLI.

In this example, a database override is applied to HTTP traffic in a protocol options profile named `web_traffic`. The flow-based database is specified.

To override the default antivirus database — CLI

```
config antivirus profile
    edit web-traffic
        config http
            set avdb flow-based
        end
    end
```

With this configuration, the flow-based database is used for antivirus scans on HTTP traffic controlled by firewall policies in which this antivirus profile is selected. Other traffic types will use the default database, as specified in *UTM Profiles > AntiVirus > Virus Database*.

Adding the antivirus profile to a security policy

This procedure is required only if your antivirus profile does not yet belong to a security policy. You need to add the antivirus profile to a policy before any antivirus profile settings can take effect.

To add the antivirus profile to a policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New* to add a new policy, or select the *Edit* icon of the security policy to which you want to add the profile.
- 3 Enable *UTM*.
- 4 Select *Enable AntiVirus* and select the antivirus profile that contains the quarantine configuration.
- 5 Select *OK* to save the security policy.

Configuring the scan buffer size

When checking files for viruses using the proxy-based scanning method, there is a maximum file size that can be buffered. Files larger than this size are passed without scanning. The default size for all FortiGate models is 10 megabytes.

Archived files are extracted and email attachments are decoded before the FortiGate unit determines if they can fit in the scan buffer. For example, a 7 megabyte ZIP file containing a 12 megabyte EXE file will be passed without scanning with the default buffer size. Although the archive would fit within the buffer, the uncompressed file size will not.

In this example, the `uncompsizelimit` CLI command is used to change the scan buffer size to 20 megabytes for files found in HTTP traffic:

```
config antivirus service http
  set uncompsizelimit 20
end
```

The maximum buffer size varies by model. Enter `set uncompsizelimit ?` to display the buffer size range for your FortiGate unit.



Flow-based scanning does not use a buffer and therefore has no file-size limit. File data is scanned as it passes through the FortiGate unit. The `uncompsizelimit` setting has no effect for flow-based scanning.

Configuring archive scan depth

The antivirus scanner will open archives and scan the files inside. Archives within other archives, or nested archives, are also scanned to a default depth of twelve nestings. You can adjust the number of nested archives to which the FortiGate unit will scan with the `uncompnestlimit` CLI command. Further, the limit is configured separately for each traffic type.

For example, this CLI command sets the archive scan depth for SMTP traffic to 5. That is, archives within archives will be scanned five levels deep.

```
config antivirus service smtp
  set uncompnestlimit 5
end
```

You can set the nesting limit from 2 to 100.

Configuring a maximum allowed file size

The protocol option profile allows you to enforce a maximum allowed file size for each of the network protocols in the profile. They are HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP. If your FortiGate unit supports SSL content scanning and inspection, you can also configure a maximum file size for HTTPS, IMAPS, POP3S, SMTPS, and FTPS.

The action you set determines what the FortiGate unit does with a file that exceeds the oversized file threshold. Two actions are available:

Block	Files that exceed the oversize threshold are dropped and a replacement message is sent to the user instead of the file.
Pass	Files exceed the oversized threshold are allowed through the FortiGate unit to their destination. Note that passed files are not scanned for viruses. File Filtering, both file pattern and file type, are applied, however.

You can also use the maximum file size to help secure your network. If you're using a proxy-based virus scan, the proxy scan buffer size limits the size of the files that can be scanned for infection. Files larger than this limit are passed without scanning. If you configure the maximum file size to block files larger than the scan buffer size, large infected files will not by-pass antivirus scanning.

In this example, the maximum file size will be configured to block files larger than 10 megabytes, the largest file that can be antivirus scanned with the default settings. You will need to configure a protocol options profile and add it to a security policy.

Create a protocol options profile to block files larger than 10 MB

- 1 Go to *Policy > Policy > Protocol Options*.
- 2 Select *Create New*.
- 3 Enter `10MB_Block` for the protocol options policy name.
- 4 For the comment, enter `Files larger than 10MB are blocked`.
- 5 Expand each protocol listed and select *Block* for the *Oversized File/Email* setting. Also confirm that the *Threshold* is set to 10.
- 6 Select *OK*.

The protocol options profile is configured, but to block files, you must select it in the firewall profiles handling the traffic that contains the files you want blocked.

To select the protocol options profile in a security policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a security policy.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select `10MB_Block` from the *Protocol Options* list.
- 6 Select *OK* to save the security policy.

Once you complete these steps, any files in the traffic handled by this policy that are larger than 10MB will be blocked. If you have multiple firewall policies, examine each to determine if you want to apply similar file blocking the them as well.

Configuring client comforting

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit scans it. If no infection is found, the file is sent along to the client. The client initiates the file transfer and nothing happens until the FortiGate finds the file clean, and releases it. Users can be impatient, and if the file is large or the download slow, they may cancel the download, not realizing that the transfer is in progress.

The client comforting feature solves this problem by allowing a trickle of data to flow to the client so they can see the file is being transferred. The default client comforting transfer rate sends one byte of data to the client every ten seconds. This slow transfer continues while the FortiGate unit buffers the file and scans it. If the file is infection-free, it is released and the client will receive the remainder of the transfer at full speed. If the file is infected, the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.

Enable and configure client comforting

- 1 Go to *Policy > Policy > Protocol Options*.
- 2 Select a protocol options profile and choose *Edit*, or select *Create New* to make a new one.
- 3 Expand HTTP, FTP, and if your FortiGate unit supports SSL content scanning and inspection, expand HTTPS and FTPS as well.
- 4 To enable client comforting, select *Comfort Clients* for each of the protocols in which you want it enabled.
- 5 Select *OK* to save the changes.
- 6 Select this protocol options profile in any security policy for it to take effect on all traffic handled by the policy.

The default values for Interval and Amount are 10 and 1, respectively. This means that when client comforting takes effect, 1 byte of the file is sent to the client every 10 seconds. You can change these values to vary the amount and frequency of the data transferred by client comforting.

Enable the file quarantine

You can quarantine blocked and infected files if you have a FortiGate unit with a local hard disk. You can view the file name and status information about the file in the *Quarantined Files* list and submit specific files and add file patterns to the *AutoSubmit* list so they will automatically be uploaded to the FortiGuard AntiVirus service for analysis.

FortiGate units can also quarantine blocked and infected files to a FortiAnalyzer unit. Files stored on the FortiAnalyzer unit can also be viewed from the *Quarantined Files* list in the FortiGate unit.

General configuration steps

The following steps provide an overview of the file quarantine configuration. For best results, follow the procedures in the order given. Note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Go to *UTM Profiles > AntiVirus > Quarantine* to configure the quarantine service and destination.
- 2 Go to *UTM Profiles > AntiVirus > Profile* and edit an existing antivirus profile or create a new one. In the *Quarantine* row, select the check boxes of the protocols for which you want the quarantine enabled. The *Quarantine* option only appears if your FortiGate unit has a local disk or if your FortiGate unit is configured to use a FortiAnalyzer unit to quarantine files.



Antivirus profiles also have a configurable feature called *Quarantine Virus Sender (to Banned User List)*. This is a different feature unrelated to the *Quarantine* option.

- 3 If you have not previously done so, go to *Policy > Policy > Policy* and add the antivirus profile to a security policy.

Configuring the file quarantine

You can configure quarantine options for HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP Traffic. If your FortiGate unit supports SSL content scanning and inspection you can also quarantine blocked and infected files from HTTPS, IMAPS, POP3S, SMTPS, and FTPS traffic.

The quarantine configuration is only available in the CLI. See the [CLI Reference](#) for a full description of the `config antivirus quarantine` command.

In this example, the quarantine is configured to use the FortiGate unit disk, save files for 24 hours, use a maximum of 500 MB, and overwrite the oldest file with a new one should the disk space limit be exceeded.

To configure the file quarantine

```
config antivirus quarantine
  set destination disk
  set agelimit 24
  set quarantine-quota 500
  set lowspace ovrw-old
end
```

Viewing quarantined files

The *Quarantined Files* list displays information about each quarantined file. You can sort the files by file name, date, service, status, duplicate count (DC), or time to live (TTL). You can also filter the list to view only quarantined files from a specific service.

To view quarantined files, go to *Log&Report > Quarantined Files*.

Downloading quarantined files

You can download any non-expired file from the quarantine. You may want to do so if it was quarantined as the result of a false positive or if you want to examine the contents.

To download a quarantined file

- 1 Go to *Log&Report > Quarantined Files*.
- 2 In the quarantine file list, find the file you want to download.

To find the file more quickly, use the *Sort by* function to change the sort order. Available sort criteria include status, services, file name, date, TTL, and duplicate count. You can also use the *Filter* function to display the files quarantined from an individual traffic type.
- 3 Select the *Download* icon to save a copy of the quarantined file on your computer.

Enable grayware scanning

Grayware programs are unsolicited software programs installed on computers, often without the user's consent or knowledge. Grayware programs are generally considered an annoyance, but they can also cause system performance problems or be used for malicious purposes.

To allow the FortiGate unit to scan for known grayware programs, you must enable both antivirus scanning and grayware detection. By default, grayware detection is disabled. To enable antivirus scanning, see ["Enable antivirus scanning" on page 900](#).

To enable grayware detection — web-based manager

- 1 Go to *UTM Profiles > AntiVirus > Virus Database*.
- 2 Select *Enable Grayware Detection*.

To enable grayware detection — CLI

```
config antivirus settings
    set grayware enable
end
```

With grayware detection enabled, the FortiGate unit will scan for grayware any time it checks for viruses.

Testing your antivirus configuration

You have configured your FortiGate unit to stop viruses, but you'd like to confirm your settings are correct. Even if you have a real virus, it would be dangerous to use for this purpose. An incorrect configuration will allow the virus to infect your network.

To solve this problem, the European Institute of Computer Anti-virus Research has developed a test file that allows you to test your antivirus configuration. The EICAR test file is not a virus. It can not infect computers, nor can it spread or cause any damage. It's a very small file that contains a sequence of characters. Your FortiGate unit recognizes the EICAR test file as a virus so you can safely test your FortiGate unit antivirus configuration.

Go to <http://www.fortiguard.com/antivirus/eicartest.html> to download the test file (eicar.com) or the test file in a ZIP archive (eicar.zip).

If the antivirus profile applied to the security policy that allows you access to the Web is configured to scan HTTP traffic for viruses, any attempt to download the test file will be blocked. This indicates that you are protected.

Antivirus examples

The following examples provide a sample antivirus configuration scenario for a fictitious company.

Configuring simple antivirus protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antivirus protection on a FortiGate unit located in a satellite office. The satellite office does not have an internal email server. To send and retrieve email, the employees connect to an external mail server.

Creating an antivirus profile

Most antivirus settings are configured in an antivirus profile. Antivirus profiles are selected in firewall policies. This way, you can create multiple antivirus profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one antivirus profile.

To create an antivirus profile — web-based manager

- 1 Go to *UTM Profiles > AntiVirus > Profile*.
- 2 Select *Create New*.
- 3 In the *Name* field, enter `basic_antivirus`.
- 4 In the *Comments* field, enter *Antivirus protection for web and email traffic*.
- 5 Select the *Virus Scan* check boxes for the *HTTP*, *IMAP*, *POP3*, and *SMTP* traffic types.
- 6 Select *OK* to save the antivirus profile.

To create an antivirus profile — CLI

```
config antivirus profile
edit basic_antivirus
set comment "Antivirus protection for web and email traffic"
config http
set options scan
end
config imap
set options scan
end
```

```
config pop3
  set options scan
end
config smtp
  set options scan
end
end
```

Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the antivirus profile in a security policy — web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a policy.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select `default` from the *Protocol Options* list.
UTM can not be enabled without selecting a protocol options profile. A default profile is provided.
- 6 Select the *Enable AntiVirus* option.
- 7 Select the `basic_antivirus` profile from the list.
- 8 Select *OK* to save the security policy.

To select the antivirus profile in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set av-profile basic_antivirus
  end
```

HTTP, IMAP, POP3, and SMTP traffic handled by the security policy you modified will be scanned for viruses. A small office may have only one security policy configured. If you have multiple policies, consider enabling antivirus scanning for all of them.

Protecting your network against malicious email attachments

Grayware is commonly delivered by email or the web. The Example.com corporation has been the victim of multiple greyware infections in the past. Now that the company has a FortiGate unit protecting its network, you (Example.com's system administrator) can configure the unit to scan email and web traffic to filter out greyware attachments.

Enabling antivirus scanning in the antivirus profile

The primary means to avoid viruses is to configure the FortiGate unit to scan email and web traffic for virus signatures. You enable virus scanning in the antivirus profile and then select the antivirus profile in firewall policies that control email traffic.

To enable antivirus scanning in the antivirus profile

- 1 Go to *UTM Profiles > AntiVirus > Profile*.

- 2 Select *Create New* to add a new antivirus profile, or select the *Edit* icon of an existing antivirus profile.
- 3 Select the *Virus Scan* check box for *HTTP* to scan web traffic for viruses.
- 4 Select the *Virus Scan* check box for *IMAP*, *POP3*, and *SMTP* to scan all email protocols for viruses.
- 5 Select *OK* to save the antivirus profile.

Enabling grayware scanning

Grayware can also threaten Example.com's network. Viruses, email messages and the web are often the means by which grayware infections are delivered.

To enable grayware scanning

- 1 Go to *UTM Profiles > AntiVirus > Virus Database*.
- 2 Select *Enable Grayware Detection*.
- 3 Select *Apply*.

When *Enable Grayware Detection* is selected, virus scanning will also include grayware scanning. Any traffic scanned for viruses will also be scanned for grayware.

Selecting the antivirus profile in a security policy

An antivirus profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an antivirus profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the antivirus profile in a security policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select the policy that controls the network traffic controlling email traffic.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select the *Enable AntiVirus* option.
- 6 Select the antivirus profile from the list.
- 7 Select *OK* to save the security policy.

AntiVirus interface reference

The following explains the antivirus options that you can configure in the Antivirus menu. When configuring a profile, you can apply an antivirus profile to a firewall policy for HTTP, FTP, IMAP, POP3, SMTP, IM, and NNTP sessions. If your unit supports SSL content scanning and inspection you can also configure antivirus protection for HTTPS, IMAPS, POP3S, and SMTPS sessions.

This topic includes the following:

- [Profile](#)
- [Virus Database](#)

Profile

From the Profile submenu, you can configure antivirus profiles that are then applied to firewall policies. A profile is specific configuration information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

You can create multiple antivirus profiles for different antivirus scanning requirements. For example, you create an antivirus profile that specifies only virus scanning for POP3 which you then apply to the out-going firewall policy. You can also choose specific protocols, such as POP3, that will be blocked and then archived by the unit. This option is available only in the CLI.

Within antivirus profiles, you can also choose specific protocols to be blocked and then archive them. This is available only in the CLI.

Antivirus profile configuration settings

The following are antivirus profile configuration settings in *UTM Profiles > Antivirus > Profile*.

Profile page Lists each individual antivirus profile that you created. On this page, you can edit, delete or create a new antivirus profile. You are redirected to this page when you select <i>View List</i> on the Edit Antivirus Profile page. Note: If you want to configure the profile to block and archive specific protocols, use the <code>options</code> value in the <code>config antivirus profile</code> command in the CLI.	
Create New	Creates a new antivirus profile. When you select <i>Create New</i> , you are automatically redirected to the New Antivirus Profile page.
Edit	Modifies settings within the antivirus profile. When you select <i>Edit</i> , you are automatically redirected to the Edit Antivirus Profile page.
Delete	Removes an antivirus profile from the list on the Profile page. To remove multiple antivirus profiles from within the list, on the Antivirus Profile page, in each of the rows of the profiles you want to remove, select the check box and then select <i>Delete</i> . To remove all antivirus profiles in the list, on the Antivirus Profile page, select the check box in the check box column, and then select <i>Delete</i> .
Name	The name of the antivirus profile.
Comments	A description for the antivirus profile.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
<p>New Antivirus Profile page</p> <p>Provides settings for configuring a new antivirus profile. This page also allows you to configure quarantine settings for including a virus sender to the Banned User List.</p> <p>This page appears when you select <i>Create New</i> on the Edit Antivirus Profile page. If you are on the Profile page, and you select <i>Create New</i>, you will be redirected to the <i>New Antivirus Profile</i> page.</p> <p>Note: Logging is enabled in the CLI.</p>	
Name	Enter a name for the profile. If you are editing an existing antivirus profile and want to change the name, enter a new name in this field. You must select <i>OK</i> to save these changes.
Comments	Enter a description for the profile; this is optional. If you are editing an existing antivirus profile and want to change the description, enter the changes in this field. You must select <i>OK</i> to save the changes.
Virus Scan and Removal	Select any of the following to have the unit scan for viruses when the available protocols are used for web (Internet activity, for example HTTP), email (for example, POP3 or POP3S), and transferring files (for example, FTP).
Quarantine	<p>Select to enable the quarantine of detected viruses.</p> <p>Quarantined information is available in <i>Log&Report > Log & Archive Access</i>.</p>

Virus Database

The unit contains multiple antivirus databases for you to choose from, so that you can get the maximum protection that you need for your network environment. The Virus Database, located in *UTM Profiles > Antivirus > Virus Database*, is used to detect viruses in network traffic. The databases are available on the Virus Database page:

- Regular Virus Database
- Extended Virus Database
- Extreme Virus Database
- Flow-based Virus Database

On the Virus Database page, you can also enable grayware detection. This grayware detection includes adware, dial, downloader, hacker tool, keylogger, RAT, and spyware.

The extended database provides “in the wild” viruses as well as a large collection of zoo viruses that have not yet been seen in current virus studies. An enhanced security environment is best suited for this type of database. The flow-based database provides “in the wild” viruses as well as some commonly seen viruses on the network. Flow-based virus scanning is an alternative to the file-based virus scanning, providing better performance but lower coverage rates than the file-based virus scan.

The extreme antivirus database allows scanning for both “in the wild” and “zoo” viruses that are no longer seen in recent studies as well as all available signatures that are currently supported. The extreme database provides flexibility, providing the maximum protection without sacrificing performance and is suited to an enhanced security environment. The extreme antivirus database is available only on models that have AMC-enabled platforms and large capacity hard drives.

The flow-based antivirus database helps to detect malware using IPS. This database includes “in the wild” viruses along with some commonly seen viruses on the network. The flow-based antivirus database provides an alternative to the file-based virus scan while also providing better performance.

The FortiGuard virus definitions are updated when the unit receives a new version of FortiGuard antivirus definitions from the FDN.



The FDN updates only the virus database selected in UTM Profiles > AntiVirus > Virus Database. If you’ve configured the use of other databases in antivirus profiles that are selected in security policies, those will also be updated by the FDN. To save bandwidth, antivirus databases that are in use are not updated.

The [FortiGuard Center Virus Encyclopedia](#) contains detailed descriptions of the viruses, worms, trojans, and other threats that can be detected and removed by your unit using the information in the FortiGuard virus definitions.

The FortiGuard AV definitions are updated automatically from the FortiGuard Distribution Network (FDN). Automatic antivirus definition updates are configured from the FDN by going to *System > Maintenance > FortiGuard*. You can also update the antivirus definitions manually from the system dashboard by going to *System > Dashboard > Status*.



If virtual domains are enabled, you must configure antivirus settings in antivirus profiles separately for each virtual domain. Grayware settings can only be enabled or disabled when running FortiOS 4.0 MR2 or higher on the unit.



Email filter

This section describes how to configure FortiGate email filtering for IMAP, POP3, and SMTP email. Email filtering includes both spam filtering and filtering for any words or files you want to disallow in email messages. If your FortiGate unit supports SSL content scanning and inspection, you can also configure spam filtering for IMAPS, POP3S, and SMTPS email traffic.

The following topics are included in this section:

- [Email filter concepts](#)
- [Enable email filter](#)
- [Configure email traffic types to inspect](#)
- [Configure the spam action](#)
- [Configure the tag location](#)
- [Configure the tag format](#)
- [Configure FortiGuard email filters](#)
- [Configure local email filters](#)
- [Email filter examples](#)

Email filter concepts

You can configure the FortiGate unit to manage unsolicited commercial email by detecting and identifying spam messages from known or suspected spam servers.

The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools, to detect and block a wide range of spam messages. Using FortiGuard Antispam email filter profile settings, you can enable IP address checking, URL checking, email checksum checking, and spam submission. Updates to the IP reputation and spam signature databases are provided continuously via the global FortiGuard Distribution Network.

From the [FortiGuard Antispam Service](#) page in the FortiGuard Center, you can find out whether an IP address is blacklisted in the FortiGuard antispam IP reputation database, or whether a URL or email address is in the signature database.

Email filter techniques

The FortiGate unit has a number of techniques available to help detect spam. Some use the FortiGuard Antispam Service and require a subscription. The remainder use your DNS servers or use lists that you must maintain.

FortiGuard IP address check

The FortiGate unit queries the FortiGuard Antispam Service to determine if the IP address of the client delivering the email is blacklisted. A match will cause the FortiGate unit to treat delivered messages as spam.

The default setting of the `smtp-spamhdrop` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the FortiGuard Antispam Service. For more information, see the [FortiGate CLI Reference](#).

FortiGuard URL check

The FortiGate unit queries the FortiGuard Antispam service to determine if any URL in the message body is associated with spam. If any URL is blacklisted, the FortiGate unit determines that the email message is spam.

Detect phishing URLs in email

The FortiGate unit sends the URL links in email messages to FortiGuard to determine if the links are associated with a known phishing site. If such a link is detected, the link is removed from the message. The URL remains, but it is no longer a selectable hyperlink.

FortiGuard email checksum check

The FortiGate unit sends a hash of an email to the FortiGuard Antispam server, which compares the hash to hashes of known spam messages stored in the FortiGuard Antispam database. If the hash results match, the email is flagged as spam.

FortiGuard spam submission

Spam submission is a way you can inform the FortiGuard AntiSpam service of non-spam messages incorrectly marked as spam. When you enable this setting, the FortiGate unit adds a link to the end of every message marked as spam. You then select this link to inform the FortiGuard AntiSpam service when a message is incorrectly marked.

IP address black/white list check

The FortiGate unit compares the IP address of the client delivering the email to the addresses in the IP address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry against all delivered email.

The default setting of the `smtp-spamhdrop` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the specified IP address black/white list. For more information, see the [FortiGate CLI Reference](#).

HELO DNS lookup

The FortiGate unit takes the domain name specified by the client in the HELO greeting sent when starting the SMTP session and does a DNS lookup to determine if the domain exists. If the lookup fails, the FortiGate unit determines that any messages delivered during the SMTP session are spam.

Email address black/white list check

The FortiGate unit compares the sender email address, as shown in the message envelope MAIL FROM, to the addresses in the email address black/white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry.

Return email DNS check

The FortiGate unit performs a DNS lookup on the reply-to domain to see if there is an A or MX record. If no such record exists, the message is treated as spam.

Banned word check

The FortiGate unit blocks email messages based on matching the content of the message with the words or patterns in the selected spam filter banned word list.

Order of spam filtering

The FortiGate unit checks for spam using various filtering techniques. The order in which the FortiGate unit uses these filters depends on the mail protocol used.

Filters requiring a query to a server and a reply (FortiGuard Antispam Service and DNSBL/ORDBL) are run simultaneously. To avoid delays, queries are sent while other filters are running. The first reply to trigger a spam action takes effect as soon as the reply is received.

Each spam filter passes the email to the next if no matches or problems are found. If the action in the filter is *Mark as Spam*, the FortiGate unit tags the email as spam according to the settings in the email filter profile.

For SMTP and SMTPS, if the action is discard, the email message is discarded or dropped.

If the action in the filter is *Mark as Clear*, the email is exempt from any remaining filters. If the action in the filter is *Mark as Reject*, the email session is dropped. Rejected SMTP or SMTPS email messages are substituted with a configurable replacement message.

Order of SMTP and SMTPS spam filtering

The FortiGate unit scans SMTP and SMTPS email for spam in the order given below. SMTPS spam filtering is available on FortiGate units that support SSL content scanning and inspection.

- 1 IP address black/white list (BWL) check on last hop IP
- 2 DNSBL & ORDBL check on last hop IP, FortiGuard Antispam IP check on last hop IP, HELO DNS lookup
- 3 MIME headers check, E-mail address BWL check
- 4 Banned word check on email subject
- 5 IP address BWL check (for IPs extracted from "Received" headers)
- 6 Banned word check on email body
- 7 Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check on public IP extracted from header.

Order of IMAP, POP3, IMAPS and POP3S spam filtering

The FortiGate unit scans IMAP, POP3, IMAPS and POP3S email for spam in the order given below. IMAPS and POP3S spam filtering is available on FortiGate units that support SSL content scanning and inspection.

- 1 MIME headers check, E-mail address BWL check
- 2 Banned word check on email subject
- 3 IP BWL check
- 4 Banned word check on email body
- 5 Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check.

Enable email filter

Unlike antivirus protection, no single control enables all email filtering. Your FortiGate unit uses many techniques to detect spam; some may not be appropriate for every situation. For this reason, when you enable email filtering, you must then choose when techniques are applied to email traffic.

To enable email traffic inspection

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Select *Enable Spam Detection and Filtering*.
- 4 Select *Apply*.

Once you allow the FortiGate unit to examine one or more types of email traffic, you can enable any of the individual email filtering techniques.

Configure email traffic types to inspect

The FortiGate unit can examine IMAP, POP3, and SMTP email traffic. If your FortiGate unit supports content inspection, it can also examine IMAPS, POP3S, and SMTPS traffic. You can select which types of email traffic are examined by each email filter profile.

To select the email traffic types to inspect

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 If *Enable Spam Detection and Filtering* is enabled, the row immediately below shows a check box for each traffic type. Select the traffic types you want the FortiGate unit to examine when using this email filter profile.
- 4 Select *Apply*.

The traffic types you enable will be examined according to the settings in the email filter profile.

Configure the spam action

When spam is detected, the FortiGate unit will deal with it according to the *Spam Action* setting in the email filter profile. Note that POP3S, IMAPS and SMTPS spam filtering is available only on FortiGate units that support SSL content scanning and inspection. POP3, IMAP, POP3S and IMAPS mail can only be tagged. SMTP and SMTPS mail can be set to *Discard* or *Tagged*:

- **Discard:** When the spam action is set to *Discard*, messages detected as spam are deleted. No notification is sent to the sender or recipient.
- **Tagged:** When the spam action is set to *Tagged*, messages detected as spam are labelled and delivered normally. The text used for the label is set in the *Tag Format* field and the label is placed in the subject or the message header, as set with the *Tag Location* option.

To configure the spam action

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 The *Spam Action* row has a drop-down selection under the SMTP and SMTPS traffic type. Select *Discard* or *Tagged*.

No selection is available for POP3, IMAP, POP3S or IMAPS traffic. *Tagged* is the only applicable action for those traffic types.

By default, the tag location for any traffic set to *Tagged* is *Subject* and the tag format is *Spam*. If you want to change these settings, continue with [“Configure the tag location” on page 917](#) and [“Configure the tag format” on page 917](#).

- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Configure the tag location

When the spam action is set to *Tagged*, the *Tag Location* setting determines where the tag is applied in the message.

To configure the tag location

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 The *Tag Location* row has two options for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag location will not be available. Select the tag location:
 - *Subject*: The FortiGate unit inserts the tag at the beginning of the message subject. For example, if the message subject is “Buy stuff!” and the tag is “[spam]”, the new message subject is “[spam] Buy stuff!” if the message is detected as spam.
 - *MIME*: The FortiGate unit inserts the tag into the message header. With most mail readers and web-based mail services, the tag will not be visible. Despite this, you can still set up a rule based on the presence or absence of the tag.
- 4 Select *Apply*.

Configure the tag format

When the spam action is set to *Tagged*, the *Tag Format* setting determines what text is used as the tag applied to the message.

To configure the tag format

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.

- 3 The *Tag Format* row has a field for each traffic type. Note that if the spam action for SMTP traffic is set to discard, the tag format will not be available. Enter the text the FortiGate unit will use as the tag for each traffic type.
- 4 Select *Apply*.

Configure FortiGuard email filters

FortiGuard email filtering techniques use FortiGuard services to detect the presence of spam among your email. A FortiGuard subscription is required to use the FortiGuard email filters.

Enabling FortiGuard IP address checking

When you enable FortiGuard IP address checking, your FortiGate unit will submit the IP address of the client to the FortiGuard service for checking. If the IP address exists in the FortiGuard IP address black list, your FortiGate unit will treat the message as spam.

To enable FortiGuard IP address checking

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *FortiGuard Spam Filtering*, select *IP Address Check*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling FortiGuard URL checking

When you enable FortiGuard IP address checking, your FortiGate unit will submit all URLs appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL black list, your FortiGate unit will treat the message as spam.

To enable FortiGuard URL checking

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *FortiGuard Spam Filtering*, select *URL Check*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling FortiGuard phishing URL detection

When you enable FortiGuard phishing URL detection, your FortiGate unit will submit all URL hyperlinks appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL phishing list, your FortiGate unit will remove the hyperlink from the message. The URL will remain in place, but it will no longer be a selectable hyperlink.

To enable FortiGuard phishing URL detection

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *FortiGuard Spam Filtering*, select *Detect Phishing URLs in Email*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling FortiGuard email checksum checking

When you enable FortiGuard email checksum checking, your FortiGate unit will submit a checksum of each email message to the FortiGuard service for checking. If a checksum exists in the FortiGuard checksum black list, your FortiGate unit will treat the message as spam.

To enable FortiGuard checksum checking

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *FortiGuard Email Filtering*, select *E-mail Checksum Check*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling FortiGuard spam submission

When you enable FortiGuard email checksum checking, your FortiGate unit will append a link to the end of every message detected as spam. This link allows email users to “correct” the FortiGuard service by informing it that the message is not spam.



Carefully consider the use of the *Spam submission* option on email leaving your network. Users not familiar with the feature may click the link on spam messages because they are curious. This will reduce the accuracy of the feature.

To enable FortiGuard Spam submission

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *FortiGuard Email Filtering*, select *Spam Submission*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Configure local email filters

Local email filtering techniques use your own resources, whether DNS checks or IP address and email address lists that you maintain.

Enabling IP address black/white list checking

When you enable IP address black/white list checking, your FortiGate unit will compare the client IP address with the IP address black/white list specified in the email filter profile. If the client IP address exists, the FortiGate unit acts according to the action configured for the IP address in the list: allow the message, reject it, or mark it as spam.

The next two topics describe adding and configuring the IP address black/white list that you will need before you can enable the checking. If you already have this list, go to [“Enabling the IP address black/white list checking” on page 921](#).

Creating an IP address black/white list

Before you can enable IP address black/white list spam filtering in the email filter profile, you must create an IP address black/white list.

To create an IP address black/white list

- 1 Go to *UTM Profiles > Email Filter > IP Address*.
- 2 Select *Create New*.
- 3 Enter a name for the IP address list.
- 4 Optionally, enter a description or comments about the list.
- 5 Select *OK* to save the IP address black/white list.

When a new IP address black/white list is created, it is empty. To perform any actions, you must add IP addresses to the list.

Adding addresses to an IP address black/white list

Each IP address black/white list contains a number of IP addresses, each having a specified action. When the FortiGate unit accepts mail from a client with an IP address on the IP address black/white list specified in the active email filter profile, it performs the action specified for the address.

To add an address to an IP address black/white list

- 1 Go to *UTM Profiles > Email Filter > IP Address*.
- 2 Select the list to which you want to add an address and choose *Edit*.
- 3 Select *Create New*.
- 4 Enter the address or netmask in the IP/netmask field.

- 5 Select the action:
 - *Mark as Clear*: Messages from clients with matching IP addresses will be allowed, bypassing further email filtering.
 - *Mark as Reject*: Messages from clients with matching IP addresses will be rejected. The FortiGate unit will return a reject message to the client. *Mark as Reject* only applies to mail delivered by SMTP. If an IP address black/white list is used with POP3 or IMAP mail, addresses configured with the *Mark as Reject* action will be marked as spam.
 - *Mark as Spam*: Messages from clients with matching IP addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see [“Configure the spam action” on page 916](#).
- 6 By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.
- 7 Select OK.

Enabling the IP address black/white list checking

Once you have created a black/white list and added the IP addresses, you can enable IP address the checking.

To enable IP address black/white list checking

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *Local Spam Filtering*, select *IP Address BWL Check*.
- 4 Select the IP address black/white list to use from the drop-down list.
- 5 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling HELO DNS lookup

Whenever a client opens an SMTP session with a server, the client sends a HELO command with the client domain name. When you enable HELO DNS lookup, your FortiGate unit will take the domain the client submits as part of the HELO greeting and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat all messages the client delivers as spam.

The HELO DNS lookup is available only for SMTP traffic.

To enable HELO DNS lookup

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *Local Spam Filtering*, select *HELO DNS Lookup*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling email address black/white list checking

When you enable email address black/white list checking, your FortiGate unit will compare the sender email address with the email address black/white list specified in the email filter profile. If the sender email address exists, the FortiGate unit acts according to the action configured for the email address in the list: allow the message or mark it as spam.

The next two topics describe adding and configuring the email address black/white list that you will need before you can enable the checking. If you already have this list, go to [“Enabling email address black/white list checking” on page 923](#).

Creating an email address black/white list

Before you can enable email address black/white list spam filtering in the email filter profile, you must create an email address black/white list.

To create an email address black/white list

- 1 Go to *UTM Profiles > Email Filter > E-mail Address*.
- 2 Select *Create New*.
- 3 Enter a name for the email address list.
- 4 Optionally, enter a description or comments about the list.
- 5 Select *OK* to save the email address black/white list.

When a new IP address back/white list is created, it is empty. To perform any actions, you must add email addresses to the list.

Adding addresses to an email address black/white list

Each email address black/white list may contain a number of email addresses, each having a specified action. When the FortiGate unit accepts an email message from a client with a reply-to address that appears in the email address black/white list specified in the active email filter profile, it performs the action specified for the email message.

To add an address to an email address black/white list

- 1 Go to *UTM Profiles > Email Filter > E-mail Address*.
- 2 Select the *Edit* icon of the list to which you want to add an address.
- 3 Select *Create New*.
- 4 Enter the email address in the *Email Address* field.
- 5 If you need to enter a pattern in the *Email Address* field, select whether to use wildcards or regular expressions to specify the pattern.

Wildcard uses an asterisk (“*”) to match any number of any character. For example, *@example.com will match all addresses ending in @example.com.

Regular expressions use Perl regular expression syntax. See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.

- 6 Select the action:
 - *Mark as Spam*: Messages with matching reply-to email addresses will be treated as spam, subject to the action configured in the applicable email filter profile. For more information, see [“Configure the spam action” on page 916](#).
 - *Mark as Clear*: Messages with matching reply-to addresses will be allowed, bypassing further email filtering.

- 7 By default, the address is enabled and the FortiGate unit will perform the action if the address is detected. To disable checking for the address, clear the *Enable* check box.
- 8 Select *OK* to save the address.

Enabling email address black/white list checking

Once you have created a black/white list and added the email addresses, you can enable the checking.

To enable email address black/white list checking

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *Local Spam Filtering*, select *E-mail Address BWL Check*.
- 4 Select the email address black/white list to use from the drop-down list.
- 5 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling return email DNS checking

When you enable return email DNS checking, your FortiGate unit will take the domain in the reply-to email address and send it to the configured DNS. If the domain does not exist, your FortiGate unit will treat the message as spam.

To enable return email DNS check

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3 Under the heading *Local Spam Filtering*, select *Return E-mail DNS Check*.
- 4 Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Enabling banned word checking

When you enable banned word checking, your FortiGate unit will examine the email message for words appearing in the banned word list specified in the email filter profile. If the total score of the banned word discovered in the email message exceeds the threshold value set in the email filter profile, your FortiGate unit will treat the message as spam.

When determining the banned word score total for an email message, each banned word score is added once no matter how many times the word appears in the message.

The next two topics describe adding and configuring the banned word list that you will need before you can enable the checking. If you already have this list, go to [“Enabling banned word checking” on page 926](#).

How content is evaluated

Every time the banned word filter detects a pattern in an email message, it adds the pattern score to the sum of scores for the message. You set this score when you create a new pattern to block content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the total score equals or exceeds the threshold, the email message is considered as spam and treated according to the spam action configured in the email filter profile. The score for each pattern is counted only once, even if that pattern appears many times in the email message. The default score for banned word patterns is 10 and the default threshold is 10. This means that by default, an email message is blocked by a single match.

A pattern can be part of a word, a whole word, or a phrase. Multiple words entered as a pattern are treated as a phrase. The phrase must appear as entered to match. You can also use wildcards or regular expressions to have a pattern match multiple words or phrases.

For example, the FortiGate unit scans an email message that contains only this sentence: "The score for each word or phrase is counted only once, even if that word or phrase appears many times in the email message."

Banned word pattern	Pattern type	Assigned score	Score added to the sum for the entire page	Comment
word	Wildcard	20	20	The pattern appears twice but multiple occurrences are only counted once.
word phrase	Wildcard	20	0	Although each word in the phrase appears in the message, the words do not appear together as they do in the pattern. There are no matches.
word*phrase	Wildcard	20	20	The wildcard represents any number of any character. A match occurs as long as "word" appears before "phrase" regardless of what is in between them.
mail*age	Wildcard	20	20	Since the wildcard character can represent any characters, this pattern is a match because "email message" appears in the message.

In this example, the message is treated as spam if the banned word threshold is set to 60 or less.

Creating a banned word list

Before you can enable IP address black/white list spam filtering in the email filter profile, you must create an IP address black/white list.

To create an IP address black/white list

- 1 Go to *UTM Profiles > Email Filter > Banned Word*.
- 2 Select *Create New*.
- 3 Enter a name for the banned word list.
- 4 Optionally, enter a description or comments about the list.
- 5 Select *OK* to save the banned word list.

When a new banned word list is created, it is empty. To perform any actions, you must add words to the list.

Adding words to a banned word list

Each banned word list contains a number of words, each having a score, and specifying whether the email FortiGate unit will search for the word in the message subject, message body, or both.

When the FortiGate unit accepts an email message containing one or more words in the banned word list specified in the active email filter profile, it totals the scores of the banned words in the email message. If the total is higher than the threshold set in the email filter profile, the email message will be detected as spam. If the total score is lower than the threshold, the message will be allowed to pass as normal.

The score of a banned word present in the message will be counted toward the score total only once, regardless of how many times the word appears in the message.

To add words to a banned word list

- 1 Go to *UTM Profiles > Email Filter > Banned Word*.
- 2 Select the list to which you want to add a word.
- 3 Select *Edit*.
- 4 Select *Create New*.
- 5 Enter the word or the pattern in the *Pattern* field.
- 6 In the *Pattern Type* field, select whether you use wildcards or regular expressions.
Wildcard uses an asterisk (“*”) to match any number of any character. For example, *re** will match all words starting with “re”.
Regular expression uses Perl regular expression syntax. See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.
- 7 In the *Language* field, select the language.
- 8 Select where the FortiGate unit will check for the banned word. The options are *Body*, *Subject*, or *All*, which combines the other two options.
- 9 Enter a score. If the word appears in the message as determined by the *Where* setting, the score is added to the scores of all the other banned words appearing in the email message. If the score total is higher than the threshold set in the email filter profile, the email message will be detected as spam. If the total score is lower than the threshold, the message will be allowed to pass as normal.

10 By default, the banned word is enabled and will appear in the list. To disable checking for the banned word, clear the *Enable* check box.

11 Select *OK* to save the banned word.

Enabling banned word checking

Once you have created a black/white list and added the email addresses, you can enable the checking.

To enable banned word checking

- 1** Go to *UTM Profiles > Email Filter > Profile*.
- 2** The default email filter profile is presented. To edit another profile, select it from the drop down in the *Edit Email Filter Profile* title bar.
- 3** Under the heading *Local Spam Filtering*, select *Banned Word Check*.
- 4** Select the banned word list to use from the drop-down list.
- 5** Enter a threshold value. If the total score of the banned words appearing in the message exceeds this threshold, the FortiGate unit treats the message as spam.
- 6** Select *Apply*.

Select the edited email filter profile in a security policy, and the traffic controlled by the security policy will be scanned according to the settings you configured. You may select the email filter profile in more than one security policy if required.

Email filter examples

Configuring simple antispam protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antispam protection on a FortiGate unit located in a satellite office.

Creating an email filter profile

Most email filter settings are configured in an email filter profile. Email filter profiles are selected in firewall policies. This way, you can create multiple email filter profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one email filter profile.

To create an email filter profile — web-based manager

- 1** Go to *UTM Profiles > Email Filter > Profile*.
- 2** Select the *Create New* icon in the Edit Email Filter Profile window title.
- 3** In the *Name* field, enter `basic_emailfilter`.
- 4** Select *Enable Spam Detection and Filtering*.
- 5** Ensure that *IMAP*, *POP3*, and *SMTP* are selected in the header row.
These header row selections enable or disable examination of each email traffic type. When disabled, the email traffic of that type is ignored by the FortiGate unit and no email filtering options are available.
- 6** Under *FortiGuard Spam Filtering*, enable *IP Address Check*.
- 7** Under *FortiGuard Spam Filtering*, enable *URL Check*.
- 8** Under *FortiGuard Spam Filtering*, enable *E-mail Checksum Check*.

- 9 Select *OK* to save the email filter profile.

To create an email filter profile — CLI

```
config spamfilter profile
edit basic_emailfilter
    set options spamfsip spamfsurl spamfschksum
end
```

Selecting the email filter profile in a security policy

An email filter profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an email filter profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the email filter profile in a security policy — web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a policy.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select the *Enable Email Filter* option.
- 6 Select the `basic_emailfilter` profile from the list.
- 7 Select *OK* to save the security policy.

To select the email filter profile in a security policy — CLI

```
config firewall policy
edit 1
    set utm-status enable
    set profile-protocol-options default
    set spamfilter-profile basic_emailfilter
end
```

IMAP, POP3, and SMTP email traffic handled by the security policy you modified will be scanned for spam. Spam messages have the text “Spam” added to their subject lines. A small office may have only one security policy configured. If you have multiple policies, consider enabling spam scanning for all of them.

Blocking email from a user

Employees of the Example.com corporation have been receiving unwanted email messages from a former client at a company called example.net. All ties between the company and the client have been severed, but the messages continue. The FortiGate unit can be configured to prevent these messages from being delivered.

To create the email address list

- 1 Go to *UTM Profiles > Email Filter > E-mail Address*.
- 2 Select *Create New*.
- 3 Enter a name for the new email address list.
- 4 Optionally, enter a descriptive comment for the email address list.
- 5 Select *OK* to create the list.
- 6 Select *Create New* to add a new entry to the email address list.
- 7 Enter `*@example.net` in the *E-mail Address* field.

- 8 Leave *Pattern Type* set to the default, *Wildcard*.
- 9 Leave *Action* as *Mark as Spam* to have the FortiGate unit mark all messages from example.net as spam.

Now that the email address list is created, you must enable the email filter in the email filter profile.

To enable Email Filter

- 1 Go to *UTM Profiles > Email Filter > Profile*.
- 2 Select the email filter profile that is used by the firewall policies handling email traffic from the email filter profile drop down list.
- 3 In the row *Tag Location*, select *Subject* for all three mail protocols.
- 4 In the row *Tag Format*, enter *SPAM:* in all three fields.
- 5 Select *Enable Spam Detection and Filtering*.
- 6 Ensure that the check boxes labeled *IMAP*, *POP3*, and *SMTP* in the header row are selected.
- 7 Under *Local Spam Filtering*, enable *E-mail Address BWL Check* and select the email address list you created in the previous procedure from the drop down list.
- 8 Select *OK*.

When this email filter profile is selected in a security policy, the FortiGate unit will add "SPAM:" to the subject of any email message from an address ending with @example.net for all email traffic handled by the security policy. Recipients can ignore the message or they can configure their email clients to automatically delete messages with "SPAM:" in the subject.

Email Filter interface reference

Reports page Provides settings for configuring a report that you generated. The information for these reports are taken from a web filter profile. You must first configure a web filter profile before you can generate a report.	
Web Filter Profile	Select the web filter profile that you want to see a report based on.
Clear report data	Removes all data within the report that you are currently viewing.
Report Type	Select the time period for the report. Choose from <i>Hour</i> , <i>Day</i> , or <i>All</i> .
Report Range	Select the time range (format is in the 24 hour clock) or day range (from six days ago to today) for the report. For example, for an "hour" report type with a range of 13 to 16, the result is a category block report for 1 pm and 4 pm today. For a "day" report type with a range of 0 to 3, the result is a category block report for three days ago from today.
Get Report	Select to generate a report.
The generated report includes the following columns that appear below the pie chart on the Reports page:	
Category	The category for which the statistic was generated.

Allowed	The number of allowed web addresses accessed in the selected time frame.
Blocked	The number of blocked web addresses accessed in the selected time frame.
Logged	The logged web filter information.
Overridden	The blocked instances where an override was allowed.

If your unit supports SSL content scanning and inspection you can also configure email filtering for IMAPS, POP3S, and SMTPS email traffic.

You can configure the unit to manage unsolicited commercial email by identifying spam messages from known or suspected spam servers.

The [FortiGuard Antispam Service](#) uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools, to detect and block a wide range of spam messages. Using FortiGuard Email filtering profile settings you can enable IP address checking, URL checking, E-mail checksum checking, and Spam submission. Updates to the IP reputation and spam signature databases are provided continuously from the global FortiGuard distribution network.

From the [FortiGuard Antispam Service](#) page in the FortiGuard center you can use IP and signature lookup to check whether an IP address is blacklisted in the FortiGuard antispam IP reputation database, or whether a URL or email address is in the signature database.

This topic contains the following:

- [Order of email filteringProfile](#)
- [Banned Word](#)
- [IP Address](#)
- [Email Filter interface reference](#)

Order of email filtering

Email filtering uses various filtering techniques. The order the unit uses these filters depends on the mail protocol that is used.

Filters requiring a query to a server and a reply (FortiGuard Antispam Service and DNSBL/ORDBL) are run simultaneously. To avoid delays, queries are sent while other filters are running. The first reply to trigger a spam action takes effect as soon as the reply is received.

Each filter passes the email to the next if no matches or problems are found. If the action in the filter is Mark as Spam, the FortiGate unit tags as spam the email according to the settings in the email filter profile.

For SMTP and SMTPS if the action is discard the email message is discarded or dropped.

If the action in the filter is *Mark as Clear*, the email is exempt from any remaining filters. If the action in the filter is *Mark as Reject*, the email session is dropped. Rejected SMTP or SMTPS email messages are substituted with a configurable replacement message.

Order of SMTP and SMTPS email filtering

SMTPS email filtering is available on units that support SSL content scanning and inspection.

- 1 IP address BWL check on last hop IP.

- 2 DNSBL & ORDBL check on last hop IP, FortiGuard Email Filtering IP address check on last hop IP, HELO DNS lookup.
- 3 MIME headers check, E-mail address BWL check.
- 4 Banned word check on email subject.
- 5 IP address BWL check (for IPs extracted from “Received” headers).
- 6 Banned word check on email body.
- 7 Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Email Filtering URL check, DNSBL & ORDBL check on public IP extracted from header.

Order of IMAP, POP3, IMAPS and POP3S email filtering

IMAPS and POP3S email filtering is available on units the support SSL content scanning and inspection.

- 1 MIME headers check, Email address BWL check.
- 2 Banned word check on email subject.
- 3 IP BWL check.
- 4 Banned word check on email body.
- 5 Return email DNS check, FortiGuard Email Filtering email checksum check, FortiGuard Email Filtering URL check, DNSBL & ORDBL check.

Profile

The Profile menu allows you to configure email filter profiles for applying to firewall policies. A profile is specific information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

Email filter profile configuration settings

The following are email filter profile configuration settings in *UTM Profiles > Email Filter > Profile*. Advanced settings are configured in the CLI.

Profile page Lists each individual email filter profile that you created. On this page, you can edit, delete or create a new email filter profile. You are redirected to this page when you select <i>View List</i> on the Edit Email Filter Profile page.	
Create New	Creates a new email filter profile. When you select <i>Create New</i> , you are automatically redirected to the New Email Filter Profile page.
Edit	Modifies settings within an email filtering profile. When you select <i>Edit</i> , you are automatically redirected to the Edit Email Filter Profile page.
Delete	Removes an email filter profile from the list. To remove multiple email filter profiles from within the list, on the Profile page, in each of the rows of the email filter profiles you want removed, select the check box and then select <i>Delete</i> . To remove all email filter profiles from the list, on the Profile page, select the check box in the check box column and then select <i>Delete</i> .

Name	The name of the email filter profile.
Comments	The description given to the email filter profile. This is an optional setting.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
<p>New Email Filter Profile page</p> <p>Provides settings for configuring multiple email filter profiles. If you are editing an email filter profile, you are automatically redirected to the Edit Email Filter Profile page.</p> <p>This page appears when you select <i>Create New</i> on the Edit Email Filter Profile page. If you are on the Profile page, and you select <i>Create New</i>, you will be redirected to the New Email Filter Profile page.</p> <p>Note: Logging is enabled in the CLI.</p>	
Name	<p>Enter a name for the email filter profile.</p> <p>If you are editing an existing email profile and want to change the name, enter the new name in the <i>Name</i> field and then select <i>Apply</i> to save the changes.</p>
Comments	<p>Enter a description about the email filter profile. This is optional.</p> <p>If you are editing an existing email profile and want to change the description, enter the new description in the <i>Comments</i> field and then select <i>Apply</i> to save the changes.</p>
Log Email Summary	<p>Select to log specific information about the spam email activity on protocols such as IMAP or web mail such as Yahoo Mail.</p> <p>This feature does not record all email filtering activity, only IMAP, POP3, SMTP, Yahoo Mail, and MSN Hotmail spam detected email messages. If you want to log all email filtering activity, you must enable it in the CLI.</p>

Enable Spam Detection and Filtering	Select to enable specific settings for detecting and filtering spam email messages for IMAP, POP3 and SMTP as well as IMAPS, POP3S and SMTPS.
Spam Action	<p>Select to either tag or discard email that the unit determines to be spam. Tagging adds the text in the <i>Tag Format</i> field to the subject line or header of an email message that is identified as spam. Discard is available only for SMTP.</p> <p>Note: When you enable virus scanning for SMTP and SMTPS in an antivirus profile, scanning in splice mode is also called streaming mode and is enabled automatically. When scanning in splice mode, the unit scans and streams the traffic to the destination at the same time, terminating the stream to the destination if a virus is selected.</p> <p>For more information about splicing behavior for SMTP, see the Knowledge Base article FortiGate Proxy Splice and Client Comforting Technical Note.</p> <p>When virus scanning is enabled for SMTP, the unit can only discard spam email if a virus is detected. Discarding immediately drops the connection. If virus scanning is not enabled, you can choose to either tag or discard SMTP spam.</p>
Tag Location	<p>Select to add the tag to the subject or MIME header of email identified as spam.</p> <p>If you select to add the tag to the subject line, the unit converts the entire subject line, including the tag, to UTF-8 format. This improves display for some email clients that cannot properly display subject lines that use more than one encoding.</p> <p>To add the tag to the MIME header, you must enable spamhdrcheck in the CLI for each protocol (IMAP, POP3 and SMTP).</p>
Tag Format	<p>Enter a word or phrase with which to tag email identified as spam. When typing a tag, use the same language as the unit's current administrator language setting. Tag text using other encodings may not be accepted. For example, when entering a spam tag that uses Japanese characters, first verify that the administrator language settings is Japanese; the unit will not accept a spam tag written in Japanese characters while the administrator language setting is English.</p> <p>Tags must not exceed 64 bytes. The number of characters constituting 64 bytes of data varies by text encoding, which may vary by the FortiGate administrator language setting.</p>

FortiGuard Spam Filtering	<p>Appears only when <i>Enable Spam Detection and Filtering</i> is enabled.</p> <p>Select to enable FortiGuard spam filtering.</p>
Local Spam Filtering	<p>Appears only when <i>Enable Spam Detection and Filtering</i> is enabled.</p> <p>Select to enable local spam filtering. Select the various check boxes beside the options you want included in the profile.</p> <p>When you enable local spam filtering, you can apply email address and IP address black and white lists, as well as a banned word list to the profile.</p>

Banned Word

Control spam by blocking email messages containing specific words or patterns. You can add words, phrases, wild cards and Perl regular expressions to match content in email messages. For information, about wild cards and Perl regular expressions, see [“Using wildcards and Perl regular expressions” on page 1145](#).

The unit checks each email message against the banned word list. The unit can sort email messages containing those banned words in the subject, body, or both. The score value of each banned word appearing in the message is added, and if the total is greater than the threshold value set in the email filter profile, the unit processes the message according to the setting in the profile. The score for a pattern is applied only once even if the word appears in the message multiple times.

Banned word configuration settings

The following are banned word configuration settings in *UTM Profiles > Email Filter > Banned Word*.

Banned Word page Lists each banned word list that you created. On this page you can edit, delete or create a new banned word.	
Create New	Creates a new banned word. When you select <i>Create New</i> , you are automatically redirected to the New List page. This page provides a name field and comment field. You must enter a name to go to the Banned Word Settings page.
Edit	Modifies the banned word list, list name, or list comment. When you select <i>Edit</i> , you are automatically redirected to the Banned Word Settings page.
Delete	<p>Removes the banned word list from the catalog. The <i>Delete</i> icon is available only if the banned word list is not selected in any email filter profiles.</p> <p>To remove multiple banned word lists from within the list, on the Banned Word page, in each of the rows of the banned word lists you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all banned word lists from the list, on the Banned Word page, select the check box in the check box column and then select <i>Delete</i>.</p>
Name	The available Email Filter banned word lists.
# Entries	The number of entries in each banned word list.
Comments	Optional description of each banned word list.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.

Banned Word Settings page	
Provides settings for configuring a word pattern or word that will be considered banned by the unit. These words and word patterns make up a banned word list which appears on the Banned Word page. If you are editing a banned word, you are automatically redirected to the Banned Word Settings page.	
Name	If you are editing an existing banned word list and you want to change the name, enter a new name in this field. You must select <i>OK</i> to save these changes.
Comments	If you are editing an existing banned word list and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save these changes.
Create New	Adds a word or phrase to the banned word list. When you select <i>Create New</i> , you are automatically redirected to the Add Banned Word page.
Edit	Modifies banned word settings. When you select <i>Edit</i> , you are automatically redirected to the Edit Banned Word page.
Delete	Removes a banned word from the list. To remove multiple banned words from within the list, on the Banned Word Settings page, in each of the rows of the words you want removed, select the check box and then select <i>Delete</i> . To remove all banned words from the list, on the Banned Word Settings page, select the check box in the check box column and then select <i>Delete</i> .
Enable	Enables a banned word within the list.
Disable	Disables a banned word within the list.
Remove All Entries	Removes all banned word entries within the list on the Banned Word Settings page.
Page Controls	Use to navigate through the information in the Banned Word menu.
Enable	A green checkmark appears if the banned word is enabled.
Pattern	The list of banned words. Select the check box to enable all the banned words in the list.
Pattern Type	The pattern type used in the banned word list entry. Choose from wildcard or regular expression. For more information, see “Using wildcards and Perl regular expressions” on page 1145 .
Language	The character set to which the banned word belongs.
Where	The location where the unit searches for the banned word: <i>Subject</i> , <i>Body</i> , or <i>All</i> .
Score	A numerical weighting applied to the banned word. The score values of all the matching words appearing in an email message are added, and if the total is greater than the Banned word check value set in the profile, the email is processed according to whether the spam action is set to <i>Discard</i> or <i>Tagged</i> in the email filter profile. The score for a banned word is counted once even if the word appears multiple times on the web page in the email. For more information, see “Email Filter interface reference” on page 928 .

Add Banned Word page	
Provides settings for configuring a banned word entry.	
Pattern	Enter the banned word pattern. A pattern can be part of a word, a whole word, or a phrase. Multiple words entered as a pattern are treated as a phrase. The phrase must appear exactly as entered to match. You can also use wildcards or regular expressions to have a pattern match multiple words or phrases.
Pattern Type	Select the pattern type for the banned word. Choose from wildcard or regular expressions. For more information, see “Using wildcards and Perl regular expressions” on page 1145 .
Language	Select the character sets for the banned word.
Where	Select where the FortiGate unit should search for the banned word, <i>Subject</i> , <i>Body</i> , or <i>All</i> .
Score	Enter a score for the pattern. Each entry in the banned word list added to the profile includes a score. When an email message is matched with an entry in the banned word list, the score is recorded. If an email message matches more than one entry, the score for the email message increases. When the total score for an email message equals or exceeds the threshold, the message is considered spam and handled according to the spam action configured in the profile.
Enable	Select to enable a disable banned word. By default, a banned word is enabled.



Perl regular expression patterns are case sensitive for banned words. To make a word or phrase case insensitive, use the regular expression `/i`. For example, `/bad language/i` will block all instances of `bad language` regardless of case. Wildcard patterns are not case sensitive.

IP Address

You can add IP address black/white lists and email address black/white lists to filter email. When performing an IP address list check, the unit compares the IP address of the message sender to the IP address list items in sequence. When performing an email list check, the unit compares the email address of the message sender to the email address list items in sequence. If a match is found, the action associated with the IP address or email address is taken. If no match is found, the message is passed to the next enabled email filter.

You can add multiple IP address lists and then select the best one for each email filter profile.

After creating an IP address list, you can add IP addresses to the list.

Enter an IP address or a pair of IP address and mask in the following formats:

- `x.x.x.x`, for example, 192.168.69.100.
- `x.x.x.x/x.x.x.x`, for example, 192.168.69.100/255.255.255.0
- `x.x.x.x/x`, for example, 192.168.69.100/24

IP address configuration settings

The following are IP address configuration settings in *UTM Profiles > Email Filter > IP Address*.

IP Address page Lists each individual IP address list that you created. On this page, you can edit, delete or create a new IP address list. An IP address list contains multiple IP addresses and this list is configured in the IP Address Settings page.	
Create New	Creates a new IP address list. When you select <i>Create New</i> , you are automatically redirected to the New List page. This page provides a name field and comment field. You must enter a name to go to the IPS Address Settings page.
Edit	Modifies settings within an IP address list. When you select <i>Edit</i> , you are automatically redirected to the IP Address Settings page.
Delete	Removes the IP address black/white list from the list. The <i>Delete</i> icon is available only if the IP address list is not selected in any profiles. To remove multiple IP address black/white lists from within the list, on the IP Address page, in each of the rows of the black/white lists you want removed, select the check box and then select <i>Delete</i> . To remove all IP address lists from the list, on the IP Address page, select the check box in the check box column and then select <i>Delete</i> .
Name	The available name of the IP address lists.
# Entries	The number of entries in each IP address list.
Comments	Optional description of each IP address list.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i></p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
IP Address Settings page Provides settings for configuring multiple IP addresses that are then grouped together to form a list of IP addresses. This list is then applied within the email filter profile. You are automatically redirected to this page from the New List page. If you are editing an IP Address, you are automatically redirected to the IP Address Settings page.	
Name	If you are editing an existing IP address list and want to change the name, enter a new name in this field. You must select <i>OK</i> to save the change.
Comments	If you are editing an existing IP address list and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save the changes.
Create New	Creates a new IP address with settings. When you select <i>Create New</i> , you are automatically redirected to the Add IP Address page.
Edit	Modifies the settings within an IP address. When you select <i>Edit</i> , you are automatically redirected to the Edit IP Address page.

Delete	<p>Removes an IP address from the list.</p> <p>To remove multiple IP addresses from within the list, on the IP Address Settings page, in each of the rows of the IP addresses you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all IP addresses from the list, on the IP Address Settings page, select the check box in the check box column and then select <i>Delete</i>.</p>
Enable	Enables an IP address within that IP address list.
Disable	Disables an IP address within that IP address list.
Move To	<p>Moves the entry to a different position in the list. When you select <i>Move To</i>, the Move IP Address window appears.</p> <p>To move an IP address, select the new position <i>Before</i> or <i>After</i>, which will place the current IP address before or after the IP address you enter in the field (<i>IP/Netmask</i>). Enter the IP address and netmask in the (<i>IP/Netmask</i>) field.</p> <p>The firewall policy executes the list from top to bottom. For example, if you have IP address 192.168.100.1 listed as spam and 192.168.100.2 listed as clear, you must put 192.168.100.1 above 192.168.100.2 for 192.168.100.1 to take effect.</p>
Remove All Entries	Removes all IP addresses from within the list on the IP Address Settings page.
Enable	Indicates that the IP address is either enabled or disabled. A green check mark indicates that it is enabled; a gray x indicates that it is disabled.
IP/Netmask	The IP address and/or netmask.
Action	The type of action the unit will take when a match is detected. For example, the <i>Action</i> is <i>Spam</i> ; when a match is found, the detected IP address is considered spam.
Add IP Address page	
Provides settings for configuring an IP address to add to the list.	
IP/Netmask	Enter the IP address or the IP address/mask pair.
Action	Select: <i>Mark as Spam</i> to apply the spam action configured in the profile, <i>Mark as Clear</i> to bypass this and remaining email filters, or <i>Mark as Reject</i> (SMTP or SMTPS) to drop the session.
Enable	Select to enable the address.

E-mail Address

The unit can filter email from specific senders or all email from a domain (such as example.net). You can add email address lists and then select the best one for each profile.

Email address configuration settings

The following are email address configuration settings in *UTM Profiles > Email Filter > E-mail Address*.

E-mail Address page	
Lists each individual email address list that you created. On this page, you can edit, delete or create a new email address list.	
Create New	Creates a new email address list. When you select <i>Create New</i> , you are automatically redirected to the New List page. This page provides a name field and comment field; you must enter a name to go to the E-mail Address Settings page.
Edit	Modifies settings within an email address list. When you select <i>Edit</i> , you are automatically redirected to the E-mail Address Settings page.
Delete	Removes the email address list from the list on the E-mail Address page. The <i>Delete</i> icon is only available if the email address list is not selected in any profiles. To remove multiple email address lists from within the list, on the E-mail Address page, in each of the rows of the lists you want removed, select the check box and then select <i>Delete</i> . To remove all email filter lists from the list, on the E-mail Address page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the email address list.
# Entries	The number of entries in each email address list.
Comments	Optional description of each email address list.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
E-mail Address Settings page Provides settings for configuring multiple email addresses that are then grouped together to form a list of email addresses. This list is then applied within the email filter profile. You are automatically redirected to this page from the New List page. If you are editing an email address list, you are automatically redirected to the E-mail Address Settings page.	
Name	If you are editing an existing email address list and want to change the name, enter a new name in this field. You must select <i>OK</i> to save the change.
Comments	If you are editing an existing email address list and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save the changes.
Create New	Creates a new email address to the email address list. When you select <i>Create New</i> , you are automatically redirected to the Add E-mail Address page.
Edit	Modifies settings within that email address in the list on the E-mail Address Settings page. When you select <i>Edit</i> , you are automatically redirected to the Edit E-mail Address page.
Delete	Removes an email address from within the list on the E-mail Address Settings page.
Enable	Enables an email address within the list.
Disable	Disables an email address within the list.

Move To	Moves the entry to a different position in the list. When you select <i>Move To</i> , the Move E-mail Address window appears. To move an email address, select the new position <i>Before</i> or <i>After</i> , which will place the current email address before or after the email address you enter in the (<i>E-mail Address</i>) field. Enter the email address in the field and then select <i>OK</i> .
Remove All Entries	Removes all email addresses within the list on the E-mail Address Settings page.
Page Controls	Use to navigate through the lists on the E-mail Address Settings page.
Enable	A green checkmark appears if an email address is enabled. A gray x appears if an email address is disabled.
E-mail Address	The email address entered.
Pattern Type	The pattern type chosen for that email address.
Action	The action that will be take when that email address is detected.
Add E-Mail Address page	
E-mail Address	Enter the email address.
Pattern Type	Select a pattern type: <i>Wildcard</i> or <i>Regular Expression</i> . For more information, see “Using wildcards and Perl regular expressions” on page 1145 .
Action	Select: <i>Mark as Spam</i> to apply the spam action configured in the profile, or <i>Mark as Clear</i> to bypass this and remaining email filters.
Enable	Select to enable the email address.



Intrusion protection

The FortiGate Intrusion Protection system combines signature detection and prevention with low latency and excellent reliability. With intrusion protection, you can create multiple IPS sensors, each containing a complete configuration based on signatures. Then, you can apply any IPS sensor to any security policy.

This section describes how to configure the FortiGate Intrusion Protection settings.

If you enable virtual domains (VDOMs) on the FortiGate unit, intrusion protection is configured separately for each virtual domain.

The following topics are included:

- [IPS concepts](#)
- [Enable IPS scanning](#)
- [Configure IPS options](#)
- [Enable IPS packet logging](#)
- [IPS examples](#)

IPS concepts

The FortiGate intrusion protection system protects your network from outside attacks. Your FortiGate unit has two techniques to deal with these attacks: anomaly- and signature-based defense.

Anomaly-based defense

Anomaly-based defense is used when network traffic itself is used as a weapon. A host can be flooded with far more traffic than it can handle, making the host inaccessible. The most common example is the denial of service (DoS) attack, in which an attacker directs a large number of computers to attempt normal access of the target system. If enough access attempts are made, the target is overwhelmed and unable to service genuine users. The attacker does not gain access to the target system, but it is not accessible to anyone else.

The FortiGate DoS feature will block traffic above a certain threshold from the attacker and allow connections from other legitimate users.

Signature-based defense

Signature-based defense is used against known attacks or vulnerability exploits. These often involve an attacker attempting to gain access to your network. The attacker must communicate with the host in an attempt to gain access and this communication will include particular commands or sequences of commands and variables. The IPS signatures include these command sequences, allowing the FortiGate unit to detect and stop the attack.

Signatures

IPS signatures are the basis of signature-based intrusion protection. Every attack can be reduced to a particular string of commands or a sequence of commands and variables. Signatures include this information so your FortiGate unit knows what to look for in network traffic.

Signatures also include characteristics about the attack they describe. These characteristics include the network protocol in which the attack will appear, the vulnerable operating system, and the vulnerable application.

To view the complete list of predefined signatures, go to *UTM Profiles > Intrusion Protection > Predefined*.

Protocol decoders

Before examining network traffic for attacks, the IPS engine uses protocol decoders to identify each protocol appearing in the traffic. Attacks are protocol-specific, so your FortiGate unit conserves resources by looking for attacks only in the protocols used to transmit them. For example, the FortiGate unit will only examine HTTP traffic for the presence of a signature describing an HTTP attack.

To view the protocol decoders, go to *UTM Profiles > Intrusion Protection > Protocol Decoder*.

IPS engine

Once the protocol decoders separate the network traffic by protocol, the IPS engine examines the network traffic for the attack signatures.

IPS sensors

The IPS engine does not examine network traffic for all signatures, however. You must first create an IPS sensor and specify which signatures are included. Add signatures to sensors individually using signature entries, or in groups using IPS filters.

To view the IPS sensors, go to *UTM Profiles > Intrusion Protection > IPS Sensor*.

IPS filters

IPS sensors contain one or more IPS filters. A filter is a collection of signature attributes that you specify. The signatures that have all of the attributes specified in a filter are included in the IPS filter.

For example, if your FortiGate unit protects a Linux server running the Apache web server software, you could create a new filter to protect it. By setting *OS* to *Linux*, and *Application* to *Apache*, the filter will include only the signatures that apply to both Linux and Apache. If you wanted to scan for all the Linux signatures and all the Apache signatures, you would create two filters, one for each.

To view the filters in an IPS sensor, go to *UTM Profiles > Intrusion Protection > IPS Sensor*, select the IPS sensor containing the filters you want to view, and choose *Edit*.

Custom/predefined signature entries

Signature entries allow you to add an individual custom or predefined IPS signature. If you need only one signature, adding a signature entry to an IPS sensor is the easiest way. Signature entries are also the only way to include custom signatures in an IPS sensor.

Another use for signature entries are to change the settings of individual signatures that are already included in a filter within the same IPS sensor. Add a signature entry with the required settings above the filter, and the signature entry will take priority.

Policies

To use an IPS sensor, you must select it in a security policy or an interface policy. An IPS sensor that is not selected in a policy will have no effect on network traffic.

IPS is most often configured as part of a security policy. Unless stated otherwise, discussion of IPS sensor use will be in regards to firewall policies in this document.

Enable IPS scanning

Enabling IPS scanning involves two separate parts of the FortiGate unit:

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day. Firewall policies can also be used to deny traffic, but those policies do not apply to IPS scanning.
- The IPS sensor contains filters, signature entries, or both. These specify which signatures are included in the IPS sensor.

When IPS is enabled, an IPS sensor is selected in a security policy, and all network traffic matching the policy will be checked for the signatures in the IPS sensor.

General configuration steps

For best results in configuring IPS scanning, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create an IPS sensor.
- 2 Add filters and/or predefined signatures and custom signatures to the sensor. The filters and signatures specify which signatures the IPS engine will look for in the network traffic.
- 3 Select a security policy or create a new one.
- 4 In the security policy, enable UTM protection, select *Enable IPS*, and choose the IPS sensor from the list.

All the network traffic controlled by this security policy will be processed according to the settings in the policy. These settings include the IPS sensor you specify in the policy.

Creating an IPS sensor

You need to create an IPS sensor and save it before configuring it with filters and entries.

To create a new IPS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select the *Create New* icon in the top of the Edit IPS Sensor window.
- 3 Enter the name of the new IPS sensor.
- 4 Optionally, you may also enter a comment. The comment will appear in the IPS sensor list and serves to remind you of the details of the sensor.
- 5 Select *OK*.

The IPS sensor is created and the sensor configuration window appears. A newly created sensor is empty and contains no filters or signatures. You need to add one or more filters or signatures before the sensor can take effect.

Creating an IPS filter

While individual signatures can be added to a sensor, a filter allows you to add multiple signatures to a sensor by specifying the characteristics of the signatures to be added.

To create a new IPS filter

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select the IPS sensor to which you want to add the filter using the drop-down list in the top row of the Edit IPS Sensor window.
- 3 Select the *Create New* drop-down and choose *Filter*.
- 4 Configure the filter that you require. Signatures matching all of the characteristics you specify in the filter will be included in the filter.

Severity	<p>Select <i>Specify</i> and choose the severity levels to include in the filter.</p> <p>If you select <i>All</i>, the severity attribute will not be used to determine which signatures are included in the filter.</p>
Target	<p>Select <i>Specify</i> and choose the type of system the signature protects.</p> <p>If you select <i>All</i>, the target attribute will not be used to determine which signatures are included in the filter.</p>
OS	<p>Select <i>Specify</i> and choose the operating system the signature protects.</p> <p>If you select <i>All</i>, the OS attribute will not be used to determine which signatures are included in the filter.</p> <p>Predefined signatures listed with an OS attribute of <i>All</i> affect all operating system and are automatically included in any filter regardless of whether a single, multiple, or all operating systems are specified.</p>
Protocol	<p>Select <i>Specify</i> and choose the network protocols the signature protects by selecting an item in the Available column and using the right-arrow icon to move the item to the Selected column.</p> <p>Similarly, remove protocols by selecting an item in the Selected column and use the left-arrow icon to move the item to the Available column.</p> <p>If you select <i>All</i>, the Protocol attribute will not be used to determine which signatures are included in the filter.</p>
Application	<p>Select <i>Specify</i> and choose the applications the signature protects by selecting an item in the Available column and using the right-arrow icon to move the item to the Selected column.</p> <p>Similarly, remove applications by selecting an item in the Selected column and using the left-arrow icon to move the item to the Available column.</p> <p>If you select <i>All</i>, the Application attribute will not be used to determine which signatures are included in the filter.</p>

Tags	<p>Tags are a means by which you can apply customized labels to your IPS filters. Specified tags are displayed only within the filter itself on the Edit IPS Filter page.</p> <p>By default, the tag feature is disabled on all but the largest FortiGate models. If the Tags option is not visible, you must go to <i>System > Admin > Settings</i> and enable <i>Display Object Tagging and Coloring</i> to enable it.</p> <p>For more information about tags, see “Tag management” on page 582.</p>
Applied Tags	Displays the tags that you have applied to the filter.
Add tags	Enter a tag and then select the plus (+) icon to add the tag to the filter. This also adds the tag to the <i>Applied tags</i> list.
View Matched rules	Select view a list of all the signatures included in the filter with the current settings.
Enable All Matching Signatures	<p>All predefined signatures have an <i>Enable</i> attribute that is set to on or off. Select <i>Accept signature defaults</i> use the <i>Enable</i> default setting for each included signature. This means that if a signature included in the filter has an <i>Enable</i> setting of off, traffic matching the signature will not be detected even though the signature is included in the filter.</p> <p>Select <i>Enable all</i> to enable all the signatures included in the filter, regardless of their <i>Enable</i> setting. Similarly, you may select <i>Disable all</i> to disable all the signatures included in the filter.</p>
Action When Signature Is Triggered	<p>All predefined signatures have an <i>Action</i> attribute that is set to Pass or Drop. Select <i>Accept signature defaults</i> use the default action for each included signature. This means that if a signature included in the filter has an <i>Action</i> setting of Pass, traffic matching the signature will be detected and then allowed to continue to its destination.</p> <p>Select <i>Monitor all</i> to pass all traffic matching the signatures included in the filter, regardless of their default <i>Action</i> setting. Similarly, you may select <i>Block all</i> to drop traffic matching any the signatures included in the filter.</p>
Quarantine Attackers (to Banned Users List)	Enable this option to add the source of the offending traffic to the Banned User list.
Packet Logging	<p>Select to enable packet logging for the filter.</p> <p>When you enable packet logging on a filter, the unit saves a copy of the packets that match any signatures included in the filter. The packets can be analyzed later.</p> <p>For more information about packet filtering, see “Viewing and saving logged packets” on page 1144</p>

5 Select OK.

The filter is created and added to the filter list.

Updating predefined IPS signatures

The FortiGuard Service periodically updates the pre-defined signatures and adds new signatures to counter emerging threats as they appear.

Because the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Viewing and searching predefined IPS signatures

Go to *UTM Profiles > Intrusion Protection > Predefined* to view the list of predefined IPS signatures. You may find signatures by paging manually through the list, apply filters, or by using the search field.

Searching manually

Signatures are displayed in a paged list, with 50 signatures per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

Applying filters

You can enter criteria for one of more columns, and only the signatures matching all the conditions you specify will be listed.

To apply filters

- 1 Go to *UTM Profiles > Intrusion Protection > Predefined*.
- 2 Select Filter Settings.
- 3 Select Add New Filter.
- 4 Select column by which to filter.
- 5 Select the item or items by which to filter.
- 6 Continue to add more filters to narrow your search, if required.
- 7 Select OK.

The available options vary by column. For example, Enable allows you to choose between two options, while OS has multiple options, and you may select multiple items together. Filtering by name allows you to enter a text string and all signature names containing the string will be displayed.

Using the search field

To use the search field, located above the signature list, start typing any portion of the signature name. Signatures names matching the text you enter are displayed in a drop-down list. A maximum of ten matches are displayed at a time.

Select a signature from the drop-down list to display its signature list entry.

Creating a signature entry

Pre-defined and custom signature entries are configured and work largely the same as filters, except they define the behavior of only one signature.

You can use entries in two ways:

- To change the behavior of a signature already included in a filter.
For example, to protect a web server, you can create a filter that includes and enables all signatures related to servers. If you want to disable one of those signatures, the simplest way is to create an entry and mark the signature as disabled.
- To add an individual signature, not included in any filters, to an IPS sensor. This is the only way to add custom signatures to IPS sensors.

When a pre-defined signature is specified in an entry, the default status and action attributes of the signature are ignored. These settings must be explicitly set when creating the entry.

To create an IPS signature entry

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select the IPS sensor to which you want to add the entry from the sensor drop down list in the *Edit IPS Sensor* window title bar.
- 3 Select the Create New drop down and choose either *Pre-defined Entry* or *Custom Entry*, depending on the type of IPS signature entry you require.
- 4 For the *Action*, select *Pass*, *Block*, or *Reset*. When the entry is enabled, the action determines what the FortiGate will do with traffic containing the specified signature.
- 5 Select *Packet Log* to save the packets containing the specified signature. For more information, see [“Enable IPS packet logging” on page 967](#).
- 6 Select the *Browse* icon and choose the signature to include in the entry.
- 7 Select *Enable*.
- 8 Select *OK*.

Creating a custom IPS signature

The FortiGate predefined signatures cover common attacks. If you use an unusual or specialized application or an uncommon platform, add custom signatures based on the security alerts released by the application and platform vendors.

You can add or edit custom signatures using the web-based manager or the CLI.

To create a custom signature

- 1 Go to *UTM Profiles > Intrusion Protection > Custom*.
- 2 Select *Create New* to add a new custom signature.
- 3 Enter a *Name* for the custom signature.
- 4 Enter the *Signature*. For information about completing this field, see [“Custom signature syntax and keywords”](#).
- 5 Select *OK*.

Custom signature syntax and keywords

All custom signatures follow a particular syntax. Each begins with a header and is followed by one or more keywords. The syntax and keywords are detailed in the next two topics.

Custom signature syntax

A custom signature definition is limited to a maximum length of 512 characters. A definition can be a single line or span multiple lines connected by a backslash (\) at the end of each line.

A custom signature definition begins with a header, followed by a set of keyword/value pairs enclosed by parenthesis [()]. The keyword and value pairs are separated by a semi colon (;) and consist of a keyword and a value separated by a space. The basic format of a definition is HEADER (KEYWORD VALUE;)

You can use as many keyword/value pairs as required within the 512 character limit. To configure a custom signature, go to *UTM Profiles > Intrusion Protection > Custom* and enter the data directly into the *Signature* field, following the guidance in the next topics.

Table 58 shows the valid characters and basic structure. For details about each keyword and its associated values, see [“Custom signature keywords” on page 951](#).

Table 58: Valid syntax for custom signature fields

Field	Valid Characters	Usage
HEADER	F-SBID	The header for an attack definition signature. Each custom signature must begin with this header.
KEYWORD	Each keyword must start with a pair of dashes (--), and consist of a string of 1 to 19 characters. Normally, keywords are an English word or English words connected by an underscore (_). Keywords are case insensitive.	The keyword is used to identify a parameter. See “Custom signature keywords” on page 951 for tables of supported keywords.
VALUE	Double quotes (") must be used around the value if it contains a space and/or a semicolon (;). If the value is NULL, the space between the KEYWORD and VALUE can be omitted. Values are case sensitive. Note: If double quotes are used for quoting the value, the double quotes are not considered as part of the value string.	The value is set specifically for a parameter identified by a keyword.

Custom signature keywords

Table 59: Information keywords

Keyword and value	Description
<code>--attack_id <id_int>;</code>	<p>Use this optional value to identify the signature. It cannot be the same value as any other custom rules. If an attack ID is not specified, the FortiGate automatically assigns an attack ID to the signature. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same attack ID for signatures in different VDOMs.</p> <p>An attack ID you assign must be between 1000 and 9999.</p> <p>Example:</p> <pre>--attack_id 1234;</pre>
<code>--name <name_str>;</code>	<p>Enter the name of the rule. A rule name must be unique. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same rule name for signatures in different VDOMs.</p> <p>The name you assign must be a string greater than 0 and less than 64 characters in length.</p> <p>Example:</p> <pre>--name "Buffer_Overflow";</pre>

Table 60: Session keywords

Keyword and value	Description
<code>--flow {from_client[,reversed] from_server[,reversed] bi_direction };</code>	<p>Specify the traffic direction and state to be inspected. They can be used for all IP traffic.</p> <p>Example:</p> <pre>--src_port 41523; --flow bi_direction;</pre> <p>The signature checks traffic to and from port 41523.</p> <p>If you enable “quarantine attacker”, the optional <code>reversed</code> keyword allows you to change the side of the connection to be quarantined when the signature is detected.</p> <p>For example, a custom signature written to detect a brute-force log in attack is triggered when “Login Failed” is detected <code>from_server</code> more than 10 times in 5 seconds. If the attacker is quarantined, it is the server that is quarantined in this instance. Adding <code>reversed</code> corrects this problem and quarantines the actual attacker.</p> <p>Previous FortiOS versions used <code>to_client</code> and <code>to_server</code> values. These are now deprecated, but still function for backwards compatibility.</p>

Table 60: Session keywords

<pre>--service {HTTP TELNET FTP DNS SMTP POP3 IMAP SNMP RADIUS LDAP MSSQL RPC SIP H323 NBSS DCERPC SSH SSL};</pre>	<p>Specify the protocol type to be inspected.</p> <p>This keyword allows you to specify the traffic type by protocol rather than by port. If the decoder has the capability to identify the protocol on any port, the signature can be used to detect the attack no matter what port the service is running on. Currently, HTTP, SIP, SSL, and SSH protocols can be identified on any port based on the content.</p>
--	--

Table 61: Content keywords

Keyword and value	Description
<pre>--byte_jump <bytes_to_convert>, <offset>[, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct] [, align];</pre>	<p>Use the <code>byte_jump</code> option to extract a number of bytes from a packet, convert them to their numeric representation, and jump the match reference up that many bytes (for further pattern matching or byte testing). This keyword allows relative pattern matches to take into account numerical values found in network data.</p> <p>The available keyword options include:</p> <ul style="list-style-type: none"> • <code><bytes_to_convert></code>: The number of bytes to examine from the packet. • <code><offset></code>: The number of bytes into the payload to start processing. • <code>relative</code>: Use an offset relative to last pattern match. • <code>big</code>: Process the data as big endian (default). • <code>little</code>: Process the data as little endian. • <code>string</code>: The data is a string in the packet. • <code>hex</code>: The converted string data is represented in hexadecimal notation. • <code>dec</code>: The converted string data is represented in decimal notation. • <code>oct</code>: The converted string data is represented in octal notation. • <code>align</code>: Round up the number of converted bytes to the next 32-bit boundary.

Table 61: Content keywords (Continued)

Keyword and value	Description
<pre>--byte_test <bytes_to_convert>, <operator>, <value>, <offset>[, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct];</pre>	<p>Use the <code>byte_test</code> keyword to compare a byte field against a specific value (with operator). This keyword is capable of testing binary values or converting representative byte strings to their binary equivalent and testing them.</p> <p>The available keyword options include:</p> <ul style="list-style-type: none"> • <code><bytes_to_convert></code>: The number of bytes to compare. • <code><operator></code>: The operation to perform when comparing the value (<, >, =, !, &). • <code><value></code>: The value to compare the converted value against. • <code><offset></code>: The number of bytes into the payload to start processing. • <code>relative</code>: Use an offset relative to last pattern match. • <code>big</code>: Process the data as big endian (default). • <code>little</code>: Process the data as little endian. • <code>string</code>: The data is a string in the packet. • <code>hex</code>: The converted string data is represented in hexadecimal notation. • <code>dec</code>: The converted string data is represented in decimal notation. • <code>oct</code>: The converted string data is represented in octal notation.
<pre>--depth <depth_int>;</pre>	<p>Use the <code>depth</code> keyword to search for the contents within the specified number of bytes after the starting point defined by the <code>offset</code> keyword. If no <code>offset</code> is specified, the <code>offset</code> is assumed to be equal to 0.</p> <p>If the value of the <code>depth</code> keyword is smaller than the length of the value of the <code>content</code> keyword, this signature will never be matched.</p> <p>The <code>depth</code> must be between 0 and 65535.</p>
<pre>--distance <dist_int>;</pre>	<p>Use the <code>distance</code> keyword to search for the contents within the specified number of bytes relative to the end of the previously matched contents. If the <code>within</code> keyword is not specified, continue looking for a match until the end of the payload.</p> <p>The <code>distance</code> must be between 0 and 65535.</p>

Table 61: Content keywords (Continued)

Keyword and value	Description
<pre>--content [!]"<content_str>;</pre>	<p>Deprecated, see <code>pattern</code> and <code>context</code> keywords.</p> <p>Use the <code>content</code> keyword to search for the content string in the packet payload. The content string must be enclosed in double quotes.</p> <p>To have the FortiGate search for a packet that does not contain the specified context string, add an exclamation mark (!) before the content string.</p> <p>Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe () character.</p> <p>The double quote ("), pipe sign() and colon(:) characters must be escaped using a back slash if specified in a content string.</p> <p>If the value of the <code>content</code> keyword is greater than the length of the value of the <code>depth</code> keyword, this signature will never be matched.</p>
<pre>--context {uri header body host};</pre>	<p>Specify the protocol field to look for the pattern. If context is not specified for a pattern, the FortiGate unit searches for the pattern anywhere in the packet buffer. The available context variables are:</p> <ul style="list-style-type: none"> • <code>uri</code>: Search for the pattern in the HTTP URI line. • <code>header</code>: Search for the pattern in HTTP header lines or SMTP/POP3/SMTP control messages. • <code>body</code>: Search for the pattern in HTTP body or SMTP/POP3/SMTP email body. • <code>host</code>: Search for the pattern in HTTP HOST line. <p>Example:</p> <pre>--pattern "GET " --context uri --pattern "yahoo.com" --context host --no_case --pcre "/DESCRIBE\s+\/\s+RTSP\/\//i" --context header</pre>
<pre>--no_case;</pre>	<p>Use the <code>no-case</code> keyword to force the FortiGate unit to perform a case-insensitive pattern match.</p>

Table 61: Content keywords (Continued)

Keyword and value	Description
<code>--offset <offset_int>;</code>	<p>Use the <code>offset</code> keyword to look for the contents after the specified number of bytes into the payload. The specified number of bytes is an absolute value in the payload. Follow the <code>offset</code> keyword with the <code>depth</code> keyword to stop looking for a match after a specified number of bytes. If no <code>depth</code> is specified, the FortiGate unit continues looking for a match until the end of the payload.</p> <p>The <code>offset</code> must be between 0 and 65535.</p>
<code>--pattern [!]"<pattern_str>;</code>	<p>The FortiGate unit will search for the specified pattern.</p> <p>A <code>pattern</code> keyword normally is followed by a <code>context</code> keyword to define where to look for the pattern in the packet. If a <code>context</code> keyword is not present, the FortiGate unit looks for the pattern anywhere in the packet buffer.</p> <p>To have the FortiGate search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI.</p> <p>Example:</p> <pre>--pattern "/level/" --pattern " E8 D9FF FFFF /bin/sh" --pattern "! 20 RTSP/"</pre>

Table 61: Content keywords (Continued)

Keyword and value	Description
<pre>--pcre [!]"(/<regex>/ m<delim> <regex><delim>) [ismxAEGRUB]";</pre>	<p>Similarly to the <code>pattern</code> keyword, use the <code>pcre</code> keyword to specify a pattern using Perl-compatible regular expressions (PCRE). A <code>pcre</code> keyword can be followed by a <code>context</code> keyword to define where to look for the pattern in the packet. If no <code>context</code> keyword is present, the FortiGate unit looks for the pattern anywhere in the packet buffer.</p> <p>For more information about PCRE syntax, go to http://www.pcre.org.</p> <p>The switches include:</p> <ul style="list-style-type: none"> • <code>i</code>: Case insensitive. • <code>s</code>: Include newlines in the dot metacharacter. • <code>m</code>: By default, the string is treated as one big line of characters. <code>^</code> and <code>\$</code> match at the beginning and ending of the string. When <code>m</code> is set, <code>^</code> and <code>\$</code> match immediately following or immediately before any newline in the buffer, as well as the very start and very end of the buffer. • <code>x</code>: White space data characters in the pattern are ignored except when escaped or inside a character class. • <code>A</code>: The pattern must match only at the start of the buffer (same as <code>^</code>). • <code>E</code>: Set <code>\$</code> to match only at the end of the subject string. Without <code>E</code>, <code>\$</code> also matches immediately before the final character if it is a newline (but not before any other newlines). • <code>G</code>: Invert the “greediness” of the quantifiers so that they are not greedy by default, but become greedy if followed by <code>?</code>. • <code>R</code>: Match relative to the end of the last pattern match. (Similar to <code>distance:0</code>). • <code>U</code>: Deprecated, see the <code>context</code> keyword. Match the decoded URI buffers.

Table 61: Content keywords (Continued)

Keyword and value	Description
<code>--uri [!]"<uri_str>;</code>	<p>Deprecated, see pattern and context keywords.</p> <p>Use the <code>uri</code> keyword to search for the URI in the packet payload. The URI must be enclosed in double quotes (").</p> <p>To have the FortiGate unit search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI.</p> <p>Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe () character.</p> <p>The double quote ("), pipe sign () and colon (:) characters must be escaped using a back slash (\) if specified in a URI string.</p>
<code>--within <within_int>;</code>	<p>Use this together with the <code>distance</code> keyword to search for the contents within the specified number of bytes of the payload.</p> <p>The <code>within</code> value must be between 0 and 65535.</p>

Table 62: IP header keywords

Keyword and Value	Description
<code>--dst_addr [!]<ipv4>;</code>	<p>Use the <code>dst_addr</code> keyword to search for the destination IP address.</p> <p>To have the FortiGate search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address.</p> <p>You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.</p> <p>Example: <code>dst_addr [172.20.0.0/16, 10.1.0.0/16, 192.168.0.0/16]</code></p>
<code>--ip_id <field_int>;</code>	Check the IP ID field for the specified value.
<code>--ip_option {rr eol nop ts sec lsrr ssrr satid any};</code>	<p>Use the <code>ip_option</code> keyword to check various IP option settings. The available options include:</p> <ul style="list-style-type: none"> <code>rr</code>: Check if IP RR (record route) option is present. <code>eol</code>: Check if IP EOL (end of list) option is present. <code>nop</code>: Check if IP NOP (no op) option is present. <code>ts</code>: Check if IP TS (time stamp) option is present. <code>sec</code>: Check if IP SEC (IP security) option is present. <code>lsrr</code>: Check if IP LSRR (loose source routing) option is present. <code>ssrr</code>: Check if IP SSRR (strict source routing) option is present. <code>satid</code>: Check if IP SATID (stream identifier) option is present. <code>any</code>: Check if IP any option is present.
<code>--ip_tos <field_int>;</code>	Check the IP TOS field for the specified value.
<code>--ip_ttl [< >] <ttl_int>;</code>	Check the IP time-to-live value against the specified value. Optionally, you can check for an IP time-to-live greater-than (>) or less-than (<) the specified value with the appropriate symbol.
<code>--protocol {<protocol_int> tcp udp icmp};</code>	<p>Check the IP protocol header.</p> <p>Example:</p> <pre>--protocol tcp;</pre>
<code>--src_addr [!]<ipv4>;</code>	<p>Use the <code>src_addr</code> keyword to search for the source IP address.</p> <p>To have the FortiGate unit search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address.</p> <p>You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.</p> <p>Example: <code>src_addr 192.168.13.0/24</code></p>

Table 63: TCP header keywords

Keyword and Value	Description
<code>--ack <ack_int>;</code>	Check for the specified TCP acknowledge number.
<code>--dst_port</code> <code>[!]{<port_int> </code> <code>:<port_int> </code> <code><port_int>: </code> <code><port_int>:<port_int>;</code>	<p>Use the <code>dst_port</code> keyword to specify the destination port number.</p> <p>You can specify a single port or port range:</p> <ul style="list-style-type: none"> • <code><port_int></code> is a single port. • <code>:<port_int></code> includes the specified port and all lower numbered ports. • <code><port_int>:</code> includes the specified port and all higher numbered ports. • <code><port_int>:<port_int></code> includes the two specified ports and all ports in between.
<code>--seq <seq_int>;</code>	Check for the specified TCP sequence number.
<code>--src_port</code> <code>[!]{<port_int> </code> <code>:<port_int> </code> <code><port_int>: </code> <code><port_int>:<port_int>;</code>	<p>Use the <code>src_port</code> keyword to specify the source port number.</p> <p>You can specify a single port or port range:</p> <ul style="list-style-type: none"> • <code><port_int></code> is a single port. • <code>:<port_int></code> includes the specified port and all lower numbered ports. • <code><port_int>:</code> includes the specified port and all higher numbered ports. • <code><port_int>:<port_int></code> includes the two specified ports and all ports in between.

Table 63: TCP header keywords (Continued)

Keyword and Value	Description
<pre>--tcp_flags <SAFRUP120>[! * +] [, <SAFRUP120>];</pre>	<p>Specify the TCP flags to match in a packet.</p> <ul style="list-style-type: none"> • S: Match the SYN flag. • A: Match the ACK flag. • F: Match the FIN flag. • R: Match the RST flag. • U: Match the URG flag. • P: Match the PSH flag. • 1: Match Reserved bit 1. • 2: Match Reserved bit 2. • 0: Match No TCP flags set. • !: Match if the specified bits are not set. • *: Match if any of the specified bits are set. • +: Match on the specified bits, plus any others. <p>The first part if the value (<SAFRUP120>) defines the bits that must be present for a successful match. For example:</p> <pre>--tcp_flags AP</pre> <p>only matches the case where both A and P bits are set.</p> <p>The second part ([, <SAFRUP120>]) is optional, and defines the additional bits that can be present for a match. For example:</p> <pre>tcp_flags S,12</pre> <p>matches the following combinations of flags: S, S and 1, S and 2, S and 1 and 2.</p> <p>The modifiers !, * and + cannot be used in the second part.</p>
<pre>--window_size [!]<window_int>;</pre>	<p>Check for the specified TCP window size.</p> <p>You can specify the window size as a hexadecimal or decimal integer. A hexadecimal value must be preceded by 0x.</p> <p>To have the FortiGate search for the absence of the specified window size, add an exclamation mark (!) before the window size.</p>

Table 64: UDP header keywords

Keyword and Value	Description
<pre>--dst_port [!]{<port_int> :<port_int> <port_int>: <port_int>:<port_int>;</pre>	<p>Specify the destination port number.</p> <p>You can specify a single port or port range:</p> <ul style="list-style-type: none"> • <port_int> is a single port. • :<port_int> includes the specified port and all lower numbered ports. • <port_int>: includes the specified port and all higher numbered ports. • <port_int>:<port_int> includes the two specified ports and all ports in between.
<pre>--src_port [!]{<port_int> :<port_int> <port_int>: <port_int>:<port_int>;</pre>	<p>Specify the source port number.</p> <p>You can specify a single port or port range:</p> <ul style="list-style-type: none"> • <port_int> is a single port. • :<port_int> includes the specified port and all lower numbered ports. • <port_int>: includes the specified port and all higher numbered ports. • <port_int>:<port_int> includes the two specified ports and all ports in between.

Table 65: ICMP keywords

Keyword and Value	Usage
--icmp_code <code_int>;	Specify the ICMP code to match.
--icmp_id <id_int>;	Check for the specified ICMP ID value.
--icmp_seq <seq_int>;	Check for the specified ICMP sequence value.
--icmp_type <type_int>;	Specify the ICMP type to match.

Table 66: Other keywords

Keyword and Value	Description
<pre>--data_size {<size_int> <<size_int> ><size_int> <port_int><><port_int>;</pre>	<p>Test the packet payload size. With data_size specified, packet reassembly is turned off automatically. So a signature with data_size and only_stream values set is wrong.</p> <ul style="list-style-type: none"> • <size_int> is a particular packet size. • <<size_int> is a packet smaller than the specified size. • ><size_int> is a packet larger than the specified size. • <size_int><><size_int> is a packet within the range between the specified sizes.

Table 66: Other keywords (Continued)

Keyword and Value	Description
<code>--data_at <offset_int>[, relative];</code>	Verify that the payload has data at a specified offset, optionally looking for data relative to the end of the previous content match.
<code>--rate <matches_int>,<time_int>;</code>	<p>Instead of generating log entries every time the signature is detected, use this keyword to generate a log entry only if the signature is detected a specified number of times within a specified time period.</p> <ul style="list-style-type: none"> • <code><matches_int></code> is the number of times a signature must be detected. • <code><time_int></code> is the length of time in which the signature must be detected, in seconds. <p>For example, if a custom signature detects a pattern, a log entry will be created every time the signature is detected. If <code>--rate 100,10;</code> is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds.</p> <p>Use this command with <code>--track</code> to further limit log entries to when the specified number of detections occur within a certain time period involving the same source or destination address rather than all addresses.</p>
<code>--rpc_num <app_int>[, <ver_int> *][, <proc_int> *];</code>	Check for RPC application, version, and procedure numbers in SUNRPC CALL requests. The * wildcard can be used for version and procedure numbers.
<code>--same_ip;</code>	Check that the source and the destination have the same IP addresses.

Table 66: Other keywords (Continued)

Keyword and Value	Description
<code>--track {client server}[,block_int];</code>	<p>When used with <code>--rate</code>, this keyword narrows the custom signature rate totals to individual addresses.</p> <ul style="list-style-type: none"> <code>client</code> has the FortiGate unit maintain a separate count of signature matches for each source address. <code>server</code> has the FortiGate unit maintain a separate count of signature matches for each destination address. <code>block_int</code> has the FortiGate unit block connections for the specified number of seconds, from the client or to the server, depending on which is specified. <p>For example, if <code>--rate 100,10</code> is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. The FortiGate unit maintains a single total, regardless of source and destination address.</p> <p>If the same custom signature also includes <code>--track client</code>; matches are totalled separately for each source address. A log entry is added when the signature is detected 100 times in 10 seconds within traffic from the same source address.</p> <p>The <code>--track</code> keyword can also be used without <code>--rate</code>. If an integer is specified, the client or server will be blocked for the specified number of seconds every time the signature is detected.</p>

IPS processing in an HA cluster

IPS processing in an HA cluster is no different than with a single FortiGate unit, from the point of view of the network user. The difference appears when a secondary unit takes over from the primary, and what happens depends on the HA mode.

Active-passive

In an active-passive HA cluster, the primary unit processes all traffic just as it would in a stand-alone configuration. Should the primary unit fail, a secondary unit will assume the role of the primary unit and begin to process network traffic. By default, the state of active communication sessions are not shared with secondary units and will not survive the fail-over condition. Once the sessions are reestablished however, traffic processing will continue as normal.

If your network requires that active sessions are taken over by the new primary unit, select *Enable Session Pick-up* in your HA configuration. Because session information must be sent to all subordinate units on a regular basis, session pick-up is a resource-intensive feature and is not enabled by default.

Active-active

The fail-over process in an active-active cluster is similar to an active-passive cluster. When the primary unit fails, a secondary unit takes over and traffic processing continues. The load-balancing schedule used to distribute sessions to the cluster members is used by the new primary unit to redistribute sessions among the remaining subordinate units. If session pick-up is not enabled, the sessions active on the failed primary are lost, and the sessions redistributed among the secondary units may also be lost. If session pick-up is enabled, all sessions are handled according to their last-known state.

For more information about HA options and settings, see [“High Availability” on page 1983](#).

Configure IPS options

There are a number of CLI commands that influence how IPS functions.

Configuring the IPS engine algorithm

The IPS engine is able to search for signature matches in two ways. One method is faster but uses more memory, the other uses less memory but is slower. Use the `algorithm` CLI command to select one method:

```
config ips global
    set algorithm {high | low | engine-pick}
end
```

Specify `high` to use the faster more memory intensive method or `low` for the slower memory efficient method. The default setting is `engine-pick`, which allows the IPS engine to choose the best method on the fly.

Configuring the IPS engine-count

FortiGate units with multiple processors can run more than one IPS engine concurrently. The `engine-count` CLI command allows you to specify how many IPS engines are used at the same time:

```
config ips global
    set engine-count <int>
end
```

The recommended and default setting is 0, which allows the FortiGate unit to determine the optimum number of IPS engines.

Configuring fail-open

If the IPS engine fails for any reason, it will fail open by default. This means that traffic continues to flow without IPS scanning. If IPS protection is more important to your network than the uninterrupted flow of network traffic, you can disable this behavior using the `fail-open` CLI command:

```
config ips global
    set fail-open {enable | disable}
end
```

The default setting is `enable`.

Configuring the session count accuracy

The IPS engine can keep track of the number of open session in two ways. An accurate count uses more resources than a less accurate heuristic count.

```
config ips global
  set session-limit-mode {accurate | heuristic}
end
```

The default is heuristic.

Configuring the IPS buffer size

Set the size of the IPS buffer.

```
config ips global
  set socket-size <int>
end
```

The acceptable range is from 1 to 64 megabytes. The default size varies by model.

Configuring protocol decoders

The FortiGate Intrusion Protection system uses protocol decoders to identify the abnormal traffic patterns that do not meet the protocol requirements and standards. For example, the HTTP decoder monitors traffic to identify any HTTP packets that do not meet the HTTP protocol standards.

To view the decoders and the port numbers that each protocol decoder monitors, go to *UTM > Intrusion Protection > Protocol Decoder*. The port or ports monitored by each decoder are listed. Many decoders are able to recognize traffic by type rather than by port. These decoders have their port listed as *auto* because the traffic will be recognized automatically, regardless of the port.

To change the ports a decoder examines, you must use the CLI. In this example, the ports examined by the DNS decoder are changed from the default 53 to 100, 200, and 300.

```
config ips decoder dns_decoder
  set port_list "100,200,300"
end
```

You cannot assign specific ports to decoders that are set to *auto* by default. These decoders can detect their traffic on any port. Specifying individual ports is not necessary.

Configuring security processing modules

FortiGate Security Processing Modules, such as the CE4, XE2, and FE8, can increase overall system performance by accelerating some security and networking processing on the interfaces they provide. They also allow the FortiGate unit to offload the processing to the security module, thereby freeing up its own processor for other tasks. The security module performs its own IPS and firewall processing, but you can configure it to favor IPS in hostile high-traffic environments.

If you have a security processing module, use the following CLI commands to configure it to devote more resources to IPS than firewall. This example shows the CLI commands required to configure a security module in slot 1 for increased IPS performance.

```
config system amc-slot
  edit sw1
    set optimization-mode fw-ips
    set ips-weight balanced
    set ips-p2p disable
```

```
set ips-fail-open enable
set fp-disable none
set ipsec-inb-optimization enable
set syn-proxy-client-timer 3
set syn-proxy-server-timer 3
end
```

In addition to offloading IPS processing, security processing modules provide a hardware accelerated SYN proxy to defend against SYN flood denial of service attacks. When using a security module, configure your DoS sensor `tcp_syn_flood` anomaly with the *Proxy* action. The *Proxy* action activates the hardware accelerated SYN proxy.



Because DoS sensors are configured before being applied to an interface, you can assign a DoS sensor with the *Proxy* action to an interface that does not have hardware SYN proxy support. In this circumstance, the *Proxy* action is invalid and a *Pass* action will be applied.

Enable IPS packet logging

Packet logging saves the network packets containing the traffic matching an IPS signature to the attack log. The FortiGate unit will save the logged packets to wherever the logs are configured to be stored, whether memory, internal hard drive, a FortiAnalyzer unit, or the FortiGuard Analysis and Management Service.

You can enable packet logging in signature entries or in filters. Use caution in enabling packet logging in a filter. Filters configured with few restrictions can contain thousands of signatures, potentially resulting in a flood of saved packets. This would take up a great deal of space, require time to sort through, and consume considerable system resources to process. Packet logging is designed as a focused diagnostic tool and is best used with a narrow scope.



Although logging to multiple FortiAnalyzer units is supported, packet logs are not sent to the secondary and tertiary FortiAnalyzer units. Only the primary unit receives packet logs.

To enable packet logging for a signature

- 1 Create either a pre-defined or custom entry in an IPS sensor. For more information, see [“Creating a signature entry” on page 948](#).
- 2 Before saving the entry, select *Packet Log*.
- 3 Select the IPS sensor in the security policy that allows the network traffic the FortiGate unit will examine traffic for the signature.

To enable packet logging for a filter

- 1 Create a filter in an IPS sensor. For more information, see [“Creating an IPS filter” on page 946](#).
- 2 Before saving the filter, select *Enable All for Packet Logging*.
- 3 Select the IPS sensor in the security policy that allows the network traffic the FortiGate unit will examine for the signature.

For information on viewing and saving logged packets, see [“Viewing and saving logged packets” on page 1144](#).

IPS examples

Configuring basic IPS protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable IPS protection on a FortiGate unit located in a satellite office. The satellite office contains only Windows clients.

Creating an IPS sensor

Most IPS settings are configured in an IPS sensor. IPS sensors are selected in firewall policies. This way, you can create multiple IPS sensors, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one IPS sensor.

To create an IPS sensor— web-based manager

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select the *Create New* icon in the top of the Edit IPS Sensor window.
- 3 In the *Name* field, enter `basic_ips`.
- 4 In the *Comments* field, enter `IPS protection for Windows clients`.
- 5 Select *OK*.
- 6 Select the *Create New* drop-down and choose *Filter*.
- 7 For *Target*, select *Specify* and *Client*.
- 8 For *OS*, select *Specify* and *Windows*.
- 9 Select *OK* to save the filter.
- 10 Select *OK* to save the IPS sensor.

To create an IPS sensor — CLI

```
config ips sensor
  edit basic_ips
    set comment "IPS protection for Windows clients"
  config filter
    edit 1
      set location client
      set os windows
    end
  end
end
```

Selecting the IPS sensor in a security policy

An IPS sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an IPS sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the IPS sensor in a security policy — web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a policy.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.

- 5 Select the *Enable IPS* option.
- 6 Select the `basic_ips` profile from the list.
- 7 Select *OK* to save the security policy.

To select the IPS sensor in a security policy — CLI

```
config firewall policy
edit 1
set utm-status enable
set ips-sensor basic_ips
end
```


All traffic handled by the security policy you modified will be scanned for attacks against Windows clients. A small office may have only one security policy configured. If you have multiple policies, consider enabling IPS scanning for all of them.

Using IPS to protect your web server

Many companies have web servers and they must be protected from attack. Since web servers must be accessible, protection is not as simple as blocking access. IPS is one tool your FortiGate unit has to allow you to protect your network.

In this example, we will configure IPS to protect a web server. As shown in [Figure 87 on page 969](#), a FortiGate unit protects a web server and an internal network. The internal network will have its own policies and configuration but we will concentrate on the web server in this example.

Figure 87: A simple network configuration



The FortiGate unit is configured with:

- a virtual IP to give the web server a unique address accessible from the Internet.
- a security policy to allow access to the web server from the Internet using the virtual IP.

To protect the web server using intrusion protection, you need to create an IPS sensor, populate it with filters, then enable IPS scanning in the security policy.

To create an IPS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor* and select *Create New*.
- 2 Enter `web_server` as the name of the new IPS sensor.
- 3 Select *OK*.

The new IPS sensor is created but it has no filters, and therefore no signatures are included.

The web server operating system is Linux, so you need to create a filter for all Linux server signatures.

To create the Linux server filter

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor* and select the `web_server` IPS sensor and select the *Edit* icon.
- 2 Select *Add Filter*.
- 3 Enter `Linux Server` as the name of the new filter.
- 4 For *Target*, select *Specify* and choose *server*.
- 5 For *OS*, select *Specify* and choose *Linux*.
- 6 Select *OK*.

The filter is saved and the IPS sensor page reappears. In the filter list, find the *Linux Server* filter and look at the value in the *Count* column. This shows how many signatures match the current filter settings. You can select the *View Rules* icon to see a listing of the included signatures.

The web server software is Apache, so you need to create a second filter for all Apache signatures.

To create the Apache filter

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor* and select the `web_server` IPS sensor and select the *Edit* icon.
- 2 Select *Add Filter*.
- 3 Enter `Apache` as the name of the new filter.
- 4 For *Application*, select *Specify* and choose *Apache* from the *Available* list.
- 5 Select the right-arrow to move *Apache* to the *Selected* list.
- 6 Select *OK*.

The filter is saved and the IPS sensor page reappears.

It might seem that you can skip a step and create one filter that specifies both Linux server and Apache signatures. However, this would include a smaller number of filters. It would not include signatures to detect attacks against the operating system directly, for example.

You have created the IPS sensor and the two filters that include the signatures you need. To have it start scanning traffic, you must edit the security policy.

To edit the security policy

- 1 Go to *Policy > Policy > Policy*, select security policy that allows access to the web server, and select the *Edit* icon.

- 2 Enable *UTM*.
- 3 Select the Enable IPS option and choose the `web_server` IPS sensor from the list.
- 4 Select *OK*.

Since IPS is enabled and the `web_server` IPS sensor is specified in the security policy controlling the web server traffic, the IPS sensor examines the web server traffic for matches to the signatures it contains.

Create and test a packet logging IPS sensor

In this example, you create a new IPS sensor and include a filter that detects the EICAR test file and saves a packet log when it is found. This is an ideal first experience with packet logging because the EICAR test file can cause no harm, and it is freely available for testing purposes.

Create an IPS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select *Create New*.
- 3 Name the new IPS sensor `EICAR test`.
- 4 Select *OK*.

Create an entry

- 1 Select the Create New drop down menu and choose *Pre-defined Entry*.
- 2 Select the signature browse icon.
- 3 Rather than search through the signature list, use the name filter by selecting the filter icon in the header of the *Name* column.
- 4 In the *Filters* list, select *Name*.
- 5 Select *Enable*.
- 6 In the Field selection, choose *Contains*.
- 7 Enter `EICAR` in the Text field.
- 8 Select *OK*.
- 9 Select the *EICAR.AV.Test.File.Download* signature.
- 10 Select *OK*.
- 11 Select *Enable*, *Logging*, and *Packet Log*.
- 12 Select *OK*.
- 13 Select *Block* as the *Action*.
- 14 Select *OK* to save the IPS sensor.

You are returned to the IPS sensor list. The `EICAR test` sensor appears in the list.

Add the IPS sensor to the security policy allowing Internet access

- 1 Go to *Policy > Policy > Policy*.
- 2 Select the security policy that allows you to access the Internet.
- 3 Select the *Edit* icon.
- 4 Enable *Log Allowed Traffic*.
- 5 Enable *UTM*.
- 6 Select *Enable IPS*.

- 7 Choose `EICAR test` from the available IPS sensors.
- 8 Select *OK*.

With the IPS sensor configured and selected in the security policy, the FortiGate unit blocks any attempt to download the EICAR test file.

Test the IPS sensor

- 1 Using your web browser, go to http://www.eicar.org/anti_virus_test_file.htm.
- 2 Scroll to the bottom of the page and select *eicar.com* from the row labeled as using the standard HTTP protocol.
- 3 The browser attempts to download the requested file and,
 - If the file is successfully downloaded, the custom signature configuration failed at some point. Check the custom signature, the IPS sensor, and the firewall profile.
 - If the download is blocked with a high security alert message explaining that you're not permitted to download the file, the EICAR test file was blocked by the FortiGate unit antivirus scanner before the IPS sensor could examine it. Disable antivirus scanning and try to download the EICAR test file again.
 - If no file is downloaded and the browser eventually times out, the custom signature successfully detected the EICAR test file and blocked the download.

Viewing the packet log

- 1 Go to *Log&Report > Log & Archive Access > UTM Log*.
- 2 Locate the log entry that recorded the blocking of the EICAR test file block. The Message field data will be `tools: EICAR.AV.Test.File.Download`.
- 3 Select the *View Packet Log* icon in the *Packet Log* column.
- 4 The packet log viewer is displayed.

Creating a custom signature to block access to example.com

In this first example, you will create a custom signature to block access to the *example.com* URL.

This example describes the use of the custom signature syntax to block access to a URL. To create the custom signature entry in the FortiGate unit web-based manager, see [“Creating a custom IPS signature” on page 949](#).

- 1 Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

- 2 Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before adding any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.example.com"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3 Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; )
```

The signature will now detect the `example.com` URL appearing in network traffic. The custom signature should only detect the URL in HTTP traffic, however. Any other traffic with the URL should be allowed to pass. For example, an email message to or from `example.com` should not be stopped.

4 Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";  
--service HTTP; )
```

The FortiGate unit will limit its search for the pattern to the HTTP protocol. Even though the HTTP protocol uses only TCP traffic, the FortiGate will search for HTTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5 Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";  
--service HTTP; --protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore UDP and ICMP network traffic.

6 Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to `Example.com`, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";  
--service HTTP; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7 Limit pattern scans to only traffic sent from the client

The `--flow` command can be used to further limit the network traffic being scanned to only that sent by the client or by the server.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";  
--service HTTP; --no_case; --flow from_client; )
```

Web servers do not contact clients until clients first open a communication session. Therefore, using the `--flow from_client` command will force the FortiGate to ignore all traffic from the server. Since the majority of HTTP traffic flows from the server to the client, this will save considerable system resources and still maintain protection.

8 Specify the context

When the client browser tries to contact example.com, a DNS is first consulted to get the example.com server IP address. The IP address is then specified in the URL field of the HTTP communication. The domain name will still appear in the host field, so this custom signature will not function without the `--context host` keyword/value pair.

```
F-SBID( --name "Block.example.com"; --pattern "example.com";
        --service HTTP; --no_case; --flow from_client;
        --context host; )
```

Creating a custom signature to block the SMTP “vrfy” command

The SMTP “vrfy” command can be used to verify the existence of a single email address or to list all of the valid email accounts on an email server. A spammer could potentially use this command to obtain a list of all valid email users and direct spam to their inboxes.

In this example, you will create a custom signature to block the use of the vrfy command. Since the custom signature blocks the vrfy command from coming through the FortiGate unit, the administrator can still use the command on the internal network.

This example describes the use of the custom signature syntax to block the vrfy command. To create the custom signature entry in the FortiGate unit web-based manager, see [“Creating a custom IPS signature” on page 949](#).

1 Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2 Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before you add any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3 Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; )
```

The signature will now detect the vrfy command appearing in network traffic. The custom signature should only detect the command in SMTP traffic, however. Any other traffic with the pattern should be allowed to pass. For example, an email message discussing the vrfy command should not be stopped.

4 Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
        --service SMTP; )
```

The FortiGate unit will limit its search for the pattern to the SMTP protocol.

Even though the SMTP protocol uses only TCP traffic, the FortiGate will search for SMTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5 Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
        --service SMTP; --protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore the pattern in UDP and ICMP network traffic.

6 Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
        --service SMTP; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7 Specify the context

The SMTP vrfy command will appear in the SMTP header. The `--context host` keyword/value pair allows you to limit the pattern search to only the header.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy";
        --service SMTP; --no_case; --context header; )
```

Configuring a Fortinet Security Processing module

The Example Corporation has a web site that is the target of SYN floods. While they investigate the source of the attacks, it's very important that the web site remain accessible. To enhance the ability of the company's FortiGate-620B to deal with SYN floods, the administrator will install an ASM-CE4 Fortinet Security Processing module and have all external access to the web server come through it.

The security processing modules not only accelerate and offload network traffic from the FortiGate unit's processor, but they also accelerate and offload security and content scanning. The ability of the security module to accelerate IPS scanning and DoS protection greatly enhances the defense capabilities of the FortiGate-620B.

Assumptions

As shown in other examples and network diagrams throughout this document, the Example Corporation has a pair of FortiGate-620B units in an HA cluster. To simplify this example, the cluster is replaced with a single FortiGate-620B.

An ASM-CE4 is installed in the FortiGate-620B.

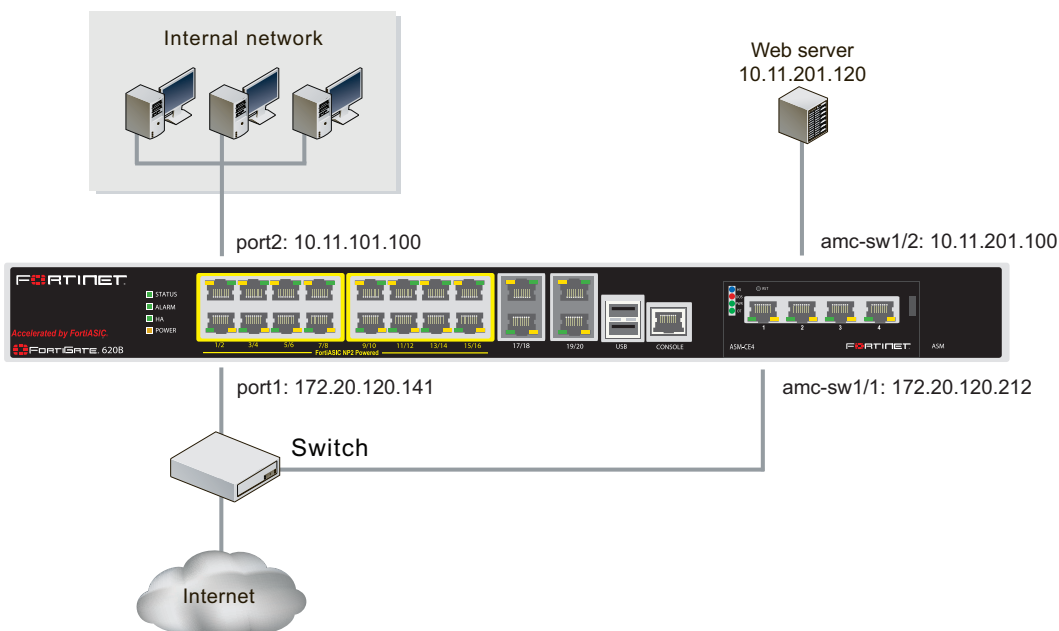
The network is configured as shown in [Figure 88](#).

Network configuration

The Example Corporation network needs minimal changes to incorporate the ASM-CE4. Interface amc-sw1/1 of the ASM-CE4 is connected to the Internet and interface amc-sw1/1 is connected to the web server.

Since the main office network is connected to port2 and the Internet is connected to port1, a switch is installed to allow both port1 and amc-sw1/1 to be connected to the Internet.

Figure 88: The FortiGate-620B network configuration



The switch used to connect port1 and amc-sw1/1 to the Internet must be able to handle any SYN flood, all of the legitimate traffic to the web site, and all of the traffic to and from the Example Corporation internal network. If the switch can not handle the bandwidth, or if the connection to the service provider can not provide the required bandwidth, traffic will be lost.

Security module configuration

The Fortinet security modules come configured to give equal priority to content inspection and firewall processing. The Example Corporation is using a ASM-CE4 module to defend its web server against SYN flood attacks so firewall processing is a secondary consideration.

Use these CLI commands to configure the security module in ASM slot 1 to devote more resources to content processing, including DoS and IPS, than to firewall processing.

```
config system amc-slot
edit sw1
set optimization-mode fw-ips
set ips-weight balanced
set ips-p2p disable
set ips-fail-open enable
set fp-disable none
set ipsec-inb-optimization enable
set syn-proxy-client-timer 3
```

```
set syn-proxy-server-timer 3
end
```

These settings do not disable firewall processing. Rather, when the security module nears its processing capacity, it will choose to service content inspection over firewall processing.

DoS sensor configuration

Defend against anomaly-based attacks using a DoS sensor. For the SYN floods launched against the Example Corporation web site, the *tcp_syn_flood* anomaly is the best defense.

Create a DoS sensor for SYN flood protection

- 1 Go to *UTM Profiles > Intrusion Protection > DoS Sensor*.
- 2 Select *Create New*.
- 3 Enter *Web site SYN protection* for the DoS sensor name.
- 4 Select *OK* to create the sensor.

The default *tcp_syn_flood* threshold is 2000. This means that the configured action will be triggered when the number of TCP packets with the SYN flag set exceeds 2000 per second.

For some applications, this value will be too high, while for others it will be too low. One way to find the correct values for your environment is to set the action to *Pass* and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your network will be from DoS attacks, but your network traffic will also be more likely to generate false alarms.

Configure a DoS sensor for SYN flood protection

- 1 Go to *UTM Profiles > Intrusion Protection > DoS Sensor*.
- 2 Select the *Web site SYN protection* sensor and select the *Edit* icon.
- 3 Select *Enable* and *Logging* for the *tcp_syn_flood* anomaly.
- 4 Select the *Proxy* action for the *tcp_syn_flood* anomaly.
- 5 Enter the threshold value for the *tcp_syn_flood* anomaly.
- 6 Select *OK*.

With the action configured as *Proxy*, TCP packets with the SYN flag set will be passed until the threshold value is exceeded. At that point, TCP packets with the SYN flag set until their numbers fall below the threshold value.

The ASM-CE4 security module will intercept the packet, and reply to the client with a TCP packet that has the SYN and ACK flags set. If the connection request is legitimate, the client will reply with a packet that has the ACK flag set. The ASM-CE4 will then 'replay' this exchange to the server and allow the client and server to communicate directly.

If the client does not reply with the expected packet, the ASM-CE4 will close the connection. Therefore, if the security module receives a flood of SYN packets, they will be blocked. Only the legitimate connections will be allowed through to the server.

DoS policy configuration

Before the DoS sensor can begin examining network traffic, you must create and configure a DoS policy and specify the DoS sensor.

Create a DoS policy

- 1 Go to *Policy > Policy > DoS Policy*.
- 2 Select *Create New*.
- 3 Select *amc-sw1/1* for *Source Interface/Zone*.
- 4 Select *all* for *Source Address*.
- 5 Select *all* for *Destination Address*.
- 6 Select *ANY* for *Service*.
- 7 Enable *DoS Sensor* and select the *Web site SYN protection* sensor from the list.
- 8 Select *OK*.

Virtual IP configuration

Traffic destined for the web server will arrive at the *amc-sw1/1* interface. You must create a virtual IP mapping to have the ASM-CE4 direct the traffic to the web server.

Create a virtual IP mapping

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.
- 3 In the *Name* field, enter *web_server*.
- 4 Select *amc-sw1/1* as the *External Interface*.
- 5 Enter *172.20.120.212* as the *External IP Address/Range*.
- 6 Enter *10.11.201.120* as the *Mapped IP Address/Range*.
- 7 Select *OK*.

Security policy configuration

A security policy is required to allow traffic through to the web server. Further, the security policy must include the virtual IP so the traffic is directed to the web server.

Create a security policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Select *amc-sw1/1* for the *Source Interface/Zone*.
- 4 Select *all* for the *Source Address*.
- 5 Select *amc-sw1/2* for the *Destination Interface/Zone*.
- 6 Select *web_server* for the *Destination Address*.
- 7 Select *Enable NAT*.
- 8 Select *OK*.

Attempts to connect to 172.20.120.212 will be forwarded to the web server with this security policy in place.

View proxy statistics

With a FortiGate security module installed, a CLI command displays the current proxy statistics.

At the CLI prompt, type `execute npu-cli /dev/ce4_0 showsynproxy`. The last nine lines will list the proxy statistics:

```
Total Proxied TCP Connections:      434055223
Working Proxied TCP Connections:    515699
Retired TCP Connections:            433539524
Valid TCP Connections:              0
Attacks, No Ack From Client:        433539524
No SynAck From Server:              0
Rst By Server (service not supported): 0
Client timeout setting:              3 Seconds
Server timeout setting:              3 Seconds
```

Total Proxied TCP Connections	The number of proxied TCP connection attempts since the FortiGate unit was restarted. This value is the sum of the working and retired connection totals.
Working Proxied TCP Connections	The number of TCP connection attempts currently being proxied.
Retired TCP Connections	The number of proxied TCP connection attempts dropped or allowed. These connection attempts are no-longer being serviced. This value is the sum of the valid and attacks totals.
Valid TCP Connections	The number of valid proxied TCP connection attempts.
Attacks, No Ack From Client	The number of proxied TCP connection attempts in which the client did not reply. These are typically attacks.
No SynAck From Server	The number of valid client connection attempts in which the server does not reply.
Rst By Server (service not supported)	The number of valid client connection attempts in which the server resets the connection.
Client timeout setting	The client time-out duration.
Server timeout setting	The server time-out duration.

Intrusion Protection interface reference

The Intrusion Protection system combines signature and anomaly detection and prevention with low latency and excellent reliability. With Intrusion Protection, you can create multiple IPS sensors, each containing a complete configuration based on signatures. Then, you can apply any IPS sensor to a firewall policy. You can also create DoS sensors to examine traffic for anomaly-based attacks.

This topic contains the following:

- [IPS Sensor](#)
- [DoS sensor](#)
- [Predefined](#)
- [Custom](#)
- [Protocol Decoder](#)

IPS Sensor

You can group signatures into IPS sensors for easy selection when applying to firewall policies. You can define signatures for specific types of traffic in separate IPS sensors, and then select those sensors in profiles designed to handle that type of traffic. For example, you can specify all of the web-server related signatures in an IPS sensor, and that sensor can then be applied to a firewall policy that controls all of the traffic to and from a web server protected by the unit.

The FortiGuard Service periodically updates the pre-defined signatures, with signatures added to counter new threats. Since the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Each IPS sensor consists of two parts: filters and overrides. Overrides are always checked before filters.

Each filter consists of a number of signatures attributes. All of the signatures with those attributes, and only those attributes, are checked against traffic when the filter is run. If multiple filters are defined in an IPS Sensor, they are checked against the traffic one at a time, from top to bottom. If a match is found, the unit takes the appropriate action and stops further checking.

A signature override can modify the behavior of a signature specified in a filter. A signature override can also add a signature not specified in the sensor's filters. Custom signatures are included in an IPS sensor using overrides.

The signatures in the overrides are first compared to network traffic. If the IPS sensor does not find any matches, it then compares the signatures in each filter to network traffic, one filter at a time, from top to bottom. If no signature matches are found, the IPS sensor allows the network traffic.

The signatures included in the filter are only those matching every attribute specified. When created, a new filter has every attribute set to *all* which causes every signature to be included in the filter. If the severity is changed to high, and the target is changed to server, the filter includes only signatures checking for high priority attacks targeted at servers.

IPS sensor configuration settings

The following are IPS sensor configuration settings in *UTM Profiles > Intrusion Protection > IPS Sensor*.

IPS Sensor page Lists each individual IPS sensor, either default or ones that you created. On this page you can edit, delete or create a new IPS sensor. If you want to configure tags, and the tag options are not available on the web-based manager, you must enable tag options by going to <i>System > Admin > Settings</i> . Note: You can configure IPS signatures to not be triggered until the threshold is met; however, this is configured only in the CLI. You must use <code>config override in config ips sensor</code> command to configure this option.	
Create New	Creates a new IPS sensor. When you select <i>Create New</i> , you are automatically redirected to the New IPS Sensor page. This page provides a name field and comment field. You must enter a name to go the IPS Sensor Settings page.
Delete	Removes the IPS sensor from the list. To remove multiple IPS sensors from within the list, on the IPS Sensor page, in each of the rows of the sensors you want removed, select the check box and then select <i>Delete</i> . To remove all IPS sensors from the list, on the IPS Sensor page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies settings within an IPS sensor. When you select <i>Edit</i> , you are automatically redirected to the Edit IPS Sensor page.
Name	The name of each IPS sensor.
Comments	An optional description of the IPS sensor.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
all_default (default)	Includes all signatures. The sensor is set to use the default enable status and action of each signature.
all_default_pass (default)	Includes all signatures. The sensor is set to use the default enable status of each signature, but the action is set to pass.
protect_client (default)	Includes only the signatures designed to detect attacks against clients and uses the default enable status and action of each signature.
protect_email_server (default)	Includes only the signatures designed to detect attacks against servers and the SMTP, POP3, or IMAP protocols and uses the default enable status and action of each signature.
protect_http_server (default)	Includes only the signatures designed to detect attacks against servers and the HTTP protocol and uses the default enable status and action of each signature.
Edit IPS Sensor page <p>Provides settings for configuring multiple filters and overrides that make up an IPS sensor. The Edit IPS Sensor Settings page contains two sections, one called Filters where you can configure filters and the other called Overrides where you can configure either predefined or custom overrides.</p> <p>Note: Logging is enabled in the CLI.</p>	
Name	If you are editing an existing IPS sensor and you want to change the name, enter a new name in the field. You must select <i>Apply</i> to save the change.

Comments	If you are editing an existing IPS sensor and you want to change the description, enter the changes in the field. You must select <i>Apply</i> to save the change.
Create New	Creates a new filter. You can also use the Insert icon to create a new filter for an IPS sensor. When you select <i>Create New</i> , you are automatically redirected to the Edit IPS Filter page. If you want to create a predefined override or a custom override, select the down arrow beside <i>Create New</i> . From the list, you can select <i>Pre-defined Entry</i> , which creates a predefined override, or you can select <i>Custom Entry</i> , which creates a custom override. For more information about overrides, see “Pre-defined overrides and custom overrides configuration settings” on page 985 .
Edit	Modifies settings within a filter. When you select <i>Edit</i> , you are automatically redirected to the Edit IPS Filter page.
Delete	Removes a filter from the list within the Filters section of the IPS Sensor Settings page. To remove multiple filter lists from within the list, in the Filters section, in each of the rows of the filters you want removed, select the check box and then select <i>Delete</i> . To remove all filters from the list, in the Filters section, select the check box in the check box column and then select <i>Delete</i> .
Insert	Inserts a new filter in filter list in the list in the Filters section. When you select <i>Insert</i> , you are automatically redirected to the Edit IPS Filter page.
Move To	Moves a filter to any position within the list in the Filters section. You must select the check box in the row of the filter you want moved so that filter will be moved within the list. When you select <i>Move To</i> , the following appears: <code>Please enter the destination filter position.</code> Enter the number for the filter's new position within the list, for example, 5, to place the first entry in the fifth position Select <i>OK</i> .
View Rules	View the rules of a filter. When you select <i>View Rules</i> , the Matched Rules window appears. Scroll through the list to see all the rules within that filter.
ID	The identification number of the entry you created.
Severity	The severity level of the filter.
Target	The target specified for that filter.
Protocol	The type of protocol for that filter.
OS	The type of operating system.
Application	The software application, such as Adobe.

Enable	A green check mark appears if you select <i>Enable all</i> within the filter's settings. If you select <i>Disable all</i> , a gray x appears. Note: For the default IPS sensor (called "default") the word <i>Default</i> displays instead of the check mark. It cannot be disabled.
Action	The type of action the unit will take. This action can be <i>Block</i> , <i>Pass</i> , or <i>Reset</i> .
Packet Logging	A green checkmark appears if you select <i>Enable all</i> within the filter's settings. A gray x appears if you select <i>Disable all</i> .
Matched Signatures	The number of signatures included in the filter. Overrides are not included in the total.

Filters configuration settings

A filter is a collection of signature attributes that you specify. The signatures that have all of the attributes specified in a filter are included in the IPS signature. An IPS sensor can contain multiple IPS filters. The following are the available options when configuring filters.

The following are filter configuration settings for an IPS sensor in *UTM Profiles > Intrusion Protection > IPS Sensors*.

New IPS Filter page Provides settings for configuring a filter. You are automatically redirected to this page when you select <i>Create New</i> in the <i>Filters</i> section of the IPS Sensor Settings page.	
Name	Enter a name for the filter.
Severity	Select a severity level. You must specify a severity level if you do not want to all severity levels.
Target	Select the type of system targeted by the attack.
OS	Select to specify the type of operating system, or select <i>All</i> to include all operating systems. The operating system available include BSD and Solaris. Signatures with an OS attack attribute of <i>All</i> affect all operating system and these signatures are automatically included in any filter regardless of whether a single, multiple, or all operating systems are specified.
Protocol	Select to choose multiple protocols or all available protocols. To select specific protocols, select <i>Specify</i> , and then move each protocol that you want from the <i>Available</i> column to the <i>Selected</i> column using the -> arrow. To remove a protocol from the <i>Selected</i> column, select the protocol and then use the <- arrow to move the protocol back to the <i>Available</i> column.
Application	Select to choose multiple applications or all available applications. To select specific applications, select <i>Specify</i> , and then move each application that you want from the <i>Available</i> list to the <i>Selected</i> list using the -> arrow. To remove an application from the <i>Selected</i> list, select the protocol and then use the <- arrow to move the application back to the <i>Available</i> list.

Tags	Applies tags to the file filter only. These tags do not display within the Filters section, only within the filter itself on the Edit IPS Filter page. If you do not see <i>Tags</i> and its available settings, go to <i>System > Admin > Settings</i> to enable them on the web-based manager.
Applied tags	Displays the tags that you have added to the filter.
Add tags	Enter the tag in the field and then select the plus (+) sign to add the tag to the filter. This also adds the tag to the <i>Applied tags</i> list.
Enable All Matching Signatures	Select to accept the defaults, or enable all signatures, or disable all.
Action When Signature Is Triggered	Select to indicate to the unit what action to take when a signature is triggered. When you select <i>Accept signature defaults</i> , and then select to enable <i>Quarantine Attackers (to Banned Users List)</i> , the <i>Method</i> and <i>Expires</i> options display. The <i>Quarantine Attackers (to Banned Users List)</i> appears only when <i>Accept signature defaults</i> is selected.
Quarantine Attackers (to Banned Users List)	Select if you want to add an attacker to the <i>Banned Users List</i> .
Method	Select <i>Attacker's IP Address</i> to block all traffic sent from the attacker's IP address. Traffic from the attacker's IP address is blocked because the attacker's IP address is in the Banned Users List. Select <i>Attacker and Victim IP Addresses</i> to block all traffic sent from the attacker IP address to the target (victim) IP address. Traffic from the attacker IP address to addresses other than the victim IP address is allowed. The attacker and target IP addresses are added to the banned user list as one entry. Select <i>Attack's Incoming Interface</i> to block all traffic from connecting to the FortiGate interface that received the attack. The interface is added to the banned user list.
Expires	You can select whether the attacker is banned indefinitely or for a specified number of days, hours, or minutes. Select <i>Indefinitely</i> if you do not want the information to expire. Select <i>After</i> and enter the number of hours, minutes or days.
Packet Logging	Select to enable packet logging on the filter. When you select to enable packet logging on a filter, the unit saves a copy of the packets that triggered an attack detection for future analysis or for audit purposes. Note: You must enable logging for both filter and sensor so that logging occurs. The logging option is available only in the CLI.

Pre-defined overrides and custom overrides configuration settings

Pre-defined and custom overrides are configured and work mainly in the same way as filters. Unlike filters, each override defines the behavior of one signature.

Overrides can be used in two ways:

- Change the behavior of a signature already included in a filter. For example, to protect a web server, you could create a filter that includes and enables all signatures related to servers. If you wanted to disable one of those signatures, the simplest way would be to create an override and mark the signature as disabled.
- Add an individual signature that is not included in any filters to an IPS sensor. This is the only way to add custom signatures to IPS sensors.

When a pre-defined signature is specified in an override, the default status and action attributes have no effect. These settings must be explicitly set when creating the override.

When configuring either a pre-defined override or a custom override, the following options are available regardless which override you are configuring.

Predefined and custom overrides are configured in the IPS Sensor itself, located in *UTM Profiles > Intrusion Protection > IPS Sensors*. The following are configuration settings for both predefined and custom overrides.



Before an override can affect network traffic, you must add it to a filter, and you must select the IPS sensor and then apply it to a policy. An override does not have the ability to affect network traffic until these steps are taken.

Configure IPS Entry page

Provides settings for configuring both predefined overrides and custom overrides. You are automatically redirected to this page after selecting either *Pre-defined Entry* or *Custom Entry*, which is accessed when you select the down arrow beside *Create New*.

Note: Logging is enabled in the CLI.

Signature	Select the browse icon to view the list of available signatures. From this list, select a signature the override will apply to and then select <i>OK</i> .
Enable	Select to enable the signature override.
Action	Select <i>Pass</i> , <i>Block</i> or <i>Reset</i> . When the override is enabled, the action determines what the unit will do with traffic containing the specified signature.
Packet Log	Select to save packets that trigger the override to the unit's hard drive for later examination.
Quarantine Attackers (to Banned Users List)	Select to enable NAC quarantine for this override. The unit deals with the attack according to the IPS sensor or DoS sensor configuration regardless of this setting.

Method	<p>Select <i>Attacker's IP address</i> to block all traffic sent from the attackers IP address. The attackers IP address is also added to the banned user list. The target address is not affected.</p> <p>Select <i>Attacker and Victim IP Addresses</i> to block all traffic sent from the attacker IP address to the target (victim) IP address. Traffic from the attacker IP address to addresses other than the victim IP address is allowed. The attacker and target IP addresses are added to the banned user list as one entry.</p> <p>Select <i>Attack's Incoming Interface</i> to block all traffic from connecting to the FortiGate interface that received the attack. The interface is added to the banned user list.</p>
Expires	You can select whether the attacker is banned indefinitely or for a specified number of days, hours, or minutes.
Exempt IP	Enter IP addresses to exclude from the override. The override will then apply to all IP addresses except those defined as exempt. The exempt IP addresses are defined in pairs, with a source and destination, and traffic moving from the source to the destination is exempt from the override.
Source	The exempt source IP address. Enter 0.0.0.0/0 to include all source IP addresses.
Destination:	The exempt destination IP address. Enter 0.0.0.0/0 to include all destination IP addresses.
Add	Select to add other exempt IP addresses to the list in the table below Add.
#	The number identifying the order of the item in the list.
Source	The source IP address and netmask entered.
Destination	The destination IP address and netmask entered.
Delete	Select to remove an item in the list.

DoS sensor

IPS uses a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and behavior. For example, one type of flooding is the denial of service (DoS) attack that occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system so legitimate users can no longer use it. This type of attack gives the DoS sensor its name, although it is capable of detecting and protecting against a number of anomaly attacks.

You can enable or disable logging for each traffic anomaly, and configure the detection threshold and action to take when the detection threshold is exceeded.

You can create multiple DoS sensors. Each sensor consists of 12 anomaly types that you can configure. When a sensor detects an anomaly, it applies the configured action. One sensor can be selected for use in each DoS policy, allowing you to configure the anomaly thresholds separately for each interface. Multiple sensors allow great granularity in detecting anomalies because each sensor can be configured for the specific needs of the interface it is attached to by the DoS policy.

The traffic anomaly detection list can be updated only when the firmware image is upgraded on the unit.

Since an improperly configured DoS sensor can interfere with network traffic, no DoS sensors are present on a factory default unit. You must create your own and then select them in a DoS policy before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.



It is important to know normal and expected network traffic before changing the default anomaly thresholds. Setting the thresholds too low could cause false positives, and setting the thresholds too high could allow otherwise avoidable attacks.



If virtual domains are enabled on the unit, the Intrusion Protection settings must be configured separately in each VDOM. All sensors and custom signatures will appear only in the VDOM in which they were created.

DoS sensor configuration settings

The following are DoS sensor configuration settings in *UTM Profiles > Intrusion Protection > DoS Sensor*.

DoS Sensor page

Lists each default DoS sensor and each DoS sensor that you created. On this page, you can create, edit or delete a DoS sensor.

Create New	Creates a new DoS sensor. When you select <i>Create New</i> , you are automatically redirected to the New DoS Sensor page. The New DoS Sensor page provides a name field and a comment file. You must enter a name to go to the Edit DoS Sensor page.
Edit	Modifies the settings within a DoS sensor. You can modify the following information: Action, Severity, and Threshold. When you select <i>Edit</i> , you are automatically redirected to the Edit DoS Sensor page.
Delete	Removes a DoS sensor from the list on the DoS Sensor page. To remove multiple DoS sensors from within the list, on the DoS Sensor page, in each of the rows of the sensors you want removed, select the check box and then select <i>Delete</i> . To remove all DoS sensors from the list, on the DoS Sensor page, select the check box in the check box column and then select <i>Delete</i> .
Name	The DoS sensor name.
Comments	An optional description of the DoS sensor.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
Edit DoS Sensor page Provides settings for configuring the action type, threshold amount, and if logging should be enabled for the anomaly. There are twelve default anomalies to configure settings for. If you are editing a DoS Sensor, you are redirected to this page.	
Name	If you are editing an existing DoS sensor setting and want to change the name, enter a new name in this field. You must select <i>OK</i> to save the change.
Comments	If you are editing an existing DoS sensor setting and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save these changes.
Anomalies Configuration	
Name	The name of the anomaly.
Enable	Select the check box to enable the DoS sensor to detect when the specified anomaly occurs. Selecting the check box in the header row will enable all anomalies.
Logging	Select the check box to enable the DoS sensor to log when the anomaly occurs. Selecting the check box in the header row will enable logging for all anomalies. Anomalies that are not enabled are not logged.

Action	Select <i>Pass</i> to allow anomalous traffic to pass when the unit detects it, or set <i>Block</i> to prevent the traffic from passing.
Threshold	Displays the number of sessions/packets that must show the anomalous behavior before the FortiGate unit triggers the anomaly action (pass or block). If required, change the number. Range 1 to 2 147 483 647. For more information about how these settings affect specific anomalies, see Table 68 on page 990 and “SYN threshold (preventing SYN floods using a DoS sensor)” on page 990.

SYN proxy

FortiGate units with Fortinet security processing modules installed offer a third action for the tcp_syn_flood threshold when a module is installed. Instead of Block and Pass, you can choose to Proxy the incomplete connections that exceed the threshold value.

When the tcp_syn_flood threshold action is set to proxy, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the unit receives an ACK response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

SYN threshold (preventing SYN floods using a DoS sensor)

The preferred primary defense against any type of SYN flood is the DoS sensor tcp_syn_flood threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to Pass, the unit will allow the SYN packets that exceed the threshold. If the action is set to Block, the unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet with a SYN+ACK response if not received from the server. These tools will not “retry”. Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to Block.

Understanding the anomalies

Each DoS sensor offers four configurable statistical anomaly types for each of the TCP, UDP, and ICMP protocols.

Table 67: The four statistical anomaly types.

For each of the TCP, UDP, and ICMP protocols, DoS sensors offer four statistical anomaly types. The result is twelve configurable anomalies, which are shown in [Table 68](#).

Table 68: The twelve individually configurable anomalies

Anomaly	Description
tcp_syn_flood	If the SYN packet rate, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.

Table 68: The twelve individually configurable anomalies (Continued)

Anomaly	Description
<code>tcp_port_scan</code>	If the SYN packets rate, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
<code>tcp_src_session</code>	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.
<code>tcp_dst_session</code>	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.
<code>udp_flood</code>	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
<code>udp_scan</code>	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
<code>udp_src_session</code>	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.
<code>udp_dst_session</code>	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.
<code>icmp_flood</code>	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
<code>icmp_sweep</code>	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
<code>icmp_src_session</code>	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.
<code>icmp_dst_session</code>	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.

Predefined

The Intrusion Protection system can use signatures once you have grouped the required signatures in an IPS sensor. If required, you can override the default settings of the signatures specified in an IPS sensor. The unit provides a number of pre-built IPS sensors, but you should check their settings before using them, to ensure they meet your network requirements.

By using only the signatures you require, you can improve system performance and reduce the number of log messages and alert email messages that the IPS sensor generates. For example, if the unit is not protecting a web server, web server signatures are not included.

The predefined signature list, located in *UTM Profiles > Intrusion Protection > Predefined*, includes signatures that are currently in the [FortiGuard Center Vulnerability Encyclopedia](#). This encyclopedia also includes additional signatures not found in the Predefined menu. Each signature name is a link to the vulnerability encyclopedia entry for the signature. The vulnerability encyclopedia describes the attack detected by the signature and provides recommended actions and links for more information.

The predefined signature list also includes characteristics such as the severity level of the attack, protocol, and applications affected for each signature. These characteristics give you a quick reference to what the signature is for. You can also use these characteristics to sort the signature list, grouping signatures by common characteristics. The signature list also displays the default action, the default logging status, and whether the signature is enabled by default. The signatures are sorted by name, which is default. The predefined signature list table allows you to view these characteristics in detail. The table is located at the bottom of the page; however, you can customize the table so that it appears on the right side of the page or hidden.

Viewing predefined signatures

You can view predefined signatures in *UTM Profiles > Intrusion Protection > Predefined*.

When you are viewing signatures, you can view each one in detail, from the IPS signatures viewer table. The IPS Signatures Viewer table appears, by default, at the bottom of the page; however, you can choose to hide the table or position the table at the right side of the page.



If virtual domains are enabled on the unit, the Intrusion Protection settings are configured separately in each VDOM. All sensors and custom signatures will appear only in the VDOM in which they were created.

Predefined page

Lists each predefined signature that is currently on your unit. When you select the name of the signature, you are automatically redirected to that signature's detailed definition in the FortiGuard Center Vulnerability Encyclopedia. This page also indicates which signatures are enabled and which are disabled.

Tip: To determine what effect IPS protection will have on your network traffic, enable the required signatures, set the *Action* to *Pass*, and enable logging. Traffic will not be interrupted, and you will be able to examine, in detail, which signatures were detected.

Tags	<p>Select to add or remove tags to the predefined signature.</p> <p>Note: If tag settings are not available on the web-based manager, they must be enabled in <i>System > Admin > Settings</i> to then appear on the Predefined page.</p> <p>When you select the down arrow beside <i>Tags</i>, you can add tags or remove tags.</p> <p>To add tags to a predefined signature, select the signature first, select the down arrow beside <i>Tags</i>, and then select <i>Add Tags</i>. The Add Tags window appears. Enter the tag in the <i>Add Tag</i> field and then select the plus (+) sign; repeat until all tags are in the <i>Tags to apply</i> list.</p> <p>To remove tags, select the signature first, select the down arrow beside <i>Tags</i>, and then select <i>Remove Tags</i>. The Remove Tags window appears. Select the tags that you want removed in the <i>Applied Tags</i> row; repeat until all the tags are in the <i>Tags to remove</i> row. The tags will automatically be put in the <i>Tags to remove</i> row after being selected in the <i>Applied Tags</i> row.</p> <p>Note: When you select a signature, the Signature Viewer Table appears. The Signature Viewer Table is the same as the Log Viewer Table, providing detailed information about a signature.</p> <p>If there are tags that you want to add that have been configured for another object, you can add those tags as well to signatures. To apply these other object tags, select the signature first, select the down arrow beside <i>Tags</i>, and then select <i>Add Tags</i>. The Add Tags window appears. Select the tags you want to add in the <i>Click tag to add</i> row. The tags automatically appear in the <i>Tags to apply</i> row. Select <i>OK</i> to add those tags to the application.</p>
Column Settings	<p>Select to customize the signature information displayed in the table. You can also readjust the column order.</p>

Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings. To modify settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filter settings, select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p> <p>The <i>Filters Settings</i> on the <i>Predefined</i> page contains <i>Copy to Sensor</i>, which allows you to copy filter settings and apply them to a IPS sensor.</p> <p>To apply existing filter settings to a sensor, select the down arrow beside <i>Filter Settings</i>, and then select <i>Copy to Sensor</i>. The <i>Select Object</i> window appears. Select the sensor that you want to apply the settings to from the drop-down list. Select <i>OK</i>.</p>
View Details	Select the down arrow to choose to either hide the predefined signature viewer table or position the table on the right side of the page.
Search	Enter search criteria in the field provided. Select the <i>Clear All</i> icon beside the field to clear the criteria for a new search.
Name	The name of the signature. Each name is also a link to the description of the signature in the FortiGuard Center Vulnerability Encyclopedia .
Severity	The severity rating of the signature. The severity levels, from lowest to highest, are <i>Information</i> , <i>Low</i> , <i>Medium</i> , <i>High</i> , and <i>Critical</i> . These levels appear as bars in this column; each bar is explained in the predefined signature viewer table.
Target	The target of the signature: servers, clients, or both.
Protocols	The protocol the signature applies to.
OS	The operating system the signature applies to.
Applications	The applications the signature applies to.
Tags	The tags that are applied to the predefined signature.
Enable	The default status of the signature. A green check mark indicates the signature is enabled. A gray x indicates the signature is not enabled.
Action	<p>The default action for the signature:</p> <ul style="list-style-type: none"> • <i>Pass</i> – allows the traffic to continue without any modification. • <i>Drop</i> – prevents the traffic with detected signatures from reaching its destination. <p>If logging is enabled, the action appears in the status field of the log message generated by the signature.</p>
Page Controls	Use to navigate through the list on the page.

Custom

Custom signatures provide the power and flexibility to customize the Intrusion Protection system for diverse network environments. The predefined signatures represent common attacks. If you use an unusual or specialized application or an uncommon platform, you can add custom signatures based on the security alerts released by the application and platform vendors.

You can also create custom signatures to help you block P2P protocols.

After creating custom signatures, you need to specify them in IPS sensors that were created to scan traffic.

Use custom signatures to block or allow specific traffic. For example, to block traffic containing profanity, add custom signatures similar to the following:

```
set signature 'F-SBID (--protocol tcp; --flow bi_direction; --
pattern "bad words"; --no_case) '
```

Custom signatures must be added to a signature override in an IPS filter to have any effect. Creating a custom signature is a necessary step, but a custom signature does not affect traffic simply by being created.



Custom signatures are an advanced feature. This document assumes the user has previous experience creating intrusion detection signatures.

Custom signature configuration settings

The following are custom signature configuration settings in *UTM Profiles > Intrusion Protection > Custom*.



If virtual domains are enabled on the unit, the Intrusion Protection settings are configured separately in each VDOM. All sensors and custom signatures will appear only in the VDOM in which they were created.

Custom page

Lists each custom signature that you created. On this page you can edit, delete or create a new custom signature.

Create New	Creates a new custom signature. When you select <i>Create New</i> , you are automatically redirected to the New Custom Signature page.
Edit	Modifies the name or signature of a custom signature. When you select <i>Edit</i> , you are automatically redirected to the Edit Custom Signature page.
Delete	Removes a custom signature from the list on the Custom page. To remove multiple custom signatures from within the list, on the Custom page, in each of the rows of the custom signatures you want removed, select the check box and then select <i>Delete</i> . To remove all custom signatures from the list, on the Custom page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the custom signature.
Signature	The signature itself.

New Custom Signature page Provides settings for configuring a new custom signature.	
Name	Enter a name for the custom signature.
Signature	Enter the signature.

Protocol Decoder

The Intrusion Protection system uses protocol decoders to identify the abnormal traffic patterns that do not meet the protocol requirements and standards. For example, the HTTP decoder monitors traffic to identify any HTTP packets that do not meet the HTTP protocol standards.


The decoder list is provided for your reference and can be configured using the CLI.

You can view protocol decoders in *UTM Profiles > Intrusion Protection > Protocol Decoder*.

Protocol Decoder page Displays a list of the current protocol decoders that are on your unit. The unit automatically updates this list by contacting the FDN. This lists includes the port number that the protocol decoder monitors.	
Protocols	The protocol decoder name.
Ports	The port number or numbers that the decoder monitors.

Upgrading the IPS protocol decoder list

The Intrusion Protection system protocol decoders are upgraded automatically through the FortiGuard Distribution Network (FDN) if existing decoders are modified or new decoders added. The FDN keeps the protocol decoder list up-to-date with protection against new threats such as the latest versions of existing IM/P2P as well as against new applications.



Web filter

This section describes FortiGate web filtering for HTTP traffic. The three main parts of the web filtering function, the Web Content Filter, the URL Filter, and the FortiGuard Web Filtering Service interact with each other to provide maximum control over what the Internet user can view as well as protection to your network from many Internet content threats. Web Content Filter blocks web pages containing words or patterns that you specify. URL filtering uses URLs and URL patterns to block or exempt web pages from specific sources. FortiGuard Web Filtering provides many additional categories you can use to filter web traffic.

This section describes the Web Content Filter and URL Filter functions. For information on FortiGuard Web Filtering, see [“FortiGuard Web Filter” on page 1023](#).

The following topics are included in this section:

- [Web filter concepts](#)
- [Web content filter](#)
- [URL filter](#)
- [SafeSearch](#)
- [Advanced web filter configuration](#)
- [Web filtering example](#)

Web filter concepts

Web filtering is a means of controlling the content that an Internet user is able to view. With the popularity of web applications, the need to monitor and control web access is becoming a key component of secure content management systems that employ antivirus, web filtering, and messaging security. Important reasons for controlling web content include:

- lost productivity because employees are accessing the web for non-business reasons
- network congestion — when valuable bandwidth is used for non-business purposes, legitimate business applications suffer
- loss or exposure of confidential information through chat sites, non-approved email systems, instant messaging, and peer-to-peer file sharing
- increased exposure to web-based threats as employees surf non-business-related web sites
- legal liability when employees access/download inappropriate and offensive material
- copyright infringement caused by employees downloading and/or distributing copyrighted material.

As the number and severity of threats increase on the World Wide Web, the risk potential increases within a company's network as well. Casual non-business related web surfing has caused many businesses countless hours of legal litigation as hostile environments have been created by employees who download and view offensive content. Web-based attacks and threats are also becoming increasingly sophisticated. Threats and web-based applications that cause additional problems for corporations include:

- spyware/grayware
- phishing
- pharming
- instant messaging
- peer-to-peer file sharing
- streaming media
- blended network attacks.

Spyware, also known as grayware, is a type of computer program that attaches itself to a user's operating system. It does this without the user's consent or knowledge. It usually ends up on a computer because of something the user does such as clicking on a button in a pop-up window. Spyware can track the user's Internet usage, cause unwanted pop-up windows, and even direct the user to a host web site. For further information, visit the [FortiGuard Center](#).

Some of the most common ways of grayware infection include:

- downloading shareware, freeware, or other forms of file-sharing services
- clicking on pop-up advertising
- visiting legitimate web sites infected with grayware.

Phishing is the term used to describe attacks that use web technology to trick users into revealing personal or financial information. Phishing attacks use web sites and email that claim to be from legitimate financial institutions to trick the viewer into believing that they are legitimate. Although phishing is initiated by spam email, getting the user to access the attacker's web site is always the next step.

Pharming is a next generation threat that is designed to identify and extract financial, and other key pieces of information for identity theft. Pharming is much more dangerous than phishing because it is designed to be completely hidden from the end user. Unlike phishing attacks that send out spam email requiring the user to click to a fraudulent URL, pharming attacks require no action from the user outside of their regular web surfing activities. Pharming attacks succeed by redirecting users from legitimate web sites to similar fraudulent web sites that have been created to look and feel like the authentic web site.

Instant messaging presents a number of problems. Instant messaging can be used to infect computers with spyware and viruses. Phishing attacks can be made using instant messaging. There is also a danger that employees may use instant messaging to release sensitive information to an outsider.

Peer-to-peer (P2P) networks are used for file sharing. Such files may contain viruses. Peer-to-peer applications take up valuable network resources and may lower employee productivity but also have legal implications with the downloading of copyrighted or sensitive company material.

Streaming media is a method of delivering multimedia, usually in the form of audio or video to Internet users. Viewing streaming media impacts legitimate business by using valuable bandwidth.

Blended network threats are rising and the sophistication of network threats is increasing with each new attack. Attackers learn from each previous successful attack and enhance and update attack code to become more dangerous and fast spreading. Blended attacks use a combination of methods to spread and cause damage. Using virus or network worm techniques combined with known system vulnerabilities, blended

threats can quickly spread through email, web sites, and Trojan applications. Examples of blended threats include Nimda, Code Red, Slammer, and Blaster. Blended attacks can be designed to perform different types of attacks, which include disrupting network services, destroying or stealing information, and installing stealthy backdoor applications to grant remote access.

Different ways of controlling access

The methods available for monitoring and controlling Internet access range from manual and educational methods to fully automated systems designed to scan, inspect, rate and control web activity.

Common web access control mechanisms include:

- establishing and implementing a well-written usage policy in the organization on proper Internet, email, and computer conduct
- installing monitoring tools that record and report on Internet usage
- implementing policy-based tools that capture, rate, and block URLs.

The final method is the focus of this topic. The following information shows how the filters interact and how to use them to your advantage.

Order of web filtering

The FortiGate unit applies web filters in a specific order:

- 1 URL filter
- 2 FortiGuard Web Filter
- 3 web content filter
- 4 web script filter
- 5 antivirus scanning.

If you have blocked a FortiGuard Web Filter category but want certain users to have access to URLs within that pattern, you can use the *Override* within the FortiGuard Web Filter. This will allow you to specify which users have access to which blocked URLs and how long they have that access. For example, if you want a user to be able to access www.example.com for one hour, you can use the override to set up the exemption. Any user listed in an override must fill out an online authentication form that is presented when they try to access a blocked URL before the FortiGate unit will grant access to it. For more information, see [“FortiGuard Web Filter” on page 1023](#).

Web content filter

You can control web content by blocking access to web pages containing specific words or patterns. This helps to prevent access to pages with questionable material. You can also add words, phrases, patterns, wild cards and Perl regular expressions to match content on web pages. You can add multiple web content filter lists and then select the best web content filter list for each web filter profile.

Enabling web content filtering involves three separate parts of the FortiGate configuration.

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day.
- The web filter profile specifies what sort of web filtering is applied.
- The web content filter list contains blocked and exempt patterns.

The web content filter feature scans the content of every web page that is accepted by a security policy. The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the scores of banned words and phrases in the page. If the sum is higher than a threshold set in the web filter profile, the FortiGate unit blocks the page.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create a web content filter list.
- 2 Add patterns of words, phrases, wildcards, and regular expressions that match the content to be blocked or exempted.

You can add the patterns in any order to the list. You need to add at least one pattern that blocks content.
- 3 In a web filter profile, enable the web content filter and select a web content filter list from the options list.

To complete the configuration, you need to select a security policy or create a new one. Then, in the security policy, enable *UTM* and select the appropriate web filter profile from the list.

Creating a web filter content list

You can create multiple content lists and then select the best one for each web filter profile. Creating your own web content lists can be accomplished only using the CLI.

This example shows how to create a web content list called inappropriate language, with two entries, offensive and rude.

To create a web filter content list

```
config webfilter content
  edit 3
    set name "inappropriate language"
    config entries
      edit offensive
        set action block
        set lang western
        set pattern-type wildcard
        set score 15
        set status enable
      next
      edit rude
        set action block
        set lang western
        set pattern-type wildcard
        set score 5
        set status enable
      end
    end
  end
end
```

See the [CLI Reference](#) for a complete description of all the web filter content list commands and options.

How content is evaluated

Every time the web content filter detects banned content on a web page, it adds the score for that content to the sum of scores for that web page. You set this score when you create a new pattern to block the content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the sum of scores equals or exceeds the threshold score, the web page is blocked. The default score for web content filter is 10 and the default threshold is 10. This means that by default a web page is blocked by a single match. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

Banned words or phrases are evaluated according to the following rules:

- The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.
- The score for any word in a phrase without quotation marks is counted.
- The score for a phrase in quotation marks is counted only if it appears exactly as written.

The following table describes how these rules are applied to the contents of a web page. Consider the following, a web page that contains only this sentence: “The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.”

Table 69: Banned Pattern Rules

Banned pattern	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
word	20	20	20	Appears twice but only counted once. Web page is blocked.
word phrase	20	40	20	Each word appears twice but only counted once giving a total score of 40. Web page is blocked
word sentence	20	20	20	“word” appears twice, “sentence” does not appear, but since any word in a phrase without quotation marks is counted, the score for this pattern is 20. Web page is blocked.

Table 69: Banned Pattern Rules (Continued)

Banned pattern	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
"word sentence"	20	0	20	"This phrase does not appear exactly as written. Web page is allowed.
"word or phrase"	20	20	20	This phrase appears twice but is counted only once. Web page is blocked.

Enabling the web content filter and setting the content threshold

When you enable the web content filter, the web filter will block any web pages when the sum of scores for banned content on that page exceeds the content block threshold. The threshold will be disregarded for any exemptions within the web filter list.

To enable the web content filter and set the content block threshold

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the *Create New* icon on the Edit Web Filter Profile window title bar.
- 3 In the *Name* field, enter the name of the new web filter profile.
- 4 Optionally, you may also enter a comment. The comment can remind you of the details of the sensor.
- 5 Select the *Inspection Method*.
 Proxy-based detection involves buffering the file and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.
 Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.
- 6 Expand the *Advanced Filter* heading.
- 7 Enable *Web Content Filter*.
- 8 Select the required web filter content list from the *Web Content Filter* drop-down list.
- 9 Select *Apply*.

The web filter profile configured with web content filtering is ready to be added to a firewall profile.

URL filter

You can allow or block access to specific URLs by adding them to the URL filter list. You add the URLs by using patterns containing text and regular expressions. The FortiGate unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message instead.



URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to ftp://ftp.example.com. Instead, use firewall policies to deny ftp connections.

When adding a URL to the URL filter list, follow these rules:

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and file name to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls access to the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.



URLs with an action set to exempt or pass are not scanned for viruses. If users on the network download files through the FortiGate unit from a trusted web site, add the URL of this web site to the URL filter list with an action to pass it so the FortiGate unit does not virus scan files downloaded from this URL.

URL filter actions

You can select one of four actions for URL patterns you include in URL filter lists.

Block

Attempts to access any URLs matching the URL pattern are denied. The user will be presented with a replacement message.

Allow

Any attempt to access a URL that matches a URL pattern with an allow action is permitted. The traffic is passed to the remaining antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning.

Allow is the default action. If a URL does not appear in the URL list, it is permitted.

Pass

Traffic to, and reply traffic from, sites matching a URL pattern with a pass action will bypass all antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning.

Make sure you trust the content of any site you pass.

Exempt

Exempt is similar to Pass in that it allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently. In general, if you're not certain that you need to use the Exempt action, use Pass.

HTTP 1.1 connections are persistent unless declared otherwise. This means the connections will remain in place until closed or the connection times out. When a client loads a web page, the client opens a connection to the web server. If the client follows a link to another page on the same site before the connection times out, the same connection is used to request and receive the page data.

When you add a URL pattern to a URL filter list and apply the Exempt action, traffic sent to and replies traffic from sites matching the URL pattern will bypass all antivirus proxy operations, as with the Pass action. The difference is that the connection itself inherits the exemption. This means that all subsequent reuse of the existing connection will also bypass all antivirus proxy operations. When the connection times out, the exemption is cancelled.

For example, consider a URL filter list that includes `example.com/files` configured with the Exempt action. A user opens a web browser and downloads a file from the URL `example.com/sample.zip`. This URL does not match the URL pattern so it is scanned for viruses. The user then downloads `example.com/files/beautiful.exe` and since this URL does match the pattern, the connection itself inherits the exempt action. The user then downloads `example.com/virus.zip`. Although this URL does not match the exempt URL pattern, a previously visited URL did, and since the connection inherited the exempt action and was re-used to download a file, the file is not scanned.

If the user next goes to an entirely different server, like `example.org/photos`, the connection to the current server cannot be reused. A new connection to `example.org` is established. This connection is not exempt. Unless the user goes back to `example.com` before the connection to that server times out, the server will close the connection. If the user returns after the connection is closed, a new connection to `example.com` is created and it is not exempt until the user visits a URL that matches the URL pattern.

Web servers typically have short time-out periods. A browser will download multiple components of a web page as quickly as possible by opening multiple connections. A web page that includes three photos will load more quickly if the browser opens four connections to the server and downloads the page and the three photos at the same time. A short time-out period on the connections will close the connections faster, allowing the server to avoid unnecessarily allocating resources for a long period. The HTTP session time-out is set by the server and will vary with the server software, version, and configuration.

Using the exempt action can have unintended consequences in certain circumstances. You have a web site at `example.com` and since you control the site, you trust the contents and configure `example.com` as exempt. But `example.com` is hosted on a shared server with a dozen other different sites, each with a unique domain name. Because of the shared hosting, they also share the same IP address. If you visit `example.com`, your connection your site becomes exempt from any antivirus proxy operations. Visits to any of the 12 other sites on the same server will reuse the same connection and the data you receive is exempt from scanned.

Use of the exempt action is not suitable for configuration in which connections through the FortiGate unit use an external proxy. For example, you use proxy.example.net for all outgoing web access. Also, as in the first example, URL filter list that includes a URL pattern of example.com/files configured with the Exempt action. Users are protected by the antivirus protection of the FortiGate unit until a user visits a URL that matches the of example.com/files URL pattern. The pattern is configured with the Exempt action so the connection to the server inherits the exemption. With a proxy however, the connection is from the user to the proxy. Therefore, the user is entirely unprotected until the connection times out, no matter what site he visits.

Ensure you are aware of the network topology involving any URLs to which you apply the Exempt action.

Examples using exempt and pass actions

These examples illustrate the differences between the exempt and pass actions.

The URL filter list in use has a single entry: `www.example.com/files/content/`

- With an exempt action, the user downloads the file `www.example.com/files/content/eicar.com`. This URL matches the URL filter list entry so the file is not scanned. Further, the connection itself inherits the exemption. The user next downloads `www.example.com/virus/eicar.com`. Although this does not match the URL filter list entry, the existing connection to example.com will be used so the file is not scanned.
- With a pass action, the user downloads the file `www.example.com/files/content/eicar.com`. This URL matches the URL filter list entry so the file is not scanned. The user next downloads `www.example.com/virus/eicar.com`. This does not match the URL filter entry so it will be scanned. The pass action does not affect the connection and every URL the user accesses is checked against the URL filter list.

The URL filter list in use has a single entry: `www.domain.com/files/content/`. The user's browser is configured to use an external web proxy. All user browsing takes advantage of this proxy.

- With an exempt action, the user downloads `www.example.com/files/content/eicar.com` through the proxy. This matches the URL filter list entry so the file is not scanned. Further, the connection to the proxy inherits the exemption. The user next downloads `www.eicar.org/virus/eicar.com` through the proxy. Although this does not match the URL filter list entry, the existing connection to the proxy will be used so the file is not scanned. In fact, until the user stops browsing long enough for the connection to time out, all the user web traffic is exempt and will not be scanned.
- With a pass action, the user downloads `www.example.com/files/content/eicar.com` through the proxy. This matches the URL filter list entry so the file is not scanned. The user next downloads `www.eicar.org/virus/eicar.com` through the proxy. This does not match the URL filter entry so it will be scanned. The pass action does not affect the connection and every URL the user accesses is checked against the URL filter list.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create a URL filter list.
- 2 Add URLs to the URL filter list.

- 3 Select a web filter profile or create a new one.
- 4 In the web filter profile, create a URL List filter.

To complete the configuration, you need to select a security policy or create a new one. Then, in the security policy, enable *UTM* and select the appropriate web filter profile from the list.

Creating a URL filter list

To create a URL Filter list

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select *Create New*.
- 3 Enter a *Name* for the new URL filter list.
- 4 Enter optional comments to describe it.
- 5 Select *OK*.

Configuring a URL filter list

Each URL filter list can have up to 5000 entries. For this example, the URL `www.example*.com` will be used. You configure the list by adding one or more URLs to it.

To add a URL to a URL filter list

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select an existing list and choose *Edit*.
- 3 Select *Create New*.
- 4 Enter the URL, without the “http”, for example: `www.example*.com`.
- 5 Select a *Type*: *Simple*, *Wildcard* or *Regular Expression*.
In this example, select *Wildcard*.
- 6 Select the *Action* to take against matching URLs: *Exempt*, *Block*, *Allow*, or *Pass*.
- 7 Select *Enable*.
- 8 Select *OK*.

SafeSearch

SafeSearch is a feature of popular search sites that prevents explicit web sites and images from appearing in search results. Although SafeSearch is a useful tool, especially in educational environments, the resourceful user may be able to simply turn it off. Enabling SafeSearch for the supported search sites enforces its use by rewriting the search URL to include the code to indicate the use of the SafeSearch feature.

Three search sites are supported:

Google	Enforce the strict filtering level of safe search protection for Google search results by adding <code>&safe=on</code> to search URL requests. Strict filtering removes both explicit text and explicit images from the search results.
Yahoo!	Enforce the strict filtering level of safe search protection for Yahoo! search results by adding <code>&vm=r</code> to search URL requests. Strict filtering removed adult web, video, and images from search results.
Bing	Enforce the strict filtering level of safe search protection for Bing search results by adding <code>adlt=strict</code> to search URL requests. Strict filtering removes explicit text, images, and video from the search results.

Enabling SafeSearch – CLI

```
config webfilter profile
  edit default
    config web
      set safe-search bing google yahoo
    end
  end
```

This enforces the use of SafeSearch in traffic controlled by the firewall policies using the web filter you configure.

Advanced web filter configuration

The *Advanced Filter* section of the web filter profile provides a number of advanced filtering options. The FortiGuard Web Filter options in the advance filter section are detailed in the FortiGuard Web Filter section, in [“Advanced FortiGuard Web Filter configuration” on page 1029](#).

ActiveX filter

Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.

Cookie filter

Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.

Java applet filter

Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.

Web resume download block

Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.

This prevents the unintentional download of viruses hidden in fragmented files.

Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.

Block Invalid URLs

Select to block web sites when their SSL certificate CN field does not contain a valid domain name.

FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:

- If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.
- If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.

HTTP POST action

Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.

The available actions include:

Normal	Allow use of the HTTP POST command as normal.
Comfort	Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic. The client comforting settings used are those defined in the protocol options profile selected in the security policy. For more information, see “Configuring client comforting” on page 904 .
Block	Block the HTTP POST command. This will limit users from sending information and files to web sites. When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.

Web filtering example

Web filtering is particularly important for protecting school-aged children. There are legal issues associated with improper web filtering as well as a moral responsibility not to allow children to view inappropriate material. The key is to design a web filtering system in such a way that students and staff do not fall under the same web filter profile in the FortiGate configuration. This is important because the staff may need to access websites that are off-limits to the students.

School district

The background for this scenario is a school district with more than 2300 students and 500 faculty and staff in a preschool, three elementary schools, a middle school, a high school, and a continuing education center. Each elementary school has a computer lab and the high school has three computer labs with connections to the Internet. Such easy access to the Internet ensures that every student touches a computer every day.

With such a diverse group of Internet users, it was not possible for the school district to set different Internet access levels. This meant that faculty and staff were unable to view websites that the school district had blocked. Another issue was the students' use of proxy sites to circumvent the previous web filtering system. A proxy server acts as a go-between for users seeking to view web pages from another server. If the proxy server has not been blocked by the school district, the students can access the blocked website.

When determining what websites are appropriate for each school, the district examined a number of factors, such as community standards and different needs of each school based on the age of the students.

The district decided to configure the FortiGate web filtering options to block content of an inappropriate nature and to allow each individual school to modify the options to suit the age of the students. This way, each individual school was able to add or remove blocked sites almost immediately and have greater control over their students' Internet usage.

In this simplified example of the scenario, the district wants to block any websites with the word **example** on them, as well as the website `www.example.com`. The first task is to create web content filter lists for the students and the teachers.

To create a web content filter list for the students

```
config webfilter content
  edit 5
    set name "Student Web Content List"
    config entries
      edit example
        set action block
        set status enable
      end
    end
  end
```

It might be more efficient if the Teacher Web Content List included the same blocked content as the student list. From time to time a teacher might have to view a blocked page. It would then be a matter of changing the *Action* from *Block* to *Allow* as the situation required.

To create a web content filter list for the teachers

```
config webfilter content
  edit 5
    set name "Teacher Web Content List"
    config entries
      edit example
        set action exempt
        set status enable
      end
    end
  end
```

URL filter lists with filters to block unwanted web sites must be created for the students and teachers. For this example the URL `www.example.com` will be used.

To create a URL filter for the students

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select *Create New*.
- 3 Enter `Student URL List` as the URL filter *Name*.
- 4 Enter optional comments to describe the contents of the list.
- 5 Select *OK*.
The URL filter for the students has been created. Now it must be configured.
- 6 Select *Create New*.
- 7 Enter `example.com` in the URL field.
- 8 Select *Simple* from the *Type* list.
- 9 Select *Block* from the *Action* list.
- 10 Select *Enable*.
- 11 Select *OK*.
- 12 Select *OK*.

The teachers should be able to view the students' blocked content, however, so an addition URL filter is needed.

To create a URL filter for the teachers

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select *Create New*.
- 3 Enter `Teacher URL List` as the URL filter *Name*.
- 4 Enter optional comments to describe the list.
- 5 Select *OK*.
The URL filter for the students has been created. Now it must be configured.
- 6 Select *Create New*.
- 7 Enter `www.example.com` in the *URL* field.
- 8 Select *Simple* from the *Type* list.
- 9 Select *Exempt* from the *Action* list.
- 10 Select *Enable*.
- 11 Select *OK*.
- 12 Select *OK*.

A web filter profile must be created for the students and the teachers.

To create a web filter profile for the students

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the *Create New* icon in the Edit Web Filter window title bar.
- 3 Enter `Students` as the *Profile Name*.
- 4 Enter optional comments to identify the profile.
- 5 Expand the *Advanced Filter* heading.
- 6 Enable *Web Content Filter*.
- 7 Select *Student Web Content List* from the *Web Content Filter* drop-down list.
- 8 Enable *Web URL Filter*.

9 Select *Student URL List* from the *Web URL Filter* drop-down list.

10 Enable *Web Resume Download Block*.

Selecting this setting will block downloading parts of a file that have already been downloaded and prevent the unintentional download of virus files hidden in fragmented files. Note that some types of files, such as PDFs, are fragmented to increase download speed, and that selecting this option can cause download interruptions with these types.

11 Select *OK*.

To create a security policy for the students

1 Go to *Policy > Policy > Policy*.

2 Select *Create New*.

3 Enable *UTM*.

4 Select *Enable Web Filter*.

5 Select *Students* from the web filter drop-down list.

6 Enter optional comments.

7 Select *OK*.

To create a web filter profile for the teachers

1 Go to *UTM Profiles > Web Filter > Profile*.

2 Select the *Create New* icon in the Edit Web Filter window title bar.

3 Enter *Teachers* as the *Profile Name*.

4 Enter optional comments to identify the profile.

5 Expand the *Advanced Filter* heading.

6 Enable *Web Content Filter*.

7 Select *Teacher Web Content List* from the *Web Content Filter* drop-down list.

8 Enable *Web URL Filter*.

9 Select *Teacher URL List* from the *Web URL Filter* drop-down list.

10 Enable *Web Resume Download Block*.

11 Select *OK*.

To create a security policy for Teachers

1 Go to *Policy > Policy > Policy*.

2 Select *Create New*.

3 Enable *UTM*.

4 Select *Enable Web Filter*.

5 Select *Teachers* from the web filter drop-down list.

6 Enter optional comments.

7 Select *OK*.

Web Filter interface reference

The following explains the web filtering options in the Web Filtering menu. If your unit supports SSL content scanning and inspection you can also configure web filtering for HTTPS traffic.

If you want to configure advanced settings, such as web content filter, you must configure them within the CLI. Advanced settings also includes overrides.

This topic includes the following:

- [Profile](#)
- [Browser cookie-based FortiGuard Web Filtering overrides](#)
- [URL Filter](#)
- [Local Ratings](#)

Profile

The Profile menu allows you to configure a web filter profile to apply to a firewall policy. A profile is specific information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

Web profile configuration settings

The following are web filter profile configuration settings in *UTM Profiles > Web Filter > Profile*. If you want to configure advanced settings, such as FortiGuard web filtering overrides, you must configure these settings within the CLI.

Profile page Lists each web filter profile that you created. On this page, you can edit, delete or create a new web filter profile. You are redirected to this page when you select <i>View List</i> on the Edit Web Filter Profile page. Note: Web filtering overrides are profile-based, allowing a rule to be created that changes the web filter profile that applies to a user. An override link appears in all related blocked pages. This is available only in the CLI.	
Create New	Creates a new web filter profile. When you select <i>Create New</i> , you are automatically redirected to the New Web Filter Profile page.
Edit	Modifies settings within a web filter profile. When you select <i>Edit</i> , you are automatically redirected to the Edit Web Filter Profile page.
Delete	Removes a web filter profile from within the list on the Profile page. To remove multiple web filter profiles from within the list, on the Profile page, in each of the rows of the file filter lists you want removed, select the check box and then select <i>Delete</i> . To remove all web filter profiles from the list, on the Profile page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the web filter profile.
Comments	A description given to the web filter profile. This is an optional setting.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
<p>New Web Filter Profile page</p> <p>Provides settings for configuring a web filter profile. Advanced features, such as web content filtering and FortiGuard web filtering, is configured in the CLI.</p> <p>This page appears when you select <i>Create New</i> on the Edit Web Filter Profile page. If you are on the Profile page, and you select <i>Create New</i>, you will be redirected to the New Web Filter Profile page.</p> <p>Note: Logging is enabled in the CLI.</p>	
Name	<p>Enter a name for the web filter profile.</p> <p>If you want to edit the name at any time, select the profile and enter a new name in the <i>Name</i> field. Select <i>Apply</i> to save the change.</p>
Comments	<p>Enter a description for the web filter profile. This is optional.</p> <p>If you want to edit the description at any time, select the profile and enter the new description in the <i>Comments</i> field. Select <i>Apply</i> to save the change.</p>
Inspection mode	<p>Select to enable either flow-based web filtering or proxy-based.</p> <p>Flow-based web filtering is a non-proxy solution, which provides high concurrent session, high session rate, and low-latency web filtering service.</p>

FortiGuard Categories	A list of FortiGuard category groups and categories that are used to rate web sites. Selecting a category group will automatically select all of the categories within the group. For example, if you select Security Risk, you can see that all of the categories within are selected if you expand the group. You can however, select or deselect categories within groups as required.
Show	Select an action to view all of the categories that are currently configured with the selected action.
Change Action for Selected Categories to	Select an action, and all of the selected categories will have the selected action applied. Selected category groups will have the action applied to all categories within the group.
Quota on Categories	<p>Users can have their web browsing time limited by category through the use of quotas. Quotas can be applied only to categories that are configured with the Monitor action.</p> <p>If you create a quota for a single category, every authenticated user subject to the security policy in which the web filter profile is applied is limited in browsing web sites in the category to the duration you specify. If you create a single quota that includes multiple categories, the quota will apply to the categories as a whole.</p> <p>Quotas are ignored for unauthenticated users. To enforce quotas, configure the security policy to require authentication.</p>
Enable Safe Search (Support Search Engines: Google, Yahoo and Bing)	When enabled, the supported search engines exclude offensive material from search results.
HTTPS Scanning	<p>Available only on models that support HTTPS.</p> <p>Select to have all of the web filtering specified in the web filter profile to HTTPS traffic as well as HTTP traffic.</p>
Advanced Filter	Expand this heading for advanced web filtering options.
Web URL Filter	Enable to block access to URLs listed in the selected URL list.
Web Resume Download Block	<p>Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.</p> <p>This prevents the unintentional download of viruses hidden in fragmented files.</p> <p>Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.</p>

Block Invalid URLs	<p>Select to block web sites when their SSL certificate CN field does not contain a valid domain name.</p> <p>FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:</p> <ul style="list-style-type: none"> • If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name. • If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.
HTTP POST Action	<p>Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.</p> <p>The available actions include:</p> <ul style="list-style-type: none"> • Normal: Allow use of the HTTP POST command as normal. • Comfort: Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic. The client comforting settings used are those defined in the protocol options profile selected in the security policy. • Block: Block the HTTP POST command. This will limit users from sending information and files to web sites. When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.
Remove Java Applet Filter	Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.
Remove ActiveX Filter	Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.
Remove Cookie Filter	Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.
Search Engine Keyword Filter	Enter the keywords that you want to monitor when users enter those same or similar keywords during a search within the supported search engines.
Web Content Filter	Enable to block access to web pages that include the words included in the selected web content filter list.
Provide Details for Blocked HTTP 4xx and 5xx Errors	Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.

Rate Images by URL (Blocked images will be replaced with blanks)	<p>Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.</p> <p>Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.</p>
Allow Websites When a Rating Error Occurs	<p>Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.</p> <p>If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.</p>
Strict Blocking	<p>This setting determines when the FortiGate unit blocks a site. Enable strict blocking to deny access to a site if any category or classification assigned to the site is set to Block. Disable strict blocking to deny access to a site only if all categories and classifications assigned to the site are set to Block.</p> <p>All rated URLs are assigned one or more categories. URLs may also be assigned a classification. If Rate URLs by domain and IP address is enabled, the site URL and IP address each carry separately assigned categories and classifications. Depending on the FortiGuard rating and the FortiGate configuration, a site could be assigned to at least two categories and up to two classifications.</p>
Rate URLs by Domain and IP Address	<p>Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.</p> <p>FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.</p>
Block HTTP Redirects by Rating	<p>Enable to block HTTP redirects.</p> <p>Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.</p> <p>This option is not supported for HTTPS.</p>

Browser cookie-based FortiGuard Web Filtering overrides

By using browser cookie-based FortiGuard Web Filtering overrides, you can identify users according to their web browser cookie instead of their IP address and then to use this identification to apply FortiGuard Web Filtering overrides to individual users.

This feature uses the dynamic profile feature to assign a web filter profile that includes FortiGuard Web Filtering to a communication session. Just like normal FortiGuard Web Filtering overrides, when FortiGuard Web Filtering blocks access to a web page, the user can authenticate to override FortiGuard Web Filtering. However, with Browser cookie-based overrides enabled, the browser cookie is used to identify the user instead of the user's IP address.

To allow browser based FortiGuard Web Filtering overrides in a user group, go to *User > User Group*, edit a firewall or directory service user group. Select *Allow to create FortiGuard Web Filtering overrides* and make sure *Browser (Cookie) Override* is set to *Allow*.

You can also go to *UTM Profiles > Web Filter > Configuration* and configure the following browser cookie-based override settings.



Additional browser cookie-based configuration settings are available from the CLI using the `config webfilter cookie-ovrd` command.

Cookie (Browser Based) Override Configuration page

Provides settings for configuring the browser cookie-based override.

Override Validation Hostname	Enter the override validation hostname in the field.
Override Validation Port	Enter the port number in the field.

How browser cookie-based FortiGuard Web Filtering overrides work

The following steps occur when a user's session that can use browser cookie-based FortiGuard Web Filtering overrides is received:

- 1 The Dynamic Profile applies a profile to the user session in the normal way.
- 2 The user issues a request to a remote site blocked by FortiGuard Web Filtering.
For example, `http://www.example.com`.
- 3 FortiGuard Web Filtering blocks the page and provides an override link.
- 4 The user selects the override option and successfully authenticates.
- 5 The unit sends a cookie to the remote site that seems to come from the *Override Validation Hostname*.
- 6 The unit creates a second cookie to the user's browser for the domain of the remote site.

For example, the domain could be `example.com`.

The rest of the communication between the user and the remote site is authorized with the unit by these cookies

URL Filter

Allow or block access to specific URLs by adding them to the URL filter list. Add patterns using text and regular expressions (or wildcard characters) to allow or block URLs. The unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message.

You can add multiple URL filter lists and then select the best URL filter list for each profile.

You can add the following to block or exempt URLs:

- complete URLs
- IP addresses
- partial URLs to allow or block all sub-domains

Each URL filter list can have up to 5000 entries.

URL filter configuration settings

The following are URL filter configuration settings in *UTM Profiles > Web Filter > URL Filter*.



URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to `ftp://ftp.example.com`. Instead, use firewall policies to deny FTP connections.

URL Filter page	
Lists each URL filter that you created. On this page, you can edit, delete or create a new URL filter.	
Create New	Creates a new URL filter list. When you select <i>Create New</i> , you are automatically redirected to the New List page. This page provides a name field and comment field. You must enter a name to go to the URL Filter Settings page.
Edit	Modifies settings within a URL filter list. When you select <i>Edit</i> , you are automatically redirected to the URL Filter Settings page.
Delete	<p>Removes the URL filter list from the list on the URL Filter page. The <i>Delete</i> icon is only available if the URL filter list is not selected in any profiles.</p> <p>To remove multiple URL filter list from within the list, on the URL Filter page, in each of the rows of the file filter lists you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all URL filter list from the list, on the URL Filter page, select the check box in the check box column and then select <i>Delete</i>.</p>
Name	The available URL filter lists.
# Entries	The number of URL patterns in each URL filter list.
MMS Profiles (FortiOS Carrier only)	The name of the MMS profile
Comments	Optional description of each URL filter list.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
URL Filter Settings page Provides settings for configuring URLs that make up the URL filter, and also lists the URLs that you created. You are automatically redirected to this page from the New List Page. If you are editing a URL filter, you are automatically redirected to this page.	
Name	If you are editing an existing URL filter setting and want to change the name, enter a new name in this field. You must select <i>OK</i> to save the change.
Comments	If you are editing an existing URL filter setting and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save these changes.
Create New	Adds a URL address and filter settings to the list. When you select <i>Create New</i> , you are automatically redirected to the New URL Filter list.
Edit	Modifies the settings within a URL filter.
Delete	<p>Removes an entry from the list.</p> <p>To remove multiple URL filters from within the list, on the URL Filter Settings page, in each of the rows of the filters you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all URL filters from the list, on the URL Filter Settings page, select the check box in the check box column and then select <i>Delete</i>.</p>
Enable	Enables a filter in the list.
Disable	Disables a filter in the list.
Move To	<p>Moves the URL to any position in the list. When you select <i>Move To</i>, the Move URL Filter window appears.</p> <p>To move a URL, select the new position <i>Before</i> or <i>After</i>, which will place the current URL entry before or after the entry you enter in the (<i>URL</i>) field. For example, 1example.com is being moved after 3example.com, so 3example.com is entered in the (<i>URL</i>) field.</p>

Remove All Entries	Removes all filter entries within the list on the URL Filter Settings page.
Enable	Indicates whether the URL is enable or disabled. A green check mark indicates that the URL is enabled; a gray check mark indicates that the URL is disabled.
URL	The URL address.
Action	The type of action the unit will take when there is a match.
Type	The type of URL. For example, the type of URL is <i>Regex</i> .
New URL Filter page Provides settings for configuring a URL to add to the filter list.	
URL	Enter the URL. Do not include http://. For details about URL formats, see “URL formats” on page 1020 .
Type	Select a type from the drop-down list: <i>Simple</i> , <i>Regex</i> (regular expression), or <i>Wildcard</i> .
Action	<p>Select an action the unit will take.</p> <ul style="list-style-type: none"> • <i>Allow</i> – any attempt to access a URL that matches a URL pattern with an allow action is permitted. • <i>Exempt</i> – similar to <i>Pass</i> in that it allows trusted traffic to bypass the antivirus proxy operations, but it functions slightly differently; ensure you are aware of the network topology involving URLs that you applied the Exemption action. Additional information about the Exempt action is found in the UTM chapter of the FortiOS Handbook. • <i>Block</i> – attempts to access any URLs matching the URL pattern are denied; user is presented with a replacement message. • <i>Pass</i> – traffic to, and replay traffic from sites that match a URL pattern with a pass action will bypass all antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters, and antivirus scanning. Make sure you trust the content of any site you pass, otherwise there may be a security risk.
Enable	Select to enable the URL. By default, the URL is enabled.



Type a top-level domain suffix (for example, “com” without the leading period) to block access to all URLs with this suffix.

URL formats

When adding a URL to the URL filter list, follow these rules:

How URL formats are detected when using HTTPS

If your unit does not support SSL content scanning and inspection or if you have selected the *URL filtering* option in web content profile for *HTTPS content filtering mode* under *Protocol Recognition*, filter HTTPS traffic by entering a top level domain name, for example, `www.example.com`. HTTPS URL filtering of encrypted sessions works by extracting the CN from the server certificate during the SSL negotiation. Since the CN only contains the domain name of the site being accessed, web filtering of encrypted HTTPS sessions can only filter by domain names.

If your unit supports SSL content scanning and inspection and if you have selected Deep Scan, you can filter HTTPS traffic in the same way as HTTP traffic.

How URL formats are detected when using HTTP

URLs with an action set to exempt are not scanned for viruses. If users on the network download files through the unit from trusted web site, add the URL of this web site to the URL filter list with an action set to exempt so the unit does not virus scan files downloaded from this URL.

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and filename to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns created using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.

FortiGate URL filtering supports standard regular expressions.



If virtual domains are enabled on the unit, web filtering features are configured globally. To access these features, select *Global Configuration* on the main menu.

Local Ratings

You can configure user-defined categories and then specify the URLs that belong to the category. This allows users to block groups of web sites on a per profile basis. The ratings are included in the global URL list with associated categories and compared in the same way the URL block list is processed.

Local ratings configuration settings

The following are local ratings configuration settings in *UTM Profiles > Web Filter > Local Ratings*.

Local Ratings page	
Lists each individual local rating that you created. On this page, you can edit, delete or create a new local rating. You can also disable or enable a local rating, as well as remove all local ratings from the page.	
Create New	Creates a new local rating. When you select <i>Create New</i> , you are automatically redirected to the New Local Rating page.
Edit	Modifies settings within a local rating. When you select <i>Edit</i> , you are automatically redirected to the Edit Local Rating page.
Delete	Select to remove a local rating from the list.
Enable	Enables a local rating within the list on the Local Ratings page.

Disable	Disables a local rating within the list on the Local Ratings page.
Remove All Entries	Removes all local ratings within the list on the Local Ratings page.
Search	Enter a word or name to search for the local rating within the list. Select Go to start the search.
#	The number identifying the order of the item in the list.
Enable	A green checkmark appears if the local rating is enabled. A gray x appears if the local rating is disabled.
URL	The URL address of the local rating.
Category	The category that was selected for the local rating.
Page controls	Use to navigate through the list of local ratings.
<i>New Local Rating page</i>	
Provides settings for configuring the URL address that belongs to a category and classification rating. When editing a local rating, you are automatically redirected to the Edit Local Rating page which contains the same settings.	
URL	Enter the URL address.
Category Rating	Select the ratings for the URL.



FortiGuard Web Filter

This section describes FortiGuard Web Filter for HTTP and HTTPS traffic.

FortiGuard Web Filter is a managed web filtering solution available by subscription from Fortinet. FortiGuard Web Filter enhances the web filtering features supplied with your FortiGate unit by sorting billions of web pages into a wide range of categories users can allow or block. The FortiGate unit accesses the nearest FortiGuard Web Filter Service Point to determine the category of a requested web page, and then applies the security policy configured for that user or interface.

FortiGuard Web Filter includes over 45 million individual ratings of web sites that apply to more than two billion pages. Pages are sorted and rated into several dozen categories administrators can allow or block. Categories may be added or updated as the Internet evolves. To make configuration simpler, you can also choose to allow or block entire groups of categories. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

FortiGuard Web Filter ratings are performed by a combination of proprietary methods including text analysis, exploitation of the web structure, and human raters. Users can notify the FortiGuard Web Filter Service Points if they feel a web page is not categorized correctly, so that the service can update the categories in a timely fashion.

The following topics are discussed in this section:

- [Before you begin](#)
- [FortiGuard Web Filter and your FortiGate unit](#)
- [Enable FortiGuard Web Filter](#)
- [Advanced FortiGuard Web Filter configuration](#)
- [Add or change FortiGuard Web Filter ratings](#)
- [Create FortiGuard Web Filter overrides](#)
- [Customize categories and ratings](#)
- [FortiGuard Web Filter examples](#)

Before you begin

Before you follow the instructions in this section, you should have a FortiGuard Web Filter subscription and your FortiGate unit should be properly configured to communicate with the FortiGuard servers. For more information about FortiGuard services, see the [FortiGuard Center](#) web page. You should also have a look at [“Web filter concepts”](#) on page 997.

FortiGuard Web Filter and your FortiGate unit

When FortiGuard Web Filter is enabled in a web filter profile, the setting is applied to all firewall policies that use this profile. When a request for a web page appears in traffic controlled by one of these firewall policies, the URL is sent to the nearest FortiGuard server. The URL category is returned. If the category is blocked, the FortiGate unit provides a replacement message in place of the requested page. If the category is not blocked, the page request is sent to the requested URL as normal.


Order of web filtering

The FortiGate unit applies web filters in a specific order:

- 1 URL filter
- 2 FortiGuard Web Filter
- 3 web content filter
- 4 web script filter
- 5 antivirus scanning.

The flowchart in [Figure 89 on page 1025](#) shows the steps involved in FortiGuard Web Filtering. Most features are included but some of the advanced options, including overrides, are not. The features appearing in the flowchart are described in this section.

Figure 89: FortiGuard Web Filter sequence of events



Enable FortiGuard Web Filter

FortiGuard Web Filter is enabled and configured within web filter profiles. Overrides, local categories, and local ratings are configured in *UTM Profiles > Web Filter*.

General configuration steps

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the *Edit* icon of the web filter profile in which you want to enable FortiGuard Web Filter, or select *Create New* to add a new web filter profile.
- 3 Create a category filter in the profile.
- 4 The categories allow you to block or allow access to general or more specific web site categories. Configure access as required.
- 5 Save the filter and web filter profile.
- 6 To complete the configuration, you need to select the security policy controlling the network traffic you want to restrict. Then, in the security policy, enable *UTM* and select *Enable Web Filter* and select the appropriate web filter profile from the list.

Configuring FortiGuard Web Filter settings

FortiGuard Web Filter includes a number of settings that allow you to determine various aspects of the filtering behavior.

To configure FortiGuard Web Filter settings

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
- 3 The category groups are listed in a table. You can expand each category group to view and configure every category within the groups. If you change the setting of a category group, all categories within the group inherit the change.
- 4 Select the category groups and categories to which you want to apply an action.
- 5 Select an action from the *Change Action for Selected Categories* drop-down list immediately below the category table. Five actions are available:
 - *Allow* permits access to the sites within the category.
 - *Monitor* permits and logs access to sites in the category. You may also enable user quotas when enabling the monitor action.
 - *Warning* presents the user with a message, allowing them to continue if they choose.
 - *Authenticate* requires a user authenticate with the FortiGate unit before being allowed access to the category or category group.
 - *Block* prevents access to sites within the category. Users attempting to access a blocked site will receive a replacement message explaining that access to the site is blocked.
- 6 Select *OK*.

Configuring FortiGuard Web Filter usage quotas

In addition to using category and classification blocks and overrides to limit user access to URLs, you can set a daily timed access quota by category, category group, or classification. Quotas allow access for a specified length of time, calculated separately for each user. Quotas are reset every day at midnight.

Users must authenticate with the FortiGate unit. The quota is applied to each user individually so the FortiGate must be able to identify each user. One way to do this is to configure a security policy using the identity based policy feature. Apply the web filter profile in which you have configured FortiGuard Web Filter and FortiGuard Web Filter quotas to such a security policy.



The use of FortiGuard Web Filter quotas requires that users authenticate to gain web access. The quotas are ignored if applied to a security policy in which user authentication is not required.

When a user first attempts to access a URL, they're prompted to authenticate with the FortiGate unit. When they provide their username and password, the FortiGate unit recognizes them, determines their quota allowances, and monitors their web use. The category and classification of each page they visit is checked and FortiGate unit adjusts the user's remaining available quota for the category or classification.



Editing the web filter profile resets the quota timers for all users.

To configure the FortiGuard Web Filter categories

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the web filter profile in which you want to enable FortiGuard Web Filter from the drop down list in the Edit Web Filter Profile window title bar, or select *Create New* to add a new web filter profile.
- 3 Select *Create New*.
- 4 Select a *Filter Type* of *Category*.
- 5 Select the required category groups. You may also expand the category groups to select individual categories.
- 6 Select the *Monitor* action.
- 7 Enable *Enforce Quota* to activate the quota for the selected categories and category groups.
- 8 Select *Hours*, *Minutes*, or *Seconds* and enter the number of hours, minutes, or seconds. This is the daily quota allowance for each user.
- 9 Select *OK*.
- 10 Select *Apply*.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.

Quota hierarchy

You can apply quotas to categories and category groups. Only one quota per user can be active at any one time. The one used depends on how you configure the FortiGuard Web Filter.

When a user visits a URL, the FortiGate unit queries the FortiGuard servers for the category of the URL. From highest to lowest, the relative priority of the quotas are:

- 1 Category
- 2 Category group

So for example, the *General Interest - Business* category group contains the *Information and Computer Security* category. When a user visits a page in the *Information and Computer Security* category, the FortiGate unit will check for quotas in sequence:

- Is there is a quota set for the *Information and Computer Security* category? If there is, the category quota timer is started and the user is allowed access to the URL. If no time remains in the category quota, the URL is blocked and the user cannot access it for the remainder of the day.
- If no quota is set for the category, is there a quota set for the *General Interest - Business* category group? If there is, the category group quota timer is started and the user is allowed access to the URL. If no time remains in the category group quota, the URL is blocked and the user cannot access it for the remainder of the day.
- If there is no category group quota, the user is allowed to access the URL. Getting to this point means there are no quotas set for the page. The FortiGate unit will stop any running quota timer because the current URL has no quota.

Only one quota timer can be running at any one time for a single user. Whenever a quota timer is started or a page is blocked, the timer running because of the previous URL access, is stopped. Similarly, a URL with no quotas will stop a quota timer still running because of the URL the user previously accessed.

Quota exempt

The quota checking sequence occurs for every URL the user accesses. This is true for every web page, and every element of the web page that is loaded. For example, if a user loads a web page, the quota is checked for the web page as soon as it is loaded. If there is a photo on the page, it is also checked and the quota is adjusted accordingly.

This can cause unexpected behavior. For example, if the web page a user loads is in the *Information and Computer Security* category and it has a quota, the quota timer is started. The web page includes a number of graphics, so as these are loaded, each is checked and the appropriate quota is started. If they all share the same category rating, which they often will, there is no problem. However, if the last graphic or page element loaded comes from another site, the quota may not work as you expect. If the last graphic is an ad, loaded from a site categorized as *Advertising*, the *Information and Computer Security* category quota timer will stop almost as soon as it is started because the FortiGate unit sees the ad URL and finds that it belongs to the *Advertising* category. If *Advertising* has a quota, its timer will start. If it is blocked or allowed, the *Information and Computer Security* category quota timer is stopped and the user can view the page without using the quota set to limit the *Information and Computer Security* category.

To solve this problem, you can configure a categories and category groups as exempt. This effectively allows the quota system to ignore it entirely. Any quota timer running when an exempt URL is encountered continues to run. An exempt category or category group, or classification can not have a quota. This may sound the same as simply disabling the quota and setting the FortiGuard Web Filter action to allow, but there is a difference. This difference is that the allow and block actions stop an already running quota timer, while the exempt action does not.

The exempt action is generally used for commonly accessed web pages that load elements from other sites that have different category ratings. Pages that load ads from advertising sites are the most common example.

To set a category or category group, see the [CLI Reference](#) for the `exempt-quota` webfilter command.

Checking quota usage

With quotas enabled, the FortiGate unit keeps track of quota usage for each user in each web filter profile. You can check the amount of quota usage for each user and their remaining time for each individual quota on the FortiGuard Quota page.

To view FortiGuard Web Filter quota usage

- 1 Go to *UTM Profiles > Monitor > FortiGuard Quota*.
- 2 The table shows the users who have used some or all of their quota allowance. The total time used is listed by web filter profile for each user.
- 3 Select the *View* icon in any row to view the remaining quota for each category, category group, and classification. A category, category group, or classification displayed in bold type indicates the quota currently in use.

Quotas are reset every day at midnight.

Advanced FortiGuard Web Filter configuration

The *Advanced Filter* section of the web filter profile provides a number of advanced filter options. The web filter options in the advance filter section unrelated to FortiGuard Web Filter are detailed in the web filter section, in [“Advanced web filter configuration” on page 1007](#).

Provide Details for Blocked HTTP 4xx and 5xx Errors

Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to circumvent web filtering.

Rate Images by URL (blocked images will be replaced with blanks)

Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.

Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.

Allow Websites When a Rating Error Occurs

Enable to allow access to web pages that return a rating error from the FortiGuard Web Filter service.

If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines what access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.

Strict Blocking

This setting determines when the FortiGate unit blocks a site. Enable strict blocking to deny access to a site if any category or classification assigned to the site is set to *Block*. Disable strict blocking to deny access to a site only if all categories and classifications assigned to the site are set to *Block*.

All rated URLs are assigned one or more categories. URLs may also be assigned a classification. If *Rate URLs by domain and IP address* is enabled, the site URL and IP address each carry separately assigned categories and classifications. Depending on the FortiGuard rating and the FortiGate configuration, a site could be assigned to at least two categories and up to two classifications.

Rate URLs by Domain and IP Address

Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.



FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.

Block HTTP Redirects by Rating

Enable to block HTTP redirects.

Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.

This option is not supported for HTTPS.

Daily log of remaining quota

Enable to log daily quota use.

As part of the quota reset at midnight, the FortiGate unit will record a log entry for every quota each user consumed during the day. These log entries are labeled with the sub-type `ftgd_quota`. Each entry includes the VDOM, user name, web filter profile name, category description, quota used (in seconds), and quota (in seconds). You can use log filtering to quickly limit the displayed entries to those you want, and generate reports from the logs.

Add or change FortiGuard Web Filter ratings

The FortiGuard Center web site allows you to check the current category assigned to any URL.

To check the category assigned to a URL

- 1 Using your web browser, go to the FortiGuard Center Web Filter URL Lookup & Submission page at <http://www.fortiguard.com/webfiltering/webfiltering.html>.
- 2 Enter the URL as directed.
- 3 Select *Search*.
- 4 If the URL has been rated by the FortiGuard web filter team, the category is displayed.

If a URL has not been rated, or you believe it is incorrectly rated, you can suggest the appropriate category and classification.

To add or change the category for a URL

- 1 Check the category assigned to the URL as described in the previous procedure.
- 2 Below the rating, select *Check to submit the URL*.

- 3 Enter your name, company, and email address.
- 4 Optionally, you may enter a comment.
- 5 Select the most appropriate category and classification for the URL.
- 6 Select *Submit* to send your submission to the FortiGuard web filter team.

Create FortiGuard Web Filter overrides

You can configure FortiGuard Web Filter to allow or deny access to web sites by category and classification. You may want to block a category but allow your users temporary access to one site within the blocked category. You may need to allow only some users to temporarily access one site within a blocked category. Do these things by using administrative and user overrides.

Understanding administrative and user overrides

The administrative overrides are backed up with the main configuration. The administrative overrides are not deleted when they expire and you can reuse them by extending their expiry dates. You can create administrative overrides either through the CLI or the web-based manager.

The user overrides are not backed up as part of the main configuration. These overrides are automatically deleted when they expire. You can only view and delete the user override entries. Users create user overrides using the authentication form opened from the block page when they attempt to access a blocked site, if override is enabled.

Customize categories and ratings

The FortiGuard Web Filter rating categories are general enough that virtually any web site can be accurately categorized in one of them. However, the rigid structure of the categories can create complications. For example, your company uses a web-based email provider. If you select the Web-based Email category, all sites categorized as web-based email providers, including the one your company uses, are blocked.

Local categories and local ratings allow you to assign sites to any category you choose. You can even create new categories. These settings apply only to your FortiGate unit. The changes you make are not sent to the FortiGuard Web Filter Service.

Creating local categories

Categories are labels that describe web site content. Creating your own category allows you to customize how the FortiGuard Web Filter service works.

Local categories appear in the web filter profile, under the FortiGuard Web Filter category list, in the *Local Categories* group. Local categories are empty when created. To populate local categories with web sites, see [“Customizing site ratings” on page 1032](#).

To create a local category

```
config webfilter ftgd-local-cat
  edit "My local category"
end
```

The new local category is added to the list, but will remain empty until you add a web site to it.

Customizing site ratings

You may find it convenient to change the rating of a site. For example, if you want to block all the sites in a category except one, you can move the one site to a different category.

To customize a site rating

- 1 Go to *UTM Profiles > Web Filter > Local Ratings*.
- 2 Select *Create New*.
- 3 In the *URL* field, enter the URL of the site you want to change.
- 4 In the *Category Rating* table, select the category or categories to apply to the site.
If you created any local categories, a *Local Categories* group will appear.
- 5 Select *OK*.

FortiGuard Web Filter examples

FortiGuard Web Filter can provide more powerful filtering to your network because you can use it to restrict access to millions of sites by blocking the categories they belong to.

Configuring simple FortiGuard Web Filter protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable FortiGuard Web Filter protection on a FortiGate unit located in a satellite office.

Creating a web filter profile

Most FortiGuard Web Filter settings are configured in a web filter profile. Web filter profiles are selected in firewall policies. This way, you can create multiple web filter profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one web filter profile.

To create a web filter profile — web-based manager

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select *Create New* in the Edit Web Filter Profile window title bar
- 3 In the *Name* field, enter `basic_FGWF`.
- 4 Select *OK*.
- 5 Select *Add Filter*.
- 6 Select a *Filter type* of *Category*.
- 7 the *FortiGuard Web Filtering* check boxes for the *HTTP* and *HTTPS* traffic types.
- 8 Select the *FortiGuard Web Filtering* expand arrow.
- 9 Select the *Potentially Liable*, *Controversial*, and *Security Risk* categories.
- 10 Select an *Action* of *Block*.
- 11 Select *OK*.
- 12 Enable *HTTPS Scanning*.
- 13 Select *Apply*.

Applying the web filter profile to a security policy

A web filter profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When a web filter profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the web filter profile in a security policy — web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a policy and choose the *Edit* icon.
- 3 Enable *UTM*.
- 4 Select the *Enable Web Filter* option.
- 5 Select the `basic_FGWF` profile from the list.
- 6 Select *OK* to save the security policy.

To select the web filter profile in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set webfilter-profile basic_FGWF
  end
```

HTTP and HTTPS traffic handled by the security policy you modified will be monitored for attempts to access to the blocked sites. A small office may have only one security policy configured. If you have multiple policies, consider enabling web filter scanning for all outgoing policies.



If you have multiple policies, consider enabling web filter scanning for all outgoing policies.

School district

Continuing with the example in the Web filter section, you can use FortiGuard Web Filter to protect students from inappropriate material. For the first part of this example, see [“Web filtering example” on page 1008](#).

To enable FortiGuard Web Filter

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the web filter profile named *Students*.
- 3 Enable *HTTPS Scanning*.
- 4 Select *OK*.

The *Students* web filter profile has no FortiGuard Web Filtering filter. You must create and configure a filter to use FortiGuard Web Filter.

To configure the sites to block

- 1 Go to *UTM Profiles > Web Filter > Profile*.
- 2 Select the web filter profile named *Students*.
- 3 In the FortiGuard Categories table, select these categories: *Potentially Liable*, *Controversial*, and *Bandwidth Consuming*.

- 4 Select *Block* from the *Change Action for Selected Categories to* drop-down list.
- 5 Select *Apply* to save the web filter profile.

The students will not be able to access any of the web sites in those three general categories or the categories within them.



Data leak prevention

The FortiGate data leak prevention (DLP) system allows you to prevent sensitive data from leaving your network. When you define sensitive data patterns, data matching these patterns will be blocked, or logged and allowed, when passing through the FortiGate unit. You configure the DLP system by creating individual filters based on file type, file size, a regular expression, an advanced rule, or a compound rule, in a DLP sensor and assign the sensor to a security policy.

Although the primary use of the DLP feature is to stop sensitive data from leaving your network, it can also be used to prevent unwanted data from entering your network and to archive some or all of the content passing through the FortiGate unit.

This section describes how to configure the DLP settings.

The following topics are included:

- [Data leak prevention concepts](#)
- [Enable data leak prevention](#)
- [DLP document fingerprinting](#)
- [File filter](#)
- [Advanced rules](#)
- [Compound rules](#)
- [DLP archiving](#)
- [DLP examples](#)

Data leak prevention concepts

Data leak prevention examines network traffic for data patterns you specify. You define whatever patterns you want the FortiGate unit to look for in network traffic. The DLP feature is broken down into a number of parts.

DLP sensor

A DLP sensor is a package of filters. To use DLP, you must enable it in a security policy and select the DLP sensor to use. The traffic controlled by the security policy will be searched for the patterns defined in the filters contained in the DLP sensor. Matching traffic will be passed or blocked according to how you configured the filters.

DLP filter

Each DLP sensor has one or more filters configured within it. Filters can examine traffic for known files using DLP fingerprints, for files of a particular type or name, for files larger than a specified size, for data matching a specified regular expression, or for traffic matching an advanced rule or compound rule.

You can configure the action taken when a match is detected. The actions include Log Only, Block, Exempt, and Quarantine User, IP address, or Interface.

Logging is enabled by default, but you can also choose to archive matching traffic or generate an archive summary.

Fingerprint

Fingerprint scanning allows you to create a library of files for the FortiGate unit to examine. It will create checksum fingerprints so each file can be easily identified. Then, when files appear in network traffic, the FortiGate will generate a checksum fingerprint and compare it to those in the fingerprint database. A match triggers the configured action.

File filter

File filters use file filter lists to examine network traffic for files that match either file names or file types. For example, you can create a file filter list that will find files called secret.* and also all JPEG graphic files. You can create multiple file filter lists and use them in filters in multiple DLP sensors as required.

File size

This filter-type checks for files exceeding a configured size. All files larger than the specified size are subject to the configured action.

Regular expression

The FortiGate unit checks network traffic for the regular expression specified in a regular expression filter. Matching traffic is subject to the configured action.

Advanced rule

Each advanced rule includes a single condition and the type of traffic in which the condition is expected to appear. For more information about the supplied advanced rules, see [“Understanding the default advanced rules” on page 1044](#).

Compound rule

Compound rules combine multiple advanced rules and require that all the conditions in the included advanced rules are true before the compound rule is triggered. For more information about the supplied compound rules, see [“Understanding the default compound rules” on page 1045](#)

Enable data leak prevention

DLP examines your network traffic for data patterns you specify. You must configure DLP in sequence.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create a DLP sensor.

New DLP sensors are empty. You must create one or more filters in a sensor before it can examine network traffic.

- 2 Add one or more filters to the DLP sensor.

Each filter searches for a specific data pattern. When a pattern in the active DLP sensor appears in the traffic, the FortiGate unit takes the action configured in the matching filter.

- 3 Add the DLP sensor to one or more firewall policies that control the traffic to be examined.

Creating a DLP sensor

DLP sensors are collections of filters. You must also specify an action for the filter when you create it in a sensor. Once a DLP sensor is configured, you can select it a security policy profile. Any traffic handled by the security policy will be examined according to the DLP sensor configuration.

To create a DLP sensor

- 1 Go to *UTM Profiles > Data Leak Prevention > Sensor*.
- 2 Select the *Create New* icon on the Edit DLP Sensor window title bar.
- 3 In the *Name* field, enter the name of the new DLP sensor.
- 4 Optionally, you may also enter a comment. The comment appears in the DLP sensor list and can remind you of the details of the sensor.
- 5 Select the *Inspection Method*.

Proxy-based detection involves buffering the file or message and examining it as a whole. Advantages of proxy-based detection include a more thorough examination of attachments, especially archive formats and nesting.

Flow-based detection examines the file as it passes through the FortiGate unit without any buffering. Advantages of flow-based detection include speed and no interruption of detection during conserve mode.

- 6 Select *OK*.

The DLP sensor is created and the sensor configuration window appears.

- 7 Select *OK*.

A newly created sensor is empty, containing no filters. Without filters, the DLP sensor will do nothing.

Adding filters to a DLP sensor

Once you have created a DLP sensor, you need to add filters.

To add filters to a DLP sensor

- 1 Go to *UTM Profiles > Data Leak Prevention> Sensor*.
- 2 Select the Sensor in the Edit DLP Sensor window title bar drop-down list.
- 3 Select *Create New*.
- 4 Enter a filter name.

- 5 Select the type of filter from the *Filter By* drop-down list. The filter you choose determines the options available to you.

Fingerprint	<p>A fingerprint filter checks files in traffic against those in the FortiGate unit document fingerprint database. A match triggers the configured action.</p> <p>You must configure a document source or uploaded documents to the FortiGate unit for fingerprint scanning to work. For more information about document fingerprinting, see “DLP document fingerprinting” on page 1040.</p>
File Type	Files in the network traffic are filtered by filename and file type.
File Pattern	<p>Select a file filter list that includes the file patterns and file types the network traffic will be examined for. Files matching the types or patterns in the selected list are treated according to the selected action.</p> <p>To create a file filter list, see “Creating a file filter list” on page 1043.</p>
File Size	Files in the network traffic are filtered by size.
Maximum Size	Enter a file size in kilobytes. Files larger than the specified size are treated according to the selected action.
Regular Expression	Network traffic is examined for the pattern described by the regular expression.
Regular Expression	<p>Enter a regular expression. Traffic matching the regular expression is treated according to the selected action.</p> <p>For details about regular expression syntax, see “Using wildcards and Perl regular expressions” on page 1145.</p>
Advanced Rule	Advanced rules detect a variety of data patterns in network traffic.
Advanced Rule	Select an advanced rule. Traffic matching the selected advanced rule is treated according to the selected action.
Compound Rule	Compound rules detect a variety of data patterns in network traffic. Compound rules are made of multiple advanced rules and all of the conditions must match for the compound rule to take effect.
Advanced Rule	Select a compound rule. Traffic matching the selected compound rule is treated according to the selected action.

- 6 Select the *Action* the FortiGate unit will take against network traffic matching the rule. A number of actions are available:

Log Only	The FortiGate unit will take no action on network traffic matching a rule with this action. The filter match is logged, however. Other matching filters in the same sensor may still operate on matching traffic.
Block	Traffic matching a filter with the block action will not be delivered. The matching message or download is replaced with the data leak prevention replacement message.
Quarantine User	<p>If the user is authenticated, this action blocks all traffic to or from the user using the protocol that triggered the rule and adds the user to the Banned User list.</p> <p>If the user is not authenticated, this action blocks all traffic of the protocol that triggered the rule from the user's IP address.</p> <p>If the banned user is using HTTP, FTP, or NNTP (or HTTPS if the FortiGate unit supports SSL content scanning and inspection) the FortiGate unit displays the "Banned by data leak prevention" replacement message for the protocol. If the user is using IM, the IM and P2P "Banned by data leak prevention" message replaces the banned IM message and this message is forwarded to the recipient. If the user is using IMAP, POP3, or SMTP (or IMAPS, POP3S, SMTPS if your FortiGate unit supports SSL content scanning and inspection) the Mail "Banned by data leak prevention" message replaces the banned email message and this message is forwarded to the recipient. These replacement messages also replace all subsequent communication attempts until the user is removed from the banned user list.</p>
Quarantine IP Address	This action blocks access for any IP address that sends traffic matching a filter with this action. The IP address is added to the Banned User list. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts from this IP address until the IP address is removed from the banned user list.
Quarantine Interface	This action blocks access to the network for all users connecting to the interface that received traffic matching a filter with this action. The FortiGate unit displays the "NAC Quarantine DLP Message" replacement message for all connection attempts to the interface until the interface is removed from the banned user list.
Exempt	The exempt action prevents any filters from taking action on matching traffic. This action overrides the action assigned to any other matching filters.

Quarantine User, *Quarantine IP*, and *Quarantine Interface* provide functionality similar to NAC quarantine. However, these DLP actions block users and IP addresses at the application layer while NAC quarantine blocks IP addresses and interfaces at the network layer.



If you have configured DLP to block IP addresses and if the FortiGate unit receives sessions that have passed through a NAT device, all traffic from that NAT device — not just traffic from individual users — could be blocked. You can avoid this problem by implementing authentication.



To view or modify the replacement message text, go to *System > Config > Replacement Message*.

- 7 Select how traffic matching the rule will be archived.

Disable	Do not archive network traffic matching the rule.
Summary Only	Archive a summary of matching network traffic. For example, if applied to a rule governing email, the information archived includes the sender, recipient, message subject, message size, and other details.
Full	Archive the matching network traffic in addition to the summary information. For example, full archiving of email traffic includes the email messages and any attached files.



Archiving requires a FortiAnalyzer device or a subscription to the FortiGuard Analysis and Management Service.

- 8 Select *OK*.

The filter is added to the sensor. You may select *Create New* to add more filters, if required.

DLP document fingerprinting

One of the DLP techniques to detect sensitive data is fingerprinting (also called document fingerprinting). Most DLP techniques rely on you providing a characteristic of the file you want to detect, whether it's the file type, the file name, or part of the file contents. Fingerprinting is different in that you provide the file itself. The FortiGate unit then generates a checksum fingerprint and stores it. The FortiGate unit generates a fingerprint for all files detected in network traffic, and it is compared to all of the fingerprints stored in its fingerprint database. If a match is found, the configured action is taken.

The document fingerprint feature requires a FortiGate unit with internal storage. The document fingerprinting menu item does not appear on models without internal storage.

Any type of file can be detected by DLP fingerprinting and fingerprints can be saved for each revision of your files as they are updated.

To use fingerprinting you select the documents to be fingerprinted and then add fingerprinting filters to DLP sensors and add the sensors to firewall policies that accept the traffic to which to apply fingerprinting.

Fingerprinted Documents

The FortiGate unit must have access to the documents for which it generates fingerprints. One method is to manually upload documents to be fingerprinted directly to the FortiGate unit. The other is to allow the FortiGate unit to access a network share that contains the documents to be fingerprinted.

If only a few documents are to be fingerprinted, a manual upload may be the easiest solution. If many documents require fingerprinting, or if the fingerprinted documents are frequently revised, using a network share makes user access easier to manage.

To configure manual document fingerprints

- 1 Go to *UTM Profiles > Data Leak Prevention > Document Fingerprinting*.
- 2 In the Manual Document Fingerprints section, select *Create New*.
- 3 Select the file to be fingerprinted.
- 4 Choose a security level.
- 5 If the file is an archive containing other files, select *Process files inside archive* if you also want the individual files inside the archive to have fingerprints generated in addition to the archive itself.
- 6 Select *OK*.

The file is uploaded and a fingerprint generated.

To configure a fingerprint document source

- 1 Go to *UTM Profiles > Data Leak Prevention > Document Fingerprinting*.
- 2 In the Document Sources section, select *Create New*.
- 3 Configure the settings:

Name	Enter a descriptive name for the document source.
Server Address	Enter the IP address of the server.
User Name Password	Enter the user name and password of the account the FortiGate unit uses to access the server network share.
Path	Enter the path to the document folder.
Filename Pattern	You may enter a filename pattern to restrict fingerprinting to only those files that match the pattern. To fingerprint all files, enter an asterisk ("*").
Severity Level	Select a severity level. The severity is a tag for your reference that is included in the log files. It does not change how fingerprinting works.
Scan Periodically	To have the files on the document source scanned on a regular basis, select this option. This is useful if files are added or changed regularly. Once selected, you can choose Daily, Weekly, or Monthly update options, and enter the time of day the files are fingerprinted.
Advanced	Expand the Advanced heading for additional options.
Fingerprint files in subdirectories	By default, only the files in the specified path are fingerprinted. Files in subdirectories are ignored. Select this option to fingerprint files in subdirectories of the specified path.

Remove fingerprints for deleted files	Select this option to retain the fingerprints of files deleted from the document source. If this option is disabled, fingerprints for deleted files will be removed when the document source is rescanned.
Keep previous fingerprints for modified files	Select this option to retain the fingerprints of previous revisions of updated files. If this option is disabled, fingerprints for previous version of files will be deleted when a new fingerprint is generated.

4 Select OK.

File filter

File filter is a DLP option that allows you to block files based on their file name or their type.

- **File patterns** are a means of filtering based purely on the names of files. They may include wildcards (*). For example, blocking *.scr will stop all files with an scr file extension, which is commonly used for Windows screen saver files. Files trying to pass themselves off as Windows screen saver files by adopting the file-naming convention will also be stopped.

Files can specify the full or partial file name, the full or partial file extension, or any combination. File pattern entries are not case sensitive. For example, adding *.exe to the file pattern list also blocks any files ending with .EXE.

Files are compared to the enabled file patterns from top to bottom, in list order.

In addition to the built-in patterns, you can specify more file patterns to block. For details, see [“Creating a file filter list” on page 1043](#).

- **File types** are a means of filtering based on an examination of the file contents, regardless of the file name. If you block the file type *Archive (zip)*, all zip archives are blocked even if they are renamed with a different file extension. The FortiGate examines the file contents to determine what type of file it is and then acts accordingly.

The FortiGate unit can take either of the following actions toward the files that match a configured file pattern or type:

- **Block:** the file is blocked and a replacement message is sent to the user. If both file pattern filtering and virus scan are enabled, the FortiGate unit blocks files that match the enabled file filter and does not scan these files for viruses.
- **Allow:** the file is allowed to pass.

The FortiGate unit also writes a message to the UTM log and sends an alert email message if configured to do so.



File filter does not detect files within archives. You can use file filter to block or allow the archives themselves, but not the contents of the archives.

General configuration steps

The following steps provide an overview of file filter configuration. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create a file filter list.

- 2 Create one or more file patterns or file types to populate the file filter list.
- 3 Enable the file filter list by adding it to a filter in a DLP sensor.
- 4 Select the DLP sensor in a security policy.

Creating a file filter list

Before your FortiGate unit can filter files by pattern or type, you must create a file filter list.

To create a file filter list

- 1 Go to *UTM Profiles > Data Leak Prevention > File Filter*.
- 2 Select *Create New*.
- 3 Enter a *Name* for the new file filter list.
- 4 Select *OK*.

The new list is created and the edit file filter list window appears. The new list is empty. You need to populate it with one or more file patterns or file types.

Creating a file pattern

A file pattern allows you to block or allow files based on the file name. File patterns are created within file filter lists.

To create a file pattern

- 1 Go to *UTM Profiles > Data Leak Prevention > File Filter*.
- 2 Select a file filter list.
- 3 Select the *Edit* icon.
- 4 Select *Create New*.
- 5 Select *File Name Pattern* as the *Filter Type*.
- 6 Enter the pattern in the *Pattern* field. The file pattern can be an exact file name or can include wildcards (*). The file pattern is limited to a maximum of 80 characters.
- 7 Select the action the FortiGate unit will take when it discovers a matching file: *Allow* or *Block*.
- 8 The filter is enabled by default. Clear the *Enable* check box if you want to disable the filter.
- 9 Select *OK*.

Creating a file type

A file type allows you to block or allow files based on the kind of file. File types are created within file filter lists.

To create a file type

- 1 Go to *UTM Profiles > Data Leak Prevention > File Filter*.
- 2 Select the *Edit* icon of the file filter list to which you will add the file type.
- 3 Select *Create New*.
- 4 Select *File Type* as the *Filter Type*.
- 5 Select the kind of file from the *File Type* list.

- 6 Select the action the FortiGate unit will take when it discovers a matching file: *Allow* or *Block*.
- 7 The filter is enabled by default. Clear the *Enable* check box if you want to disable the filter.
- 8 Select *OK*.

Advanced rules

A number of advanced rules are provided with your FortiGate unit. If you require additional rules, you can create your own.

Understanding the default advanced rules

You can use the default advanced rules as provided, or modify them to fit your needs.



These rules affect only unencrypted traffic types. If you are using a FortiGate unit that can decrypt and examine encrypted traffic, you can enable those traffic types in these rules to extend their functionality if required.



Before using the rules, examine them closely to ensure you understand how they will affect the traffic on your network.

All-Email, All-FTP, All-HTTP, All-IM, All-NNTP	These rules will detect all email, FTP, HTTP, instant messaging, and NNTP traffic.
Email-AmEx, Email-Canada-SIN, Email-US-SSN, Email-Visa-Mastercard	These four rules detect American Express numbers, Canadian Social Insurance Numbers, U.S. Social Security Numbers, or Visa and Mastercard numbers within the message bodies of SMTP, POP3, and IMAP email traffic.
HTTP-AmEx, HTTP-Canada-SIN, HTTP-US-SSN, HTTP-Visa-Mastercard	These four rules detect American Express numbers, Canadian Social Insurance Numbers, U.S. Social Security Numbers, or Visa and Mastercard numbers sent using the HTTP POST command. The HTTP POST command is used to send information to a web server. As written, these rules are designed to detect data the user is sending to web servers. This rule does not detect the data retrieved with the HTTP GET command, which is used to retrieve web pages.
Large-Attachment	This rule detects files larger than 5MB attached to SMTP, POP3, and IMAP email messages.
Large-FTP-Put	This rule detects files larger than 5MB sent using the FTP PUT command. Files received using FTP GET are not examined.

Large-HTTP-Post	This rule detects files larger than 5MB sent using the HTTP POST command. Files received using HTTP GET are not examined.
Email-Not-Webex, HTTP-Post-Not-Webex	These rules detect all traffic that is not generated by the WebEx web conference service. While not very useful on their own, these rules are used in the default compound rules.

Creating advanced rules

Most DLP functions can be accomplished using file filtering, file size, or regular expression features, but advanced rules do offer filtering features unavailable elsewhere in DLP. Creating your own advanced rules can be accomplished only using the CLI.

The general procedure when creating a DLP advanced rule is to specify the protocol, sub-protocol (if any), the field, and then the remaining options as required.

This example shows one way to create an advanced rule called email-confidential to detect the word “confidential” in the body of any email message. The rule can detect the word regardless of whether it is uppercase, lowercase, or a mixture of both.

To create a DLP advanced rule

```
config dlp rule
edit email-confidential
set protocol email
set sub-protocol imap pop3 smtp
set field body
set regexp /confidential/i
set description "Detect the word Confidential in email"
end
```

See the [CLI Reference](#) for a complete list of the DLP advanced rule commands and options.

Compound rules

A number of compound rules are provided with your FortiGate unit. If you require additional compound rules, you can create your own.

Understanding the default compound rules

You can use the default compound rules as provided, or modify them to fit your needs.



These rules affect only unencrypted traffic types. If you are using a FortiGate unit that can decrypt and examine encrypted traffic, you can enable those traffic types in these rules to extend their functionality if required.



Before using the rules, examine them closely to ensure you understand how they will affect the traffic on your network.

Email-SIN	This rule combines the Email-Canada-SIN and Email-Not-Webex advanced rules. If you use the Webex web conferencing service, the Email-Canada-SIN advanced rule can be mistakenly triggered by Webex traffic. Combining the advanced rules in a compound rules ensures that the Email-Canada-SIN rule will be triggered only by non-Webex network traffic.
HTTP-Post-SIN	This rule combines the HTTP-Canada-SIN and HTTP-Post-Not-Webex advanced rules. If you use the Webex web conferencing service, the HTTP-Canada-SIN advanced rule can be mistakenly triggered by Webex traffic. Combining the advanced rules in a compound rules ensures that the HTTP-Canada-SIN rule will be triggered only by non-Webex network traffic.

Creating compound rules

Compound rules are a very powerful feature to detect a specific set of circumstances. Because of this, you may need to create your own compound rules. Creating your own compound rules can be accomplished only using the CLI.

This example shows one way to create a compound rule called email-confidential-attachment to detect email messages that have the word “confidential” in the body of any email message and a file attachment of at least 5 MB in size.

To create a DLP compound rule

```
config dlp compound
edit email-confidential-attachment
set protocol email
set sub-protocol imap pop3 smtp
set member email-confidential
set member Large-Attachment
end
```

See the [CLI Reference](#) for a complete list of the DLP compound rule commands and options.

DLP archiving

DLP is typically used to prevent sensitive information from getting out of your company network, but it can also be used to record network use. This is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Enabling archiving for rules when you add them to sensors directs the FortiGate unit to record all occurrences of these traffic types when they are detected by the sensor.

Since the archive setting is configured for each rule in a sensor, you can have a single sensor that archives only the things you want.

DLP archiving comes in two forms: *Summary Only*, and *Full*.

Summary archiving records information about the supported traffic types. For example, when an email message is detected, the sender, recipient, message subject, and total size are recorded. When a user accesses the Web, every URL the user visits recorded. The result is a summary of all activity the sensor detected.

For more detailed records, full archiving is necessary. When an email message is detected, the message itself, including any attachments, is archived. When a user accesses the Web, every page the user visits is archived. Far more detailed than a summary, full DLP archives require more storage space and processing.

Because both types of DLP archiving require additional resources, DLP archives are saved to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service (subscription required).

Two sample DLP sensors are provided with DLP archiving capabilities enabled. If you select the `Content_Summary` sensor in a security policy, it will save a summary DLP archive of all traffic the security policy handles. Similarly, the `Content_Archive` sensor will save a full DLP archive of all traffic handled the security policy you apply it to. These two sensors are configured to detect all traffic of the supported types and archive them.

DLP examples

Blocking sensitive email messages

Someone in the Example.com corporation has been sending copies of the company president's monthly update email messages to the press. These messages have included the full header. Rather than try to block them, the IT department at Example.com will find out who is sending the messages using DLP.

All messages include the text `From: president@example.com` and `Subject: XYZ Monthly Update` where XYZ is the month the update applies to.

You will create a rule for the email address and a rule for the subject, combine them in a compound rule, and add the compound rule to a DLP sensor. You will then add the DLP sensor to the security policy that controls outgoing email traffic.

To create the "address" rule

```
config dlp rule
  edit "President address"
    set description "Finds president@example.com in email"
    set protocol email
    set sub-protocol imap pop3 smtp
    set field body
    set regexp-wildcard enable
    set regexp president@example.com
  end
```

To create the "subject" rule

```
config dlp rule
  edit "President subject"
    set description "Finds XYZ Monthly Update in email subject"
    set protocol email
    set sub-protocol imap pop3 smtp
    set field subject
    set regexp-wildcard enable
    set regexp "* Monthly Update"
  end
```

The asterisk (*) can represent any characters so the rule will match any monthly update.

Adding these two rules to a DLP sensor may generate a large number of false positives because any rule in a sensor will trigger the action. If the action were to log messages matching the address and subject rules in this example, then left as individual rules, the DLP sensor would log Monthly Updates from any employee and log all the president's email messages. In this case, you only want to know when both rules are true for a single message. To do this, you must first add the rules to a compound rule.

To create the “president + subject” compound rule.

```
config dlp compound
edit "President + Subject"
set protocol email
set sub-protocol imap pop3 smtp
set member "President address"
set member "President subject"
end
```

To create the “president” DLP sensor

- 1 Go to *UTM Profiles > Data Leak Prevention > Sensor*.
- 2 Select the *Create New* icon on the Edit DLP Sensor window title bar.
- 3 In the *Name* field, enter `president`.
- 4 In the *Comments* field, enter Finds “`president@example.com`” and “XYZ Monthly Update” in email.
- 5 Select *OK* to save the new sensor.
- 6 Select *Create New* to add a rule to the sensor.
- 7 Enter `President + Subject` for the *Filter Name*.
- 8 Set *Filter By* to *Compound Rule*.
- 9 Set *Advanced Rule* to *President + Subject*.
- 10 Set *Archive* to *Full*.
- 11 Select *OK*.

With the DLP sensor ready for use, you need to select it in the security policy.

To select the DLP sensor in the security policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select the security policy that controls outgoing email traffic.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select the *Enable DLP Sensor* option.
- 6 Select the `president` sensor from the list.
- 7 Select *OK* to save the security policy.

With the DLP sensor specified in the correct security policy, any email message with both `president@example.com` and Monthly Update in the subject will trigger the sensor and the email message will be archived.

Data Leak Prevention interface reference

You can use the Data Leak Prevention (DLP) system to prevent sensitive data from leaving or entering your network. You can define sensitive data patterns, and data matching these patterns will be blocked and/or logged or archived when passing through the unit. The DLP system is configured by creating individual rules, combining the rules into DLP sensors, and then assigning a sensor to a firewall policy.

Although the primary use of the DLP feature is to stop sensitive data from leaving your network, it can also be used to prevent unwanted data from entering your network and to archive some or all of the content passing through the unit.

This topic includes the following:

- [Sensor](#)
- [Document Fingerprinting](#)
- [File Filter](#)
- [DLP archiving](#)

Sensor

DLP sensors are simply collections of DLP rules and DLP compound rules. The DLP sensor also includes settings such as action, archive, and severity for each rule or compound rule. Once a DLP sensor is configured, it can be specified in a firewall policy. Any traffic handled by the policy in which the DLP sensor is specified will enforce the DLP sensor configuration.

You can create a new DLP sensor and configure it to include the DLP rules and DLP compound rules required to protect the traffic leaving your network.

A DLP sensor must be created before it can be configured by adding rules and compound rules.



Before use, examine the sensors and rules in the sensors closely to ensure you understand how they will affect the traffic on your network.

DLP sensor configuration settings

The following are DLP sensor configuration settings in *UTM Profiles > Data Leak Prevention > Sensor*.

Sensor page Lists each individual DLP sensor that you created, as well as the default DLP sensors. On this page, you can edit DLP sensors (default or ones that you created), delete or create new DLP sensors. You are redirected to this page when you select <i>View List</i> on the Edit DLP Sensor page.	
Create New	Creates a new DLP sensor. When you select <i>Create New</i> , you are automatically redirected to the New DLP List page. This page provides a name field and comment field. You must enter a name to go to the Sensor Settings page.
Edit	Modifies settings within a DLP sensor. When you select <i>Edit</i> , you are automatically redirected to the Edit DLP Sensor page.
Delete	Removes a DLP sensor from the list. To remove multiple DLP sensors from within the list, on the Sensor page, in each of the rows of the sensors you want removed, select the check box and then select <i>Delete</i> . To remove all DLP sensors from the list, on the Sensor page, select the check box in the check box column and then select <i>Delete</i> .
Clone	Select to use an existing DLP sensor's settings as the basis for a new DLP sensor's settings.
Name	The DLP sensor name. There are six default sensors. The following default DLP sensors are provided with your unit. You can use these as provided, or modify them as required.
Comments	The optional description of the DLP sensor.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref</i>., and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
Content_Archive (default)	<p>DLP archive all email (POP3, IMAP, and SMTP), FTP, HTTP, and IM traffic. For each rule in the sensor, <i>Archive</i> is set to <i>Full</i>. No blocking or quarantine is performed. See “DLP archiving” on page 1060.</p> <p>You can add the <i>All-Session-Control</i> rule to also archive session control content.</p> <p>If you have a unit that supports SSL content scanning and inspection, you can edit the <i>All-Email</i> rule to archive POP3S, IMAPS, and SMTPS traffic. You can also edit the <i>All-HTTP</i> rule to archive HTTPS traffic.</p>
Content_Summary (default)	<p>DLP summary archive all email (POP3, IMAP, and SMTP), FTP, HTTP, and IM traffic. For each rule in the sensor, <i>Archive</i> is set to <i>Summary Only</i>. No blocking or quarantine is performed. See “DLP archiving” on page 1060.</p> <p>You can add the <i>All-Session-Control</i> rule to also archive session control content.</p> <p>If you have a unit that supports SSL content scanning and inspection, you can edit the <i>All-Email</i> rule to archive POP3S, IMAPS, and SMTPS traffic.</p> <p>You can also edit the <i>All-HTTP</i> rule to archive HTTPS traffic.</p>

Credit-Card (default)	<p>The number formats used by American Express, Visa, and Mastercard credit cards are detected in HTTP and email traffic.</p> <p>As provided, the sensor is configured not to archive matching traffic and an action of <i>None</i> is set. Configure the action and archive options as required.</p>
Large-File (default)	<p>Files larger than 5MB will be detected if attached to email messages or if send using HTTP or FTP.</p> <p>As provided, the sensor is configured not to archive matching traffic and an action of <i>None</i> is set. Configure the action and archive options as required.</p>
SSN-Sensor (default)	<p>The number formats used by U.S. Social Security and Canadian Social Insurance numbers are detected in email and HTTP traffic.</p> <p>As provided, the sensor is configured not to archive matching traffic and an action of <i>None</i> is set. Configure the action and archive options as required.</p>
<p>Edit DLP Sensor page</p> <p>Provides settings for configuring rules that are added to DLP sensors. When you select <i>Create New</i> to create a new sensor, you are automatically redirected to the New DLP Sensor page. You must enter a name for the sensor in the <i>Name</i> field to continue configuring the sensor, at which time you are redirected to the Sensor Settings page. When you select <i>Create New</i> on this page, you are redirected to the New DLP Sensor Rule page.</p> <p>Note: Enabling logging or NAC quarantine is available in the CLI.</p>	
Name	If you are editing an existing sensor and want to change the name, enter a name in this field. You must select <i>OK</i> to save the change.
Comment	If you are editing an existing sensor and want to change or add a description, enter the new text in this field. You must select <i>OK</i> to save these changes.
Inspection Method	Select the type of DLP inspection method.
Flow-based Detection	Select to enable flow-based DLP scanning. Flow-based is a non-proxy solution which provides high concurrent session, high session rate, and low-latency DLP service.
Proxy-based Detection (Extended)	Select to enable a proxy-based detection scanning method.
Create New	<p>Adds a new rule or compound rule to the sensor. When you select a specific type of member, either <i>Compound rule</i> or <i>Rule</i>, different options become available.</p> <p>When you select <i>Create New</i>, you are automatically redirected to the New DLP Sensor Rule page.</p>
Edit	Modifies a rule or compound rule. When you select <i>Edit</i> , you are automatically redirected to Edit DLP Sensor Rule page.

Delete	<p>Removes a compound rule or a rule from the list.</p> <p>To remove multiple compound rules or rules from the list, on the DLP Sensor Settings page, in each of the rows of the compound rules or rules you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all compound rules and/or rules from the list, select the check box in the check box column and then select <i>Delete</i>.</p>
Filter Name	The name of the rule that you have created. This is often referred to as a filter because it can filter the information using the rule applied.
Type	The type of rule that was applied, such as compound rule or fingerprint.
Action	<p>The action configured for each rule. If the selected action is <i>None</i>, no action will be listed.</p> <p>Although archiving is enabled independent of the action, the <i>Archive</i> designation appears with the selected action.</p> <p>For example, if you select the <i>Block</i> action and set <i>Archive</i> to <i>Full</i> for a rule, the action displayed in the sensor rule list is <i>Block, Archive</i>.</p>
Archive	The type of archiving that is selected for that rule. For example, summary only.
New DLP Sensor Filter page Provides settings for configuring filter, such as compound rule or document finger print.	
Filter Name	Enter a name for the rule.
Filter By	<p>Select what the DLP sensor will filter by, for example compound rule. You can choose to filter by the following rules:</p> <ul style="list-style-type: none"> • Finger Print • File Type • File Size • Regular Expression • Advanced Rule • Compound Rule
Sensitivity	The sensitivity level for document fingerprinting. Appears only when <i>Filter By</i> is <i>Finger Print</i> .
Advanced Rule	Select the advance rule that you want applied to the filter. Appears when <i>Filter By</i> is either <i>Compound Rule</i> or <i>Advanced Rule</i> .
Regular Expression	Enter the regular expression in the field provided. Appears only when <i>Filter By</i> is <i>Regular Expression</i> .
Maximum Size	Enter the maximum file size in kB. Appears only when <i>Filter By</i> is <i>File Size</i> .

File Pattern	Select the type of file pattern. Appears only when <i>Filter By</i> is <i>File Type</i> .
Action	Select an action that the unit will take for that particular rule or compound rule.



DLP prevents duplicate action. Even if more than one rule in a sensor matches some content, DLP will not create more than one DLP archive entry, quarantine item, or ban entry from the same content.

Document Fingerprinting

DLP document fingerprinting allows you to better protect specific documents from leakage. Document fingerprinting, in this sense, is a method of uniquely identifying a document. This method breaks up files into chunks, taking a checksum of those chunks and using that checksum as the fingerprint. The fingerprint is then applied to a DLP filter rule within a DLP sensor which is then used during the scanning process of DLP activity.

Document Fingerprint page	
Lists the document sources as well as the document files that were manually inputted.	
Document Sources section	
Lists the configured document sources for document fingerprinting.	
Create New	Creates a new document source. When you select <i>Create New</i> , you are automatically redirected to the Document Source page.
Delete	Removes a document source from within the list. To remove multiple document sources from the list, in the Document Sources section of the page, in each of the rows of sources you want removed, select the check box and then select <i>Delete</i> . To remove all document sources from the list, select the check box in the check box column and then select <i>Delete</i> .
View	Select to view the document sources' information.
Name	The name of the document source.
Server	The source's IP address or server IP address of where the documents/files are located.
Path	The location of where the files are located on the server.
Sensitivity Level	The level of sensitivity for the documents on the server.
# Documents	The total number of documents for fingerprinting.
Manual Document Fingerprints section	
Lists each individual document or file that you uploaded to the FortiGate unit for fingerprinting.	
Create New	Uploads a file or document that you want for fingerprinting. When you select <i>Create New</i> , you are automatically redirected to the Upload File for Fingerprinting page. For more information about uploading individual documents for fingerprinting, see "Manual Document Fingerprints" on page 1056 .

Delete	Removes a document from within the list. To remove multiple documents from the list, in the Manual Document Fingerprints section of the page, in each of the rows of documents you want removed, select the check box and then select <i>Delete</i> . To remove all documents from the list, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the document.
Sensitivity Level	The level of sensitivity given to that uploaded document.

Document Sources configuration settings

The following are configuration settings for document sources in *UTM Profiles > Data Leak Prevention > Document Fingerprinting*.

Document Source page	
Provides settings for configuring the document sources, or servers, that contain the files or documents that you want to prevent being leaked from your intranet.	
Name	Enter the name of the document source.
Server Type	Select the type of server, such as a windows share server. The Windows Share server is the default server type available.
Server Address	Enter the server's IP address.
User Name	Enter the name of the user that logs in to the server.
Password	Enter the user's password that is required to log in to the server.
Path	Enter the location of where the files are located on the server.
Filename Pattern	Enter the pattern of the name of the files, such as *.pdf.
Sensitivity Level	Select the type of sensitivity level you created in the CLI, using the <code>config dlp fp-sensitivity</code> command. There are three default sensitivity levels: Private, Critical and Warning.
Scan Periodically	Select to schedule when the FortiGate unit scans the server. When you select the check box to enable this, the following options appear.
Daily	Select to schedule a daily scan. Enter the hour and minutes in the fields provided.
Weekly	Select to schedule a scan during a day of the week and at a specific time. When you select <i>Weekly</i> , <i>Weekday</i> appears. Select a day from the drop-down list in <i>Weekday</i> and then enter the hour and minutes in the fields provided. For example, a scan will occur every Monday at 5:30.

Monthly	Select the day of the month to scan. Enter the date of the month, such as 5, in the <i>Date</i> field. Enter the hour and minutes in the fields provided.
Advanced	Expand to enable or disable any of the following: <ul style="list-style-type: none"> • <i>Fingerprint files in subdirectories</i> – fingerprints files that are in subdirectories as well as in directories • <i>Remove fingerprints for deleted files</i> – removes fingerprints from files that are deleted from the source or server. • <i>Keep previous fingerprints for modified files</i> – keeps the fingerprints for files that were recently modified.

Manual Document Fingerprints

The Manual Document Fingerprints section of the DLP Fingerprint page allows you to upload a document that will be fingerprinted by the FortiGate unit.

The following are settings for uploading documents for fingerprinting in *UTM Profiles > Data Leak Prevention > Document Fingerprinting*.

Upload File for Fingerprinting page	
Uploads the individual files that you want to provide for fingerprinting.	
File	Enter the file name.
Sensitivity Level	Select the sensitivity level from the drop-down list.
Process files inside archives	Select to enable the FortiGate unit to process files that are inside archive files.

File Filter

The Filter menu allows you to configure filtering options that block specific file patterns and file types. Files are compared to the enabled file patterns and then the file types from top to bottom. If a file does not match any specified patterns or types, it is passed along to antivirus scanning (if enabled). In effect, files are passed if not explicitly blocked. The unit also writes a message to the virus log and sends an alert email message if configured to do so.

The unit can take either of these actions toward files that match a configured file pattern or type:

- Allow: the file is allowed to pass.
- Block: the file is blocked and a replacement messages will be sent to the user. If both file filter and virus scan are enabled, the unit blocks files that match the enabled file filter and does not scan these files for viruses.
- Intercept: the file will be archived to the local hard disk or the FortiAnalyzer unit. (FortiOS Carrier only)

By using the *Allow* action, this behavior can be reversed with all files being blocked unless explicitly passed. Simply enter all the file patterns or types to be passed with the allow attribute. At the end of the list, add an all-inclusive wildcard (*.*) with a block action. Allowed files continue to antivirus scanning (if enabled) while files not matching any allowed patterns are blocked by the wildcard at the end. For standard operation, you can choose to disable file filter in the profile, and enable it temporarily to block specific threats as they occur.

The unit is preconfigured with a default list of file patterns:

- executable files (*.bat, *.com, and *.exe)
- compressed or archive files (*.gz, *.rar, *.tar, *.tgz, and *.zip)
- dynamic link libraries (*.dll)
- HTML application (*.hta)
- Microsoft Office files (*.doc, *.ppt, *.xl?)
- Microsoft Works files (*.wps)
- Visual Basic files (*.vb?)
- screen saver files (*.scr)
- program information files (*.pif)
- control panel files (*.cpl)

The unit can detect the following file types:

Table 70: Supported file types

arj	activemime	aspack	base64	bat	binhex	bzip	bzip2
cab	class	cod	elf	exe	fsg	gzip	hlp
hta	html	jad	javascript	lzh	mime	msc	msoffice
petite	prc	rar	sis	tar	upx	uue	zip
unknown	ignored						



The “unknown” type is any file type that is not listed in the table. The “ignored” type is the traffic the unit typically does not scan. This includes primarily streaming audio and video.

File filter configuration

You can add multiple file filter lists to the antivirus profile. For file patterns, you can add a maximum of 5000 patterns to a list. For file types, you can select only from the supported types.

The following are file filter configuration settings in *UTM Profiles > Data Leak Prevention > File Filter*.

File Filter page Lists each individual file filter that you created. On this page, you can edit, delete or create a new file filter.	
Create New	Creates a new file filter. When you select <i>Create New</i> , you are automatically redirected first to the New List page. You must enter a name for the file filter list in the <i>Name</i> field on the New List page and then select <i>OK</i> then be redirected to the File Filter Settings page.
Delete	Removes the file filter list from the list on the File Filter page. To remove multiple file filter lists from within the list, on the File Filter page, in each of the rows of the file filter lists you want removed, select the check box and then select <i>Delete</i> . To remove all file filter lists from the list, on the File Filter page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies the settings within a file filter list. When you select <i>Edit</i> , you are redirected to the File Filter Settings page.
Name	The name of the file filter list.
# Entries	The number of file patterns or file types in each file filter list.
DLP Rule	The DLP rules in which each filter is used.
Comments	An optional description of each file filter list.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> . To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object. To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window: <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
File Filter Settings page Provides settings for configuring multiple file patterns and file types that make up a file filter. This page also lists the file patterns and file types that were created for the file filter. If you are editing a file filter, you are redirected to this page.	

Name	The name that was entered in the <i>Name</i> field on the New List page. To change the name, edit the text in this field and select <i>OK</i> .
Comment	The comment that was entered in the <i>Comment</i> field on the New List page. If you want to edit or add the description, enter the text in this field and select <i>OK</i> .
Create New	Creates a new file filter pattern or type within the list on the File Filter Settings page. When you select <i>Create New</i> , you are automatically redirected to the New File Filter page.
Edit	Modifies settings within the file pattern/type and action. When you select <i>Edit</i> , you are automatically redirected to the Edit File Filter page.
Delete	Removes the file pattern or type from the list on the File Filter Settings page. To remove multiple file filter lists from within the list, on the File Filter page, in each of the rows of the file filter lists you want removed, select the check box and then select <i>Delete</i> . To remove all file filter lists from the list, on the File Filter page, select the check box in the check box column and then select <i>Delete</i> .
Enable	Enables a disabled file pattern or type.
Disable	Disables a file pattern or type.
Move To	Moves the file pattern or type to any position in the list. When you select <i>Move To</i> , the Move AV File Filter Entry window appears. To move a file pattern or type, select the new position <i>Before</i> or <i>After</i> , which will place the current entry before or after the entry you enter in (<i>Entry</i>). Enter the entry's name in the (<i>Entry</i>) field.
Filter	The current list of file patterns and types.
Action	Files matching the file patterns and types can be set to <i>Block</i> or <i>Allow</i> .
Enable	Indicates that the file pattern or file type is either enabled or disabled. A green check mark indicates that the pattern or type is enabled; a gray x indicates that it is disabled.
New File Filter page	
Provides settings for configuring the file pattern or file type for the list. When you select <i>Create New</i> on the File Filter Settings page, you are automatically redirected to this page.	
Filter Type	Select <i>File Name Pattern</i> or <i>File Type</i> .
File Type	Select a file type from the list. Appears only when <i>File Type</i> is selected in <i>Filter Type</i> .
Pattern	Enter the file pattern. The file pattern can be an exact file name or can include wildcards. The file pattern can be 80 characters long.
Action	Select an action from the drop down list: <i>Block</i> or <i>Allow</i> .
Enable	Select to enable or disable the filter.



The default file pattern list catalog is called *builtin-patterns*.

DLP archiving

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the FortiGate configuration. The FortiGuard Analysis server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

You can configure full DLP archiving and summary DLP archiving. Full DLP archiving includes all content, for example, full email DLP archiving includes complete email messages and attachments. Summary DLP archiving includes just the meta data about the content, for example, email message summary records include only the email header.

You can archive Email, FTP, HTTP, IM, and session control content:

- Email content includes IMAP, POP3, and SMTP sessions. Email content can also include email messages tagged as spam by Email filtering. If your unit supports SSL content scanning and inspection, Email content can also include IMAPS, POP3S, and SMTPS sessions.
- HTTP content includes HTTP sessions. If your unit supports SSL content scanning and inspection HTTP content can also include HTTPS sessions.
- IM content includes AIM, ICQ, MSN, and Yahoo! sessions.
- MMS content includes MM1, MM3, MM4, and MM7 sessions. (FortiOS Carrier only)
- Session control content includes SIP, SIMPLE and SCCP sessions. Only summary DLP archiving is available for SIP and SCCP. Full and summary DLP archiving is available for SIMPLE.

You add DLP sensors to archive Email, Web, FTP, IM, and session control content. Archiving of spam email messages is configured in the DLP sensor.

In FortiOS Carrier, MMS archiving is configured in MMS profiles.

DLP archiving is enabled in the DLP sensor itself. DLP sensors are located in *UTM Profiles > Data Leak Prevention > Sensor*. You can also use either Content_Archive or Content_Summary sensors to archive DLP logs instead of creating a new DLP sensor for archiving purposes.

You can now create a session control DLP rule that includes SIP, SIMPLE or SCCP for DLP archiving within the CLI.



Application control

Using the application control UTM feature, your FortiGate unit can detect and take action against network traffic depending on the application generating the traffic. Based on FortiGate Intrusion Protection protocol decoders, application control is a user-friendly and powerful way to use Intrusion Protection features to log and manage the behavior of application traffic passing through the FortiGate unit. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic even if the traffic uses non-standard ports or protocols.

The FortiGate unit can recognize the network traffic generated by a large number of applications. You can create application control sensors that specify the action to take with the traffic of the applications you need to manage and the network on which they are active, and then add application control sensors to the firewall policies that control the network traffic you need to monitor.

This section describes how to configure the application control settings.

If you enable virtual domains (VDOMs) on the Fortinet unit, you need to configure application control separately for each virtual domain.

The following topics are included in this section:

- [Application control concepts](#)
- [Enable application control](#)
- [Application traffic shaping](#)
- [Application control monitor](#)
- [Application control packet logging](#)
- [Application considerations](#)
- [Application control examples](#)

Application control concepts

You can control network traffic generally by the source or destination address, or by the port, the quantity or similar attributes of the traffic itself in the security policy. If you want to control the flow of traffic from a specific application, these methods may not be sufficient to precisely define the traffic. To address this problem, the application control feature examines the traffic itself for signatures unique to the application generating it. Application control does not require knowledge of any server addresses or ports. The FortiGate unit includes signatures for over 1000 applications, services, and protocols.

Updated and new application signatures are delivered to your FortiGate unit as part of your FortiGuard Application Control Service subscription. Fortinet is constantly increasing the number of applications that application control can detect by adding applications to the [FortiGuard Application Control Database](#). Because intrusion protection protocol decoders are used for application control, the application control database is part of the [FortiGuard Intrusion Protection System Database](#) and both of these databases have the same version number.

To view the version of the application control database installed on your FortiGate unit, go to the *License Information* dashboard widget and find the *IPS Definitions* version.

To see the complete list of applications supported by FortiGuard Application Control go to the [FortiGuard Application Control List](#). This web page lists all of the supported applications. You can select any application name to see details about the application.

Enable application control

Application control examines your network traffic for traffic generated by the applications you want it to control.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 Create an application sensor.
- 2 Configure the sensor to include the signatures for the application traffic you want the FortiGate unit to detect. Configure each entry to allow or pass the traffic.
- 3 Enable UTM and application control in a security policy and select the application sensor.

Creating an application sensor

You need to create an application sensor before you can enable application control.

To create an application sensor

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.
- 3 In the *Name* field, enter the name of the new application sensor.
- 4 Optionally, you may also enter a comment.
- 5 Select *OK*.

The application sensor is created and the sensor configuration window appears. A newly created application sensor is empty. Without applications, the application sensor will have no effect.

Adding applications to an application sensor

Once you have created an application sensor, you need to need to define the applications that you want to control.

You can add applications using application entries and application filters. Entries allow you to choose individual applications. Filters allow you to choose application attributes and all the applications with matching attributes are included in the filter.



The sequence of the entries in the table is significant. The entries are checked in sequence, from top to bottom, and when a match is found and the action executed, further checking is stopped.

To add an application entry to an application sensor

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select an application sensor from the drop-down list in the *Edit Application Sensor* window title bar.
- 3 Select the *Create New* drop-down icon in the sensor and choose *Entry*.

- 4 Using the *Category* selection, choose the type of application you want to add. For example, if you want to add Facebook chat, choose *im*.
The *Category* selection displays only the options available in the *Application* you select. If you want to see all the applications listed, leave *Category* set to *All Categories*.
- 5 Using the *Application* selection, choose the application you want to add.
The applications available to you will be limited to those in the category you select.
- 6 Select the *Action* the FortiGate unit will take when it detects network traffic from the application:
 - *Block* will stop all traffic from the application and log all occurrences.
 - *Monitor* allows the application traffic to flow normally and log all occurrences.
 If you set the action to *Monitor*, you have the option of enabling traffic shaping for the application or applications specified in this application list entry. For more information about application control traffic shaping, see [“Application traffic shaping” on page 1067](#)
- 7 Enable *Session TTL* to specify a time-to-live value for the session, in seconds. If this option is not enabled, the TTL defaults to the setting of the CLI command `config system session-ttl`.
- 8 Select *Enable Packet Log* to have the FortiGate unit save the packets that application control used to determine the traffic came from the application.
- 9 Some applications have additional options:

IM Options (for some IM applications)	
Block Login	Select to prevent users from logging in to the selected IM system.
Block File Transfers	Select to prevent the sending and receiving of files using the selected IM system.
Block Audio	Select to prevent audio communication using the selected IM system.
Inspect Non-standard Port	Select to allow the FortiGate unit to examine nonstandard ports for the IM client traffic.
Display DLP meta-information on the system dashboard	Select to include meta-information detected for the IM system on the FortiGate unit dashboard.
Other Options	
Command	<p>Some traffic types include a command option. These include FTP.Command, NNTP.Command, POP3.Command, and SMTP.Command. Specify a command that appears in the traffic that you want to block or pass.</p> <p>For example, enter <code>GET</code> as a command in the <i>FTP.Command</i> application to have the FortiGate unit examine FTP traffic for the <code>GET</code> command. Multiple commands can be entered.</p>

Method	<p>A method option is available for HTTP, RTSP, and SIP protocols. Specify a method that appears in the traffic that you want to block or pass.</p> <p>For example, enter <code>POST</code> as a method in the <i>HTTP.Method</i> application to have the FortiGate unit examine HTTP traffic for the POST method. Multiple methods can be entered.</p>
Program Number	<p>Enter the program number appearing in Sun Remote Procedure Calls (RPC) that you want to block or pass. Multiple program numbers can be entered.</p>
UUID	<p>Enter the UUID appearing in Microsoft Remote Procedure Calls (MSRPC) that you want to block or pass. Multiple UUIDs can be entered.</p>

To add an application filter to an application sensor

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.
- 3 Select the *Create New* drop-down icon in the sensor and choose *Filter*.
- 4 Configure the filter that you require. Applications matching all of the characteristics you specify in the filter will be included in the filter.

Category	<p>Select <i>Specify</i> and choose an application category to include in the filter. All applications within the category will be included. Some categories have one or more subcategories to allow you to narrow the included applications. For example, selecting IM will include all of the instant messaging applications, but the VoIP subcategory will restrict the applications to only those related to VoIP.</p> <p>If you select <i>All</i>, the category attribute will not be used to determine which signatures are included in the filter.</p>
Vendor	<p>Select <i>Specify</i> and choose a vendor to include all of that vendor's applications in the filter.</p> <p>If you select <i>All</i>, the vendor attribute will not be used to determine which signatures are included in the filter.</p>
Behavior	<p>Select <i>Specify</i> and choose the type of application behavior to include. For example, selecting <i>Encrypted-Tunneling</i> will include all applications providing or related to encrypted tunneling.</p> <p>If you select <i>All</i>, the behavior attribute will not be used to determine which signatures are included in the filter.</p>
Technology	<p>Select <i>Specify</i> and choose the technology of the applications to include. Options include web-browser, peer-to-peer, client, and server.</p> <p>If you select <i>All</i>, the technology attribute will not be used to determine which signatures are included in the filter.</p>

Protocol	<p>Select <i>Specify</i> and choose the network protocols the applications use.</p> <p>If you select <i>All</i>, the Protocol attribute will not be used to determine which signatures are included in the filter.</p>
Tags	<p>Tags are a means by which you can apply customized labels to your application filters. Specified tags are displayed only within the filter itself on the Edit Application Filter page.</p> <p>By default, the tag feature is disabled on all but the largest FortiGate models. If the Tags option is not visible, you must go to <i>System > Admin > Settings</i> and enable <i>Display Object Tagging and Coloring</i> to enable it.</p> <p>For more information about tags, see “Tag management” on page 582.</p>
Applied Tags	Displays the tags that you have applied to the filter.
Add tags	Enter a tag and then select the plus (+) icon to add the tag to the filter. This also adds the tag to the <i>Applied tags</i> list.
View Matched rules	Select view a list of all the applications included in the filter with the current settings.
Action	<p>Select the <i>Action</i> the FortiGate unit takes when it detects network traffic from the application:</p> <ul style="list-style-type: none"> • <i>Block</i> will stop all traffic from the application and log all occurrences. • <i>Monitor</i> allows the application traffic to flow normally and log all occurrences. • <i>Reset</i> will reset the network session the application is using. <p>If you set the action to <i>Monitor</i>, you have the option of enabling traffic shaping for the application or applications specified in this application list entry. For more information about application control traffic shaping, see “Application traffic shaping” on page 1067.</p>
Session TTL	Enable <i>Session TTL</i> to specify a time-to-live value for the session, in seconds. If this option is not enabled, the TTL defaults to the setting of the CLI command <code>config system session-ttl</code> .
Packet Logging	<p>Select to enable packet logging for the filter.</p> <p>When you enable packet logging on a filter, the unit saves a copy of the packets that match any signatures included in the filter. The packets can be analyzed later.</p> <p>For more information about packet filtering, see “Viewing and saving logged packets” on page 1144</p>

5 Select OK.

The filter is created and added to the filter list.

Understanding the default application sensor

A default application sensor is provided with your FortiGate unit. You can use it as provided, or modify it as required.



Before using the default application sensor, examine it closely to ensure you understand how it works.

monitor-all

This sensor allows all application traffic and enables the application control monitoring for all traffic.

Viewing and searching the application list

Go to *UTM Profiles > Application Control > Application List* to view the list of applications the FortiGate unit recognizes. You may find applications by paging manually through the list, apply filters, or by using the search field.

Searching manually

Applications are displayed in a paged list, with 50 applications per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

Applying application list filters

You can enter criteria for one or more columns, and only the applications matching all the conditions you specify will be listed.

To apply filters

- 1 Go to *UTM Profiles > Application Control > Application List*.
- 2 Select *Filter Settings*.
- 3 Select *Add New Filter*.
- 4 Select column by which to filter.
- 5 Select the item or items by which to filter.
- 6 Continue to add more filters to narrow your search, if required.
- 7 Select *OK*.

The options available to you will vary by column. For example, Category allows you to choose one option from a list, while Behavior allows you to select multiple items. Filtering by name allows you to enter a text string and all application names containing the string will be displayed.

Using the search field

To use the search field, located above the application list, start typing any portion of the application name. Application names matching the text you enter are displayed in a drop-down list. A maximum of ten matches are displayed at a time.

Select an application from the drop-down list to display its application list entry.

Application traffic shaping

You can apply traffic shaping for application list entries you configure to pass. Traffic shaping enables you to limit or guarantee the bandwidth available to the application or applications specified in an application list entry. You can also prioritize traffic by using traffic shaping.

When the action is set to *Monitor*, two options appear: *Traffic Shaping* and *Reverse Direction Traffic Shaping*. When enabled, you can select traffic shapers configured in *Firewall Objects > Traffic Shaper*.

You can create or edit traffic shapers by going to *Firewall Objects > Traffic Shaper > Shared*. Per-IP traffic shapers are not available for use in application traffic shaping.

For more information about traffic shaping, see [“Traffic shaping methods” on page 2261](#).

Enabling application traffic shaping

Enabling traffic shaping in an application sensor involves selecting the required shaper. You can create or edit shapers in *Firewall Objects > Traffic Shaper > Shared*.

To enable traffic shaping

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.
- 3 Select the application control list entry and choose *Edit*.
- 4 Select *Traffic Shaping* and choose the required traffic shaper from the list.
If the action is set to *Block*, the traffic shaping option is not available. Only allowed traffic can be shaped.
- 5 Select *Reverse Direction Traffic Shaping* and choose the required traffic shaper from the list if traffic flowing in the opposite direction also requires shaping.
- 6 Select *OK*.

Any security policy with this application sensor selected will shape application traffic according to the applications specified in the list entry and the shaper configuration.

Reverse direction traffic shaping

To enable traffic shaping, you must set the action to *Monitor*, enable *Traffic Shaping* and then choose the shaper. This will apply the shaper configuration to the application traffic specified in the entry, but only in the direction as specified in the security policy in which the application sensor is selected. To shape traffic travelling in the opposite direction, enable *Reverse Direction Traffic Shaper*.

For example, if you find that your network bandwidth is being overwhelmed by streaming HTTP video, one solution is to limit the bandwidth by applying a traffic shaper to an application control entry that allows the HTTP.Video application. Your users access the Web using a security policy that allows HTTP traffic from the internal interface to the external interface. Firewall policies are required to initiate communication so even though web sites respond to requests, a policy to allow traffic from the external interface to the internal interface is not required for your users to access the Web. The internal to external policy allows them to open communication sessions to web servers, and the external servers can reply using the existing session.

If you enable *Traffic Shaping* and select the shaper in an application sensor specified in the security policy, the problem will continue. The reason is the shaper you select for *Traffic Shaping* is applied only to the application traffic moving in the direction stated in the security policy. In this case, that is from the internal interface to the external interface. The security policy allows the user to visit the web site and start the video, but the video itself is streamed from the server to the user, or from the external interface to the internal interface. This is the reverse of the direction specified in the security policy. To solve the problem, you must enable *Reverse Direction Traffic Shaping* and select the shaper.

Shaper re-use

Shapers are created independently of firewall policies and application sensors so you are free to reuse the same shapers in multiple list entries and policies. Shared shapers can be configured to apply separately to each security policy or across all policies. This means that if a shaper is configured to guaranteed 1000 KB/s bandwidth, each security policy using the shaper will have its own 1000 KB/s reserved, or all of the policies using the shaper will share a pool of 1000 KB/s, depending on how it is configured.

The same thing happens when a shaper is used in application sensors. If an application sensor using a shaper is applied to two separate policies, how the bandwidth is limited or guaranteed depends on whether the shaper is set to apply separately to each policy or across all policies. In fact, if a shaper is applied directly to one security policy, and it is also included in an application sensor that is applied to another security policy, the same issue occurs. How the bandwidth is limited or guaranteed depends on the shaper configuration.

If a shaper is used more than once within a single application sensor, all of the applications using the shaper are restricted to the maximum bandwidth or share the same guaranteed bandwidth.

For example, you want to limit the bandwidth used by Skype and Facebook chat to no more than 100 KB/s. Create a shaper, enable *Maximum Bandwidth*, and enter 100. Then create an application sensor with an entry for Skype and another entry for Facebook chat. Apply the shaper to each entry and select the application sensor in the security policy that allows your users to access both services.

This configuration uses the same shaper for each entry, so Skype **and** Facebook chat traffic are limited to no more than 100 KB/s in total. That is, traffic from both applications is added and the total is limited to 100 KB/s. If you want to limit Skype traffic to 100 KB/s and Facebook chat traffic to 100 KB/s, you must use separate shapers for each application control entry.

Application control monitor

The application monitor enables you to gain an insight into the applications generating traffic on your network. When monitor is enabled in an application sensor entry and the list is selected in a security policy, all the detected traffic required to populate the selected charts is logged to the SQL database on the FortiGate unit hard drive. The charts are available for display in the executive summary section of the log and report menu.



Because the application monitor relies on a SQL database, the feature is available only on FortiGate units with an internal hard drive.

While the monitor charts are similar to the top application usage dashboard widget, it offers several advantages. The widget data is stored in memory so when you restart the FortiGate unit, the data is cleared. Application monitor data is stored on the hard drive and restarting the system does not affect old monitor data.

Application monitor allows you to choose to compile data for any or all of three charts: top ten applications by bandwidth use, top ten media users by bandwidth, and top ten P2P users by bandwidth. Further, there is a chart of each type for the traffic handled by each security policy with application monitor enabled. The top application usage dashboard widget shows only the bandwidth used by the top applications since the last system restart.

Enabling application control monitor

Once you have configured and enabled application control, you can enable application monitor. There are three steps, as detailed below: enabling application monitor in an application sensor, selecting the charts in the security policy, and displaying the charts in the Executive Summary.

To enable application control monitor in an application sensor

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select an application sensor from the drop-down list in the Edit Application Sensor window title bar.
- 3 Select *Enable Monitoring*.
- 4 Select *OK*.

With application control monitoring enabled, the FortiGate unit begins collecting data for the applications specified in the application sensor from the traffic handled by all policies using the list. If you require monitoring in other application sensors, follow the same procedure to enable it in each sensor.

To configure the charts for which data is collected

- 1 Go to *Policy > Policy > Policy*.
- 2 Select the security policy in which the application sensor is selected and choose *Edit*. Note the security policy ID number.
- 3 Under *UTM*, the *Enable Application Control* selection has three new options, one for each chart type. Select one or more chart types.
- 4 Select *OK*.
- 5 If you have the application sensor specified in multiple firewall policies, repeat this procedure for each policy.

To display the application monitor charts

- 1 Go to *Log&Report > Report Access > Executive Summary*.
- 2 Select *Add Widget*.

- 3 Select the chart you want from the *Widgets* list.

The three application monitor charts correspond to the three chart selections in the security policy. They are listed in the list as:

- `top10-application-bw-X-0`
- `top10-media-user-X-0`
- `top10-p2p-user-bw-X-0`

If you have application monitor enabled in multiple firewall policies, one chart of each type per policy will be available for you to choose. The 'X' in the chart name is the security policy number.

- 4 Select a *Daily* or *Weekly* schedule. The chart will display the data collected from only the current day or current week, depending on the setting. The chart will be reset daily on the hour specified, or weekly on the hour and day specified.
- 5 Select *OK*.

Application control packet logging

Packet logging saves the network packets that application control identifies application traffic with. These packets can be used to trouble-shoot false positives or for forensic investigation. The FortiGate unit saves the logged packets to the attack log, wherever the logs are configured to be stored, whether memory, internal hard drive, a FortiAnalyzer unit, or the FortiGuard Analysis and Management Service.

You can enable packet logging in individual application list entries. Use caution in enabling packet logging. Application sensor entries configured with few restrictions can contain hundreds of applications, potentially resulting in a flood of saved packets. This would take up a great deal of space, require time to sort through, and consume considerable system resources to process. Packet logging is designed as a focused diagnostic tool and is best used with a narrow scope.



Although logging to multiple FortiAnalyzer units is supported, packet logs are not sent to the secondary and tertiary FortiAnalyzer units. Only the primary unit receives packet logs.

To enable application control packet logging

- 1 Create an entry in an application sensor. For more information, see [“Adding applications to an application sensor” on page 1062](#).
- 2 Before saving the entry, select *Packet Log*.
- 3 Select the application sensor in the security policy that allows the network traffic the FortiGate unit will examine for the application or applications.

For information on viewing and saving logged packets, see [“Viewing and saving logged packets” on page 1144](#).

Application considerations

Some applications behave differently from most others. You should be aware of these differences before using application control to regulate their use.

IM applications

Application control regulates most instant messaging applications by preventing or allowing user access to the service. Selecting *Block Login* will not disconnect users who are logged in when the change is made. Once users log out, however, they will not be able to log in again.

Skype

Based on the NAT firewall type, Skype takes advantage of several NAT firewall traversal methods, such as STUN (Simple Traversal of UDP through NAT), ICE (Interactive Connectivity Establishment) and TURN (Traversal Using Relay NAT), to make the connection.

The Skype client may try to log in with either UDP or TCP, on different ports, especially well-known service ports, such as HTTP (80) and HTTPS (443), because these ports are normally allowed in firewall settings. A client who has previously logged in successfully could start with the known good approach, then fall back on another approach if the known one fails.

The Skype client could also employ Connection Relay. This means if a reachable host is already connected to the Skype network, other clients can connect through this host. This makes any connected host not only a client but also a relay server.

Application control examples

Blocking all instant messaging

Instant messaging use is not permitted at the Example Corporation. Application control helps enforce this policy.

First you will create an application sensor with a single entry that includes all instant messaging applications. You will set the list action to block.

To create the application sensor

- 1 Go to *UTM Profiles > Application Control > Application Sensor*.
- 2 Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.
- 3 In the *Name* field, enter `no IM` for the application sensor name.
- 4 Select *OK*.
- 5 Select the *Create New* drop-down icon in the sensor and choose *Entry*.
- 6 For *Category*, select *im*.
- 7 For *Action*, select *Block*.
- 8 Select *OK* to save the new list entry.
- 9 Select *OK* to save the list.

Next you will enable application control and select the list.

To enable application control and select the application sensor

- 1 Go to *Policy > Policy > Policy*.
- 2 Select the security policy that allows the network users to access the Internet and choose *Edit*.
- 3 Enable *UTM*.
- 4 Select *Enable Application Control*.

5 Select the *no IM* application sensor.

6 Select OK.

No IM use will be allowed by the security policy. If other firewall policies handle traffic that users could use for IM, enable application control with the *no IM* application sensor for those as well.

Allowing only software updates

Some departments at Example Corporation do not require access to the Internet to perform their duties. Management therefore decided to block their Internet access. Software updates quickly became an issue because automatic updates will not function without Internet access and manual application of updates is time-consuming.

The solution is configuring application control to allow only automatic software updates to access the Internet.

To create an application sensor — web-based manager

1 Go to *UTM Profiles > Application Control > Application Sensor*.

2 Select the *Create New* icon in the title bar of the *Edit Application Sensor* window.

3 In the *Name* field, enter `Updates_Only` as the application sensor name.

4 Select OK.

5 Select the *Create New* drop-down icon in the sensor and choose *Entry*.

6 Select *update* from the *Category* list.

7 Select *Pass* from the *Action* list.

8 Select OK to save the entry.

This application list entry will allow all software update application traffic.

9 Select the *All Other Known Applications* entry.

10 Select *Edit*.

11 Select *Block* from the *Action* list.

12 Select OK.

This application list entry will block all traffic from recognized applications that are not specified in this application sensor.

13 Select the *All Other Unknown Applications* entry.

14 Select *Edit*.

15 Select *Block* from the *Action* list.

16 Select OK.

This application list entry will block all traffic from applications that are not recognized by the application control feature.

17 Select OK.

18 Select OK to save the application sensor.

To create an application sensor — CLI

```
config application list
edit Updates_Only
config entries
edit 1
set category 17
```

```
        set action pass
    end
    set other-application-action block
    set unknown-application-action block
end
```

Selecting the application sensor in a security policy

An application sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an application sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the application sensor in a security policy — web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select a policy.
- 3 Select the *Edit* icon.
- 4 Enable *UTM*.
- 5 Select the *Enable Application Control* option.
- 6 Select the *Updates_only* list.
- 7 Select *default* from the *Protocol Options* list.

Application control can not be enabled without selecting a protocol options profile. A default profile is provided.

- 8 Select *OK*.

To select the application sensor in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set application-list Updates_Only
  end
```

Traffic handled by the security policy you modified will be scanned for application traffic. Software updates are permitted and all other application traffic is blocked.

Application Control interface reference

This section describes how to configure the application control options associated with firewall policies.

By using the UTM feature's application control, the unit can detect and take action against network traffic depending on the application generating the traffic. Based on Intrusion Protection protocol decoders, application control is a more user-friendly and powerful way to use Intrusion Protection features to log and manage the behavior of application traffic passing through the unit. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic even if the traffic uses non-standard ports or protocols.

The unit can recognize the network traffic generated by a large number of applications. You can create application control black/white lists that specify the action to take with the traffic of the applications you need to manage and the network on which they are active. Add application control lists to firewall policies applied to the network traffic you need to monitor.

Fortinet is constantly increasing the list of applications that application control can detect by adding applications to the [FortiGuard Application Control Database](#). Because intrusion protection protocol decoders are used for application control, the application control database is part of the [FortiGuard Intrusion Protection System Database](#) and both of these databases have the same version number.

You can find the version of the application control database that is installed on your unit, by going to the *License Information* dashboard widget and find IPS Definitions version.

You can go to the [FortiGuard Application Control List](#) to see the complete list of applications supported by FortiGuard. This web page lists all of the supported applications. You can select any application name to see details about the application.

This topic includes the following:

- [Application Sensor](#)
- [Application List](#)



DiffServ is supported per-application and is available only in the CLI.

Application Sensor

Each application control list contains details about the application traffic to be monitored and the actions to be taken when it is detected. An application control list must be selected in a firewall policy to take effect.

There are no default application control lists provided.

The unit examines network traffic for the application entries in the listed order, one at a time, from top to bottom. Whenever a match is detected, the action specified in the matching rule is applied to the traffic and further checks for application entry matches are stopped. Because of this, you can use both actions to create a complex rule with fewer entries.

Application sensor configuration settings

The following are application sensor configuration settings in *UTM Profiles > Application Control > Application Sensor*.

Application Sensor page

Lists each individual black/white list that you created. On this page, you can edit, delete and create a new application sensors.

You are redirected to this page when you select *View List* on the Edit Application Sensor page.

Create New	Creates a new application control sensor. When you select <i>Create New</i> , you are automatically redirected to the New Application Control List page. This page provides a name field and a comment field. You must enter a name to go to the Edit Application Sensor page where you can then configure settings for the new application sensor.
Edit	Modifies settings within an application black/white list. When you select <i>Edit</i> , you are automatically redirected to the Edit Application Control List page.

Delete	Removes the application control black/white list from within the list on the page.
Name	The available application control lists.
# of Entries	The number of application rules in each application control list.
Comments	An optional description of each application control list.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
Edit Application Sensor page Provides settings for configuring the applications for the list. When you are editing a list, you are redirected to this page. Note: Logging is enabled in the CLI.	
Name	If you are editing an existing application control list and want to change the name, enter a new name in the field. You must select <i>OK</i> to save the change.
Comments	If you are editing an existing list and want to change or add a description, enter the new text in the field. You must select <i>OK</i> to save the change.

Create New	<p>Creates a new application entry. When you select <i>Create New</i>, you are automatically redirected to the New Application Filter page.</p> <p>When you select the down arrow beside <i>Create New</i>, you can choose to create either a new application filter or entry.</p> <p>When you select <i>Filter</i>, you are automatically redirected to the New Application Filter page where you can configure a filter.</p> <p>When you select <i>Entry</i>, you are automatically redirected to the New Application Entry page where you can configure an entry.</p>
Edit	<p>Modifies settings within the application control entry. When you select <i>Edit</i> you are automatically redirected to the Edit Application Entry page.</p>
Delete	<p>Removes the application control entry in the list.</p> <p>To remove multiple application control entries from within the list, on the Edit Application Sensor page, in each of the rows of the entries you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all application control entries from the list, on the Edit Application Sensor page, select the check box in the check box column and then select <i>Delete</i>.</p>
Insert	<p>Creates a new application control entry above the entry you highlighted. When you select <i>Insert</i>, you are automatically redirected to the New Application Entry page.</p>
Move To	<p>Moves the application control entry to any position in the list. When you select <i>Move To</i>, the Move Application Control Entry window appears.</p> <p>To move an application control entry, select the new position <i>Before</i> or <i>After</i>, which will place the current entry before or after the entry you enter in the <i>Application ID</i> field. Use the number found in the <i>ID</i> column when entering the new position in the <i>Application ID</i> field.</p>
View Rules	<p>Select to view the rules of an entry. When you select <i>View Rules</i>, the Matched Rules window appears. You can view all the rules associated within that entry in this window.</p>
Page Controls	<p>Use to navigate through the application control entries within an application control list.</p>
ID	<p>The identification number of the entry.</p>
Category	<p>The category indicates the scope of the applications included in the application entry if <i>Application</i> is set to <i>all</i>. For example, if <i>Application</i> is <i>all</i> and <i>Category</i> is <i>toolbar</i>, then all the toolbar applications are included in the application entry even though they are not specified individually.</p> <p>If <i>Application</i> is a single application, the value in <i>Category</i> has no effect on the operation of the application entry.</p>
Vendor	<p>The application's vendor name.</p>
Behavior	<p>The type of behavior chosen.</p>
Technology	<p>The type of technology associated with the application.</p>

Application	<p>The type of application that was chosen.</p> <p>Note: <i>Full List</i> appears in this column only when you choose all applications for a filter. When you select <i>Full List</i>, the Match Rules window appears where you can view all the applications.</p>
Action	If the unit detects traffic from the specified application, the selected action will be taken.
<p>New Application Filter page</p> <p>Provides settings for configuring an application filter to add to the application control sensor.</p> <p>This page appears when you select <i>Create New</i> on the Edit Application Sensor page. If you are on the Application Sensor page, and you select <i>Create New</i>, you will be redirected to the New Application Filter page.</p>	
Category	<p>The applications are categorized by type. If you want to choose an IM application, for example, select the <i>im</i> category, and the application control list will show only the <i>im</i> applications.</p> <p>The <i>Category</i> selection can also be used to specify an entire category of applications. To select all IM applications for example, select the <i>im</i> category, and select <i>all</i> as the application. This specifies all the IM applications with a single application control black/white list entry.</p>
Vendor	Select to either specify a vendor or enable all vendors.
Behavior	When you select <i>Specify</i> , the lists <i>Available</i> and <i>Selected</i> appear. Select the available behavior of the application (for example, encrypted-tunneling) in the <i>Available</i> list and then move it to the <i>Selected</i> list using the -> arrow. Use the same method to remove a behavior from the <i>Selected</i> list except use the <- arrow.
Technology	Select to include all technology or specify the technology. For example, web browsers.
Protocol	Select to include all protocols or specify a particular protocol.
Tags	<p>Adds and displays the tags you create.</p> <p>If tag settings are not available on the web-based manager, you must enable them in <i>System > Admin > Settings</i>.</p>
Applied tags	Displays the current tags configured for the filter.
Add tags	Enter the tag in the field and then select the plus sign (+) beside the field to add the tag to the list in <i>Applied tags</i> .
View Matched rules	Select to view the matched rules.
Action	If the unit detects traffic from the specified application, the selected action will be taken.
Monitor	Select to monitor the filter using traffic shaping or reverse direction traffic shaping or both.
Block	Select to block.

Session TTL	The application's session TTL. If this option is not enabled, the TTL defaults to the setting of the <code>config system session-ttl</code> CLI command.
Packet Logging	Select to log the occurrence of packet logs which concern application control.
New Application Entry page Provides settings for configuring an application entry, to add to the application sensor.	
Category	Select a category from the drop-down list.
Application	Select an application from the drop-down list.
Action	Select either <i>Monitor</i> or <i>Block</i> for the type of action the unit will take. When you select <i>Monitor</i> , the options for applying traffic shaping and reverse direction traffic shaping. Select the check box beside either one or both to enable these options and then select
Session TTL	Enter a number for the session's time to live.
Packet Logging	Select to enable packet logging.

Application List

The application list displays applications, which also shows their popularity and risk. You can view the details of each application by selecting the application's name; this link redirects you to the [FortiGuard Application Control List](#) where the details are given for the application. You can also filter the information that appears in *UTM Profiles > Application Control > Application List*.

You can go to the [FortiGuard Application Control List](#) to see the complete list of applications supported by FortiGuard. This web page lists all of the supported applications. You can select any application name to see details about the application.

Application lists are viewed from *UTM Profiles > Application Control > Application List*.

Application List page Lists the applications that are available on the unit, which includes their category, popularity rating and risk.	
Tags	<p>Select to add or remove tags to the applications in the list.</p> <p>Note: If <i>Tags</i> is not available on the web-based manager, you must enable it in <i>System > Admin > Settings</i>.</p> <p>When you select the down arrow beside <i>Tags</i>, you can add tags or remove tags.</p> <p>To add tags to an application, select the application first then select the down arrow to then select <i>Add Tags</i>. The Add Tags window appears. Enter the tag in the <i>Add tag</i> field and select the plus (+) sign; repeat until all tags are in the Tags to apply list.</p> <p>To remove tags, select the application first, select the down arrow beside Tags, and then select <i>Remove Tags</i>. The Remove Tags window appears. Select the tags that you want removed in the <i>Applied Tags</i> row; repeat until all the tags are in the <i>Tags to remove</i> row. The tags will automatically be put in the <i>Tags to remove</i> row after being selected in the <i>Applied Tags</i> row.</p> <p>If there are tags that you want to add that have been configured for another object, you can add those tags as well to signatures. To apply these other object tags, select the signature first, select the down arrow beside Tags, and then select <i>Add Tags</i>. The Add Tags window appears. Select the tags you want to add in the <i>Click tag to add</i> row. The tags automatically appear in the <i>Tags to apply</i> row. Select <i>OK</i> to add those tags to the application.</p>
Column Settings	Customize the column view. You can select the columns to hide or display them and specify the column display order.

Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filter settings, select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p> <p>The <i>Filters Settings</i> on the Application List page contains <i>Copy to Sensor</i>, which allows you to copy filter settings and apply them to an application sensor.</p> <p>To apply existing filter settings to a sensor, select the down arrow beside <i>Filter Settings</i>, and then select <i>Copy to Sensor</i>. The Select Object window appears. Select the sensor that you want to apply the settings to from the drop-down list. Select <i>OK</i>.</p>
Search	Enter search criteria into the field and then press Enter on your keyboard. Use the <i>Clear All</i> icon beside the field to clear the search results.
Page Controls	Use to navigate through the list to view the applications.
[Total: <maximum number>]	The maximum number of applications that are currently in the FortiGuard Application Control List.
Application Name	The name of the application.
Category	The category that the application is associated with.
Vendor	The type of vendor.
Technology	The type of technology that the application uses. For example, 56.COM using Peer-to-Peer technology.
Protocol	The type of protocol the application uses.
Behavior	The type of behavior that is associated with the application.
Tags	<p>The tags that are associated with that application.</p> <p>If no tag configuration settings display, this indicates that this feature is disabled. You can enable tag configuration settings from <i>System > Admin Settings</i>.</p>



DoS policy

Denial of Service (DoS) policies are primarily used to apply DoS sensors to network traffic based on the FortiGate interface it is entering as well as the source and destination addresses. DoS sensors are a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and behavior. A common example of anomalous traffic is the denial of service attack. A denial of service occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system, so that legitimate users can no longer use it.

This section describes how to create and configure DoS sensors and policies to protect the publicly accessible servers on your network.

The following topics are included in this section:

- [DoS policy concepts](#)
- [Enable DoS](#)
- [DoS example](#)

DoS policy concepts

DoS policies are similar to firewall policies except that instead of defining the way traffic is allowed to flow, they keep track of certain traffic patterns and attributes and will stop traffic displaying those attributes. Further, DoS policies affect only incoming traffic on a single interface. You can further limit a DoS policy by source address, destination address, and service.

DoS policies examine network traffic very early in the sequence of protective measures the FortiGate unit deploys to protect your network. Because of this early detection, DoS policies are a very efficient defence that uses few resources. Denial of service attacks, for example, are detected and its packets dropped before requiring security policy look-ups, antivirus scans, and other protective but resource-intensive operations. For more information about DoS attacks, see [“Defending against DoS attacks” on page 888](#).

Enable DoS

A DoS policy examines network traffic arriving at an interface for anomalous patterns usually indicating an attack. Enable DoS sensors to protect your FortiGate unit from attack. To apply a DoS policy, you must follow the steps below in sequence:

- 1 Create a DoS sensor.
- 2 Create a DoS policy
- 3 Apply the DoS sensor to the DoS policy.

Creating and configuring a DoS sensor

Because an improperly configured DoS sensor can interfere with network traffic, no DoS sensors are present on a factory default FortiGate unit. You must create your own and then enable them before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.



It is important to know normal and expected network traffic before changing the default anomaly thresholds. Setting the thresholds too low could cause false positives, and setting the thresholds too high could allow otherwise avoidable attacks.

To create a DoS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > DoS Sensor*.
- 2 Select *Create New*.
- 3 In the *Name* field, enter the name of the DoS sensor.
- 4 Optionally, enter a description of the DoS sensor in the *Comment* field.
- 5 Select *OK*.

The DoS sensor is created and the sensor configuration window appears. However, a newly created DoS sensor contains default values which may not be appropriate for your network. You can adjust these values by configuring the DoS sensor thresholds.

To configure a DoS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > DoS Sensor*.
- 2 Select the DoS sensor you want to configure and choose *Edit*.
- 3 The DoS sensor configuration window appears.

The *Anomalies Configuration* table lists 12 types of network anomalies.

Anomaly	Description
tcp_syn_flood	If the SYN packet rate of new TCP connections, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
tcp_port_scan	If the SYN packet rate of new TCP connections, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
tcp_src_session	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.
tcp_dst_session	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.
udp_flood	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
udp_scan	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.

udp_src_session	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.
udp_dst_session	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.
icmp_flood	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
icmp_sweep	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed. The threshold is expressed in packets per second.
icmp_src_session	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.
icmp_dst_session	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.

- 4 Select *Enable* to have the FortiGate unit examine traffic for the anomaly.
- 5 Select *Logging* to create an entry in the attack log if the anomaly is detected.
- 6 Select an *Action* for the anomaly. By default, the action is *Pass*, which allows the traffic containing the anomaly to pass uninterrupted. If set to *Block*, the anomalous traffic is blocked and will not flow through the FortiGate unit.

With a Fortinet security processing module installed, FortiGate units that support these modules offer a third action for the `tcp_syn_flood` threshold. In addition to *Block* and *Pass*, you can choose to *Proxy* connect attempts when their volume exceeds the threshold value. When the `tcp_syn_flood` threshold action is set to *Proxy*, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets with a hardware accelerated SYN proxy to determine whether the connection attempts are legitimate or a SYN flood attack. Legitimate connections are allowed while an attack is blocked.



Because DoS sensors are configured before being applied to an interface, you can assign a DoS sensor with the *Proxy* action to an interface that does not have hardware SYN proxy support. In this circumstance, the *Proxy* action is invalid and a *Pass* action will be applied.

- 7 Set the *Threshold* value for the anomaly. See the table in step 3 for details about the threshold values for each anomaly.
- 8 Select *OK*.

Creating a DoS policy

DoS policies examine network traffic entering an interface. The DoS sensor specified in the DoS policy allows you to limit certain anomalous traffic to protect against attacks.

To create a DoS policy

- 1 Go to *Policy > Policy > DoS Policy* and select *Create New*.

- 2 For *Source Interface/Zone*, select the interface on which the DoS policy will examine incoming traffic.
- 3 For *Source Address*, select the address or address group that defines the source addresses of the traffic the DoS policy will examine. Network traffic from addresses not included in the selected address group is ignored by this DoS policy.
- 4 For *Destination Address*, select the address or address group that defines the destination addresses of the traffic the DoS policy will examine. Network traffic to addresses not included in the selected address group is ignored by this DoS policy.
- 5 For *Service*, select the type of network traffic the DoS policy will examine. Protocols not included in the selected service or service group are ignored by this DoS policy.
- 6 Select the *DoS Sensor* check box and choose the required sensor from the list.
- 7 Select *OK*.

Apply an IPS sensor to a DoS policy

Although IPS sensors are usually applied to firewall policies, you can also apply them to DoS policies by using CLI commands. There are two reasons you might want to apply an IPS sensor to a DoS policy:

- If you want to have all traffic coming into one FortiGate unit interface checked for the signatures in an IPS sensor, it is simpler to apply the IPS sensor once to a DoS policy. In a complex configuration, there could be many policies controlling the traffic coming in on a single interface.
- The operations in a DoS policy occur much earlier in the sequence of operations performed on incoming traffic. This means that IPS examination of traffic occurs much sooner if the IPS sensor is applied to a DoS policy. Fewer system resources are used because signatures set to block traffic will take effect before security policy checking and all of the scans specified in the security policy.

The CLI command for configuring DoS policies is `config firewall interface-policy`. The following command syntax shows how to add an example IPS sensor called `all_default_pass` to a DoS policy with policy ID 5 that was previously added from the web-based manager.

```
config firewall interface-policy
edit 5
    set ips-sensor-status enable
    set ips-sensor all_default_pass
end
```

DoS example

The Example.com corporation installed a web server and connected it to Port5 on its FortiGate unit. To protect against denial of service attacks, you will configure and apply a DoS sensor to protect the web server.

To create the DoS sensor

- 1 Go to *UTM Profiles > Intrusion Protection > DoS Sensor*.
- 2 Select *Create New*.
- 3 Enter `Web Server` in the *Name* field.
- 4 In the *Anomalies Configuration* table, select the *Enable* check box in the table heading. This enables all the anomalies with a single selection.

- 5 Select *OK* to save the new DoS policy.

As suggested in “[Defending against DoS attacks](#)” on page 888, the IT administrators will run the DoS policy with logging enabled and the anomaly actions set to *Pass* until they determine the correct threshold values for each anomaly.

To create a DoS policy

- 1 Go to *Policy > Policy > DoS Policy*.
- 2 Select *Create New*.
- 3 In the *Source Interface/Zone* field, select *Port1* which is the interface connected to the Internet.
- 4 In the *Source Address* field, select *all*.
- 5 In the *Destination Address* field, select *all*.
If there were more than one publicly accessible server connected to the FortiGate unit, you would specify the address of the web server in this field.
- 6 In the *Service* field, select *ANY*.
- 7 Select the *DoS Sensor* check box and choose *Web Server* from the list.
- 8 Select *OK* to save the DoS policy.

The DoS policy will monitor all network traffic entering Port1 and log the violations if the thresholds in the *Web Server* DoS sensor are exceeded.

DoS Policy interface reference

The DoS security policy list displays the DoS security policies in their order of matching precedence for each interface, source/destination address pair, and service.

If virtual domains are enabled on the unit, DoS security policies are configured separately for each virtual domain; you must access the VDOM before you can configure its security policies.

You can add, delete, edit, and re-order security policies in the DoS policy list. DoS policy order affects security policy matching. As with security policies, DoS security policies are checked against traffic in the order in which they appear in the DoS policy list, one at a time, from top to bottom. When a matching security policy is discovered, it is used and further checking for DoS policy matches are stopped.

The DoS policy configuration allows you to specify the interface, a source address, a destination address, and a service. All of the specified attributes must match network traffic to trigger the security policy.

DoS policies configuration settings

The following are DoS policy configuration settings in *Policy > Policy > DoS Policy*.

DoS Policy page Lists each individual DoS policy that you created. On this page, you can edit, delete or create a new DoS policy.	
Create New	<p>Adds a new DoS policy. Select the down arrow beside <i>Create New</i> to add a new section to the list to visually group the security policies. When you select <i>Create New</i>, (or the <i>Policy</i> option from the down arrow's drop-down list), you are automatically redirected to the New Policy page.</p> <p>When you select <i>Section Title</i>, the Section Title window appears. Enter the section title name in the <i>Name</i> field and then enter the DoS policy ID number of the security policy that you want the section title to come before in the <i>Starting Policy ID</i> field. For example, <i>branch_office_dos</i> comes before the DoS policy ID 2, so 2 is entered in the <i>Starting Policy ID</i> field.</p>
Edit	<p>Modifies settings within the security policy. When you select <i>Edit</i>, you are automatically redirected to the Edit Policy page.</p> <p>When you select the down arrow beside <i>Edit</i>, you can select to modify the security policy (<i>Edit Policy</i>), disable a security policy (<i>Disable</i>), or enable a security policy (<i>Enable</i>).</p>
Delete	<p>Removes a security policy from the list on the DoS Policy page.</p> <p>To remove multiple DoS security policies from within the list, on the DoS Policy page, in each of the rows of the security policies you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all DoS security policies from the list, on the DoS Policy page, select the check box in the check box column, and then select <i>Delete</i>.</p>
Move To	<p>Moves the corresponding security policy before or after another security policy in the list.</p> <p>When you select <i>Move To</i>, the Move Policy window appears. To move a security policy, select the new position either <i>Before</i> or <i>After</i>, which will place the current entry before or after the security policy number that you enter in the (<i>Policy ID</i>) field. Enter the security policy ID number and then select OK. For example, policy ID 2 is moved after policy ID 5.</p>
Insert	<p>Inserts a new security policy above the corresponding security policy. When you select Insert, the New Policy window appears.</p>

Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filter settings, select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p>
Column Settings	Customize the table view. You can select the columns to hide or display and specify the column displaying order in the table.
Section View	Select to display security policies organized by interface.
Global View	Select to list all security policies in order according to a sequence number.
ID	A unique identifier for each security policy. Policies are numbered in the order they are created.
Source	The source address or address group to which the security policy applies.
Destination	The destination address or address group to which the security policy applies.
Service	The service to which the security policy applies.
DoS Sensor	The DoS sensor selected in this security policy.
Interface	The interface to which this security policy applies.
Status	When selected, the DoS security policy is enabled. Clear the check box to disable the security policy.
New Policy page Provides settings for configuring a DoS security policy. When you select <i>Create New</i> on the DoS Policy page, you are automatically redirected to this page.	
Source Interface/Zone	The interface or zone to be monitored.
Source Address	Select an address, address range, or address group to limit traffic monitoring to network traffic sent from the specified address or range. Select <i>Multiple</i> to include multiple addresses or ranges. You can also select <i>Create New</i> to add a new address or address group.
Destination Address	Select an address, address range, or address group to limit traffic monitoring to network traffic sent to the specified address or range. Select <i>Multiple</i> to include multiple addresses or ranges. You can also select <i>Create New</i> to add a new address or address group.

Service	Select a firewall pre-defined service or a custom service to limit traffic monitoring to only the selected service or services. You can also select <i>Create New</i> to add a custom service.
DoS Sensor	Select and specify a DoS sensor to have the FortiGate unit apply the sensor to matching network traffic. You can also select <i>Create New</i> to add a new DoS Sensor. See “Creating and configuring a DoS sensor” on page 1082 .



Endpoint Control and monitoring

This section describes the Endpoint Control feature and how to configure it.

The following topics are included in this section:

- [Endpoint Control overview](#)
- [Configuring FortiClient required version and download location](#)
- [About application detection and control](#)
- [Creating an endpoint control profile](#)
- [Enabling Endpoint Control in firewall policies](#)
- [Monitoring endpoints](#)
- [Modifying Endpoint Security replacement pages](#)
- [Example](#)

Endpoint Control overview

Endpoint Control ensures that workstation computers (endpoints) meet security requirements, otherwise they are not permitted access. Endpoint Control can enforce

- use of FortiClient Endpoint Security
- use of a licensed version of FortiClient Endpoint Security
- use of FortiClient firewall
- use of FortiClient antivirus protection
- use of FortiClient web content filtering
- use of up-to-date FortiClient antivirus signatures
- installation or running of specific applications
- absence or non-use of specific applications

Non-compliant endpoints can be either warned or blocked.

Of the features listed above, enforcement of FortiClient licensing and FortiClient web content filtering can be configured only through the CLI using the `config endpoint-control profile` command.

Endpoint Control settings are grouped into one or more Endpoint Control profiles. You enable Endpoint Security in firewall policies and select an Endpoint Control profile.

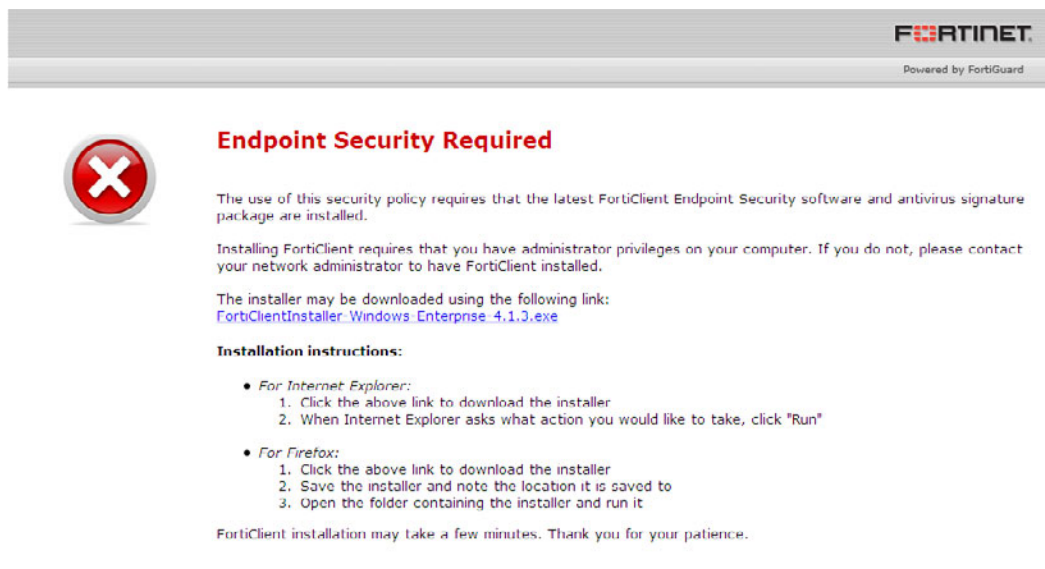
User experience

Endpoint Control applies to users attempting to make a connection that is controlled by a firewall policy with Endpoint Security enabled. The user of a non-compliant endpoint using a web browser receives a replacement message HTML page from the FortiGate unit. The message explains the non-compliance. Depending on the endpoint profile, the user may be allowed to continue or is blocked from further access. For information about modifying these replacement pages, see [“Modifying Endpoint Security replacement pages” on page 1100](#).

FortiClient version non-compliance

If the *FortiClient* application detection entry in the Endpoint Control profile has either the Warn or Block action selected, the user sees a message like this:

Figure 90: Default FortiClient non-compliance message



If there is a FortiClient installer available for the user's endpoint computer, a link is provided to download the installer from the location defined in *UTM Profiles > Endpoint Control > Client Installers*. If there is no installer available, the user is asked to contact the network administrator.

If the action on the *FortiClient* application detection entry is Warn, there is a link at the bottom of the page to enable the user to continue to the requested web site without installing FortiClient Endpoint Security. Otherwise, the same message will be displayed for every connection attempt where Endpoint Security is in effect until the user installs FortiClient Endpoint Security.

Blocked user - FortiClient features not compliant

If the Endpoint Control profile has the *Block* action selected for FortiClient features such as antivirus or firewall, the FortiGate unit sends a message like this to the user's browser.:

Figure 91: Endpoint blocked message



The user needs to resolve the listed issues and retry the connection.

Configuration overview

Endpoint Control requires that all hosts using the firewall policy have the FortiClient Endpoint Security application installed. Make sure that all hosts affected by this policy are able to install this application. Currently, FortiClient Endpoint Security is available for Microsoft Windows (2000 and later) only.

To set up Endpoint Control, you need to

- Enable Central Management by the FortiGuard Analysis & Management Service if you will use FortiGuard Services to update the FortiClient application or antivirus signatures. You do not need to enter account information. See “Centralized Management” in the System Administration chapter of this Handbook.
- Configure the minimum required version of FortiClient and the location from which non-compliant endpoints can download the FortiClient installer. See “[Configuring FortiClient required version and download location](#)” on page 1091.
- Create an endpoint control profile or use a predefined profile. See “[About predefined profiles](#)” on page 1094.
- If needed, modify the profile’s predefined FortiClient application detection rules. You can select the action to take on endpoints that do not have FortiClient Endpoint Security and you can set conditions and actions regarding FortiClient features.
- Configure application detection rules for other applications that are required, allowed, or not allowed on endpoints. See “[About application detection and control](#)” on page 1093.
- Enable Endpoint Security in firewall policies, selecting the appropriate Endpoint Control profile.



You cannot enable Endpoint in firewall policies if *Redirect HTTP Challenge to a Secure Channel (HTTPS)* is enabled in *User > User > Authentication*.

- Optionally, modify the inactivity timeout for endpoints. The default is 5 minutes. After that time period, the FortiGate unit rechecks the endpoint for Endpoint compliance. To change the timeout, adjust the `compliance-timeout` value in the `config endpoint-control settings` CLI command.
- Optionally, modify the *Endpoint NAC Download Portal* and the *Endpoint NAC Recommendation Portal* replacement messages.

Configuring FortiClient required version and download location

The Endpoint Control feature can set a minimum FortiClient version that endpoints are required to run. To make this policy easy for users, you can configure a download source for the FortiClient installer.

Configuring FortiClient requirement and download location - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Client Installers*.

Figure 92: Configuring FortiClient version requirements and installer source

FortiClient Endpoint Security

Information

FortiGuard Availability	✓
FortiClient Endpoint Versions	
Windows Installer	4.1.3 (Updated 2010-04-07) [Download]
AV Signature Package	11.670 (Updated 2010-04-07)
Application Signature Package	1.169 (Updated 2010-04-07)
FortiClient Downloads	0

[Update Now](#)

FortiClient Installer Download Location

The remediation portal will direct users to download from:

☒ FortiGuard Distribution Network
☐ This FortiGate
☐ Custom URL:

FortiClient Version

☒ Enforce Minimum Version (If enabled in Endpoint Profiles)

[Apply](#)

2 Do one of the following:

- Select *FortiGuard Distribution Network*. FortiGuard must be configured on the FortiGate unit.
- Select *This FortiGate*. Users can download a FortiClient installer file from this FortiGate unit. This option is available only on FortiGate models that support upload of FortiClient installer files.
- Select *Custom URL*. Enter the URL from which users can download the FortiClient installer.



Select *This FortiGate* or *Custom URL* if you want to provide a customized FortiClient application. This is required if a FortiManager unit will centrally manage FortiClient applications. For information about customizing the FortiClient application, see the [FortiClient Administration Guide](#).

3 Optionally, select *Enforce Minimum Version* and select the minimum acceptable version number or *Latest Available* for the FortiClient Endpoint Security application.

The list contains the FortiClient versions available from the selected *FortiClient Installer Download Location*.

Fortinet recommends that administrators wait for a reasonable period of time after deploying a FortiClient version update before updating the minimum version required to the most recent version. This gives users some time to install the update.

Configuring FortiClient requirement and download location - CLI

In this example, users are required to have FortiClient version 4.1.3 or later. FortiGuard provides the FortiClient installer.

```
config endpoint-control settings
  set enforce minimum-version enable
  set version-check minimum
  set version 4.1.3
  set download-location fortiguard
end
```

About application detection and control

In firewall policies you can select an endpoint control profile. The application detection list within the endpoint control profile allows or denies endpoint access to the network based on the applications that are installed or running on the endpoint.

The application detection list contains rules specific to individual applications, application vendors, and application categories. A rule tests for a particular condition of the application on the endpoint, which can be any of the following:

- *Installed* — application is installed and may or may not be currently running
- *Not Installed* — application is not installed
- *Running* — application is installed and currently running
- *Not Running* — application is not currently running or is not installed

The rule determines the action to take when the specified application matches the condition. The possible actions are:

- *Allow* — Allow the endpoint to connect.
- *Block* — Block the endpoint.
- *Warn* — Warn the endpoint, but then allow the user to connect.
- *Monitor* — Allow the endpoint to connect and include this endpoint's information in statistics and logs on the Endpoint Monitor page.

FortiClient application rules

There are three application rules for FortiClient that are present in every endpoint control profile: FortiClient, FortiClient AV, and FortiClient Firewall. You can edit, but not delete, these entries.

- **FortiClient** — **Select** the Block, Warn, or Monitor action to apply if the FortiClient application is not installed or not running on the endpoint.
- **FortiClient AV** — **Select** whether to allow or block endpoints that are not running the antivirus feature of the FortiClient application. Optionally, you can also apply this rule to endpoints with an outdated FortiClient antivirus database.
- **FortiClient Firewall** — **Select** whether to allow or block endpoints that are not running the firewall feature of the FortiClient application.

Other application rules

Application detection rules (entries) are based on application signatures provided by FortiGuard Services. You create your application detection list entries by selecting applications from FortiGuard-supplied lists of categories, vendors, and application names. To view application information from FortiGuard services, go to *UTM Profiles > Endpoint Control > Application Database*.

An application detection rule checks applications against the database from the top down until it finds a match. Specific entries, such as those that list one particular application, should precede more general entries, such as those that match all applications of a particular category.

The All application rule

At the bottom of every application detection list is the All rule. This specifies the action to apply to an endpoint with an application installed that does not match a rule higher in the list. The default action is Monitor. If you select Block, endpoints can have only a specific set of applications installed and are denied access if any other applications are installed.

About predefined profiles

Each VDOM has the following default endpoint profiles:

Enforce_FortiClient_AV — **blocks** endpoints without FortiClient Endpoint Security or not running FortiClient antivirus protection. Endpoints with other applications installed are monitored.

P2P_application_detection — **blocks** endpoints running peer-to-peer file sharing. Users whose endpoint does not have FortiClient Endpoint Security installed receive a warning.

Recommend_FortiClient — **monitors** endpoints and presents a warning to users whose endpoint does not have FortiClient Endpoint Security installed.

You can modify or delete these endpoint profiles.

Creating an endpoint control profile

An endpoint profile defines requirements for FortiClient Endpoint Security and other applications on endpoints. The profile is selected in firewall policies and applies to all users of the firewall policy.

To create an endpoint control profile - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and select *Create New*.
- 2 Enter a *Name* and optionally *Comments* for the profile, then select *OK*.
The profile opens, showing the application detection list. There are default entries for FortiClient Endpoint Security with default settings.
- 3 Configure FortiClient-related entries as needed.
- 4 Create additional application detection entries as needed.
- 5 Select *OK*.

To create an endpoint control profile - CLI

```
config endpoint-control profile
  edit profile1
end
```

Setting endpoint FortiClient requirements

Every endpoint control profile requires that FortiClient Endpoint Security must be installed on all endpoints. You can choose whether non-compliant endpoints are blocked, warned, or simply monitored. By default they are blocked.

You can choose whether to require the use of FortiClient antivirus, firewall, or web filtering features. By default, endpoints are not required to use these features.

You can also choose whether to require each endpoint's FortiClient application to be licensed. This is not required by default.

To set the action for endpoints not running FortiClient - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open the profile for editing.

- 2 Edit the *FortiClient* application detection entry and select the required action: *Block*, *Warn*, or *Monitor*.

For information about these actions, see [“About application detection and control” on page 1093](#).

To set the action for endpoints not running FortiClient - CLI

In this example, endpoints that do not have FortiClient Endpoint Security installed will be blocked. The other options for `recommendation-disclaimer` are `enable` to warn the endpoint and `skip` to simply monitor the endpoint.

```
config endpoint-control profile
edit profile1
set recommendation-disclaimer disable
end
```

To require endpoints to use FortiClient antivirus protection - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open the profile for editing.
- 2 Edit the *FortiClient AV* application detection entry.
- 3 In *Condition*, select *Not Running* to enforce use the antivirus feature, To also require use of an up-to-date AV database, select *Not Running or Up-to-date*.
- 4 In *Action*, select *Block*.
- 5 Select *OK*.

To require endpoints to use FortiClient antivirus protection - CLI

In this example, endpoint control is configured to require that FortiClient AV is enabled and its database is up-to-date.

```
config endpoint-control profile
edit profile1
set feature-enforcement enable
set require-av enable
set require-av-uptodate enable
end
```

To require endpoints to use FortiClient firewall protection - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open the profile for editing.
- 2 Edit the *FortiClient Firewall* application detection entry.
- 3 In *Action*, select *Block*.
- 4 Select *OK*.

To require endpoints to use FortiClient firewall protection - CLI

```
config endpoint-control profile
edit profile1
set feature-enforcement enable
set require-firewall enable
end
```

To require endpoints to use FortiClient web filtering

This is a CLI-only configuration. In the following example, profile1 is configured to require endpoints to use web filtering.

```
config endpoint-control profile
```

```
edit profile1
    set feature-enforcement enable
    set require-webfilter enable
end
```

To require endpoints to use a licensed FortiClient application

This is a CLI-only configuration. In the following example, profile1 is configured to require endpoints to use a licensed copy of FortiClient Endpoint Security.

```
config endpoint-control profile
    edit profile1
        set feature-enforcement enable
        set require-license enable
    end
```

Optionally, you can set `require-license` to `warn` to warn rather than block users of unlicensed FortiClient software.

Setting the default action for applications

To set the default action for applications - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open the profile for editing.
- 2 Select the last application detection list entry, *All*, and edit it. Select the *Action* to take for any applications **not** included in this application detection list:
 - **Block** — Endpoints must have only the applications you specify and are denied access if any other applications are installed.
 - **Warn** — Endpoints are warned, but not blocked, if they have any applications other than those included in this application detection list.
 - **Allow or Monitor** — Endpoints can have any application installed, and are denied access only if they have an application for which you created a specific Block rule.
- 3 Select *OK*.

To set the default action for applications - CLI

```
config endpoint-control app-detect rule-list
    edit profile1.list
        set other-application-action allow
    end
```

Adding application detection entries

You need to add an application detection entry for any application that requires a different action than the predefined *All* entry.

To create an application detection entry - web-based manager

- 1 With the endpoint control profile open, select *Create New*.
- 2 Select the application *Category*.
- 3 In *Application*, do one of the following:
 - Select *All*.
 - Select *Specify* and then select the application.

- 4 Select the *Action* and *Condition*, depending on the type of rule you are creating:

Application detection rule	Action	Condition
Application is allowed	Allow	N/A
Application must be installed and running	Block or Warn	Not Running
Application must be installed	Block or Warn	Not Installed
Application must not be running	Block or Warn	Running
Application must not be installed	Block or Warn	Installed
Monitor endpoint with this application running	Monitor	Running
Monitor endpoint with this application installed	Monitor	Installed

The Warn option permits users to connect after viewing a warning.

- 5 Select *OK*.
- 6 To create additional application detection entries, repeat steps 1 through 5.

To create an application detection list - CLI

This example creates an application sensor that denies access to endpoints with peer-to-peer file sharing applications installed. All other applications are allowed.

```
config endpoint-control app-detect rule-list
  edit "applist1"
    config entries
      edit 1
        set application 0
        set category 15
        set vendor 0
        set status installed
        set action deny
      end
      set other-application-action allow
    end
  end
config endpoint-control profile
  edit profile1
    set application-detection enable
    set application-detection-rule-list profile1.list
  end
```

Viewing the application database

You can view the application list provided by FortiGuard Services. Go to *UTM Profiles > Endpoint Control > Application Database*.

Figure 93: Endpoint Control Predefined application list

Category	Name	Vendor	ID	Group
Anti-Malware Software	ActiveSc Application	Null	3077	Security
Anti-Malware Software	Agnitum Outpost Service	Agnitum Ltd.	3078	Security
Anti-Malware Software	Agnitum Outpost	Agnitum Ltd.	3079	Security
Anti-Malware Software	Kingsoft Internet Security	Kingsoft	3080	Security
Anti-Malware Software	SuiMainExe	AhnLab Inc.	3081	Security
Anti-Malware Software	VSCORE.14.1.0.496.x86	McAfee, Inc.	1802	Security
Anti-Malware Software	AhnLab MyKeyDefense 2.5	AhnLab Inc.	3082	Security
Anti-Malware Software	360杀毒	360.CN	1803	Security
Anti-Malware Software	VSCORE.14.1.0.447.x86	McAfee Inc.	3083	Security
Anti-Malware Software	TrendSecure Common Platform	Trend Micro Inc.	1804	Security
Anti-Malware Software	360rpt	360.CN	2316	Security

1 / 61 [Total Applications: 3037] [Column Settings] [Clear All Filters]

The list contains the following information. You can select the name of any column to sort the data by that field. You can also create filters on each column.

Application Database page	
Lists all the applications that are provided by FortiGuard Services	
Column Settings	Select the columns to display in the list. You can also determine the order in which they appear.
Filter Settings	Set and clear column display filters.
Category	The type of application. Example: Document Viewers
Name	The name of the application.
Vendor	The vendor that the application is associated with. For example, the Adobe Reader is associated with the vendor, Adobe Systems Incorporated.
ID	Unique application ID.
Group	Another categorization of the applications. Groups are not used in application sensor rules.
Page controls	Shows the current page number in the list. Select the left and right arrows to display the first, previous, next or last page of known endpoints.
[Total Signatures: <number>]	The total number of application signatures currently in the database.


Enabling Endpoint Control in firewall policies

Endpoint Control is applied to any traffic where the controlling firewall policy has Endpoint Security enabled. The selected Endpoint Control profile determines the conditions that govern network access.

You can also enable Endpoint Security in combination with identity-based firewall policies. Users must authenticate and their computers must meet the requirements of the Endpoint Control profile.

To enable Endpoint Control - web-based manager

- 1 Go to *Policy > Policy > Policy* and edit the firewall policy where you want to enable Endpoint Control.
- 2 Select Enable Endpoint Security and select the Endpoint Control profile.

Figure 94: Enabling Endpoint Security in a firewall policy

3 Select OK.

To configure the firewall policy - CLI

In this example, the LAN connects to Port 2 and the Internet is connected to Port 1. An Endpoint Control profile is applied.

```
config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr LANusers
    set dstaddr all
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set endpoint-check enable
    set endpoint-profile "our_profile"
  end
```

Monitoring endpoints

You can view statistical information about the endpoints that were subject to endpoint control. Data is gathered every 15 minutes and the display statistics are for the past 24 hours.

Endpoint status

Endpoint Status displays a pie chart showing the proportions of the endpoint population that:

- do not have FortiClient Endpoint Security installed
- have FortiClient Endpoint Security installed, but do not completely comply with the endpoint profile
- are fully compliant with the endpoint security profile.

To view this display, go to *UTM Profiles > Monitor > Endpoint Monitor* and in *Report by* select *Status*.

You can click on the chart for a detailed list of endpoints.

Figure 95: Endpoint Monitor detailed endpoints list

Refresh Filter Settings Return									
Compliant	Host Name	IP Address	User	OS Version	FortiClient Version	AV Signature	Detected Applications	CPU Model	System Uptime
Compliant	DareDevil	192.168.100.125	Administrator	Microsoft Windows XP Professional Service Pack 3 (build 2600)	4.2.3	13.161	<ul style="list-style-type: none"> FortiClient System Tray Controller Messenger Microsoft® Windows® Operating System VMware Tools TPAutoConnect More Applications... 	Intel(R) Xeon(R) CPU S130 @ 2.00GHz	0 day(s) 0 hour(s) 50 minute(s)

Endpoint Application Usage

Endpoint Application Usage displays a bar chart of the top ten applications by traffic volume. The counts include all endpoints with FortiClient Endpoint Security installed that are subject to endpoint control. You can select any of the chart bars to see a list of the top ten endpoints contributing to the data volume for that application.

To view this display, go to *UTM Profiles > Monitor > Endpoint Monitor* and in *Report by* select *Application Usage*.

Endpoint Traffic

Endpoint Traffic displays a bar chart of the top ten endpoints by traffic volume. The counts include all endpoints with FortiClient Endpoint Security installed that are subject to endpoint control. You can select any of the chart bars to see a list of the top ten applications on that endpoint by traffic volume.

To view this display, go to *UTM Profiles > Monitor > Endpoint Monitor* and in *Report by* select *Traffic*.

Modifying Endpoint Security replacement pages

The FortiGate unit sends one of the following HTML pages to a non-compliant user who attempts to use a firewall policy in which Endpoint Security is enabled:

- *Endpoint NAC Block Page* — The endpoint has FortiClient Connect installed, rather than FortiClient Endpoint Security. In the profile's FortiClient application detection entry, the *Block* action is selected. The user should check the FortiClient Connect console to view the exact reason why the endpoint is blocked.
- *Endpoint NAC Recommendation Block Page* — This is the warning version of the *Endpoint NAC Block Page*. The user can select the *Continue to* link to access their desired destination. In the profile's FortiClient application detection entry, the *Warn* action is selected.
- *Endpoint NAC Download Portal* — The endpoint does not have FortiClient Endpoint Security installed. In the profile's FortiClient application detection entry, the *Block* action is selected. The user must install the FortiClient application to proceed. The page includes a download link for the FortiClient installer.

If you modify this replacement message, be sure to retain the `%%LINK%%` tag which provides the download URL for the FortiClient installer.

- *Endpoint NAC Recommendation Portal* — This is the warning version of the *Endpoint NAC Download Portal*. In the profile's FortiClient application detection entry, the *Warn* action is selected. The user can select the *Continue to* link to access their desired destination, without installing FortiClient Endpoint Security. The page also includes a download link for the FortiClient installer.

If you modify this replacement message, be sure to retain both the `%%LINK%%` tag which provides the download URL for the FortiClient installer and the `%%DST_ADDR%%` link that contains the URL that the user requested.

- *Endpoint NAC Feature Block Page* — The endpoint does not have all of the required FortiClient features running. In the profile, the FortiClient feature application detection entry for the missing feature - AV, Firewall, or webfilter (CLI only) - has the *Block* action selected. The message lists the non-compliances. The user must correct the FortiClient settings and try again.

- *Endpoint NAC Recommendation Feature Block Page* — This is the warning version of the *Endpoint NAC Feature Block Page*. In this case, the FortiClient feature application detection entry for the missing feature - AV, Firewall, or webfilter (CLI only) - has the *Warn* action selected. The message lists the non-compliances. The user can select the *Continue* to link to access their desired destination without correcting the non-compliances.

To modify endpoint security replacement messages - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open the profile for editing.
- 2 Select *Customize Endpoint Messages* and then select the *Edit* icon.
The *Edit Message* window opens.
- 3 In the *Message* box, select the message that you want to modify.
- 4 In the *Message HTML* box, modify the message HTML code.
- 5 Select *OK*.
- 6 Select *OK* to save the changes to the profile.

Example

The Example company has the following requirements for employee computers:

- must run FortiClient Endpoint Security version 4.1.3 with firewall enabled
- must have OpenOffice installed
- cannot have any peer-to-peer file sharing applications installed
- must not have any games running
- all other applications are allowed

Configuring FortiClient download source and required version

FortiGuard Services will provide the FortiClient installer.

To configure FortiClient requirements - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Client Installers*.
- 2 Check that *FortiGuard Availability* shows a green checkmark icon.
If you see a red 'X' icon, check your FortiGuard configuration.
- 3 Under *FortiGuard Installer Download Location*, select *FortiGuard Distribution Network*.
- 4 Select *Enforce Minimum Version* and then select 4.1.3 from the list.
- 5 Select *Apply*.

To configure FortiClient requirements - CLI

```
config endpoint-control settings
  set download-location fortiguard
  set version-check minimum
  set version 4.1.3
  set compliance-timeout 5
end
```

Creating an endpoint control profile

The endpoint control rules for FortiClient Endpoint Security and other applications are contained within an endpoint control profile.

To create an endpoint control profile

- 1 Go to *UTM Profiles > Endpoint Control > Profile*.
- 2 Select *Create New*, enter a *Name* for the profile, and then select *OK*.

To create an endpoint control profile

In this example, profile1 is created.

```
config endpoint-control profile
  edit profile1
end
```

Configuring FortiClient application detection entries

The FortiClient application detection entry is configured by default to block endpoints that do not have FortiClient installed and running. The FortiClient AV and FortiClient Firewall entries by default allow endpoints access even if they are not using these features. Only the FortiClient Firewall entry needs to be modified.

To configure the FortiClient Firewall application detection entry

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open your profile for editing.
- 2 Open the *FortiClient Firewall* entry for editing.
- 3 In *Action*, select *Block*.
- 4 Select *OK*.
- 5 Select *OK* to save the change to your profile.

To configure the FortiClient Firewall application detection entry - CLI

```
config endpoint-control profile
  edit profile1
    set feature-enforcement enable
    set require-firewall enable
  end
```

Configuring application detection entries for other applications

You need to create application detection rules for OpenOffice, P2P applications, and games.

To configure the application sensor - web-based manager

- 1 Go to *UTM Profiles > Endpoint Control > Profile* and open your profile for editing.
- 2 Select *Create New*, enter the following information, and then select *OK*:

This creates a rule requiring the OpenOffice suite.

Category	Office
Application	Specify
Browse	Open Office
Filter by Vendor	Disabled

Action	Block
Condition	Not Installed

- 3 Select *Create New*, enter the following information, and then select *OK*:
This creates a rule denying users with P2P applications installed.

Category	P2P File Sharing
Application	All
Action	Block
Status	Installed

- 4 Select *Create New*, enter the following information, and then select *OK*:
This creates a rule denying users with games applications running.

Category	Games
Application	All
Action	Block
Status	Running

- 5 Select the last application detection list entry, *All*, and edit it. In *Action* select *Allow*.
6 Select *OK*.
7 Select *OK* to save the changes to your profile.

To configure the application detection list and endpoint control profile - CLI

By convention, the application detection entries in the CLI are contained in a rule list prefixed with the profile name. For example, the rule list for profile1 is profile1.list.

The three application detection entries are entered in the same order as for the web-based manager, above. To find codes, use the '?'. For example, `set vendor ?` lists the vendor codes.

```
config endpoint-control app-detect rule-list
edit profile1.list
config entries
edit 1
set application 77
set category 31
set status not-installed
set action deny
next
edit 2
set category 15
set status installed
set action deny
next
edit 3
set category 20
set status running
set action deny
end
set other-application-action allow
end
```

```
config endpoint-control profile
  edit "our_profile"
    set application-detection enable
    set application-detection-rule-list profile1.list
  end
```

Configuring the firewall policy

The firewall policy enables access to the Internet, but requires hosts to meet the Endpoint Control requirements configured in the Endpoint Control profile that you configured earlier.

To configure the firewall policy - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*:

Source Interface/Zone	Select the interface which connects to the LAN.
Source Address	Select the LAN address range.
Destination Interface/Zone	Select the interface which connects to the Internet.
Destination Address	All
Schedule	as required
Service	ANY
Action	ACCEPT
NAT	Enable NAT
Enable Endpoint Security	Select the Endpoint Control profile that you configured earlier.

To configure the firewall policy - CLI

In this example, the LAN connects to Port 2 and the Internet is connected to Port 1.

```
config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr LANusers
    set dstaddr all
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set endpoint-check enable
    set endpoint-profile profile1
  end
```

Endpoint Control interface reference

The Endpoint Control menu helps you to configure profiles, application sensors and databases, including network monitoring.

Endpoint control requires that all hosts using the security policy have the FortiClient Endpoint Security application installed. Make sure that all hosts affected by this policy are able to install this application. Currently, FortiClient Endpoint Security is available for only Microsoft Windows 2000 and later.

To set up endpoint control, you need to:

- Enable Central Management if you are going to use FortiGuard Services to update the FortiClient application or antivirus signatures. You do not need to enter account information.
- Configure the minimum required version of FortiClient and the source of FortiClient installer downloads for non-compliant endpoints. See [“Configuring FortiClient required version and download location” on page 1091](#).
- Define application detection lists to specify which applications are allowed or not allowed. Optionally, you can deny access to endpoints that have applications installed that are not on the detection list. See [“FortiGuard” on page 407](#).
- Configure endpoint profiles which specify the FortiClient enforcement settings and the application detection list to apply.
- In firewall policies, enable Endpoint Security and select the Endpoint profile to use.
- Optionally, modify the inactivity timeout for endpoints. The default is 5 minutes. After that time period, the FortiGate unit rechecks the endpoint for Endpoint compliance. To change the timeout, adjust the `compliance-timeout` value in the `config endpoint settings` CLI command.

You can also modify the appearance of the *Endpoint Download Portal* and the *Endpoint Recommendation Portal* by making changes within their replacement messages which are Endpoint Download Portal and Endpoint Recommendation Portal.

This topic includes the following:

- [Profile](#)
- [Application Database](#)
- [Configuring FortiClient required version and download location](#)

Profile

An endpoint control profile contains FortiClient enforcement settings and can specify an application detection list. Security policies can apply an endpoint profile to the traffic they handle.

The following are endpoint profile configuration settings in *UTM Profiles > Endpoint Control > Profile*.

Profile Settings page	
Provides settings for configuring an endpoint control profile. When you edit an existing endpoint profile, you are automatically redirected to this page.	
Name	The name that was entered in the <i>Name</i> field on the New Detection List page. To change the name, edit the text in this field and then select OK.
Comments	The comment that was entered in the <i>Comment</i> field on the New Detection List page. If you want to edit or add a comment, enter the text in this field and then select OK.

Customize Endpoint Messages	Select to allow modifying the Endpoint NAC replacement messages. Select <i>Edit</i> to modify a specific replacement message. When you select <i>Edit</i> , the Edit Message window appears where you can modify each endpoint NAC replacement message.
Create New	Creates a new application entry. When you select <i>Create New</i> , you are automatically redirected to the New Application Detection Entry page.
Edit	Modifies settings within an application entry. When you select <i>Edit</i> , you are automatically redirected to the Edit Application Detection Entry page.
Delete	Removes an entry from within the list. To remove multiple entries from within the list, on the Profile Settings page, in each of the rows of the entries you want removed, select the check box and then select <i>Delete</i> . To remove all entries in the list, on the Profile Settings page, select the check box in the check box column, and then select <i>Delete</i> .
Insert	Inserts a new application detection entry in the list. When you select <i>Insert</i> , you are automatically redirected to the New Application Detection Entry page.
Move To	Moves the entry to another position in the list. When you select <i>Move To</i> , the Move Application Detection Entry window appears. To move an entry, select the new position <i>Before</i> or <i>After</i> , which will place the current entry before or after then entry you enter in the (<i>Entry ID</i>) field. Enter the entry ID number in the field and then select <i>OK</i> .
ID	The identification number for that application detection entry. This number is used when moving an entry within the list.
Application or Category	The category and vendor. If you selected <i>All Applications</i> , then <i>All</i> will appear at the end. For example, Avaya Inc. - Games - All.
Action	The type of action that the unit will take when a match is found.
Condition	The status of the application or category, for example if the FortiClient software is not installed and not running.

New/Edit Application Detection Entry page

You can edit or create a new application detection entry from the UTM Profiles > Endpoint Control > Profile page.

Category	<p>Select the software category for the entry. For example, Instant Messaging.</p> <p>This is not available for the preconfigured FortiClient application detection entries.</p>
Application	<p>Select whether to include all applications in the application detection entry or specify an application. You cannot specify multiple applications per application detection entry.</p> <p>When you select <i>Specify</i>, you can use the <i>Search</i> field to find a specific application or select one from the <i>Browse</i> list. You can also select the check box beside <i>Filter By Vendor</i> to filter the applications by vendor as well.</p>
Tags	<p>You can add tags to an application detection entry. If tag configuration settings are not available, they are disabled on the web-based manager. You must enable them in <i>System > Admin > Settings</i>. See “General Settings” on page 359.</p> <p>To add a tag, enter the tag name in the <i>Add tags</i> field and select the plus sign. To add multiple tags, enter the tags and separate each with a comma and then select the plus sign. The tags that you create are shown in the <i>Applied tags</i> row.</p> <p>This is not available for the preconfigured FortiClient entries.</p>
Action	<p>Select what to do if the application is running on the endpoint:</p> <ul style="list-style-type: none"> • <i>Allow</i> – allows the endpoint to connect • <i>Block</i> – quarantines the endpoint • <i>Monitor</i> – includes this endpoint’s information in statistics and logs on the Endpoint Monitor page. • <i>Warn</i> – displays a block page, but contains a button to let the user continue at his or her discretion. Information is then sent back to the client.
Condition	<p>Select the status of the application.</p> <ul style="list-style-type: none"> • <i>Installed</i> – application is installed but no currently running. • <i>Running</i> – application is currently running • <i>Not Installed</i> – application is not currently installed • <i>Not Running</i> – application is not currently running

Application Database

The Application Database page allows you to view the applications which are sorted by category. For example, iKey is found under the Encryption PKI category. This page also allows you to apply tags the applications within the list, search to find a particular application or category, as well as filter the applications within the list.

Lists the applications. On this page, you can create tags for applications, search, as well as filter and customize columns.

Tags	<p>Adds or removes tags to the applications in the application database list. If tag configuration settings are not available, they are disabled on the web-based manager. You must enable them in <i>System > Admin > Settings</i>. See “Tag management” on page 582.</p> <p>When you select the down arrow beside Tags, you can add or removed tags.</p> <p>To add tags to an application in the list, select the application first and then select the down arrow beside <i>Tags</i> to select <i>Add Tags</i>. The Add Tags window appears. Enter the tag in the <i>Add tag</i> field and select the plus sign to add the tag <i>Tags</i> to apply. Select <i>OK</i> to permanently add the tags to the application.</p> <p>To remove a tag in application in the list, select the application first and then select the down arrow beside <i>Tags</i> to select <i>Remove Tags</i>. Select the tag in the <i>Applied tags</i> row; it automatically moves to the <i>Tags</i> to remove row. Repeat until all tags are listed in the <i>Tags to remove</i> row. Select <i>OK</i> to permanently remove the tags from the application.</p> <p>If there are tags that you want to add to an application that have been configured for another object, such as an application within <i>UTM Profiles > Application Control > Application List</i>, you can add those tags as well to the application in the application database.</p> <p>To apply these other object tags, select the application and then select the down arrow beside <i>Tags</i>; select <i>Add Tags</i>. In the <i>Add Tags</i> window, select each of the tags you want to add from within the <i>Click tag to add</i> row. Each tag is automatically added to the <i>Tags to apply</i> row. Select <i>OK</i> to add them to the application in the application database.</p>
Column Settings	<p>Customize the column view. You can select the columns to hide or display them and specify the column display order.</p>

Filter Settings	<p>Select to filter the information on the Application Database page. Filters appears automatically after selecting <i>Filter Settings</i>, below the column headings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filter settings, select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the filter settings for that column.</p> <p>The <i>Filter Settings</i> on the Application Database page contains the option <i>Copy to Profile</i>, which allows you to copy the filter settings and apply them to an endpoint profile.</p> <p>To apply existing filter settings to an endpoint profile, select the down arrow beside <i>Filter Settings</i> and then select <i>Copy to Profile</i>. In the <i>Select Object window</i>, select the endpoint profile from the <i>Endpoint Profile's</i> drop-down list. Select <i>OK</i>.</p>
Search	Enter a search criteria into the field and then press Enter on your keyboard. Use the <i>Clear All</i> icon beside the field to clear the search results.
Category	<p>The type of category an application is included in. For example, Anti-Mailware Software includes the Billy The Goat application.</p> <p>This column can display the application database list information in ascending or descending order. Select the green arrow beside Category; a green down arrow means the information is in descending order and an green up arrow means the information is in ascending order.</p>
Name	The name of the application.
Vendor	The name of the vendor.
Tags	The tag or tags that were created for the application.
Page Controls	Use to navigate through the list.
[Total Applications:]	The maximum number of applications that are currently shown in the list on the Application Database page.

Client Installers

You can set the minimum FortiClient version that endpoints are required to run from *UTM Profiles > Endpoint Control > Client Installers*. The FortiClient Endpoint Security page also configures the download source for the FortiClient installer.

Information section	
Indicates the FortiGuard availability and current versions of antivirus and application signatures packages. This section also allows you to update your antivirus and application signature packages, as well as downloading a Windows Installer.	
FortiGuard Availability	FortiGuard Services is available if the indicator is green.
FortiClient Endpoint Versions	FortiClient software versions available from FortiGuard Services are listed. Select the <i>Download</i> link to download the installer for either a Mac computer or Windows computer.
AV Signature Package	The latest AV signature package available from FortiGuard Services.
Application Signature Package	The latest application signature package available from FortiGuard Services.
FortiClient Downloads	The number of FortiClient software downloads through this FortiGate unit.
Update Now	Retrieve the latest information from FortiGuard Services.
FortiClient Installer Download Location section	
Select one of the following options to determine the link that the FortiClient Download Portal provides to non-compliant users to download the FortiClient installer.	
FortiGuard Distribution Network	<p>The FortiClient application is provided by the FortiGuard Distribution Network. The FortiGate unit must be able to access the FortiGuard Distribution Network.</p> <p>If the FortiGate unit contains a hard disk drive, the files from FortiGuard Services are cached to more efficiently serve downloads to multiple end points.</p>
This FortiGate	<p>Users download a FortiClient installer file from this unit.</p> <p>This option is available only on FortiGate models that support upload of FortiClient installer files. Upload your FortiClient installer file using the <code>execute restore forticlient</code> CLI command. For more information, refer to the FortiGate CLI Reference.</p>
Custom URL	Specify a URL from which users can download the FortiClient installer. You can use this option to provide custom installer files even if your unit does not have storage space for them.

FortiClient Version section	
Enforce Minimum Version ... (If enabled in Endpoint Profiles)	<p>From the list select either <i>Latest Available</i> or a specific FortiClient version as the minimum requirement for endpoints.</p> <p>The list contains the FortiClient versions available from the selected <i>FortiClient Installer Download Location</i>.</p> <p>Fortinet recommends that administrators deploy a FortiClient version update to their users or ask users to install the update and then wait a reasonable period of time for the updates to be installed before updating the minimum version required to the most recent version.</p>



Select *Custom URL* if you want to provide a customized FortiClient application. This is required if a FortiManager unit will centrally manage FortiClient applications. For information about customizing the FortiClient application, see the [FortiClient Administration Guide](#).



Vulnerability Scan

The Network Vulnerability Scan helps you to protect your network assets (servers and workstations) by scanning them for security weaknesses. You can scan on-demand or on a scheduled basis. Results are viewable on the FortiGate unit, but results are also sent to an attached FortiAnalyzer unit. The FortiAnalyzer unit can collect the results of vulnerability scans from multiple FortiGate units at different locations on your network, compiling a comprehensive report about network security.

This section describes how to configure a single FortiGate unit for network scanning and how to view the results of the scan.

The following topics are included in this section:

- [Overview](#)
- [Selecting assets to scan](#)
- [Configuring scans](#)
- [Viewing scan results](#)

Overview

Network vulnerability scanning has three main parts:

- Select the assets to scan
- Schedule scans or initiate them manually
- View the scan results

Selecting assets to scan

An asset is a server or workstation computer on your network. You can specify assets individually, but it is easier to use the network vulnerability scan feature's asset discovery function. The discovery function searches a specified IP address range and populates the asset list. You then select the assets to include in network vulnerability scans.

Asset discovery scans the following ports:

- TCP: 21-23, 25, 53, 80, 88, 110-111, 135, 139, 443, 445
- UDP: 53, 111, 135, 137, 161, 500

Discovering assets

The simplest way to build the Asset list is to perform a discovery scan on the range of IP addresses where your network assets are installed.

To discover assets - web-based manager

- 1 Go to *UTM Profiles > Vulnerability Scan > Asset Definition* and select *Create New*.
- 2 Enter a *Name* for this scan.
- 3 In *Type*, select *Range* and then enter the IP address *Range* to scan.

4 Select *OK*.

This creates an entry in the Asset list.

5 Select the asset list entry that you just created and then select *Discover Assets*.

Above the table header, on the top right, the status of the current scan is shown. Depending on the number of computers to be discovered, the scan can take several minutes, until the web-based manager reports “Scan completed.” The number of assets discovered is listed to the left of the *Discover Assets* button.

6 Select *Assets Found* and then select *Import*.

The discovered assets are added to the Asset list. By default, all are enabled for scanning.

7 Unless you want to discover assets on every scan, clear the Enable check box for this Asset discovery only asset.

You might want to add authentication credentials to some of your assets. To edit an entry in the Asset list, select its check box (at the left side of the list) and then select *Edit*. For more information about individual asset settings, see [“Adding assets manually”](#), below.

To discover assets - CLI

This example discovers assets in the range 10.11.101.10 to 10.11.101.200.

1 First configure the asset range to scan:

```
config netscan assets
  edit 0
    set name "office_discovery"
    set addr-type range
    set start-ip 10.11.101.10
    set end-ip 10.11.101.200
    set mode discovery
    set status enable
  end
```

2 Execute the discovery scan:

```
execute netscan start discover
```

3 Check the status of the discovery scan:

```
execute netscan status
```

Repeat periodically until status is “scan complete”.

4 Optionally, view a list of the discovered assets:

```
execute netscan list
```

5 Add the discovered assets to the asset list:

```
execute netscan import
```

Adding assets manually

There is no need to perform a discovery scan if you know the IP address of the computer that you want to scan, or you know that you want to scan all of the computers in a particular IP address range.

If you create an asset with an IP address range, any authentication credentials you enter will apply to all devices in the range. If this is not appropriate, you need to create individual entries for each computer instead.

To add an asset - web-based manager

- 1 Go to *UTM Profiles > Vulnerability Scan > Asset Definition* and select *Create New*.
- 2 Enter the following information and select *OK*:

Name	Enter a name for this asset.
Type	Select <i>Host</i> to configure a single IP address. Select <i>Range</i> to configure a range of IP addresses to scan.
IP Address	Enter the IP address of the asset. (<i>Type is Host.</i>)
Range	Enter the start and end of the IP address range. (<i>Type is Range.</i>)
Scan Type	Select <i>Vulnerability Scan</i> .
Windows Authentication	Select to use authentication on a Windows operating system. Enter the username and password in the fields provided. For more information, see “Requirements for authenticated scanning” on page 1116 .
Unix Authentication	Select to use authentication on a Unix operating system. Enter the username and password in the fields provided. For more information, see “Requirements for authenticated scanning” on page 1116 .

To add an asset - CLI

This example adds a single computer to the Asset list:

```
config netscan assets
edit 0
    set name "server1"
    set addr-type ip
    set start-ip 10.11.101.20
    set mode scan
    set auth-windows enable
    set win-username admin
    set win-password zxcvbnm
    set status enable
end
```

This example adds an address range to the Asset list. Authentication is not used:

```
config netscan assets
edit 0
    set name "fileservers"
    set addr-type range
    set start-ip 10.11.101.160
    set end-ip 10.11.101.170
    set mode scan
    set status enable
end
```

Requirements for authenticated scanning

The effectiveness of an authenticated scan is determined by the level of access the FortiGate unit obtains to the host operating system. Rather than use the system administrator's account, it might be more convenient to set up a separate account for the exclusive use of the vulnerability scanner with a password that does not change.

Microsoft Windows hosts - domain scanning

The user account provided for authentication must

- have administrator rights
- be a Security type of account
- have global scope
- belong to the Domain Administrators group
- meet the Group Policy requirements listed below:

Group Policy - Security Options

In the Group Policy Management Editor, go to Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options.

Setting	Value
Network access: Sharing and security model for local accounts	Classic
Accounts: Guest account status	Disabled
Network access: Let Everyone permissions apply to anonymous users	Disabled

Group Policy - System Services

In the Group Policy Management Editor, go to Computer Configuration > Windows Settings > Security Settings > System Services.

Setting	Value
Remote registry	Automatic
Server	Automatic
Windows Firewall	Automatic

Group Policy - Administrative Templates

In the Group Policy Management Editor, go to Computer Configuration > Administrative Templates > Network > Network Connections > Windows Firewall > Domain Profile.

Setting	Value
Windows Firewall: Protect all network connections	Disabled

or

Setting	Value
Windows Firewall: Protect all network connections	Enabled
Windows Firewall: Allow remote administration exception Allow unsolicited messages from ¹	Enabled *
Windows Firewall: Allow file and printer sharing exception	Enabled

Allow unsolicited messages from ¹	*
Windows Firewall: Allow ICMP exceptions	Enabled
Allow unsolicited messages from ¹	*

¹Windows prompts you for a range of IP addresses. Enter either "*" or the IP address of the Fortinet appliance that is performing the vulnerability scan.

Microsoft Windows hosts - local (non-domain) scanning

The user account provided for authentication must

- be a local account
- belong to the Administrators group

The host must also meet the following requirements:

- Server service must be enabled. (Windows 2000, 2003, XP)
- Remote Registry Service must be enabled.
- File Sharing must be enabled.
- Public folder sharing must be disabled. (Windows 7)
- Simple File Sharing (SFS) must be disabled. (Windows XP)

Windows firewall settings

- Enable the *Remote Administration Exception* in Windows Firewall. (Windows 2003, Windows XP)
- Allow *File and Print sharing* and *Remote Administration* traffic to pass through the firewall. Specify the IP address or subnet of the Fortinet appliance that is performing the vulnerability scan. (Windows Vista, 2008)
- For each of the active *Inbound Rules* in the *File and Printer Sharing* group, set the *Remote IP address* under *Scope* to either *Any IP address* or to the IP address or subnet of the Fortinet appliance that is performing the vulnerability scan. (Windows 7)

Unix hosts

The user account provided for authentication must be able at a minimum to execute these commands:

- The account must be able to execute "uname" in order to detect the platform for packages.
- If the target is running Red Hat, the account must be able to read /etc/redhat-release and execute "rpm".
- If the target is running Debian, the account must be able to read /etc/debian-version and execute "dpkg".

Configuring scans

You can configure regular network scans on a daily, weekly, or monthly basis. There are three scan modes. Full scan checks every TCP and UDP port and takes the most time. Standard scan checks the ports used by most known applications. Quick scan checks only the most commonly used ports. For a detailed list of the TCP and UDP ports examined by each scan mode, see [Table 71 on page 1119](#). Also, the `get netscan settings` CLI command lists the TCP and UDP ports scanned in the current scan mode. See the `tcp-ports` and `udp-ports` fields.

You can also initiate the configured scan manually.

Table 71: Ports scanned in each scan mode

Standard Scan	<p>TCP: 1-3, 5, 7, 9, 11, 13, 15, 17-25, 27, 29, 31, 33, 35, 37-39, 41-223, 242-246, 256-265, 280-282, 309, 311, 318, 322-325, 344-351, 363, 369-581, 587, 592-593, 598, 600, 606-620, 624, 627, 631, 633-637, 666-674, 700, 704-705, 707, 709-711, 729-731, 740-742, 744, 747-754, 758-765, 767, 769-777, 780-783, 786, 799-801, 860, 873, 886-888, 900-901, 911, 950, 954-955, 990-993, 995-1001, 1008, 1010-1011, 1015, 1023-1100, 1109-1112, 1114, 1123, 1155, 1167, 1170, 1207, 1212, 1214, 1220-1222, 1234-1236, 1241, 1243, 1245, 1248, 1269, 1313-1314, 1337, 1344-1625, 1636-1774, 1776-1815, 1818-1824, 1901-1909, 1911-1920, 1944-1951, 1973, 1981, 1985-2028, 2030, 2032-2036, 2038, 2040-2049, 2053, 2065, 2067, 2080, 2097, 2100, 2102-2107, 2109, 2111, 2115, 2120, 2140, 2160-2161, 2201-2202, 2213, 2221-2223, 2232-2239, 2241, 2260, 2279-2288, 2297, 2301, 2307, 2334, 2339, 2345, 2381, 2389, 2391, 2393-2394, 2399, 2401, 2433, 2447, 2500-2501, 2532, 2544, 2564-2565, 2583, 2592, 2600-2605, 2626-2627, 2638-2639, 2690, 2700, 2716, 2766, 2784-2789, 2801, 2908-2912, 2953-2954, 2998, 3000-3002, 3006-3007, 3010-3011, 3020, 3047-3049, 3080, 3127-3128, 3141-3145, 3180-3181, 3205, 3232, 3260, 3264, 3267-3269, 3279, 3306, 3322-3325, 3333, 3340, 3351-3352, 3355, 3372, 3389, 3421, 3454-3457, 3689-3690, 3700, 3791, 3900, 3984-3986, 4000-4002, 4008-4009, 4080, 4092, 4100, 4103, 4105, 4107, 4132-4134, 4144, 4242, 4321, 4333, 4343, 4443-4454, 4500-4501, 4567, 4590, 4626, 4651, 4660-4663, 4672, 4899, 4903, 4950, 5000-5005, 5009-5011, 5020-5021, 5031, 5050, 5053, 5080, 5100-5101, 5145, 5150, 5190-5193, 5222, 5236, 5300-5305, 5321, 5400-5402, 5432, 5510, 5520-5521, 5530, 5540, 5550, 5554-5558, 5569, 5599-5601, 5631-5632, 5634, 5678-5679, 5713-5717, 5729, 5742, 5745, 5755, 5757, 5766-5767, 5800-5802, 5900-5902, 5977-5979, 5997-6053, 6080, 6103, 6110-6112, 6123, 6129, 6141-6149, 6253, 6346, 6387, 6389, 6400, 6455-6456, 6499-6500, 6515, 6558, 6588, 6660-6670, 6672-6673, 6699, 6767, 6771, 6776, 6831, 6883, 6912, 6939, 6969-6970, 7000-7021, 7070, 7080, 7099-7100, 7121, 7161, 7174, 7200-7201, 7300-7301, 7306-7308, 7395, 7426-7431, 7491, 7511, 7777-7778, 7781, 7789, 7895, 7938, 7999-8020, 8023, 8032, 8039, 8080-8082, 8090, 8100, 8181, 8192, 8200, 8383, 8403, 8443, 8450, 8484, 8732, 8765, 8886-8894, 8910, 9000-9001, 9005, 9043, 9080, 9090, 9098-9100, 9400, 9443, 9535, 9872-9876, 9878, 9889, 9989-10000, 10005, 10007, 10080-10082, 10101, 10520, 10607, 10666, 11000, 11004, 11223, 12076, 12223, 12345-12346, 12361-12362, 12456, 12468-12469, 12631, 12701, 12753, 13000, 13333, 14237-14238, 15858, 16384, 16660, 16959, 16969, 17007, 17300, 18000, 18181-18186, 18190-18192, 18194, 18209-18210, 18231-18232, 18264, 19541, 20000-20001, 20011, 20034, 20200, 20203, 20331, 21544, 21554, 21845-21849, 22222, 22273, 22289, 22305, 22321, 22555, 22800, 22951, 23456, 23476-23477, 25000-25009, 25252, 25793, 25867, 26000, 26208, 26274, 27000-27009, 27374, 27665, 29369, 29891, 30029, 30100-30102, 30129, 30303, 30999, 31336-31337, 31339, 31554, 31666, 31785, 31787-31788, 32000, 32768-32790, 33333, 33567-33568, 33911, 34324, 37651, 40412, 40421-40423, 42424, 44337, 47557, 47806, 47808, 49400, 50505, 50766, 51102, 51107, 51112, 53001, 54321, 57341, 60008, 61439, 61466, 65000, 65301, 65512</p> <p>UDP: 7, 9, 13, 17, 19, 21, 37, 53, 67-69, 98, 111, 121, 123, 135, 137-138, 161, 177, 371, 389, 407, 445, 456, 464, 500, 512, 514, 517-518, 520, 555, 635, 666, 858, 1001, 1010-1011, 1015, 1024-1049, 1051-1055, 1170, 1243, 1245, 1434, 1492, 1600, 1604, 1645, 1701, 1807, 1812, 1900, 1978, 1981, 1999, 2001-2002, 2023, 2049, 2115, 2140, 2801, 3024, 3129, 3150, 3283, 3527, 3700, 3801, 4000, 4092, 4156, 4569, 4590, 4781, 5000-5001, 5036, 5060, 5321, 5400-5402, 5503, 5569, 5632, 5742, 6073, 6502, 6670, 6771, 6912, 6969, 7000, 7300-7301, 7306-7308, 7778, 7789, 7938, 9872-9875, 9989, 10067, 10167, 11000, 11223, 12223, 12345-12346, 12361-12362, 15253, 15345, 16969, 20001, 20034, 21544, 22222, 23456, 26274, 27444, 30029, 31335, 31337-31339, 31666, 31785, 31789, 31791-31792, 32771, 33333, 34324, 40412, 40421-40423, 40426, 47262, 50505, 50766, 51100-51101, 51109, 53001, 61466, 65000</p>
----------------------	---

Full Scan	All TCP and UDP ports (1-65535)
Quick Scan	<p>TCP: 11, 13, 15, 17, 19-23, 25, 37, 42, 53, 66, 69-70, 79-81, 88, 98, 109-111, 113, 118-119, 123, 135, 139, 143, 220, 256-259, 264, 371, 389, 411, 443, 445, 464-465, 512-515, 523-524, 540, 548, 554, 563, 580, 593, 636, 749-751, 873, 900-901, 990, 992-993, 995, 1080, 1114, 1214, 1234, 1352, 1433, 1494, 1508, 1521, 1720, 1723, 1755, 1801, 2000-2001, 2003, 2049, 2301, 2401, 2447, 2690, 2766, 3128, 3268-3269, 3306, 3372, 3389, 4100, 4443-4444, 4661-4662, 5000, 5432, 5555-5556, 5631-5632, 5634, 5800-5802, 5900-5901, 6000, 6112, 6346, 6387, 6666-6667, 6699, 7007, 7100, 7161, 7777-7778, 8000-8001, 8010, 8080-8081, 8100, 8888, 8910, 9100, 10000, 12345-12346, 20034, 21554, 32000, 32768-32790</p> <p>UDP: 7, 13, 17, 19, 37, 53, 67-69, 111, 123, 135, 137, 161, 177, 407, 464, 500, 517-518, 520, 1434, 1645, 1701, 1812, 2049, 3527, 4569, 4665, 5036, 5060, 5632, 6502, 7778, 15345</p>

To configure scanning - web-based manager

- 1 Go to *UTM Profiles > Vulnerability Scan > Scan*.
- 2 Enter the following information and select *Apply*.

Scan Mode	<p>Quick — check only the most commonly used ports</p> <p>Standard — check the ports used by most known applications</p> <p>Full — check all TCP and UDP ports</p> <p>For a detailed list of the TCP and UDP ports examined by each scan mode, see Table 71 on page 1119.</p>
Schedule	<p>Manually – perform scan on request only</p> <p>Schedule – use the following fields to configure a schedule</p>
Recurrence	<p>Select <i>Daily</i>, <i>Weekly</i>, or <i>Monthly</i>.</p> <p>If you select <i>Weekly</i>, the Day of Week drop-down list appears. If you select <i>Monthly</i>, the Day of Month drop-down list appears.</p>
Time	Select the time of day to start the scan, in the format HH:MM.
Day of Week	For a weekly scan, select the day of the week.
Day of Month	For a monthly scan, select the day of the month.

To configure scanning - CLI

To configure, for example, a standard scan to be performed every Sunday at 2:00am, you would enter:

```
config netscan settings
  set scan-mode standard
  set schedule enable
  set time 02:00
  set recurrence weekly
  set day-of-week sunday
end
```

To perform a vulnerability scan manually - web-based manager

- 1 Go to *UTM Profiles > Vulnerability Scan > Asset*.

- 2 Select the Enable check box for each asset you want to scan.
- 3 Select *Start Scan*.

Above the table header, on the top right, the status of the current scan is shown.

Depending on the number of computers to be discovered, the scan can take several minutes, until the web-based manager reports “Scan completed.”

To perform a vulnerability scan manually - CLI

You must have some assets configured with `mode` set to `scan`.

- 1 Execute the discovery scan:
- 2 Check the status of the discovery scan:

```
execute netscan start scan
```

```
execute netscan status
```

Repeat periodically until status is “scan complete”.

Viewing scan results

The results of network scanning are available as summary graphs and log entries.

Viewing scan logs

To view network scan logs, go to *Log&Report > Log & Archive Access > Vulnerability Scan Log*.

Figure 96: Network scan logs

Refresh Clear All Filters Column Settings Disk Memory HA Cluster: FG600B3908600705 Formatted Raw					
64	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
65	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
66	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
67	2010-06-01 10:13:45	notice	vulnerability	4090	10.11.101.20
68	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
69	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
70	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
71	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
72	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
73	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
74	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
75	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
76	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
77	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
78	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
79	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
80	2010-06-01 10:13:45	notice	vulnerability	4098	10.11.101.20
81	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
82	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
83	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
84	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
85	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
86	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
87	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
88	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
89	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
90	2010-06-01 10:13:45	notice	discovery	4100	10.11.101.20
91	2010-06-01 10:13:45	notice	discovery	4099	10.11.101.20
92	2010-06-01 10:12:16	notice	vulnerability	4096	Linux 2.6.17 - 2.6.27
93	2010-06-01 10:12:16	notice	vulnerability	4098	10.11.101.20
94	2010-06-01 10:12:16	notice	vulnerability	4098	10.11.101.20

Select any log entry to view log details.

Figure 97: Network scan log details

The screenshot shows the FortiGate Log & Report interface. On the left, a list of log entries is displayed with columns for ID, Date, Time, Level, Sub Type, ID, Virtual Domain, Action, IP Address, Protocol, Port, and Vulnerability. The entry with ID 65 is selected. On the right, a 'Log Details' pop-up window shows the details for this entry, including Date (2010-06-01), Time (10:13:45), Level (notice), Sub Type (vulnerability), ID (4098), Virtual Domain (root), Action (vuln-detection), IP Address (10.11.101.20), Protocol (tcp), Port (80), Vulnerability (Apache mod_proxy_ftp 2.0.x/2.2.x Denial of Service Vulnerability), Vuln Category (web server), Vuln ID (18412), Reference (N/A), and Severity (high).

ID	Date	Time	Level	Sub Type	ID	Virtual Domain	Action	IP Address	Protocol	Port	Vulnerability	Vuln Category	Vuln ID	Reference	Severity
64	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	SSH Server type								
65	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
66	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
67	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
68	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
69	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
70	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
71	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
72	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
73	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
74	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
75	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
76	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
77	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
78	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
79	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
80	2010-06-01	10:13:45	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
81	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
82	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
83	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
84	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
85	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
86	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
87	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
88	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
89	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
90	2010-06-01	10:13:45	notice	discovery	4100	10.11.101.20	Apache mod_proxy_ftp								
91	2010-06-01	10:13:45	notice	discovery	4099	10.11.101.20	Linux 2.6.17 - 2.6.27								
92	2010-06-01	10:12:16	notice	vulnerability	4096		SSH Server type								
93	2010-06-01	10:12:16	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								
94	2010-06-01	10:12:16	notice	vulnerability	4098	10.11.101.20	Apache mod_proxy_ftp								

Viewing Executive Summary graphs

To view summary graphs, go to *Log&Report > Report Access > Executive Summary*. You might need to add the following widgets to the page to view the summaries you require.

Table 72: Executive summary widgets for network scan

Chart	Widget name
Vulnerabilities by Category	vulner-by-category-last24h
Vulnerabilities by Severity	vulner-by-severity-last24h
Top Vulnerable Operating Systems Detected	top-vulner-os-last24h
Top Vulnerable Services Detected	top-vulner-service-last24h
Top Vulnerable TCP Services Detected	top-vulner-tcp-service-last24h
Top Vulnerable UDP Services Detected	top-vulner-udp-service-last24h

Creating reports

You can use the FortiGate unit's Log&Report features to generate reports on the results of network vulnerability scanning.

To create a report of scanning results

- 1 Go to *Log&Report > Report Config > Layout* and select *Create New*.
- 2 Enter a *Name* for the report.

- 3 Optionally select a *Report Theme*.
- 4 Enter a *Title* to appear on the report.
- 5 Choose each *Option* and *Output Format* that you require.
- 6 If you want to have the report generated on a regular basis, create a *Schedule*.
- 7 Select the *Report Components*.

The components are listed in order below the Report Components heading. For a network vulnerability scan report, you will need to select Chart components from the Vulnerability category. The following vulnerability charts are available:

Table 73: Executive summary widgets for network scan

Chart	Component name
Vulnerabilities by Category	vulner-by-category-last24h
Vulnerabilities by Severity	vulner-by-severity-last24h
Top Vulnerable Operating Systems Detected	top-vulner-os-last24h
Top Vulnerable Services Detected	top-vulner-service-last24h
Top Vulnerable TCP Services Detected	top-vulner-tcp-service-last24h
Top Vulnerable UDP Services Detected	top-vulner-udp-service-last24h

- 8 Optionally select other components, such as headings and text.
- 9 Select OK.

The report will be generated at the scheduled time.

To generate a report manually

- 1 Go to *Log&Report > Report Config > Layout*.
- 2 Select the required report.
- 3 Select *Run*.

Viewing reports

Go to *Log&Report > Report Access > Disk* to view generated reports.

Figure 98: List of reports

Delete					
<input type="checkbox"/>	Report File	Started	Finished	Size (bytes)	Other Formats
<input checked="" type="checkbox"/>	On-Demand-netscanreport-2010-05-03-101952	2010-05-03 10:19:52	2010-05-03 10:19:53	10380	PDF(9.1k)

If HTML output was enabled, you can select the Report File name to view the report in a separate browser window.

If PDF output was enabled, you can select the link in the Other Formats column to view the report.

Vulnerability Scan interface reference

The *Vulnerability Scan* menu enables you to configure scanning of your network, a feature similar to the vulnerability scan features on FortiAnalyzer or FortiScan units.

This topic includes the following:

- [Selecting assets to scan](#)
- [Configuring scans](#)
- [Vulnerability Result](#)

Asset Definition

In the Asset Definition menu, you can configure multiple asset lists that are used for scanning purposes. On the Asset Definition page, you can use an asset list to discover assets, start a scan or view the discovered assets.

The following are asset definition configuration settings in *UTM Profiles > Vulnerability Scan > Asset Definition*.

Asset Definition page	
Lists each individual asset that you created. On this page, you can edit, delete or create a new asset.	
Create New	Creates a new asset. When you select <i>Create New</i> , you are automatically redirected to the Asset Settings page.
Edit	Modifies an asset. When you select <i>Edit</i> , you are automatically redirected to the Asset Settings page.
Delete	Removes an asset from the list on the page.
View	Displays discovered hosts that were found during the scanning process. When you select <i>View</i> , the Discovered Hosts window appears. Select <i>Return</i> to go back to the Asset Definition page.
Discover Assets	Scans to discover assets. When you select <i>Discover Assets</i> , the options <i>Pause</i> and <i>Stop</i> appear. Select <i>Pause</i> to pause the scanning process or select <i>Stop</i> to stop the scanning process altogether.
Start Scan	Starts the scanning process. When you start scanning, the <i>Pause</i> and <i>Stop</i> options appear. Select <i>Pause</i> to pause the scanning process or select <i>Stop</i> to stop the scanning process altogether.
Name	The name of the asset.
IP Address/Range	If <i>Host</i> was chosen as the type for the asset, then the IP address of the host displays. If <i>Range</i> was chosen as the type for the asset, the IP address range appears.
Scheduled Vulnerability Scan	Indicates that there is a scheduled scan.
# Assets Discovered	Indicates how many assets were discovered during the scanning process.
Scan Activity	Indicates the activity of the scan.
Asset Settings page	
Provides settings for configuring an asset.	
Name	Enter a name for the asset that you are creating.
Type	Select <i>Host</i> to configure the host's IP address. Select <i>Range</i> to configure the IP address range.

IP Address (Range)	Enter the IP address of the host, or the IP address range. This depends on what type you selected in <i>Type</i> . If you select <i>Range</i> in <i>Type</i> , then the name <i>Range</i> appears and there are two IP address fields for you to enter the IP address range in.
Enable Scheduled Vulnerability Scanning	Select to schedule a scan.
Windows Authentication	Select to use authentication on a Windows operating system. Enter the username and password in the fields provided. The fields appear after selecting <i>Windows Authentication</i> .
Unix Authentication	Select to use authentication on a Unix operating system. Enter the username and password in the fields provided. The fields appear after selecting <i>Unix Authentication</i> .

Scan Schedule

From the Scan Schedule page, you can view the status of a vulnerability scan, and adjust the times and type of future scans.

The following are configuration settings for scheduling scans in *UTM Profiles > Vulnerability Scan > Scan Schedule*.

Status section	
Indicates the status of the previous scan, as well as when the next scan will occur. You can also start a scan using <i>Start Scan</i> .	
Scan Status	Indicates if a scan is currently running. If you select <i>Start Scan</i> , a scan will immediately begin. A progress bar appears, along with <i>Pause</i> and <i>Stop</i> , and the start time, estimated rating and completion time period. Select <i>Pause</i> to pause the scan or <i>Stop</i> to stop the scan altogether.
Last Scan Start Time	Indicates the last previous scan's start time. The format is <month> <day>, <year> - <hour>:<minute> <AM/PM>. For example, October 12, 2010 - 04:30 PM. The time is in 24-hour format.
Last Scan End Time	Indicates the last previous scan's end time. The format is <month> <day>, <year> - <hour>:<minute> <AM/PM>. For example, October 12, 2010 - 06:30 PM. The time is in 24-hour format.
Last Scan Duration	Indicates how long the scan lasted, in seconds.
Next Scheduled Scan	Indicates when the next scheduled scan will begin.

Schedule section	
Configure the time and day (or date if you choose to schedule a scan on a monthly basis) as well as enable to suspend a scan between a specified time period.	
Recurrence	<p>Select when the schedule should occur, such as on a daily basis or monthly basis.</p> <p>If you select <i>Weekly</i>, you must select the day of the week that the scan will occur on, as well as the hour and minutes.</p> <p>If you select <i>Monthly</i>, you must select the day within the current month as well as the hour and minutes. For example, in the current month of October, <i>Recurrence</i> is set to <i>Monthly</i>, with <i>Day of Month</i> being 20 and <i>Hour</i> being 12, and <i>Minutes</i> being 00.</p>
Suspend Scan between	Select to suspend the unit from scanning the network during a specified time period. Specify the hours and minutes using the drop-down lists.
Vulnerability Scan Mode section	
Select a mode that will be used when scanning network activity.	
Quick	Examines the most commonly used ports for vulnerabilities.
Standard	Examines a large number of application ports which cover many known applications.
Full	Examines the full port range, 1-65535, looking for applications that are running on non-standard ports.
Advanced section	
Provides options for specifying scanning only TCP or UDP ports, as well as operating systems or services.	
Enable TCP Port Scan	Select to scan only TCP ports.
Enable Service Detection	Select to scan for the detection of services on ports.
Enable OS Detection	Select to scan for the detection of operating systems on ports.
Enable UDP Port Scan	Select to scan only UDP ports.

Vulnerability Result

The Vulnerability Scan Results page provides a graphical and tabular representation of the information the FortiGate unit gathered from scanning the network.

The Summary table lists recent scans.

The two bar graphs, *Vulnerabilities by Severity* and *Vulnerabilities by Category* show the number of vulnerabilities that were found by severity and by category.



Sniffer policy

Sniffer policies are used to configure a physical interface on the FortiGate unit as a one-arm intrusion detection system (IDS). Traffic sent to the interface is examined for matches to the configured IPS sensor and application control list. Matches are logged and then all received traffic is dropped. Sniffing only reports on attacks. It does not deny or otherwise influence traffic.

This section describes how to configure your network topology to use the FortiGate unit as a one-arm intrusion detection system. It also describes how to configure and enable a sniffer policy.

The following topics are included in this section:

- [Sniffer policy concepts](#)
- [Before you begin](#)
- [Enable one-arm sniffing](#)
- [Sniffer example](#)

Sniffer policy concepts

Using the one-arm intrusion detection system (IDS), you can configure a FortiGate unit to operate as an IDS appliance by sniffing network traffic for attacks without actually processing the packets.

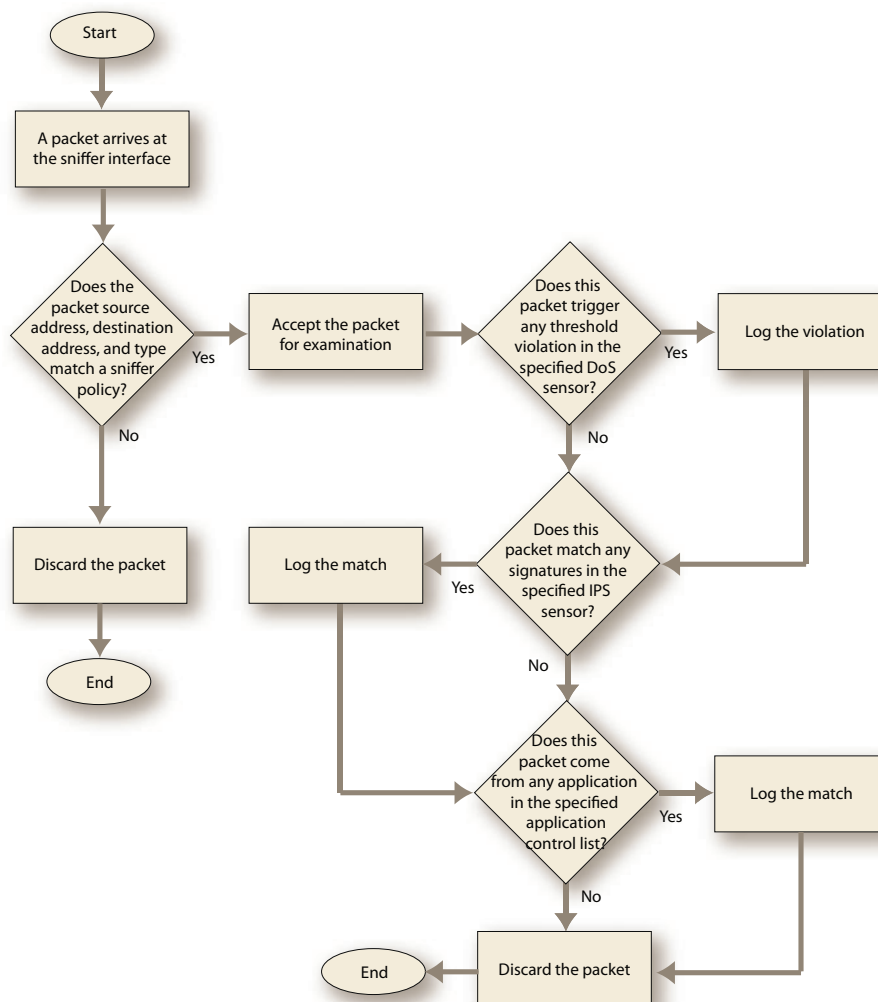
To configure one-arm IDS, you enable sniffer mode on a FortiGate interface and connect the interface to a hub or to the SPAN port of a switch that is processing network traffic. Then you add DoS policies for that FortiGate interface. Each policy can include a DoS sensor, an IPS sensor, and an application control list to detect attacks and application traffic in the network traffic that the FortiGate interface receives from the hub or switch SPAN port.

The sniffer policy list

The sniffer policy list shows all of the sniffer policies you have created. The policies are listed by sniffer interface. This is important because multiple sniffer policies can be applied to sniffer interfaces. Traffic entering a sniffer interface is checked against the sniffer policies for matching source and destination addresses and for service. This check against the policies occurs in listed order, from top to bottom. The first sniffer policy matching all three attributes then examines the traffic. Once a policy matches the attributes, checks for policy matches stop. If no sniffer policies match, the traffic is dropped without being examined.

Once a policy match is detected, the matching policy compares the traffic to the contents of the DoS sensor, IPS sensor, and application list specified in the policy. If any matches are detected, the FortiGate unit creates an entry in the log of the matching sensor/list. If the same traffic matches multiple sensors/lists, it is logged for each match. When this comparison is complete, the network traffic is dropped.

[Figure 99](#) illustrates this process.

Figure 99: How the intrusion detection system uses sniffer policies to examine traffic

Before you begin

Traffic entering an interface in sniffer mode is examined for DoS sensor violations, IPS sensor matches, and application control matches. After these checks, all network traffic is dropped. To avoid losing data, you must direct a copy of the network traffic to the FortiGate unit interface configured to sniff packets.


The easiest way to do this is to either use a hub or a switch with a SPAN port.

A hub is the easiest solution to implement but carries a downside. Connecting the FortiGate unit interface configured with the sniffer policy to a hub will deliver all traffic passing through the hub to the interface. However, if the network carries a heavy traffic load, the hub could slow the network because every hub interface sends out all the traffic the hub received on every interface.

A better solution is a switch with a SPAN port. Network switches receive traffic on all interfaces but they only send traffic out on the interface connected to the destination. Network slowdowns are less common when using switches instead of hubs.

Connecting the sniffer interface to a regular switch interface will not work because no traffic is addressed to the sniffer interface. A SPAN port is a special-purpose interface that mirrors all the traffic the switch receives. Traffic is handled normally on every other switch interface, but the SPAN port sends a copy of everything. If you connect your FortiGate unit sniffer interface to the switch SPAN port, all the network traffic will be examined without any being lost because of the examination.

Figure 100: A network configured for intrusion detection using a sniffer policy



Enable one-arm sniffing

Sniffer policies examine network traffic for anomalous patterns that usually indicate an attack. Since all traffic entering a sniffer interface is dropped, you need to first add a switch or hub to your network as described in [“Before you begin” on page 1128](#). The following steps are based on the assumption that you have added the switch or hub.

General configuration steps

The interface first must be designated as the sniffer interface, then the sniffer policy can be configured to use the sniffer interface.

- 1 Add a switch or hub to your network as described in [“Before you begin” on page 1128](#). This configuration will send a copy of your network traffic to the sniffer interface.



When an interface is configured as a sniffer interface, all traffic received by the interface is dropped after being examined by the sniffer policy.

- 2 Designate a physical interface as a sniffer interface.
- 3 Create a sniffer policy that specifies the sniffer interface.
- 4 Specify a DoS sensor, IPS sensor, application control list, or any combination of the three to define the traffic you want logged.

Designating a sniffer interface

An interface must be designated as a sniffer interface before it can be used with a sniffer policy. Once an interface is designated as a sniffer interface, it functions differently from a regular network interface in two ways:

- A sniffer mode interface accepts all traffic and drops it. If a sniffer policy is configured to use the sniffer interface, traffic matching the attributes configured in the policy will be examined before it is dropped. No traffic entering a sniffer mode interface will exit the FortiGate unit from any interface.
- A sniffer mode interface will be the only available selection in sniffer policies. The sniffer interface will not appear in firewall policies, routing tables, or anywhere else interfaces can be selected.

Designating a sniffer interface

- 1 Go to *System > Network > Interface*.
- 2 Select the interface.
- 3 Select the *Edit* icon.



When an interface is configured as a sniffer interface, all traffic received by the interface is dropped after being examined by the sniffer policy.

- 4 Select the *Enable one-arm sniffer* check box.
If the check box is not available, the interface is in use. Ensure that the interface is not selected in any firewall policies, routes, virtual IPs or other features in which a physical interface is specified.
- 5 Select *OK*.

Creating a sniffer policy

Sniffer interfaces accept all traffic. To examine the traffic before it is dropped, a sniffer policy is required.

To create a sniffer policy

- 1 Go to *Policy > Policy > Sniffer Policy* and select *Create New*.
- 2 For *Source Interface/Zone*, select the interface configured as the sniffer interface. If no interfaces are available for selection, no interfaces have been defined as sniffer interfaces. For more information, see [“Designating a sniffer interface” on page 1130](#).
- 3 For *Source Address*, select the address or address group that defines the source addresses of the traffic the sniffer policy will examine. Network traffic from addresses not included in the selected address group is ignored by this sniffer policy.
- 4 For *Destination Address*, select the address or address group that defines the destination addresses of the traffic the sniffer policy will examine. Network traffic to addresses not included in the selected address group is ignored by this sniffer policy.
- 5 For *Service*, select the type of network traffic the sniffer policy will examine. Protocols not included in the selected service or service group are ignored by this sniffer policy.
- 6 To have the sniffer policy log violations specified in a DoS sensor, select the *DoS Sensor* check box and choose the sensor from the list.
- 7 To have the sniffer policy log signatures appearing in an IPS sensor, select the *IPS Sensor* check box and choose the sensor from the list.

- 8 To have the sniffer policy log traffic from applications specified in an application control list, select the *Application Black/White List* check box and choose the application control list.
- 9 Select OK.

DoS sensors, IPS sensors, and application control lists all allow you to choose actions and log traffic. When included in a sniffer sensor, these settings are ignored. Actions in these other settings do not apply, and all matches are logged regardless of the logging setting.

Sniffer example

An IDS sniffer configuration

The Example.com Corporation uses a pair of FortiGate-620B units to secure the head office network. To monitor network attacks and create complete log records of them, the network administrator has received approval to install a FortiGate-82C to record all IPS signature matches in incoming and outgoing network traffic using a sniffer policy. This example details the set-up and execution of this network configuration.


Although this example uses a separate FortiGate unit for sniffer-mode operation, the sniffer traffic can be sent to the FortiGate unit protecting the network. The switch must still be configured to create a copy of the data because the sniffer interface drops all incoming traffic. In this case, the administrator requested a FortiGate-82C for this purpose because sniffer-mode operation is resource intensive, and using a separate FortiGate unit frees the FortiGate-620B cluster from this task. The FortiGate-82C unit also has four internal hard drives, making it ideal for storing large log files.

Configuring the network

Connect the Port1 interface of the FortiGate-82C to the Port8 interface of the switch.

You must configure your network to deliver a copy of the traffic to be examined to the sniffer interface because all network traffic entering a sniffer interface is dropped after examination.

Since the corporate network uses a pair of FortiGate units in an HA cluster, a switch is already in place connecting the Internet to Port1 of both FortiGate units.

Figure 101: Switch configuration

The company Internet feed is connected to Port1 of the switch. The FortiGate units are connected to Port2 and Port3 of the switch. Since they are configured as an HA cluster, they must both have access to the Internet in the event of a failure.

To allow a FortiGate unit sniffer interface to examine the network traffic, the switch must be configured to create a copy of all network traffic entering or leaving Port2 and Port3 and send it out Port8. When configured this way, the switch port sending the duplicate traffic is called a mirror port or a SPAN port.

Consult the switch documentation for instructions on how to configure a SPAN port.



The traffic between Port1 and Port2/Port3 is not modified or diverted in any way by the creation of a SPAN port. The traffic is duplicated with the copy being sent out of the SPAN port.

Configuring the FortiGate sniffer interface

No sniffer interfaces are included in the default configuration of any FortiGate unit. A copy of all of the network traffic is being sent to Port1 of the FortiGate-82C so you must configure Port1 as a sniffer-mode interface.



When an interface is configured as a sniffer interface, all traffic received by the interface is dropped after being examined by the sniffer policy.

To configure the sniffer mode interface — web-based manager

- 1 Log in to the FortiGate-82C web-based manager.
- 2 Go to *System > Network > Interface*.
- 3 Select the Port1 interface.
- 4 Select *Edit*.
- 5 Select *Enable one-arm sniffer*.
- 6 Select *OK*.

To configure the sniffer mode interface — CLI

```
config system interface
  edit port1
    set ips-sniffer-mode enable
  end
```

Creating an IPS sensor

A sniffer policy allows you to select an IPS sensor, a DOS sensor, and an application control list. Any conditions these sensors and list are configured to detect and log are saved to the appropriate log.

For this example, create an IPS sensor that detects and logs the occurrence of all the predefined IPS signatures.

To create an IPS sensor — web-based manager

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor*.
- 2 Select *Create New*.
- 3 In the *Name* field, enter `IPS_sniffer`.
- 4 In the *Comments* field, enter `IPS sensor for use in the sniffer policy`.
- 5 In the *Filters* section, select *Create New*.
- 6 In the name field, enter `All signatures, logged`.
- 7 For the *Logging* setting under *Signatures Settings*, select *Enable all*.
- 8 Select *OK* to save the filter.
- 9 Ensure *Enable Logging* is selected in the sensor.
- 10 Select *OK* to save the IPS sensor.

To create an IPS sensor — CLI

```
config ips sensor
  edit IPS_sniffer
    set comment "IPS sensor for use in the sniffer policy."
  config filter
    edit "All signatures, logged"
      set log enable
    end
  end
```

Creating the sniffer policy

The sniffer policy allows us to choose

To create the sniffer policy — web-based manager

- 1 Go to *Policy > Policy > Sniffer Policy*.
- 2 Select *Create New*.
- 3 Select *Port1* for the *Source Interface/Zone*.
- 4 Enable *IPS Sensor* and select the `IPS_sniffer` sensor.
- 5 Select *OK* to save the sniffer policy.

To create the sniffer policy — web-based manager

```
config firewall sniff-interface-policy
```

```

edit 0
  set interface port1
  set srcaddr all
  set dstaddr all
  set service ANY
  set ips-sensor-status enable
  set ips-sensor IPS_sniffer
end

```

With this configuration, all traffic entering the sniffer port is checked for matching signatures. Matches are logged and the traffic is dropped.

To examine the network traffic for more issues, you can create a DoS sensor and select it in the sniffer policy to log traffic anomalies. You can also create an application list with the specific application you'd like to check for and select it in the sniffer policy.

Sniffer Policy interface reference

The sniffer security policy list displays sniffer security policies in their order of matching precedence for each interface, source/destination address pair, and service.

If virtual domains are enabled on the unit, sniffer security policies are configured separately for each virtual domain; you must access the VDOM before you can configure its security policies.

You can add, delete, edit, and re-order security policies in the sniffer policy list. Sniffer policy order affects policy matching. As with security policies and DoS security policies, sniffer security policies are checked against traffic in the order in which they appear in the sniffer policy list, one at a time, from top to bottom. When a matching policy is discovered, it is used and further checking for sniffer policy matches are stopped. If no match is found the packet is dropped.

Sniffer policy configuration settings

The following are sniffer policy configuration settings in *Policy > Policy > Sniffer Policy*.

Sniffer Policy page	
Lists each individual sniffer policy that you created. On this page, you can edit, delete and create a new sniffer policy. You can also move a policy or insert a new policy on the page.	
Note: All of the specified attributes must match network traffic to trigger the policy.	
Create New	<p>Creates a new sniffer policy. Select the down arrow beside <i>Create New</i> to add a new section to the list to visually group the security policies.</p> <p>When you select <i>Create New</i> (or an option from the down arrow's drop-down list), you are automatically redirected to the New Policy page.</p>
Edit	<p>Modifies settings within the sniffer policy. When you select <i>Edit</i>, you are automatically redirected to the Edit Sniffer Policy page.</p> <p>You can also edit a sniffer policy by selecting the down arrow beside <i>Edit</i> and then selecting <i>Edit Policy</i>. If you want to disable or enable a sniffer policy, select the down arrow beside <i>Edit</i>, and then select either <i>Enable</i> or <i>Disable</i>.</p>

Delete	<p>Removes a policy from the list on the Sniffer Policy page.</p> <p>To remove multiple security policies from within the list, on the Sniffer Policy page, in each of the rows of the security policies you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all security policies from the list, on the Sniffer Policy page, select the check box in the check box column, and then select <i>Delete</i>.</p>
Move To	<p>Moves the corresponding policy before or after another policy in the list. When you select <i>Move To</i>, the Move Policy window appears.</p> <p>To move a sniffer policy, select the new position <i>Before</i> or <i>After</i>, which will place the current policy before or after the policy you enter in the field (<i>Policy ID</i>). Enter the policy ID number in the field and then select <i>OK</i>.</p>
Insert	<p>Adds a new policy above the corresponding policy (the New Policy screen appears). See “New Policy page” on page 1136.</p>
Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filter settings, select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p>
Column Settings	<p>Select when you want to Include or remove columns. You can select the columns to hide or display and specify the column displaying order in the table.</p>
Section View	<p>Select to display security policies organized by interface.</p>
Global View	<p>Select to list all security policies in order according to a sequence number.</p>
ID	<p>A unique identifier for each policy. security policies are numbered in the order they are created.</p>
Source	<p>The source address or address group to which the policy applies.</p>
Destination	<p>The destination address or address group to which the policy applies.</p>
Service	<p>The service to which the policy applies.</p>
DoS Sensor	<p>The DoS sensor selected in this policy.</p>

IPS Sensor	The IPS sensor selected in this policy.
Application Control List	The application sensor that is selected in this policy.
Status	When selected, the DoS policy is enabled. Clear the check box to disable the policy.
Interface	The interface used by the policy.
New Policy page Provides settings for configuring a new sniffer policy. When you select <i>Create New</i> on the Sniffer Policy page, you are automatically redirected to this page.	
Source Interface/Zone	The interface or zone to be monitored.
Source Address	Select an address, address range, or address group to limit traffic monitoring to network traffic sent from the specified address or range. Select <i>Multiple</i> to include multiple addresses or ranges. You can also select <i>Create New</i> to add a new address or address group.
Destination Address	Select an address, address range, or address group to limit traffic monitoring to network traffic sent to the specified address or range. Select <i>Multiple</i> to include multiple addresses or ranges. You can also select <i>Create New</i> to add a new address or address group.
Service	Select a firewall pre-defined service or a custom service to limit traffic monitoring to only the selected service or services. You can also select <i>Create New</i> to add a custom service.
DoS sensor	Select and specify a DoS sensor to have the unit apply the sensor to matching network traffic. You can also select <i>Create New</i> within the drop-down list to add a new DoS sensor.
Enable IPS	Select an IPS sensor from the drop-down list. You can also select <i>Create New</i> within the drop-down list to add a new IPS sensor.
Enable Application Control	Select an application sensor from the drop-down list. You can also select <i>Create New</i> within the drop-down list to add a new application sensor.
Enable Antivirus	Select an antivirus profile from the drop-down list. You can also select <i>Create New</i> within the drop-down list to add a new antivirus profile.
Enable Web Filter	Select a web filter profile from the drop-down list. You can also select <i>Create New</i> within the drop-down list to add a new web filter profile.
Enable DLP sensor	Select a DLP sensor from the drop-down list. You can also select <i>Create New</i> within the drop-down list to add a new DLP sensor.



Other UTM considerations

The following topics are included in this section:

- UTM and Virtual domains (VDOMs)
- Conserve mode
- SSL content scanning and inspection
- Viewing and saving logged packets
- Using wildcards and Perl regular expressions

UTM and Virtual domains (VDOMs)

If you enable virtual domains (VDOMs) on your FortiGate unit, all UTM configuration is limited to the VDOM in which you configure it.

While configuration is not shared, the various databases used by UTM features are shared. The FortiGuard antivirus and IPS databases and database updates are shared. The FortiGuard web filter and spam filter features contact the FortiGuard distribution network and access the same information when checking email for spam and web site categories and classification.

Conserve mode

FortiGate units perform all UTM processing in physical RAM. Since each model has a limited amount of memory, conserve mode is activated when the remaining free memory is nearly exhausted or the AV proxy has reached the maximum number of sessions it can service. While conserve mode is active, the AV proxy does not accept new sessions.

The AV proxy

Most content inspection the FortiGate unit performs requires that the files, email messages, URLs, and web pages be buffered and examined as a whole. The AV proxy performs this function, and because it may be buffering many files at the same time, it uses a significant amount of memory. Conserve mode is designed to prevent all the component features of the FortiGate unit from trying to use more memory than it has. Because the AV proxy uses so much memory, conserve mode effectively disables it in most circumstances. As a result, the content inspection features that use the AV proxy are also disabled in conserve mode.

All of the UTM features use the AV proxy with the exception of IPS, application control, DoS as well as flow-based antivirus, DLP, and web filter scanning. These features continue to operate normally when the FortiGate unit enters conserve mode.

Entering and exiting conserve mode

A FortiGate unit will enter conserve mode because it is nearly out of physical memory, or because the AV proxy has reached the maximum number of sessions it can service. The memory threshold that triggers conserve mode varies by model, but it is about 20% free memory. When memory use rises to the point where less than 20% of the physical memory is free, the FortiGate unit enters conserve mode.

The FortiGate unit will leave conserve mode only when the available physical memory exceeds about 30%. When exiting conserve mode, all new sessions configured to be scanned with features requiring the AV proxy will be scanned as normal, with the exception of a unit configured with the one-shot option.

Conserve mode effects

What happens when the FortiGate unit enters conserve mode depends on how you have `av-failopen` configured. There are four options:

off

The off setting forces the FortiGate unit to stop all traffic that is configured for content inspection by UTM features that use the AV proxy. New sessions are not allowed but current sessions continue to be processed normally unless they request more memory. Sessions requesting more memory are terminated.

For example, if a security policy is configured to use antivirus scanning, the traffic it permits is blocked while in conserve mode. A policy with IPS scanning enabled continues as normal. A policy with both IPS and antivirus scanning is blocked because antivirus scanning requires the AV proxy.

Use the off setting when security is more important than a loss of access while the problem is rectified.

pass

The pass setting allows traffic to bypass the AV proxy and continue to its destination. Since the traffic is bypassing the proxy, no UTM scanning that requires the AV proxy is performed. UTM scanning that does not require the AV proxy continues normally.

Use the pass setting when access is more important than security while the problem is rectified.

Pass is the default setting.

one-shot

The one-shot setting is similar to pass in that traffic is allowed when conserve mode is active. The difference is that a system configured for one-shot will force new sessions to bypass the AV proxy even after it leaves conserve mode. The FortiGate unit resumes use of the AV proxy only when the `av-failopen` setting is changed or the unit is restarted.

idledrop

The idledrop setting will recover memory and session space by terminating all the sessions associated with the host that has the most sessions open. The FortiGate may force this session termination a number of times, until enough memory is available to allow it to leave conserve mode.

The idledrop setting is primarily designed for situations in which malware may continue to open sessions until the AV proxy cannot accept more new sessions, triggering conserve mode. If your FortiGate unit is operating near capacity, this setting could cause the termination of valid sessions. Use this option with caution.

Configuring the av-failopen command

You can configure the av-failopen command using the CLI.

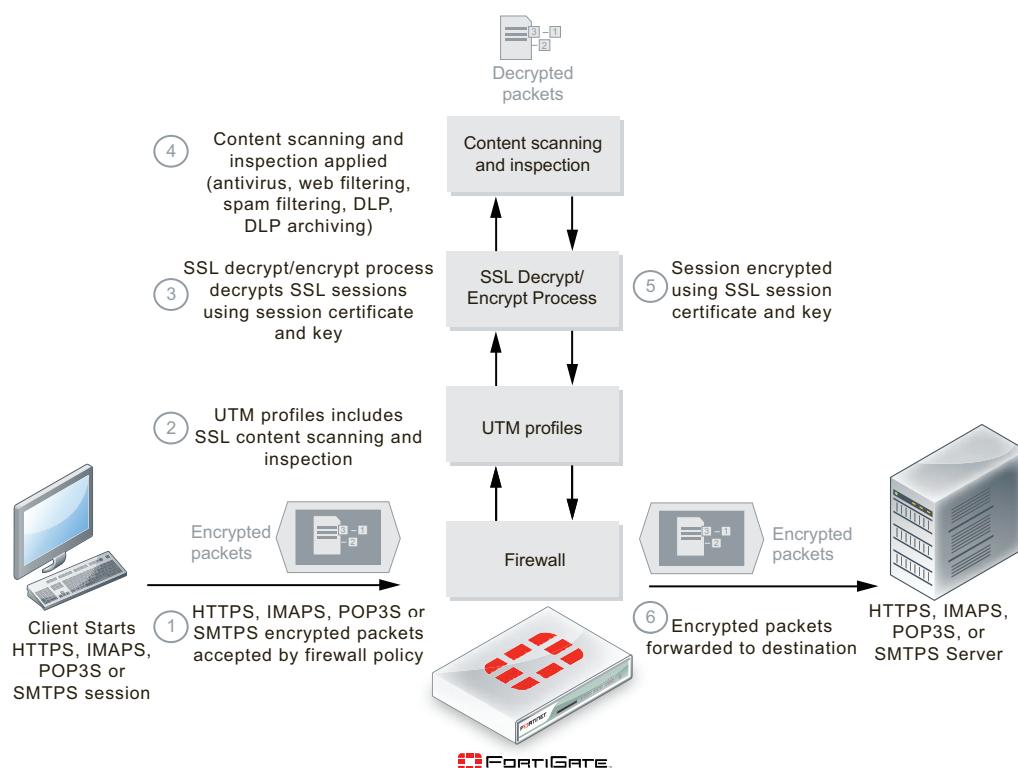
```
config system global
    set av-failopen {off | pass | one-shot | idledrop}
end
```

The default setting is pass.

SSL content scanning and inspection

If your FortiGate model supports SSL content scanning and inspection, you can apply antivirus scanning, web filtering, FortiGuard Web Filtering, and email filtering to encrypted traffic. You can also apply DLP and DLP archiving to HTTPS, IMAPS, POP3S, and SMTPS traffic. To perform SSL content scanning and inspection, the FortiGate unit does the following:

- intercepts and decrypts HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions between clients and servers (FortiGate SSL acceleration speeds up decryption)
- applies content inspection to decrypted content, including:
 - HTTPS, IMAPS, POP3S, and SMTPS Antivirus, DLP, and DLP archiving
 - HTTPS web filtering and FortiGuard web filtering
 - IMAPS, POP3S, and SMTPS email filtering
- encrypts the sessions and forwards them to their destinations.

Figure 102: FortiGate SSL content scanning and inspection packet flow

Setting up certificates to avoid client warnings

To use SSL content scanning and inspection, you need to set up and use a certificate that supports it. FortiGate SSL content scanning and inspection intercepts the SSL keys that are passed between clients and servers during SSL session handshakes and then substitutes spoofed keys. Two encrypted SSL sessions are set up, one between the client and the FortiGate unit, and a second one between the FortiGate unit and the server. Inside the FortiGate unit the packets are decrypted.

While the SSL sessions are being set up, the client and server communicate in clear text to exchange SSL session keys. The session keys are based on the client and server certificates. The FortiGate SSL decrypt/encrypt process intercepts these keys and uses a built-in signing CA certificate named `Fortinet_CA_SSLProxy` to create keys to send to the client and the server. This signing CA certificate is used only by the SSL decrypt/encrypt process. The SSL decrypt/encrypt process then sets up encrypted SSL sessions with the client and server and uses these keys to decrypt the SSL traffic to apply content scanning and inspection.

Some client programs (for example, web browsers) can detect this key replacement and will display a security warning message. The traffic is still encrypted and secure, but the security warning indicates that a key substitution has occurred.

You can stop these security warnings by importing the signing CA certificate used by the server into the FortiGate unit SSL content scanning and inspection configuration. Then the FortiGate unit creates keys that appear to come from the server and not the FortiGate unit.



You can add one signing CA certificate for SSL content scanning and inspection. The CA certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported for SSL content scanning and encryption.

You can replace the default signing CA certificate, Fortinet_CA_SSLProxy, with another signing CA certificate. To do this, you need the signing CA certificate file, the CA certificate key file, and the CA certificate password.

All SSL content scanning and inspection uses the same signing CA certificate. If your FortiGate unit is operating with virtual domains enabled, the same signing CA certificate is used by all virtual domains.

To add a signing CA certificate for SSL content scanning and inspection

- 1 Obtain a copy of the signing CA certificate file, the CA certificate key file, and the password for the CA certificate.
- 2 Go to *System > Certificates > Local Certificates* and select *Import*.
- 3 Set *Type* to *Certificate*.
- 4 For *Certificate file*, use the *Browse* button to select the signing CA certificate file.
- 5 For *Key file*, use the *Browse* button to select the CA certificate key file.
- 6 Enter the CA certificate *Password*.
- 7 Select *OK*.

The CA certificate is added to the *Local Certificates* list. In this example the signing CA certificate name is *Example_CA*. This name comes from the certificate file and key file name. If you want the certificate to have a different name, change these file names.

- 8 Add the imported signing CA certificate to the SSL content scanning and inspection configuration. Use the following CLI command if the certificate name is *Example_CA*.

```
config firewall ssl setting
  set caname Example_CA
end
```

The *Example_CA* signing CA certificate will now be used by SSL content scanning and inspection for establishing encrypted SSL sessions.

SSL content scanning and inspection settings

If SSL content scanning and inspection is available on your FortiGate unit, you can configure SSL settings. The following table provides an overview of the options available and where to find further instruction:

Table 74: SSL content scanning and inspection settings

Setting	Description
Predefined firewall services	The IMAPS, POP3S and SMTPS predefined services. You can select these services in a security policy and a DoS policy.

Table 74: SSL content scanning and inspection settings (Continued)

Setting	Description
Protocol recognition	<p>The TCP port numbers that the FortiGate unit inspects for HTTPS, IMAPS, POP3S, and SMTPS. Go to <i>Policy > Policy > Protocol Options</i>. Add or edit a protocol options profile, configure HTTPS, IMAPS, POP3S, SMTPS, and FTPS.</p> <p>Using <i>Protocol Options</i>, you can also configure the FortiGate unit to perform URL filtering of HTTPS or to use SSL content scanning and inspection to decrypt HTTPS so that the FortiGate unit can also apply antivirus and DLP content inspection and DLP archiving to HTTPS. Using SSL content scanning and inspection to decrypt HTTPS also allows you to apply more web filtering and FortiGuard Web Filtering options to HTTPS.</p> <p>To enable full SSL content scanning of web filtering, select <i>Enable Deep Scanning</i> under HTTPS in the protocol options profile.</p>
Antivirus	<p>Antivirus options including virus scanning and file filtering for HTTPS, IMAPS, POP3S, and SMTPS.</p> <p>Go to <i>UTM AntiVirus > Profile</i>. Add or edit a profile and configure <i>Virus Scan</i> for HTTPS, IMAPS, POP3S, and SMTPS.</p>
Antivirus quarantine	<p>Antivirus quarantine options to quarantine files in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.</p> <p>Go to <i>UTM Profiles > AntiVirus > Quarantine</i>. You can quarantine infected files, suspicious files, and blocked files found in HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.</p>
Web filtering	<p>Web filtering options for HTTPS:</p> <ul style="list-style-type: none"> • Web Content Filter • Web URL Filter • ActiveX Filter • Cookie Filter • Java Applet Filter • Web Resume Download Block • Block invalid URLs <p>Go to <i>UTM Profiles > Web Filter > Profile</i>. Add or edit a web filter profile and configure web filtering for HTTPS.</p>

Table 74: SSL content scanning and inspection settings (Continued)

Setting	Description
FortiGuard Web Filtering	<p>FortiGuard Web Filtering options for HTTPS:</p> <ul style="list-style-type: none"> • Enable FortiGuard Web Filtering • Enable FortiGuard Web Filtering Overrides • Provide Details for Blocked HTTP 4xx and 5xx Errors • Rate Images by URL (Blocked images will be replaced with blanks) • Allow Websites When a Rating Error Occurs • Strict Blocking • Rate URLs by Domain and IP Address • Block HTTP Redirects by Rating <p>Go to <i>UTM Profiles > Web Filter > Profile</i>. Add or edit a profile and configure FortiGuard Web Filtering for HTTPS.</p>
Email filtering	<p>Email filtering options for IMAPS, POP3S, and SMTPS:</p> <ul style="list-style-type: none"> • FortiGuard Email Filtering IP Address Check, URL check, E-mail Checksum Check, and Spam Submission • IP Address BWL Check • E-mail Address BWL Check • Return S-mail DNS Check • Banned Word Check • Spam Action • Tag Location • Tag Format <p>Go to <i>UTM Profiles > Email Filter > Profile</i>. Add or edit a profile and configure email filtering for IMAPS, POP3S, and SMTPS.</p>
Data Leak Prevention	<p>DLP for HTTPS, IMAPS, POP3S, and SMTPS. To apply DLP, follow the steps below:</p> <ul style="list-style-type: none"> • Go to <i>UTM Profiles > Data Leak Prevention > Sensor</i>, create a new DLP sensor or edit an existing one and then add any combination of the DLP advanced rules, DLP compound rules, file filters, a Regular Expressions, and file size limits to a DLP sensor. • Go to <i>Policy > Policy > Protocol Options</i>. Add or edit a profile and select <i>Enable Deep Scan</i> under HTTPS. • Go to <i>Policy > Policy > Policy</i>, edit the required policy, enable UTM, select <i>Enable DLP Sensor</i> and select the DLP sensor. • Go to <i>Policy > Policy > Policy</i>, edit the required policy, enable <i>Protocol Options</i> and select a profile that has <i>Enable Deep Scan</i> selected under HTTPS. Note: If no protocol options profile is selected, or if <i>Enable Deep Scan</i> is not selected within the protocol options profile, DLP rules cannot inspect HTTPS.
DLP archiving	<p>DLP archiving for HTTPS, IMAPS, POP3S, and SMTPS. Add DLP Rules for the protocol to be archived.</p>

Table 74: SSL content scanning and inspection settings (Continued)

Setting	Description
Monitor DLP content information on the system dashboard	<p>DLP archive information on the Log and Archive Statistics widget on the system dashboard for HTTPS, IMAPS, POP3S, and SMTPS.</p> <p>Go to <i>Policy > Policy > Protocol Options</i>. Add or edit a profile. For each protocol you want monitored on the dashboard, enable <i>Monitor Content Information for Dashboard</i>.</p> <p>These options display meta-information on the Statistics dashboard widget.</p>

Viewing and saving logged packets

The FortiGate unit supports packet logging for IPS and application control. The packets that trigger a signature match for IPS or application recognition for application control are saved for later viewing when packet logging is enabled.

For information on how to enable packet logging, see [“Enable IPS packet logging” on page 967](#) and [“Application control packet logging” on page 1070](#).

Once the FortiGate unit has logged packets, you can view or save them.

To view and save logged packets

- 1 Go to *Log&Report > Log Access > Attack*.
- 2 Depending on where the logs are configured to be stored, select the appropriate option:

Memory	Select if logs are stored in the FortiGate unit memory.
Disk	Select if the FortiGate unit has an internal hard disk and logs are stored there.
Remote	Select if logs are sent to a FortiAnalyzer unit or to the FortiGuard Analysis and Management Service.

- 3 Select the *Packet Log* icon of the log entry you want to view.
The *IPS Packet Log Viewer* window appears.
 - 4 Select the packet to view the packet in binary and ASCII. Each table row represents a captured packet.
 - 5 Select *Save* to save the packet data in a PCAP formatted file.
- PCAP files can be opened and examined in network analysis software such as Wireshark.

Configuring packet logging options

You can use a number of CLI commands to further configure packet logging.

Limiting memory use

When logging to memory, you can define the maximum amount of memory used to store logged packets.

```
config ips settings
    set packet-log-memory 256
end
```

The acceptable range is from 64 to 8192 kilobytes. This command affects only logging to memory.

Limiting disk use

When logging to the FortiGate unit internal hard disk, you can define the maximum amount of space used to store logged packets.

```
config ips settings
  set ips-packet-quota 256
end
```

The acceptable range is from 0 to 4294967295 megabytes. This command affects only logging to disk.

Configuring how many packets are captured

Since the packet containing the signature is sometimes not sufficient to troubleshoot a problem, you can specify how many packets are captured before and after the packet containing the IPS signature match.

```
config ips settings
  packet-log-history
  packet-log-post-attack
end
```

The `packet-log-history` command specifies how many packets are captured before and including the one in which the IPS signature is detected. If the value is more than 1, the packet containing the signature is saved in the packet log, as well as those preceding it, with the total number of logged packets equalling the `packet-log-history` setting. For example, if `packet-log-history` is set to 7, the FortiGate unit will save the packet containing the IPS signature match and the six before it.

The acceptable range for `packet-log-history` is from 1 to 255. The default is 1.



Setting `packet-log-history` to a value larger than 1 can affect the performance of the FortiGate unit because network traffic must be buffered. The performance penalty depends on the model, the setting, and the traffic load.

The `packet-log-post-attack` command specifies how many packets are logged after the one in which the IPS signature is detected. For example, if `packet-log-post-attack` is set to 10, the FortiGate unit will save the ten packets following the one containing the IPS signature match.

The acceptable range for `packet-log-post-attack` is from 0 to 255. The default is 0.

Using wildcards and Perl regular expressions

Many UTM feature list entries can include wildcards or Perl regular expressions.

For more information about using Perl regular expressions, see <http://perldoc.perl.org/perlretut.html>.

Regular expression vs. wildcard match pattern

A wildcard character is a special character that represents one or more other characters. The most commonly used wildcard characters are the asterisk (*), which typically represents zero or more characters in a string of characters, and the question mark (?), which typically represents any one character.

In Perl regular expressions, the '.' character refers to any single character. It is similar to the '?' character in wildcard match pattern. As a result:

- example.com not only matches example.com but also examplea.com, exampleb.com, examplec.com, and so on.



To add a question mark (?) character to a regular expression from the FortiGate CLI, enter Ctrl+V followed by ?. To add a single backslash character (\) to a regular expression from the CLI you must add precede it with another backslash character. For example, `example\\.com`.

To match a special character such as '.' and '*' use the escape character '\\'. For example:

- To match example.com, the regular expression should be: `example\\.com`

In Perl regular expressions, '*' means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- `exam*.com` matches `exammmmm.com` but does not match `example.com`

To match any character 0 or more times, use '.' where '.' means any character and the '*' means 0 or more times. For example, the wildcard match pattern `exam*.com` should therefore be `exam.*\\.com`.

Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression "test" not only matches the word "test" but also any word that contains "test" such as "atest", "mytest", "testimony", "atestb". The notation "\\b" specifies the word boundary. To match exactly the word "test", the expression should be `\\btest\\b`.

Case sensitivity

Regular expression pattern matching is case sensitive in the web and Email Filter filters. To make a word or phrase case insensitive, use the regular expression `/i`. For example, `/bad language/i` will block all instances of "bad language", regardless of case.

Perl regular expression formats

Table 75 lists and describes some example Perl regular expressions.

Table 75: Perl regular expression formats

Expression	Matches
<code>abc</code>	"abc" (the exact character sequence, but anywhere in the string)
<code>^abc</code>	"abc" at the beginning of the string
<code>abc\$</code>	"abc" at the end of the string
<code>a b</code>	Either "a" or "b"
<code>^abc abc\$</code>	The string "abc" at the beginning or at the end of the string
<code>ab{2,4}c</code>	"a" followed by two, three or four "b"s followed by a "c"
<code>ab{2,}c</code>	"a" followed by at least two "b"s followed by a "c"
<code>ab*c</code>	"a" followed by any number (zero or more) of "b"s followed by a "c"
<code>ab+c</code>	"a" followed by one or more b's followed by a c

Table 75: Perl regular expression formats (Continued)

ab?c	“a” followed by an optional “b” followed by a “c”; that is, either “abc” or “ac”
a.c	“a” followed by any single character (not newline) followed by a “c”
a\.c	“a.c” exactly
[abc]	Any one of “a”, “b” and “c”
[Aa]bc	Either of “Abc” and “abc”
[abc]+	Any (nonempty) string of “a”s, “b”s and “c”s (such as “a”, “abba”, “acbabcacaa”)
[^abc]+	Any (nonempty) string which does not contain any of “a”, “b”, and “c” (such as “defg”)
\d\d	Any two decimal digits, such as 42; same as \d{2}
/i	Makes the pattern case insensitive. For example, /bad language/i blocks any instance of bad language regardless of case.
\w+	A “word”: A nonempty sequence of alphanumeric characters and low lines (underscores), such as foo and 12bar8 and foo_1
100\s*mk	The strings “100” and “mk” optionally separated by any amount of white space (spaces, tabs, newlines)
abc\b	“abc” when followed by a word boundary (for example, in “abc!” but not in “abcd”)
perl\b	“perl” when not followed by a word boundary (for example, in “perlert” but not in “perl stuff”)
\x	Tells the regular expression parser to ignore white space that is neither preceded by a backslash character nor within a character class. Use this to break up a regular expression into (slightly) more readable parts.
/x	Used to add regular expressions within other text. If the first character in a pattern is forward slash '/', the '/' is treated as the delimiter. The pattern must contain a second '/'. The pattern between '/' will be taken as a regular expressions, and anything after the second '/' will be parsed as a list of regular expression options ('i', 'x', etc). An error occurs if the second '/' is missing. In regular expressions, the leading and trailing space is treated as part of the regular expression.

Examples of regular expressions

Block any word in a phrase

```
/block|any|word/
```

Block purposely misspelled words

Spammers often insert other characters between the letters of a word to fool spam blocking software.

```
/^.*v.*i.*a.*g.*r.*o.*$/i
```

```
/cr[eéeêë] [\+ \- \* = < > \. \, ; ! \? % & $ @ \^ ° \ $ £ € \{ \} () \[ \] \ | \ \ 01]dit/i
```

Block common spam phrases

The following phrases are some examples of common phrases found in spam messages.

```
/try it for free/i
```

```
/student loans/i
```

```
/you're already approved/i
```

```
/special[\+\\-\\*=<>\\.\\,;!\\?%&~#\\$@\\^\\°\\$£€\\{\\}()\\[\\]\\\\\\_1]offer/i
```

Protocol Options interface reference

The Protocol Options menu allows you to configure settings for specific protocols, which are grouped together in a protocol group, and then applied to a security policy. The default groups are scan, strict, unfiltered, and web.

Protocol options configuration settings

The following are protocol option configuration settings in *Policy > Policy > Protocol Options*.

Protocol Options page Lists each individual protocol setting that you created. On this page, you can edit, delete or create a new group of protocol settings. Note: If you want to provide information about any of the protocols, select the check box beside <i>Monitor Content Information for Dashboard</i> , which is available within each protocol section.	
Create New	Creates a new protocol option. When you select <i>Create New</i> , you are automatically redirected to the Protocol Options Settings page.
Edit	Modifies settings to a protocol setting. When you select <i>Edit</i> , you are automatically redirected to the Protocol Options page.
Delete	Removes a protocol setting from the list on the Protocol Options page. To remove multiple protocol settings from within the list, on the Protocol Options page, in each of the rows of the policies you want removed, select the check box and then select <i>Delete</i> . To remove all protocol options from the list, on the Protocol Options page, select the check box in the check box column, and then select <i>Delete</i> .
Name	The name of the protocol group. This group is the group you select when applying it to a security policy.
Comments	Describes the protocol group.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when the icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
<p>Protocol Options Settings page</p> <p>Provides settings for configuring options for each protocol which make up a protocol group.</p> <p>Note: There are similar settings within each protocol setting.</p>	
Name	Enter a name for the protocol group.
Comments	Enter a description about the protocol group. This is optional.
Enable Oversized File Log	Select to allow logging of oversized files.
Enable Invalid Certificate Log	Select to allow logging of invalid certificates.
HTTP section	Configure settings for the HTTP protocol or the HTTPS protocol.
Port (i.e. 80,88, 0-auto)	This is available for every protocol except for IM.
Comfort Clients	<p>This is available only for HTTP, FTP, and HTTPS.</p> <ul style="list-style-type: none"> • Interval (1-900 seconds) – enter the interval time in seconds. • Amount (1-10240 bytes) – enter the amount in bytes.
Oversized File/Email	<p>This is available for all protocols.</p> <ul style="list-style-type: none"> • Threshold – enter the threshold amount for an oversized email message or file in MB.

Monitor Content Information for Dashboard	Select to view the activity of the protocol from the Dashboard menu.
Enable Chunked Bypass	Select to enable the chunked bypass setting.
FTP section	Configure settings for the file transfer protocol. <i>FTP</i> and <i>HTTP</i> contain the same settings, except the FTP section does not contain the option <i>Enable Chunked Bypass</i> .
FTPS section	Configure settings for the FTPS protocol. FTPS is an extension of the FTP protocol, adding support for both the TLS and SSL cryptographic protocols. This section contains the same settings as in the <i>FTP</i> section.
IMAP section	Configure settings for the IMAP protocol.
Allow Fragmented Messages	Allows fragmented email messages to be passed.
POP3 section	Configure settings for the POP3 protocol. This section contains the same settings as are in the IMAP section.
SMTP section	Configure settings for the SMTP section.
Append Email Signature	Select to enable the option of entering a new email signature that appears in the email message.
Email Signature Text	Enter a signature for the email message, for example, Yours sincerely. Accessible only when Append Email Signature is selected.
IM section	Configure settings for the IM protocol.
NNTP section	Configure settings for the NNTP protocol.
HTTPS section	Configure settings for the HTTPS protocol.
Allow Invalid SSL Certificate	Select to allow invalid SSL certificates.
Enable Deep Scanning	Select to allow deep scanning.
IMAPS	Configure settings for the IMAPS protocol.
POP3S	Configure settings for the POP3S protocol. This section contains the same settings as <i>IMAPS</i> .
SMTPS	Configure settings for the SMTPS protocol. This section contains the same settings as <i>IMAPS</i> and <i>POP3S</i> .

ICAP interface reference

The Internet Content Adaption Protocol (ICAP) is supported on the unit. ICAP is a light-weight response/request protocol that allows you to offload specific filtering features, such as virus scanning and content filtering, to dedicated ICAP servers.

ICAP does not appear by default in the web-based manager. You must enable it in *System > Admin > Settings* to display ICAP in the web-based manager.

In an ICAP configuration with the unit, the unit is the surrogate or “middle-man” and carries the ICAP responses from the ICAP server to the ICAP client. The ICAP client then responds back, and the unit determines the action that should be taken with these responses and requests.

An ICAP profile must be configured and then applied to a firewall policy for this feature to work. The ICAP server or servers must also be configured. The ICAP servers are applied within the ICAP profile.

This topic contains the following:

- [ICAP profile](#)
- [ICAP server](#)

ICAP profile

The following are ICAP profile configuration settings in *UTM Profiles > ICAP > Profile*.

ICAP Profile page	
Lists each ICAP profile that you created. On this page, you can edit, delete and create a new ICAP profiles.	
Create New	Creates a new ICAP profile. When you select <i>Create New</i> , you are automatically redirected to the New ICAP Profile page.
Edit	Modifies settings within an ICAP profile. When you select <i>Edit</i> , you are automatically redirected to the Edit ICAP Profile page.
Delete	Removes the profile from within the list on the page.
Name	The name of the ICAP profile.
Request Processing	The status of the process for requests. A gray x appears if there is no request processing settings enabled and configured. A green check mark appears if request processing settings are enabled and configured.
Response Processing	The status of the process for responses. A gray x appears if there is no response processing settings enabled and configured. A green check mark appears if response processing settings are enabled and configured.
Bypass Streaming Media	The status of whether bypass streaming media is enabled or disabled. A gray x appears if bypass streaming media is disabled. A green check mark appears if bypass streaming media is enabled.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
New ICAP Profile page	
Provides settings for configuring the ICAP profiles.	
Note: Logging is enabled in the CLI.	
Name	Enter the name of the new ICAP profile.
Enable Request Processing	Select to enable how the request from an ICAP server will be processed by the unit.
Server	Select the ICAP server from the drop-down list. An ICAP server must be configured before you can configure an ICAP profile. See “ICAP server” on page 1153 .
Path	Enter the path name.
On Failure	Select either <i>Error</i> or <i>Bypass</i> . This instructs the unit on which action to take if a failure occurs.
Enable Response Processing	Select to enable how the response from an ICAP server will be processed by the unit.
Server	Select the ICAP server from the drop-down list. An ICAP server must be configured before you can configure an ICAP profile. See “ICAP server” on page 1153 .
Path	Enter the path name.
On Failure	Select either <i>Error</i> or <i>Bypass</i> . This instructs the unit on which action to take if a failure occurs.
Enable Streaming Media Bypass	Select to enable streaming media bypass.

ICAP server

The following are ICAP server configuration settings in *UTM Profiles > ICAP > Server*.

ICAP Server page Lists each ICAP profile that you created. On this page, you can edit, delete and create a new ICAP profiles.	
Create New	Adds an ICAP server. When you select <i>Create New</i> , you are automatically redirected to the New ICAP Server page.
Edit	Modifies settings within an ICAP server configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit ICAP Profile page.
Delete	Removes the profile from within the list on the page.
Name	The name of the ICAP profile.
Address	The IP address of the ICAP server.
Port	The port number used by the ICAP server.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.
New ICAP Server page Provides settings for configuring the profiles. Note: Logging is enabled in the CLI.	
Name	Enter the name of the new ICAP profile.
IP Type	Enter the type of IP address the ICAP server uses, either IPv4 or IPv6.

IP Address	Enter the IP address of the ICAP server.
Port	Enter the port that the ICAP server uses.

Profile Group interface reference

A profile group is a group of UTM features that include MMS profiles and replacement message groups. You can configure multiple profile groups, which are then applied to a firewall policy. This provides an easier way to apply multiple profiles and sensors to a firewall policy.

Profile Group configuration settings

The following are profile group configuration settings. If you are running FortiOS Carrier, these configuration settings are in *UTM Profiles > Carrier > Profile Group*; however, if you are running FortiOS, these configuration settings are in *UTM Profiles > Profile Group > Profile Group*.

Profile Group page	
Lists all the profile groups that you have created. On this page, you can edit, delete or create a new profile group.	
Create New	Creates a new profile group. When you select <i>Create New</i> , you are automatically redirected to the New Profile Group page.
Edit	Modifies a profile group's settings. When you select <i>Edit</i> , you are automatically redirected to the Edit Profile Group page.
Delete	Removes a profile group from the list. To remove multiple profile groups from within the list, on the Profile Group page, in each of the rows of the groups you want removed, select the check box and then select <i>Delete</i> . To remove all profile groups from the list, on the Profile Group page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the profile group.
New Profile Group page	
Provides settings for configuring a profile group. You must configure UTM features prior to configuring a profile group because profile groups require these UTM features. If you want to include a replacement message group, that group must also be configured prior to configuring a profile group. You cannot configure a UTM feature or replacement group from the New Profile Group page.	
Name	Enter a name for the profile group.
Protocol Options	Select the check box to enable this option and then select a protocol option from the drop-down list. Protocol options are configured in <i>Firewall > Policy > Protocol Options</i> .
Enable Antivirus	Select the check box to enable this option and then select an antivirus profile from the drop-down list. Antivirus profiles are configured in <i>UTM Profiles > Antivirus > Profile</i> .

Enable IPS	Select the check box to enable this option and then select the IPS sensor from the drop-down list. IPS sensors are configured in <i>UTM Profiles > Intrusion Protection > IPS Sensor</i> .
Enable Web Filter	Select the check box to enable this option and then select the web filter profile from the drop-down list. Web filter profiles are configured in <i>UTM Profiles > Web Filter > Profile</i> .
Enable Email Filter	Select the check box to enable this option and then select the email filter profile from the drop-down list. Email filter profiles are configured in <i>UTM Profiles > Email Filter > Profile</i> .
Enable DLP Sensor	Select the check box to enable this option and then select the DLP sensor from the drop-down list. DLP sensors are configured in <i>UTM Profiles > Data Leak Prevention > Sensors</i> .
Enable Application Control	Select the check box to enable this option and then select the application control list from the drop-down list. Application control lists are configured in <i>UTM Profiles > Application Control > Application Control Lists</i> .
Enable VoIP	Select the check box to enable this option and then select the VoIP profile from the drop-down list. VoIP profiles are configured in <i>UTM Profiles > VoIP > Profile</i> .
Enable MMS Profile	Select the check box to enable this option and then select the MMS profile from the drop-down list. MMS profiles are configured in <i>UTM Profiles > Carrier > MMS Profile</i> .
Enable Replacement Message Group	Select the check box to enable this option and then select the replacement message group from the drop-down list. Replacement message groups are configured in <i>System > Config > Replacement Message Group</i> .

Monitor interface reference

The Monitor submenus allow you to view the UTM activity occurring on your network. You must have UTM profiles and sensors applied to firewall policies, as well as logging enabled for the profiles and sensors, for the monitors to display any information regarding this activity.

This topic contains the following:

- [AV Monitor](#)
- [Intrusion Monitor](#)
- [Web Monitor](#)
- [Email Monitor](#)
- [Archive & Data Leak Monitor](#)
- [Application Monitor](#)

AV Monitor

The AV Monitor submenu allows you to view statistical information regarding viruses that were detected on your unit from *UTM Profiles > Monitor > AV Monitor*. The information displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.



You must have antivirus logging enabled for this within the profile itself, as well as within log settings and an antivirus profile is applied to a firewall policy.

AV Monitor page

Displays monitored information about viruses that were detected by the unit.

Tip: To view information about a specific virus, select a bar within the chart; the virus FortiGuard definition displays.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Viruses (all policies) since <yyyy-mm-dd hh:mm:ss>	The top viruses detected by the unit using all firewall policies.
#	The order that the viruses are listed in the table.
Virus Name	The name of the virus.
Last Detected	The last time that the virus was detected.
Count	The number of times the virus has been detected.

Intrusion Monitor

The Intrusion Monitor submenu allows you to view statistical information regarding attacks that were detected on your unit from *UTM Profiles > Monitor > Intrusion Monitor*. The information displays in a bar chart as well as in a table below the bar chart. The table contains detailed information.

Intrusion Monitor page

Displays monitored information about attacks that were detected by the unit.

Tip: To view information about a specific attack, select a bar within the chart; the attack FortiGuard definition displays.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Attacks (all policies) since <yyyy-mm-dd hh:mm:ss>	A bar chart displaying the top attacks detected by the unit.

#	The order that the attacks are listed in the table.
Attack Name	The name of the attack.
Last Detected	The last time that the attack was detected.
Count	The number of times the attack has been detected.

Web Monitor

The Web Monitor submenu allows you to view statistical information regarding the web activity from *UTM Profiles > Monitor > Web Monitor*. The information displays in both a pie chart and a bar chart

Web Monitor page Displays monitored information about web activity detected by the unit.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Report By	Select whether to view the web filter monitored information by web filter technique or by FortiGuard web filter category. If you choose FortiGuard web filter category, you are viewing the information that was gathered from the category settings for FortiGuard web filter from the web filter profile.
Web Monitor since <yyyy-mm-dd hh:mm:ss>	
Total Requests (HTTP)	A pie chart representing the total requests detected.
Blocked Requests (HTTP)	A bar chart representing the total blocked requests detected. The information is broken down to spam, banned words, file filter, viruses, archives, FortiGuard, URL filter, and fragmented.
Total Web Requests (HTTP): <number>	The total number of web requests over HTTP that occurred.

Email Monitor

The Email Monitor submenu allows you to view statistical information regarding email filtering from *UTM Profiles > Monitor > Email Monitor*. The information displays in both a pie chart and bar chart.

Email Monitor page Displays monitored information about email filter activity detected by the unit.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Total Emails	A pie chart representing the total number of emails scanned by the unit.

Blocked Emails	A bar chart representing the total number of blocked emails, broken down by protocol. The colors indicate the type of scanning that occurred.
Total Emails: <number>	The total number of email messages detected by the unit.

Archive & Data Leak Monitor

The Archive & Data Leak Monitor submenu allows you to view statistical information regarding log archives, as well as DLP usage. This page displays the information in a bar chart in *UTM Profiles > Monitor > Archive & Data Leak Monitor*.

Archive & Data Leak Monitor page Displays monitored information about archive and DLP activity detected by the unit.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Report By:	Select what type of DLP information you want to view. You can view DLP usage by DLP sensor, firewall policy usage, or by protocol.
Top DLP Usage by DLP Sensor <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using DLP sensor information.
Top DLP Usage by Policy <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using firewall policy traffic information.
Top DLP Usage by Protocol <yyyy-mm-dd hh:mm:ss>	The bar chart that displays DLP usage monitored using protocol information.
Total Dropped Archives: <number>	The total number of dropped DLP archives.

Application Monitor

The Application Monitor submenu allows you to view statistical information regarding application usage in *UTM Profiles > Monitor > Application Monitor*.

Application Monitor page Displays monitored information about the application usage detected by the unit. Tip: To view top source IP addresses for a specific application, select a bar in the chart to view that application's source IP addresses.	
Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.

Top Application Usage by <yyyy-mm-dd hh:mm:ss>	The bar chart that displays the top applications being used detected by the unit.
Resolve Host Name	Appears after selecting a bar for a specific application, for example SSL. Select to resolve the host name. Tip: Hover your mouse over the bar to view the address and total MB (or KB) used for that application.
Report By:	Appears after selecting a bar for a specific application, for example, SSL. Select to view the detailed information by destination address, or source address.
Display User Name	Appears after selecting <i>Source Address</i> from the drop-down list beside <i>Report By</i> . Select to display user names.

FortiGuard Quota

The FortiGuard Quota submenu allows you to view statistical information regarding quota usage by users in *UTM Profiles > Monitor > FortiGuard Quota*.

FortiGuard Quota page Lists the users and the amount of quota that they have used.	
Page Controls	Use to navigate through the list.
User Name	The user name of the user that has FortiGuard quota enabled for them.
Webfilter Profile	The web filter profile that was used for detecting users' FortiGuard quota usage.
Used Quota	The amount of used quota by a user.

Endpoint Monitor

You can view monitored endpoints in *UTM Profiles > Monitor > Endpoint Monitor*. An endpoint is added to the list when it uses a security policy that has *Endpoint Security* enabled.

Endpoint Monitor page Provides information about endpoints, such as endpoint traffic. Note: The pie chart displays information in percent and indicates which is non-compliant and which is compliant.	
Refresh	Updates the list, providing current endpoints that are being monitored.
Report By	Select to view endpoint information by traffic, status or application usage. When you select <i>Status</i> , a pie chart appears along with information about the total endpoints (<i>Total Endpoints</i>). When you select <i>Traffic</i> or <i>Application usage</i> , a bar chart appears; select a bar to view detailed information.



Chapter 7 User Authentication

This FortiOS Handbook chapter contains the following sections:

[Introduction to authentication](#) describes some basic elements and concepts of authentication.

[Authentication and User in the web-based manager](#) describes the web-based manager interface of FortiOS, specifically the Authentication and User top level menu items.

[Authentication servers](#) describes external authentication servers, where a FortiGate unit fits into the topology, and how to configure a FortiGate unit to work with that type of authentication server.

[Dynamic profiles and end points](#) describes how to set up dynamic profiles and carrier end points to identify users that are using RADIUS settings. Some parts are specific to FortiOS Carrier.

[Users and user groups](#) describes the different types of user accounts and user groups. Authenticated access to resources is based on user identities and user group membership. Two-factor authentication methods, including FortiToken, provide additional security.

[Configuring authenticated access](#) provides detailed procedures for setting up authenticated access in security policies and authenticated access to VPNs.

[FSSO integration with Windows AD or Novell](#) describes how to install and configure the Fortinet Single Sign On (FSSO) on Windows AD or Novell network domain controllers and the FortiGate unit. With FSSO, network users have single sign-on access to resources through the FortiGate unit. Earlier versions of FSSO were called FSAE.

[Certificate-based authentication](#) describes authentication by means of X.509 certificates.

[Monitoring authenticated users](#) describes FortiOS authenticated user monitor screens.

[Examples and Troubleshooting](#) provides a configuration example and troubleshooting in which Windows AD and other network users are provided authenticated access to the Internet.



Introduction to authentication

Identifying users and other computers—authentication—is a key part of network security. This section describes some basic elements and concepts of authentication.

The following topics are included in this section:

- [What is authentication?](#)
- [Methods of authentication](#)
- [Types of authentication](#)
- [User's view of authentication](#)
- [FortiGate administrator's view of authentication](#)

What is authentication?

Businesses need to authenticate people who have access to company resources. In the physical world this may be a swipe card to enter the building, or a code to enter a locked door. If a person has this swipe card or code, they have been authenticated as someone allowed in that building or room.

Authentication is the act of confirming the identity of a person or other entity. In the context of a private computer network, the identities of users or host computers must be established to ensure that only authorized parties can access the network. The FortiGate unit enables controlled network access and applies authentication to users of security policies and VPN clients.

Methods of authentication

FortiGate unit authentication is divided into three basic types: password authentication for people, certificate authentication for hosts or endpoints, and two-factor authentication for additional security beyond just passwords. An exception to this is that FortiGate units in an HA cluster and FortiManager units use password authentication.

Password authentication verifies individual user identities, but access to network resources is based on membership in user groups. For example, a security policy can be configured to permit access only to the members of one or more user groups. Any user who attempts to access the network through that policy is then authenticated through a request for their username and password.

Methods of authentication include:

- [Local password authentication](#)
- [Server-based password authentication](#)
- [Certificate-based authentication](#)
- [Two-factor authentication](#)

Local password authentication

The simplest authentication is based on user accounts stored locally on the FortiGate unit. For each account, a username and password is stored. The account also has a disable option so that you can suspend the account without deleting it.

Local user accounts work well for a single-FortiGate installation. If your network has multiple FortiGate units that will use the same accounts, the use of an external authentication server can simplify account configuration and maintenance.

You create local user accounts in the web-based manager under *User > User*. This page is also used to create accounts where an external authentication server stores and verifies the password.

Server-based password authentication

Using external LDAP, RADIUS, or TACACS+ authentication servers is desirable when multiple FortiGate units need to authenticate the same users, or where the FortiGate unit is added to a network that already contains an authentication server.

When you use an external authentication server to authenticate users, the FortiGate unit sends the user's entered credentials to the external server. The password is encrypted. The server's response indicates whether the supplied credentials are valid or not.

You must configure the FortiGate unit to access the external authentication servers that you want to use. The configuration includes the parameters that authenticate the FortiGate unit to the authentication server.

You can use external authentication servers in two ways:

- Create user accounts on the FortiGate unit, but instead of storing each user's password, specify the server used to authenticate that user. As with accounts that store the password locally, you add these users to appropriate user groups.
- Add the authentication server to user groups. Any user who has an account on the server can be authenticated and have the access privileges of the FortiGate user group. Optionally, when an LDAP server is a FortiGate user group member, you can limit access to users who belong to specific groups defined on the LDAP server.

Dynamic profiles

Managed Security Service Providers (MSSPs) and carrier service providers can use the FortiOS dynamic profile configuration to dynamically assign profile groups to customer traffic. Using the dynamic profile, FortiOS can receive RADIUS Start records from service provider accounting systems when customers connect to service provider networks. In real time, FortiOS can extract identifying information and profile group names from these RADIUS Start records and match the identifying information with the customer communication session. FortiOS can then dynamically select and apply the profile group named in the RADIUS Start record to the communication session. Some parts of dynamic profiles and end points are FortiOS Carrier-only features. See [“Dynamic profiles and end points” on page 1325](#).

Single Sign On authentication using FSSO

“Single sign on” or “Single user Sign On” means that users logged on to a computer network are authenticated for access to network resources through the FortiGate unit without having to enter their username and password again. Fortinet Single Sign On (FSSO) provides Single Sign On capability for:

- Microsoft Windows networks using either Active Directory or NTLM authentication
- Novell networks, using eDirectory

FSSO monitors user logons and sends the FortiGate unit the username, IP address, and the list of Windows AD user groups to which the user belongs. When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

For detailed information about FSSO, see [“FSSO integration with Windows AD or Novell” on page 1283](#).

Certificate-based authentication

An RSA X.509 server certificate is a small file issued by a Certificate Authority (CA) that is installed on a computer or FortiGate unit to authenticate itself to other devices on the network. When one party on a network presents the certificate as authentication, the other party can validate that the certificate was issued by the CA. The identification is therefore as trustworthy as the Certificate Authority (CA) that issued the certificate.

To protect against compromised or misused certificates, CAs can revoke any certificate by adding it to a Certificate Revocation List (CRL). Certificate status can also be checked online using Online Certificate Status Protocol (OCSP).

RSA X.509 certificates are based on public-key cryptography, in which there are two keys: the private key and the public key. Data encrypted with the private key can be decrypted only with the public key and vice versa. As the names suggest, the private key is never revealed to anyone and the public key can be freely distributed. Encryption with the recipient's public key creates a message that only the intended recipient can read. Encryption with the sender's private key creates a message whose authenticity is proven because it can be decrypted only with the sender's public key.

Server certificates contain a signature string encrypted with the CA's private key. The CA's public key is contained in a CA root certificate. If the signature string can be decrypted with the CA's public key, the certificate is genuine.

Certificate authorities

A certificate authority can be:

- an organization, such as VeriSign Inc., that provides certificate services
- a software application, such as Microsoft Certificate Services or OpenSSH

For a company web portal or customer-facing SSL VPN, a third-party certificate service has some advantages. The CA certificates are already included in popular web browsers and customers trust the third-party. On the other hand, third-party services have a cost.

For administrators and for employee VPN users, the local CA based on a software application provides the required security at low cost. You can generate and distribute certificates as needed. If an employee leaves the organization, you can simply revoke their certificate.

Certificates for users

FortiGate unit administrators and SSL VPN users can install certificates in their web browsers to authenticate themselves. If the FortiGate unit uses a CA-issued certificate to authenticate itself to the clients, the browser will also need the appropriate CA certificate.

FortiGate IPsec VPN users can install server and CA certificates according to the instructions for their IPsec VPN client software. The FortiClient Endpoint Security application, for example, can import and store the certificates required by VPN connections.

FortiGate units are also compatible with some Public Key Infrastructure systems. For an example of this type of system, see [“RSA ACE \(SecurID\) servers” on page 1217](#).

Two-factor authentication

A user can be required to provide both something they know (their username and password combination) and something they have (certificate or a random token code). Certificates are installed on the user's computer.

Two-factor authentication is available for PKI users. For more information, see [“Certificate” on page 1228](#).

Another type of two-factor authentication is to use a randomly generated token (multi-digit number) along with the username and password combination. One method is a FortiToken — a one time passcode (OTP) generator that generates a unique code every 60 seconds. Others use email or SMS text messaging to deliver the random token code to the user or administrator.

When one of these methods is configured, the user enters this code at login after the username and password have been verified. The FortiGate unit verifies the token code after as well as the password and username. For more information, see [“Two-factor authentication” on page 1228](#)

Types of authentication

FortiOS supports two different types of authentication based on your situation and needs.

Security policy authentication, or identity-based policies, is easily applied to all users logging on to a network, or network service. For example if a group of users on your network such as the accounting department who have access to sensitive data need to access the Internet, it is a good idea to make sure the user is a valid user and not someone trying to send company secrets to the Internet. Security policy authentication can be applied to as many or as few users as needed, and it supports a number of authentication protocols to easily fit with your existing network.

VPN authentication can be for both the remote VPN device as well as the VPN users. VPNs are used to communicate with locations outside the company network as if they were part of the company network. This level of trust, once a VPN is established, is easily established with authentication to verify the remote user is in fact a valid user. In this situation without authentication, anyone malicious or otherwise could connect to the company network with potentially full access.

Firewall authentication (identity-based policies)

Security policies enable traffic to flow between networks. If you want to limit which users have access to particular resources, you create identity-based policies (IBP) that allow access only to members of specific user groups. Authentication, a request for username and password, is triggered when a user attempts to access a resource for which data must pass through an identity-based policy.

The user's authentication expires if the connection is idle for too long, 30 minutes by default but that can be customized.

Identity-based policies are the mechanism for FSSO, NTLM, certificate based, and dynamic profile authentication.

FSSO

Fortinet Single Sign on (FSSO) provides seamless authentication support for Microsoft Windows Active Directory (AD) and Novell eDirectory users in a FortiGate environment.

On a Microsoft Windows or Novell network, users authenticate with the Active Directory or Novell eDirectory at logon. FSSO provides authentication information to the FortiGate unit so that users automatically get access to permitted resources. See [“Introduction to FSSO” on page 1283](#).

NTLM

The NT LAN Manager (NTLM) protocol is used when the MS Windows Active Directory (AD) domain controller can not be contacted. NTLM is a browser-based method of authentication.

The FSSO software is installed on each AD server and the FortiGate unit is configured to communicate with each FSSO client. When a user successfully logs into their Windows PC (and is authenticated by the AD Server), the FSSO client communicates the user's name, IP address, and group login information to the FortiGate unit. The FortiGate unit sets up a temporary access policy for the user, so when they attempt access through the firewall they do not need to re-authenticate. This model works well in environments where the FSSO client can be installed on all AD servers.

In system configurations where it is not possible to install FSSO clients on all AD servers, the FortiGate unit must be able to query the AD servers to find out if a user has been properly authenticated. This is achieved using the NTLM messaging features of Active Directory and Internet Explorer.

Even when NTLM authentication is used, the user is not asked again for their username and password. Internet Explorer stores the user's credentials and the FortiGate unit uses NTLM messaging to validate them in the Windows AD environment.

Note that if the authentication reaches the timeout period, the NTLM message exchange restarts. For more information on NTLM, see [“NTLM authentication” on page 1255](#) and [“NTLM authentication with FSSO” on page 1287](#), and

Certificates

Certificates can be used as part of an identity-based policy. All users being authenticated against the policy are required to have the proper certificate. See [“Certificate-based authentication” on page 1165](#)

Dynamic profile

Dynamic profile is a remote authentication method that does not require any local users to be configured, and relies on RADIUS Start records to provide the FortiGate unit with authentication information. That information identifies the user and user group, which is then matched using a security policy. See [“Dynamic profiles and end points” on page 1325](#)

FortiGuard Web Filter override authentication

Optionally, users can be allowed the privilege of overriding FortiGuard Web Filtering to view blocked web sites. Depending on the override settings, the override can apply to the user who requested it, the entire user group to which the user belongs, or all users who share the same web filter profile. As with other FortiGate features, access to FortiGuard overrides is controlled through user groups. Firewall and Directory Services user groups are eligible for the override privilege. For more information about web filtering and overrides, see the UTM chapter of this FortiOS Handbook.

VPN authentication

Authentication involves authenticating the user. In IPsec VPNs authenticating the user is optional, but authentication of the peer device is required.

This section includes:

- [Authenticating IPsec VPN peers \(devices\)](#)
- [Authenticating IPsec VPN users](#)
- [Authenticating SSL VPN users](#)
- [Authenticating PPTP and L2TP VPN users](#)

Authenticating IPsec VPN peers (devices)

A VPN tunnel has one end on a local trusted network, and the other end is at a remote location. The remote peer (device) must be authenticated to be able to trust the VPN tunnel. Without that authentication, it is possible for a malicious hacker to masquerade as a valid VPN tunnel device and gain access to the trusted local network.

The three ways to authenticate VPN peers are with a preshared key, RSA X.509 certificate, or a specific peer ID value.

The simplest way for IPsec VPN peers to authenticate each other is through the use of a preshared key, also called a shared secret. The preshared key is a text string used to encrypt the data exchanges that establish the VPN tunnel. The preshared key must be six or more characters. The VPN tunnel cannot be established if the two peers do not use the same key. The disadvantage of preshared key authentication is that it can be difficult to securely distribute and update the preshared keys. See [“Authenticating the FortiGate unit with a pre-shared key” on page 1410](#).

RSA X.509 certificates are a better way for VPN peers to authenticate each other. Each peer offers a certificate signed by a Certificate Authority (CA) which the other peer can validate with the appropriate CA root certificate. For more information about certificates, see [“Certificate-based authentication” on page 1263](#).

You can supplement either preshared key or certificate authentication by requiring the other peer to provide a specific peer ID value. The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

Authenticating IPsec VPN users

An IPsec VPN can be configured to accept connections from multiple dynamically addressed peers. You would do this to enable employees to connect to the corporate network while traveling or from home. On a FortiGate unit, you create this configuration by setting the *Remote Gateway* to *Dialup User*.

It is possible to have an IPsec VPN in which remote peer devices authenticate using a common preshared key or a certificate, but there is no attempt to identify the user at the remote peer. To add user authentication, you can do one of the following:

- require a unique preshared key for each peer
- require a unique peer ID for each peer
- require a unique peer certificate for each peer
- require additional user authentication (XAuth)

The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

Authenticating SSL VPN users

SSL VPN users can be

- user accounts with passwords stored on the FortiGate unit
- user accounts authenticated by an external RADIUS, LDAP or TACACS+ server
- PKI users authenticated by certificate

You need to create a user group for your SSL VPN. Simply create a firewall user group, enable SSL VPN access for the group, and select the web portal the users will access.

SSL VPN access requires an SSL VPN security policy that permits access to members of your user group.

Authenticating PPTP and L2TP VPN users

PPTP and L2TP are older VPN tunneling protocols that do not provide authentication themselves. FortiGate units restrict PPTP and L2TP access to users who belong to one specified user group. Users authenticate themselves to the FortiGate unit by username/password. You can configure PPTP and L2TP VPNs only in the CLI. Before you configure the VPN, create a firewall user group and add to it the users who are permitted to use the VPN. Users are authenticated when they attempt to connect to the VPN. For more information about configuring PPTP or L2TP VPNs, see the [FortiGate CLI Reference](#).

User's view of authentication

From the user's point of view, they see a request for authentication when they try to access a protected resource, such as an FTP repository of intellectual property or simply access a website on the Internet. The way the request is presented to the user depends on the method of access to that resource.

VPN authentication usually controls remote access to a private network.

Web-based user authentication

Security policies usually control browsing access to an external network that provides connection to the Internet. In this case, the FortiGate unit requests authentication through the web browser.

The user types a username and password and then selects *Continue* or *Login*. If the credentials are incorrect, the authentication screen is redisplayed with blank fields so that the user can try again. When the user enters valid credentials, access is granted to the required resource. In some cases, if a user tries to authenticate several times without success, a message appears, such as: "Too many bad login attempts. Please try again in a few minutes." This indicates the user is locked out for a period of time. This prevents automated brute force password hacking attempts. The administrator can customize these settings if required.



After a defined period of user inactivity (the authentication timeout, defined by the FortiGate administrator), the user's access expires. The default is 5 minutes. To access the resource, the user will have to authenticate again.

VPN client-based authentication

A VPN provides remote clients with access to a private network for a variety of services that include web browsing, email, and file sharing. A client program such as FortiClient negotiates the connection to the VPN and manages the user authentication challenge from the FortiGate unit.

FortiClient can store the username and password for a VPN as part of the configuration for the VPN connection and pass them to the FortiGate unit as needed. Or, FortiClient can request the username and password from the user when the FortiGate unit requests them.

SSL VPN is a form of VPN that can be used with a standard Web browser. There are two modes of SSL VPN operation (supported in NAT/Route mode only):

- web-only mode, for remote clients equipped with a web-browser only
- tunnel mode, for remote computers that run a variety of client and server applications.



After a defined period of user inactivity on the VPN connection (the idle timeout, defined by the FortiGate administrator), the user's access expires. The default is 30 minutes. To access the resource, the user will have to authenticate again.

FortiGate administrator's view of authentication

Authentication is based on user groups. The FortiGate administrator configures authentication for security policies and VPN tunnels by specifying the user groups whose members can use the resource. Some planning is required to determine how many different user groups need to be created. Individual user accounts can belong to multiple groups, making allocation of user privileges very flexible.

A member of a user group can be:

- a user whose username and password are stored on the FortiGate unit
- a user whose name is stored on the FortiGate unit and whose password is stored on a remote or external authentication server
- a remote or external authentication server with a database that contains the username and password of each person who is permitted access

The general process of setting up authentication is as follows:

- 1 If remote or external authentication is needed, configure the required servers.
- 2 Configure local and peer (PKI) user identities. For each local user, you can choose whether the FortiGate unit or a remote authentication server verifies the password. Peer members can be included in user groups for use in security policies.
- 3 Create user groups.
Add local/peer user members to each user group as appropriate. You can also add an authentication server to a user group. In this case, all users in the server's database can authenticate. You can only configure peer user groups through the CLI.
- 4 Configure security policies and VPN tunnels that require authenticated access.

For authentication troubleshooting, see the specific chapter for the topic or for general issues see [“Troubleshooting” on page 1377](#).



Authentication and User in the web-based manager

This section provides an introduction to the web-based manager User menu.

The following topics are included in this section:

- [User](#)
- [User groups](#)
- [Remote](#)
- [FortiToken](#)
- [Fortinet Single Sign On Agent \(FSSO\)](#)
- [PKI](#)
- [Monitor](#)



The word “unit” refers to the FortiGate unit. The words “FortiGate unit” are used when talking about different Fortinet products in one sentence. For example, “The Central Management menu provides the option of remotely managing your FortiGate unit by a FortiManager unit.”

User

The User menu allows you to configure authentication settings and user accounts. The User menu also allows you to configure user groups, remote servers, as well as monitor users.

A user is a user account that consists of a user name, password and in some cases, other information that can be configured on the unit or on an external authentication server. Users can access resources that require authentication only if they are members of an allowed user group. For more information about user groups, see [“User groups” on page 1176](#).

This topic contains the following:

- [Local user accounts](#)
- [IM users](#)
- [Authentication settings](#)

Local user accounts

A local user is a user configured on a unit. The user can be authenticated with a password stored on the unit (the user name and password must match a user account stored on the unit) or with a password stored on an authentication server. The user name must match a user account stored on the unit and the user name and password must match a user account stored on the authentication server associated with the user.

When configuring a user account, you can enable two-factor authentication. This feature allows you to add what is referred to as a FortiToken, which is a serial number used in log in credentials, and this FortiToken can be emailed to the user or sent to their mobile phone. Two-factor authentication is a security measure that provides an additional log in credential, in this case a FortiToken, that the user must enter so that they can log in. You must have the FortiToken device to use the FortiToken feature on your unit.

Local user accounts configuration settings

Local users are configured in *User > User > User*. Use the following table when configuring local user accounts.

User page Lists each individual local user's list that you created. On this page, you can edit, delete or create a new local users list. Note: If you want to have users always authenticate whenever their time expires, use the <code>hard-timeout</code> value in the <code>auth-type</code> command. This is available only in the CLI.	
Create New	Creates a new local user account. When you select <i>Create New</i> , you are automatically redirected to New User page.
Delete	Removes a user from the list. Removing the user name removes the authentication configured for the user. The <i>Delete</i> icon is not available if the user belongs to a user group. To remove multiple local user accounts from within the list, on the User page, in each of the rows of user accounts you want removed, select the check box and then select <i>Delete</i> . To remove all local user accounts from the list, on the User page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies a user's account settings. When you select <i>Edit</i> , you are automatically redirected to the Edit User page.
User Name	The local user name.
Type	The authentication type to use for this user. The authentication types are Local (user and password stored on FortiGate unit), LDAP, RADIUS, and TACACS+ (user and password matches a user account stored on the authentication server).
Two-factor Authentication	Indicates the status of whether two-factor authentication is configured for the user. A gray x displays if there is no two-factor authentication enabled for the user. A green check mark displays if two-factor authentication is enabled for the user.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
New User page Provides settings for configuring whether to allow or block a local user from authenticating.	
User Name	A name that identifies the user.
Disable	Select to prevent this user from authenticating.
Password	Select to authenticate this user using a password stored on the unit and then enter the password. Best practices dictate that the password be at least six characters long.
Match users on LDAP servers	Select to authenticate this user using a password stored on an LDAP server. Select the LDAP server from the list. You can select only an LDAP server that has been added to the FortiGate LDAP configuration. For more information, see “LDAP” on page 1184 .
Match users on RADIUS server	Select to authenticate this user using a password stored on a RADIUS server. Select the RADIUS server from the list. You can select only a RADIUS server that has been added to the FortiGate RADIUS configuration. For more information, see “RADIUS” on page 1182 .
Match users on TACACS+ server	Select to authenticate this user using a password stored on a TACACS server. Select the TACACS+ server from the list. You can select only a TACACS server that has been added to the FortiGate TACACS configuration. For more information, see “TACACS+” on page 1187 .

Enable Two-factor Authentication	Select to enable the two-factor authentication feature. When you enable this feature, <i>FortiToken</i> , <i>Email to</i> and <i>SMS</i> appear below. You can enable two-factor authentication for FortiGate administrators as well as users. For more information about enabling two-factor authentication for FortiGate administrators, see “Associating FortiTokens with accounts” on page 1233 .
Deliver Token Code by	When you enable two-factor authentication, you can send the FortiToken number to an email address or to a mobile phone.
FortiToken	Select the serial number of the FortiToken that will be used by the user for generating the token password code. See “FortiToken configuration settings” on page 1189 .
Email to	Enter the email address of the recipient that will be receiving the token password code.
SMS	Enter the phone number of the mobile phone that will receive the token password code. You must first enter the SMS in the CLI before you can select one from the drop-down list.

IM users

Instant Messenger (IM) protocols are gaining in popularity as an essential way to communicate between two or more individuals in real time. Some companies even rely on IM protocols for critical business applications such as Customer/Technical Support.

The most common IM protocols in use today include AOL Instant Messenger, Yahoo Instant Messenger, MSN messenger, and ICQ. FortiGate units allow you to set up IM users that either allow or block the use of applications, to determine which applications are allowed.

IM users configuration settings

IM users are configured in the CLI first, and then appear in *User > User > IM*. IM users must be configured using the `config imp2p` command in the CLI.

The following are IM user configuration settings in *User > User > IM*.

IM page Lists each individual IM user. On this page, you can edit, delete or create a new IM user. You can also filter the information by protocol or policy.	
Create New	Creates an IM user. When you select <i>Create New</i> , you are automatically redirected to the Edit User page.
Edit	Modifies an IM user's settings. When you select <i>Edit</i> , you are automatically redirected to the Edit User page.
Delete	Removes an IM user from the list on the page. To remove multiple IM users from within the list, on the IM page, in each of the rows of users you want removed, select the check box and then select <i>Delete</i> . To remove all IM users from the list, on the IM page, select the check box in the check box column and then select <i>Delete</i> .

Protocol:	<p>Filter the information on the page using a specific protocol. For example, only ICQ users appear when the protocol selected is ICQ.</p> <p>You can use both Protocol and Policy to filter the information on the page. For example, only ICQ users that are blocked appear when protocol is set to ICQ and policy is set to Block</p>
Policy:	<p>Filter the information on the page using a specific policy. For example, only policies that are blocked appear when the policy selected is Block.</p> <p>You can use both Protocol and Policy to filter the information on the page. For example, only ICQ users that are blocked appear when protocol is set to ICQ and policy is set to Block</p>
Protocol	The type of instant messaging protocol. For example, ICQ is for all ICQ instant messaging software.
Username	The name of the user.
Policy	The type of action that will be taken when the protocol is detected.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
Edit User page	
Provides settings for configuring an IM user.	
Protocol	Select an instant messaging protocol from the drop-down list.
Username	Enter the user name for accessing the IM user.
Policy	Select the type of action that will be taken when the protocol is detected, for example, the protocol AIM will be blocked.

Authentication settings

The Authentication submenu provides settings for configuring authentication timeout, protocol support, and authentication certificates.

When user authentication is enabled within a security policy, the authentication challenge is normally issued for any of the four protocols (depending on the connection protocol):

- HTTP (can also be set to redirect to HTTPS)
- HTTPS
- FTP
- Telnet.

The selections made in the *Protocol Support* list of the Authentication Settings screen control which protocols support the authentication challenge. Users must connect with a supported protocol first so they can subsequently connect with other protocols. If HTTPS is selected as a method of protocol support, it allows the user to authenticate with a customized Local certificate.

When you enable user authentication within a security policy, the security policy user will be challenged to authenticate. For user ID and password authentication, users must provide their user names and passwords. For certificate authentication (HTTPS or HTTP redirected to HTTPS only), you can install customized certificates on the unit and the users can also have customized certificates installed on their browsers. Otherwise, users will see a warning message and have to accept a default FortiGate certificate.

Authentication settings are configured in *User > User > Authentication*. Use the following table when configuring authentication settings.

Authentication Settings	
Provides settings for defining how users authenticate a session. For example, the authentication timeout is configured for 10 minutes so a user's session, if idle for 10 minutes, will log them out of the session automatically.	
Authentication Timeout	Enter a length of time in minutes, from 1 to 480. Authentication timeout controls how long an authenticated firewall connection can be idle before the user must authenticate again. The default value is 30
Protocol Support	Select the protocols to challenge during firewall user authentication.
Certificate	If using HTTPS protocol support, select the local certificate to use for authentication. Available only if HTTPS protocol support is selected.
Apply	Select to apply the selections for user authentication settings.



When you use certificate authentication, if you do not specify any certificate when you create the security policy, the global settings will be used. If you specify a certificate, the per-policy setting will overwrite the global setting.

User groups

A user group is a list of user identities. An identity can be:

- a local user account (user name and password) stored on the FortiGate unit
- a local user account with a password stored on a RADIUS, LDAP, or TACACS+ server

- a RADIUS, LDAP, or TACACS+ server (all identities on the server can authenticate)
- a user or user group defined on a Directory Service server.

Each user group belongs to one of three types: Firewall, Directory Service or SSL VPN.

In most cases, the unit authenticates users by requesting each user name and password. The unit checks local user accounts first. If the unit does not find a match, it checks the RADIUS, LDAP, or TACACS+ servers that belong to the user group. Authentication succeeds when the unit finds a matching user name and password.

This topic contains the following:

- [User Group](#)
- [Firewall user groups](#)
- [Fortinet Single Sign-On \(FSSO\) user groups](#)
- [SSL VPN user groups](#)

User Group

For each resource that requires authentication, you specify which user groups are permitted access. You need to determine the number and membership of user groups appropriate to your authentication needs.

Users that are associated with multiple groups have access to all services within all the user groups that they are associated with. This is only available in the CLI. The command used is `auth-multi-group`, which is enabled by default. This feature checks all groups a user belongs to for firewall authentication.

User Group configuration settings

The following are user group configuration settings in *User > User Group > User Group*.

User Group page	
Lists each individual user group list according to their type of group. On this page, you can edit, delete or create a new user group list.	
Create New	Creates a new user group. When you select <i>Create New</i> , you are automatically redirected to New User Group page.
Delete	<p>Removes a user group from the list on the User Group page. You cannot delete a user group that is included in a security policy, a dialup user phase 1 configuration, or a PPTP or L2TP configuration.</p> <p>To remove multiple user groups, expand the section (either <i>Firewall</i> or <i>Directory Service</i>, and in each of the rows of user groups you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all user groups, on the User Group page, select the check box in the check box column and then select <i>Delete</i>.</p>
Edit	Modifies membership and options of a group. When you select <i>Edit</i> , you are automatically redirected to the Edit User Group page.

Group Name	The name of the user group. User group names are listed by type of user group: Firewall, Directory Service and SSL VPN. For more information, see “Fortinet Single Sign-On (FSSO) user groups” on page 1180 , and “SSL VPN user groups” on page 1180 .
Members	The Local users, RADIUS servers, LDAP servers, TACACS+ servers, Directory Service users/user groups or PKI users found in the user group.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy’s settings appear within the table.
New User Group page Provides settings for configuring a list of users and/or groups. Note: The <i>Remote authentication servers</i> section does not appear unless there is a RADIUS, LDAP, or TACAS+ server configured.	
Name	Enter the name of the user group.
Type	Select the user group type. See “Firewall user groups” on page 1179 , “SSL VPN user groups” on page 1180 and “Fortinet Single Sign-On (FSSO) user groups” on page 1180 .
Firewall	Select this group in any security policy that requires Firewall authentication.
Fortinet Single Sign-On (FSSO)	Select this group in any security policy that requires FSSO authentication.

	SSL VPN	Select this group in any security policy with <i>Action</i> set to <i>SSL VPN</i> . See “SSL VPN access” on page 1236 . This option is not available in Transparent mode.
	Allow SSL-VPN Access	Select the type of SSL VPN access from the drop-down list.
	Available Users	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, Directory Service users/user groups, or PKI users that can be added to the user group. To add a member to <i>Available Users</i> list, select the name and then select the -> arrow. To remove a member from the <i>Members</i> list, select the member and then select the <- arrow. * <i>Available Members</i> displays only if the user group type is <i>Directory Service</i> .
	Members	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, Directory Service users/user groups, or PKI users that belong to the user group. To remove a member, select the name and then select the <- arrow.
Remote authentication servers The following appear only if a RADIUS, LDAP or TACACS+ server has been configured and if <i>Type</i> is <i>Firewall</i> .		
	Add	Select to add a new remote authentication server. When you select add the following appear: <ul style="list-style-type: none"> • remote server drop-down list • <i>Any</i> and <i>Specify and its field</i> • Delete icon
	Remote Server	Select a remote server from the drop-down list.
	Group Name	Select to specify the group name or to have any group name. If you select <i>Specify</i> , you must enter the name in the field provided. If you want to specify more than one, use a comma to separate each name.
	Delete	Select to remove the remote server from within the list.

Firewall user groups

A firewall user group provides access to a security policy that requires authentication and lists the user group as one of the allowed groups. The unit requests the group member's user name and password when the user attempts to access the resource that the policy protects. The reason for the name firewall user group is when creating a new user group you have the option to select Firewall or FSSO as the group type.

You can also authenticate a user by certificate if you have selected this method.

A firewall user group can also provide access to an IPsec VPN for dialup users. In this case, the IPsec VPN phase 1 configuration uses the Accept peer ID in dialup group peer option. The user's VPN client is configured with the user name as peer ID and the password as pre-shared key. The user can connect successfully to the IPsec VPN only if the user name is a member of the allowed user group and the password matches the one stored on the unit.



A user group cannot be a dialup group if any member is authenticated using a RADIUS or LDAP server.

You can also use a firewall user group to provide override privileges for FortiGuard web filtering.

Fortinet Single Sign-On (FSSO) user groups

On a network, you can configure the unit to allow access to members of FSSO user groups who have been authenticated on the network. The Fortinet Single Sign On Agent must be installed on the network domain controllers.

An FSSO user group provides access to a security policy that requires FSSO type authentication and lists the user group as one of the allowed groups. The members of the user group are FSSO users or groups that you select from a list that the unit receives from the FSSO servers that you have configured. For more information, see [“Fortinet Single Sign On Agent \(FSSO\)” on page 1191](#).



An FSSO user group cannot have SSL VPN access.

You cannot use FSSO user groups directly in FortiGate security policies. You must add FSSO groups to FortiGate user groups. An FSSO group belongs to only one FortiGate user group according to best practices. If you assign it to multiple FortiGate user groups, the unit recognizes only the last user group assignment.

You can also use an FSSO user group to provide override privileges for FortiGuard web filtering.

SSL VPN user groups

An SSL VPN user group provides access to a security policy that requires SSL VPN type authentication and lists the user group as one of the allowed groups. Local user accounts, LDAP, and RADIUS servers can be members of an SSL VPN user group. The unit requests the user's user name and password when the user accesses the SSL VPN web portal. The user group settings include options for SSL VPN features.

An SSL VPN user group can also provide access to an IPsec VPN for dialup users. In this case, the IPsec VPN phase 1 configuration uses the Accept peer ID in dialup group peer option. You configure the user's VPN client with the user name as peer ID and the password as pre-shared key. The user can connect successfully to the IPsec VPN only if the user name is a member of the allowed user group and the password matches the one stored on the unit.



A user group cannot be an IPsec dialup group if any member is authenticated using a RADIUS or LDAP server.

For information on configuring user groups, see [“User groups” on page 1176](#). For SSL VPN information, see [“SSL VPN access” on page 1236](#).



By default, the web-based manager displays Firewall options. You cannot add local users to a group that is used to authenticate administrators.

Dynamically assigning VPN client IP addresses from a user group

SSL VPN tunnel mode, dialup IPsec VPN, and PPTP VPN sessions can assign IP addresses to remote users by getting the IP address to assign to the user from the Framed-IP-Address field in the RADIUS record received when the RADIUS server confirms that the user has authenticated successfully. For more information about RADIUS fields, see RFC 2865 and RFC 2866.

For the unit to dynamically assign an IP address, the VPN users must be configured for RADIUS authentication and you must include the IP address to assign to the user in the Framed-IP-Address RADIUS field on your RADIUS server. You configure each type of VPN differently. In each case you are associating the configuration that assigns IP addresses to users with a user group.

Assigning IP addresses from a RADIUS record replaces dynamically assigning IP addresses from an address range. You cannot include an IP address range and assigning IP addresses from a RADIUS record in the same configuration.

Remote

Remote authentication is generally used to ensure that employees working offsite can remotely access their corporate network with appropriate security measures in place. In general terms, authentication is the process of attempting to verify the (digital) identity of the sender of a communication such as a login request. The sender may be someone using a computer, the computer itself, or a computer program. Since best practices dictate a computer system be used only by those who are authorized to do so, there must be a measure in place to detect and exclude any unauthorized access.

On a unit, you can control access to network resources by defining lists of authorized users, called user groups. To use a particular resource, such as a network or VPN tunnel, the user must:

- belong to one of the user groups that is allowed access
- correctly enter a user name and password to prove his or her identity, if asked to do so.

This topic contains the following:

- [Administrators](#)
- [RADIUS](#)
- [LDAP](#)
- [TACACS+](#)

Administrators

A local administrator account, admin, is configured by default on all FortiOS units. If you have remote servers configured, you can create remote administrator accounts as well.

These are accounts with local administrator privileges that use remote authentication. One version of this are wildcard administrator accounts. See [“Example — wildcard admin accounts - CLI” on page 1211](#).

Configuring a remote administrator account is the same as a local administrator account with the following differences:

- You must have a remote authentication server and user group with an associated authentication server configured before creating the administrator account.
- Even though the administrator account will be authenticated remotely, when creating the account using the web-based manager you must enter a password or the account will not be created. You are not prompted for the password if the account is created in the CLI. This password acts as a backup password in case the remote server does not respond during authentication. The exception to this is when wildcard admin is selected — for wildcard administrators, the password field is skipped over.

RADIUS

Remote Authentication and Dial-in User Service (RADIUS) servers provide authentication, authorization, and accounting functions. FortiGate units use the authentication function of the RADIUS server. You must configure the server before you configure the FortiGate users or user groups that will need it to use the RADIUS server for authentication.

If you have configured RADIUS support and a user is required to authenticate using a RADIUS server, the unit sends the user's credentials to the RADIUS server for authentication. If the RADIUS server can authenticate the user, the user is successfully authenticated with the unit. If the RADIUS server cannot authenticate the user, the unit refuses the connection. You can override the default authentication scheme by selecting a specific authentication protocol or changing the default port for RADIUS traffic.



The default port for RADIUS traffic is 1812. If your RADIUS server is using port 1645, use the CLI command, `config system global`, to change the default RADIUS port.

If you want to configure settings for UTF-8 encoding, you must enable this in the CLI using the `config vpn ssl setting` command.

RADIUS configuration settings

The following are RADIUS server configuration settings in *User > Remote > RADIUS*.

RADIUS page Lists each individual RADIUS server that you created. On this page, you can edit, delete or create a new RADIUS server.	
Create New	Creates a new RADIUS server. The maximum number is 10. When you select <i>Create New</i> , you are automatically redirected to the New RADIUS Server page.
Delete	Removes a RADIUS server from the list on the RADIUS page. Note: You cannot delete a RADIUS server that has been added to a user group. To remove multiple RADIUS servers from within the list, on the RADIUS page, in each of the rows of servers you want removed, select the check box and then select <i>Delete</i> . To remove all RADIUS servers from the list, on the RADIUS page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies settings within a RADIUS server configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit RADIUS Server page.
Name	Name that identifies the RADIUS server on the unit.
Type	The type of server, either query or dynamic start.
Server Name/IP	Domain name or IP address of the RADIUS server.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> . To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object. To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window: <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
New RADIUS Server page Provides settings for configuring a RADIUS server.	
Name	Enter the name that is used to identify the RADIUS server on the unit.

Type	Select either <i>Query</i> or <i>Dynamic Start</i> .
Primary Server Name/IP	Enter the domain name or IP address of the primary RADIUS server.
Primary Server Secret	Enter the RADIUS server secret key for the primary RADIUS server. The primary server secret key length can be up to a maximum of 16 characters. Note: Best practices dictate that the server secret key should be the maximum length for security reasons.
Secondary Server Name/IP	Enter the domain name or IP address of the secondary RADIUS server, if you have one.
Secondary Server Secret	Enter the RADIUS server secret key for the secondary RADIUS server. The secondary server secret key can be up to a maximum length of 16 characters.
Authentication Scheme	Select <i>Use Default Authentication Scheme</i> to authenticate with the default method. The default authentication scheme uses PAP, MS-CHAP-V2, and CHAP, in that order. Select <i>Specify Authentication Protocol</i> to override the default authentication method, and choose the protocol from the list: MS-CHAP-V2, MS-CHAP, CHAP, or PAP, depending on what your RADIUS server needs.
NAS IP/Called Station ID	Optionally enter the NAS IP address (RADIUS Attribute 31 outlined in RFC 2548). In this configuration, the FortiGate unit is the Network Access Server (NAS) and this is how the RADIUS server registers all valid servers that use its records. If you do not enter an IP address, the IP address that the FortiGate interface uses to communicate with the RADIUS server will be applied.
Include in every User Group	Select the check box beside <i>Enable</i> to have the RADIUS server automatically included in all user groups.

LDAP

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, and printers. An LDAP consists of a data-representation scheme, a set of defined operations, and a request/response network.

If you have configured LDAP support and require a user to authenticate using an LDAP server, the unit contacts the LDAP server for authentication. To authenticate with the unit, the user enters a user name and password. The unit sends this user name and password to the LDAP server. If the LDAP server can authenticate the user, the unit successfully authenticates the user. If the LDAP server cannot authenticate the user, the unit refuses the connection.

FortiGate LDAP supports password renewal, and these settings are configured in the CLI. There are settings for a warning that the password is going to expire, and threshold of the expiry as well.

LDAP configuration settings

The following are LDAP server configuration settings in *User > Remote > LDAP*. You can use both IPv6 and IPv 4 addresses when configuring LDAP servers.

LDAP page Lists each individual LDAP server that you created. On this page, you can edit, delete or create a new LDAP server.	
Create New	Creates a new LDAP server. The maximum number is 10. When you select <i>Create New</i> , you are automatically redirected to the New LDAP Server page.
Delete	Removes the LDAP server configuration from the list on the LDAP page. To remove multiple LDAP servers from within the list, on the LDAP page, in each of the rows of servers you want removed, select the check box and then select <i>Delete</i> . To remove all LDAP servers from the list, on the LDAP page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies settings from within an LDAP server configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit LDAP Server page.
Name	The name that identifies the LDAP server on the unit.
Server Name/IP	The domain name or IP address of the LDAP server.
Port	The TCP port used to communicate with the LDAP server.
Common Name Identifier	The common name identifier for the LDAP server. Most LDAP servers use cn. However, some servers use other common name identifiers such as uid.
Distinguished Name	The distinguished name used to look up entries on the LDAP servers use. The distinguished name reflects the hierarchy of LDAP database object classes above the common name identifier.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
New LDAP Server page Provides settings for configuring an LDAP server.	
Name	Enter the name that identifies the LDAP server on the FortiGate unit.
Server Name/IP	Enter the domain name or IP address of the LDAP server.
Server Port	Enter the TCP port used to communicate with the LDAP server. By default, LDAP uses port 389. If you use a secure LDAP server, the default port changes when you select <i>Secure Connection</i> .
Common Name Identifier	Enter the common name identifier for the LDAP server. The maximum number of characters is 20.
Distinguished Name	Enter the base distinguished name for the server using the correct X.500 or LDAP format. The unit passes this distinguished name unchanged to the server. The maximum number of characters is 512.
Query icon	View the LDAP server Distinguished Name Query tree for the LDAP server that you are configuring so that you can cross-reference to the Distinguished Name. The <i>Query</i> icon is located beside the <i>Distinguished Name</i> field. For more information, see “Using Query” on page 1187 .
Bind Type	Select the type of binding for LDAP authentication.

	Regular	Connect to the LDAP server directly with user name/password, then receive accept or reject based on search of given values. When you select <i>Regular</i> , the <i>Use DN</i> and <i>Password</i> fields appear.
	Anonymous	Connect as an anonymous user on the LDAP server, then retrieve the user name/password and compare them to given values.
	Simple	Connect directly to the LDAP server with user name/password authentication.
	User DN	Enter the Distinguished name of the user to be authenticated. Available if <i>Bind Type</i> is <i>Regular</i> .
	Password	Enter the password of the user to be authenticated. This is available if <i>Bind Type</i> is <i>Regular</i> .
	Secure Connection	Select to use a secure LDAP server connection for authentication.
	Protocol	Select a secure LDAP protocol to use for authentication. Depending on your selection, the value in <i>Server Port</i> will change to the default port for the selected protocol. This is available only if <i>Secure Connection</i> is selected. <ul style="list-style-type: none"> • <i>LDAPS</i>: port 636 • <i>STARTTLS</i>: port 389
	Certificate	Select a certificate to use for authentication from the list. The certificate list comes from CA certificates at <i>System > Certificates > CA Certificates</i> . For more information about certificates, see “Certificate-based authentication” on page 1263 .

Using Query

The LDAP Distinguished Name Query list displays the LDAP Server IP address, and all the distinguished names associated with the Common Name Identifier for the LDAP server. The tree helps you to determine the appropriate entry for the DN field. You can see the distinguished name associated with the Common Name identifier, by expanding *CN identifier* and then selecting the DN from the list. The DN you select is displayed in the *Distinguished Name* field. You must select *OK* to save your selection in the *Distinguished Name* field of the LDAP Server configuration.

If you want to see the users with the LDAP Server user group for the selected Distinguished Name, expand *Distinguished Name* in the LDAP Distinguished Name Query tree.

TACACS+

Terminal Access Controller Access-Control System (TACACS+) is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices via one or more centralized servers. TACACS+ allows a client to accept a user name and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies network access to the user. The default TCP port for a TACACS+ server is 49.

TACACS+ configuration settings

The following are TACACS+ server configuration settings in *User > Remote > TACACS+*.

TACACS+ page Lists each individual TACACS+ server that you created. On this page, you can edit, delete or create a new TACACS+ server.	
Create New	Creates a new TACACS+ server. The maximum number of TACACS+ servers that you can create is 10. When you select <i>Create New</i> , you are automatically redirected to the New TACACS+ Server page.
Delete	Removes a TACACS+ server from the list on the TACACS+ page. To remove multiple TACACS+ servers from within the list, on the TACACS+ page, in each of the rows of servers you want removed, select the check box and then select <i>Delete</i> . To remove all TACACS+ servers from the list, on the TACACS+ page, select the check box in the check box column and then select <i>Delete</i> .
Edit	Modifies settings within a TACACS+ server configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit TACACS+ Server page.
Name	The name of the TACACS+ server.
Server	The server domain name or IP address of the TACACS+ server.
Authentication Type	The supported authentication method. TACACS+ authentication methods include: Auto, ASCII, PAP, CHAP, and MSCHAP.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> . To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object. To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window: <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
New TACACS+ Server page Provides settings for configuring a TACACS+ server.	

Name	Enter the name of the TACACS+ server.
Server Name/IP	Enter the server domain name or IP address of the TACACS+ server.
Server Key	Enter the key to access the TACACS+ server. The server key can be a maximum of 16 characters in length.
Authentication Type	Select the authentication type to use for the TACACS+ server. Selection includes: Auto, ASCII, PAP, CHAP, and MSCHAP. Auto authenticates using PAP, MSCHAP, and CHAP (in that order).

FortiToken

In *User > FortiToken > FortiToken*, you can add token password codes that are then used by users to authenticate when logging in, for example the FortiGate unit. These token password codes are generated by the FortiToken device, which is a small device that generates a unique token password code. This can then be activated by the unit so that it can then be used for authentication purposes. This code provides additional security for you and users. This is part of the two-factor authentication feature, providing an extra login credential to prove that the user is who they claim to be.

For more information about FortiToken and two-factor authentication, see [“FortiToken” on page 1230](#).

FortiToken configuration settings

You can apply two-factor authentication to administrators that are logging into the FortiGate unit.

The following are FortiToken configuration settings in *User > FortiToken > FortiToken*.

FortiToken page	
Lists each individual FortiToken list that is currently being used and which user the token is being used by. A FortiToken list provides up to ten FortiToken serial numbers that one user can use to authenticate when entering log in credentials.	
Note: You can import FortiToken serial numbers from a file.	
Create New	Creates a new FortiToken list. When you select <i>Create New</i> , you are automatically redirected to the Add new FortiToken page.
Edit	Modifies the settings within the FortiToken list. When you select <i>Edit</i> , you are automatically redirected to the Edit new FortiToken page.
Delete	Removes a FortiToken list from within the list on the FortiToken page. To remove multiple token serial numbers from within the list, on the FortiToken page, in each of the rows of codes you want removed, select the check box and then select <i>Delete</i> . To remove all token codes from the list, on the FortiToken page, select the check box in the check box column and then select <i>Delete</i> .

Import	Select to import FortiToken serial numbers from a file. When you select Import, the Upload window appears. From the Upload window you can select Browse to locate the file and then select OK to import the file to the unit.
Activate	Activates the FortiToken device so that you can use the device to code so that the code can then be used by a user or FortiGate administrator.
Synchronization	Select to synchronize the FortiToken with an NTP server. This allows the unit's system time to be the same as the FortiToken's system time.
Serial Number	The serial number of the FortiToken device. This serial number is found on the back of the device.
Status	The status of the FortiToken, whether it is activated or not.
Drift	<p>The synchronization of the unit's system time and the FortiToken's time, which can "drift" away or become unsynchronized.</p> <p>When you synchronize the FortiToken, it calculates the unit's and its own drift. The drift is the number of minutes that the FortiToken has drifted away from the unit's system time. The minutes is shown in integers only.</p> <p>If drift is zero, then the FortiToken's internal time is the same as the unit's system time. If the drift is -1, the FortiToken is lagging behind the unit by one minute. If the drift is 1, then the FortiToken is ahead of the unit by one minute.</p>
User	The user that the FortiToken is currently being used by.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
Add new FortiToken page Provides settings for configuring multiple FortiToken serial numbers.	
Serial Number #<number>	Enter the FortiToken's serial number that you see on the back of the FortiToken. You can add up to ten FortiTokens.
Automatically Send Activate Request to FortiGuard	By default, this is enabled. When this is selected, the serial numbers that were inputted in the fields are automatically sent to FortiGuard after selecting <i>OK</i> .

Fortinet Single Sign On Agent (FSSO)

Windows Active Directory (AD) and Novell eDirectory provide central authentication services by storing information about network resources across a domain (a logical group of computers running versions of an operating system) in a central directory database. Each person who uses computers within a domain receives his or her own unique account/user name. This account can be assigned access to resources within the domain based on the role of the user. In a domain, the directory resides on computers that are configured as domain controllers. A domain controller is a server that manages all security-related features that affect the user/domain interactions, security centralization, and administrative functions.

FortiGate units use security policies to control access to resources based on user groups configured in the policies. Each FortiGate user group is associated with one or more Directory Service user groups. When a user logs in to the Windows or Novell domain, a Fortinet Single Sign On Agent (FSSO) sends the user's IP address and the names of the Directory Service user groups to which the user belongs to the FortiGate unit.

The Fortinet Single Sign On Agent has two components that you must install on your network:

- The domain controller (DC) agent must be installed on every domain controller to monitor user logins and send information about them to the collector agent.
- The Collector agent must be installed on at least one domain controller to send the information received from the DC agents to the FortiGate unit. Alternately a FortiAuthenticator server can take the place of the Collector agent in an FSSO polling mode configuration.

The unit uses this information to maintain a copy of the domain controller user group database. Because the domain controller authenticates users, the unit does not perform authentication. It recognizes group members by their IP address.

You must install the Fortinet Single Sign On Agent on the network and configure the unit to retrieve information from the Directory Service server. For more information about Fortinet Single Sign On Agent, see the [“Introduction to FSSO” on page 1283](#).

Fortinet Single Sign-on Agent configuration settings

The following are configuration settings for collector agents in *User > FSSO > FSSO Agent*.

FSSO Agent page Lists all the collector agents' lists that you have configured. On this page, you can create, edit or delete FSSO agents. Note: You can create a redundant configuration on your unit if you install a collector agent on two or more domain controllers. If the current (or first) collector agent fails, the FortiGate unit switches to the next one in its list of up to five collector agents.	
Create New	Creates a new agent. When you select <i>Create New</i> , you are automatically redirected to the New page.
Edit	Modifies settings within the Collector agent list. When you select <i>Edit</i> , you are automatically redirected to the Edit page. To remove multiple agents from within the list, on the FSSO Agent page, in each of the rows of servers you want removed, select the check box and then select <i>Delete</i> . To remove all agents from the list, on the FSSO Agent page, select the check box in the check box column and then select <i>Delete</i> .
Delete	Removes an agent from the list on the page.
Add User/Group	Adds a user or group to the list on the FSSO Agent page. You must know the distinguished name for the user or group. When you select <i>Add User/Group</i> , you are automatically redirected to the Add User/Group page.
Edit User/Group	Modifies users or groups from the remote server. When you select <i>Edit Users/Groups</i> , you are automatically redirected to the Edit Users/Groups page.
Refresh	Refreshes the current information on the page.
Name	The name of the FSSO agent list.
FSSO Agent IP/Name	The IP address or name of the agent or agents.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
<p>New page</p> <p>Provides settings for configuring FSSO agents. You can configured up to five agents, either Collector or Novell eDirectory.</p>	
Name	Enter a name for the FSSO agent.
FSSO Agent IP/Name	Enter the IP address or name of the Directory Service server where this collector agent is installed. The maximum number of characters is 63.
Port	Enter the TCP port used for Directory Service. This must be the same as the FortiGate listening port specified in the Fortinet Single Sign On Agent collector agent configuration.
Password	Enter the password for the collector agent. This is required only if you configured your Fortinet Single Sign On Agent collector agent to require authenticated access.
LDAP Server	Select the check box and select an LDAP server to access the Directory Service.

PKI

Public Key Infrastructure (PKI) authentication utilizes a certificate authentication library that takes a list of peers, peer groups, and/or user groups and returns authentication successful or denied notifications. Users only need a valid certificate for successful authentication—no user name or password are necessary. Firewall and SSL VPN are the only user groups that can use PKI authentication.



If you use the CLI to create a peer user, best practices dictate that you enter a value for either `subject` or `ca`. If you do not do so, and then open the user record in the web-based manager, you will be prompted to enter a subject or ca value before you can continue.

PKI configuration settings

If your unit is currently running FortiOS 4.0 MR2 or higher, you must configure a PKI user first in the CLI to enable the PKI menu in the web-based manager, then you can then configure other PKI users in *User > PKI > PKI*.

The following are PKI configuration settings in *User > PKI > PKI*.

PKI page	
Lists each individual PKI user that you have created. On this page, you can edit, delete or create a new PKI user.	
Create New	Creates a new PKI user. When you select <i>Create New</i> , you are automatically redirected to the New PKI User page.
Delete	<p>Removes a PKI user from the list on the PKI page.</p> <p>The delete icon is not available if the peer user belongs to a user group. Remove it from the user group first.</p> <p>To remove multiple PKI users from within the list, on the PKI page, in each of the rows of users you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all PKI users from the list, on the PKI page, select the check box in the check box column and then select <i>Delete</i>.</p>
Edit	Modifies settings within the PKI user configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit PKI User page.
Name	The name of the PKI user.
Subject	The text string that appears in the subject field of the certificate of the authenticating user.
CA	<p>The CA certificate that is used to authenticate this user.</p> <p>Note: PKI certificate authentication supports the extraction of the user name from within the UPN field, which allows users to log in without having to enter their user name. This is available only in the CLI.</p>

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> .</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
New PKI User page	
Provides settings for configuring a new PKI user.	
Name	Enter the name of the PKI user.
Subject	Enter the text string that appears in the subject field of the certificate of the authenticating user. This field is optional.
CA	Enter the CA certificate that must be used to authenticate this user. This field is optional.
Two-factor authentication section	
Require two-factor authentication	Require this PKI user to authenticate by password in addition to certificate authentication.
Password	Enter the password that this PKI user must enter.

Peer users and peer groups

You can define peer users and peer groups used for authentication in some VPN configurations and for PKI certificate authentication in security policies. Defining peer users and peer groups is configured in the CLI.

A peer user is a digital certificate holder that can use PKI authentication. Before using PKI authentication, you must define peer users to include in the user group that is incorporated into the firewall authentication policy.

When defining a peer user, you need:

- a peer user name
- the text from the subject field of the certificate of the authenticating peer user, or the CA certificate used to authenticate the peer user.

You can add or modify other configuration settings for PKI authentication.

Peer users are created in *User > PKI > PKI* for authentication; however, you must enter a value for at least one of the fields, *Subject* or *CA*. You can configure peer user groups only through the CLI.



PKI certificate authentication supports the extraction of the user name from within the UPN field, which allows users to log in without having to enter their user name. This is available only in the CLI.

Monitor

You can go to the Monitor menu to view lists of currently authenticated users, authenticated IM users, and banned users. For each authenticated user, the list includes the user name, user group, how long the user has been authenticated (*Duration*), how long until the user's session times out (*Time left*), and the method of authentication used. The list of IM users includes the source IP address, protocol, and last time the protocol was used. The Banned User list includes users configured by administrators in addition to those quarantined based on AV, IPS, or DLP rules.

This topic contains the following:

- [Firewall monitor list](#)
- [IM user monitor list](#)
- [The Banned User list](#)

Firewall monitor list

In some environments, it is useful to determine which users are authenticated by the FortiGate unit and allow the system administrator to de-authenticate (stop current session) users. With the Firewall monitor, you can de-authenticate all currently authenticated users, or select single users to de-authenticate. To permanently stop a user from re-authenticating, change the FortiGate configuration (disable a user account) and then use the User monitor to immediately end the user's current session.

Monitored firewall users are viewed from *User > Monitor > Firewall*.

Firewall page

Lists all authenticated firewall users that are currently authenticated by the FortiGate unit and are active. This page allows you to refresh the information on the page, as well as filter the information.

Refresh	Refresh the Firewall user monitor list.
Filter Settings	Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i> , below the column headings. Use to configure filter settings. Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.
	To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.
	To modify settings, select <i>Change</i> beside the setting and edit the settings.
	To clear all filter settings, select the icon beside <i>Clear all filters</i> .

		To use a filter icon to filter settings within a column, select the filter icon in the column; Filters appears. Within Filters, configure the settings for that column.
		Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i> , below the column headings. Use to configure filter settings. Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.
	De-authenticate All Users	Stop authenticated sessions for all users in the Firewall user monitor list. Users must re-authenticate with the firewall to resume their communication session.
	Page Controls	Use to navigate through the list.
	Column Settings	Customize the table view. You can select the columns to hide or display and specify the column displaying order in the table.
	User Name	The user names of all connected remote users. In the <i>User Name</i> column heading, there is a green arrow beside the name. This green arrow allows you to sort the list alphabetically. By default, the green arrow is up and when you select it, the green arrow is down. The green up arrow indicates that the list is in alphabetical order, starting with the letter A; when the green arrow is down, the list is the opposite.
	User Group	The user group that the remote user is part of.
	Policy ID	The policy identification number.
	Duration	Length of time since the user was authenticated.
	IP Address	The user's source IP address.
	Traffic Volume	The amount of traffic that is going through the unit, that is generated by the user.
	Method	Authentication method used for the user by the unit (authentication methods can be Fortinet Single Sign On Agent, firewall authentication, or NTLM).

IM user monitor list

User lists can be managed to allow or block certain users. Each user can be assigned a policy to allow or block activity for each IM protocol. Each IM function can be individually allowed or blocked providing the administrator the granularity to block the more bandwidth consuming features such as voice chat while still allowing text messaging. The IM user monitor list displays information about instant messaging users who are currently connected. The list can be filtered by protocol. After IM users connect through the firewall, the unit displays which users are connected. You can analyze the list and decide which users to allow or block. A policy can be configured to handle unknown users.

Active IM users are viewed from *User > Monitor > IM*.



IM users who are already logged on before changes are made to the IM user profile will not be affected until their next login. You cannot disconnect users who have already logged on by enabling logon blocking.

IM page	
Lists all active IM users that are currently active. This page allows you to view blocked users as well as users that are currently using a particular IM protocol, such as MSN.	
Protocol	Filter the list by selecting the protocol for which to display current users: AIM, ICQ, MSN, or Yahoo. All current users can also be displayed.
#	The position number of the IM user in the list.
Protocol	The protocol being used.
User Name	The name selected by the user when registering with an IM protocol. The same user name can be used for multiple IM protocols. Each user name/protocol pair appears separately in the list.
Source IP	The IP address where the user initiated the IM session from.
Last Login	The last time the current user used the protocol.
Block	Select to add the user name to the permanent black list. Each user name/protocol pair must be explicitly blocked by the administrator.

The Banned User list

The Banned User list shows all IP addresses and interfaces blocked by NAC quarantine. The list also shows all IP addresses, authenticated users, senders, and interfaces blocked by Data Leak Prevention (DLP). The system administrator can selectively release users or interfaces from quarantine or configure quarantine to expire after a selected time period.

All sessions started by users or IP addresses on the Banned User list are blocked until the user or IP address is removed from the list. All sessions to an interface on the list are blocked until the interface is removed from the list.

You can configure NAC quarantine to add users or IP addresses to the Banned User list under the following conditions:

- **Users or IP addresses that originate attacks detected by IPS** - To quarantine users or IP addresses that originate attacks, enable and configure *Quarantine Attackers* in an IPS Sensor Filter.
- **IP addresses or interfaces that send viruses detected by virus scanning** - To quarantine IP addresses that send viruses or interfaces that accept traffic containing a virus, enable *Quarantine Virus Sender* in an antivirus profile.
- **Users or IP addresses that are banned or quarantined by Data Leak Prevention** - Set various options in a DLP sensor to add users or IP addresses to the Banned User list.

For more information, see [FortiOS Handbook UTM chapter](#).

Banned users are viewed from *User > Monitor > Banned User*.

Banned User page	
Lists all banned users.	
Page Controls	Use to navigate through the list.
Clear	Removes all users and IP addresses from the Banned User list.
#	The position number of the user or IP address in the list.

Ban key	The Ban key.
Application Protocol	The protocol that was used by the user or IP address added to the Banned User list.
Cause or rule	The FortiGate function that caused the user or IP address to be added to the Banned User list. <i>Cause or rule</i> can be IPS, Antivirus, or Data Leak Prevention.
Created	The date and time the user or IP address was added to the Banned User list.
Expires	The date and time the user or IP address will be automatically removed from the Banned User list. If <i>Expires</i> is <i>Indefinite</i> , you must manually remove the user or host from the list.
Delete	Removes the selected user or IP address from the Banned User list.



Authentication servers

FortiGate units support the use of external authentication servers. An authentication server can provide password checking for selected FortiGate users or it can be added as a member of a FortiGate user group.

If you are going to use authentication servers, you must configure the servers before you configure FortiGate users or user groups that require them.



MAC OS and iOS devices, including iPhones and iPads, can perform user authentication with FortiOS units using RADIUS servers, but not with LDAP or TACACS+ servers.

This section includes the following topics:

- [FortiAuthenticator servers](#)
- [RADIUS servers](#)
- [LDAP servers](#)
- [TACACS+ servers](#)
- [FSSO servers](#)
- [RSA ACE \(SecurID\) servers](#)

FortiAuthenticator servers

FortiAuthenticator is an Authentication, Authorization, and Accounting (AAA) server, that includes a RADIUS server, an LDAP server, and can replace the FSSO Collector Agent on a Windows AD network. Multiple FortiGate units can use a single FortiAuthenticator for FSSO, remote authentication, and FortiToken management.

For more information, see the [FortiAuthenticator Administration Guide](#).

RADIUS servers

Remote Authentication and Dial-in User Service (RADIUS) is a broadly supported client-server protocol that provides centralized authentication, authorization, and accounting functions. RADIUS clients are built into gateways that allow access to networks such as Virtual Private Network servers, Network Access Servers (NAS), as well as network switches and firewalls that use authentication. FortiGate units fall into the last category.

RADIUS servers use UDP packets to communicate with the RADIUS clients on the network to authenticate users before allowing them access to the network, to authorize access to resources by appropriate users, and to account or bill for those resources that are used. RADIUS servers are currently defined by RFC 2865 (RADIUS) and RFC 2866 (Accounting), and listen on either UDP ports 1812 (authentication) and 1813 (accounting) or ports 1645 (authentication) and 1646 (accounting) requests. RADIUS servers exist for all major operating systems.

You must configure the RADIUS server to accept the FortiGate unit as a client. FortiGate units use the authentication and accounting functions of the RADIUS server.



FortiOS does not accept all characters from auto generated keys from MS Windows 2008. These keys are very long and as a result RADIUS authentication will not work. Maximum key length for MS Windows 2008 is 128 bytes. In older versions of FSAE, it was 40 bytes.

Microsoft RADIUS servers

Microsoft Windows Server 2000, 2003, and 2008 have RADIUS support built-in. Microsoft specific RADIUS features are defined in RFC 2548. The Microsoft RADIUS implementation can use Active Directory for user credentials.

For details on Microsoft RADIUS server configurations, refer to Microsoft documentation.

RADIUS user database

The RADIUS user database is commonly an SQL or LDAP database, but can also be any combination of:

- usernames and passwords defined in a configuration file
- user account names and passwords configured on the computer where the RADIUS server is installed.

If users are members of multiple RADIUS groups, then the user group authentication timeout value does not apply. See [“Membership in multiple groups” on page 1239](#).

RADIUS authentication with a FortiGate unit

To use RADIUS authentication with a FortiGate unit

- configure one or more RADIUS servers on the FortiGate unit
- assign users to a RADIUS server

When a configured user attempts to access the network, the FortiGate unit will forward the authentication request to the RADIUS server which will match the username and password remotely. Once authenticated the RADIUS server passes the authorization granted message to the FortiGate unit which grants the user permission to access the network.

The RADIUS server uses a “shared secret” key along with MD5 hashing to encrypt information passed between RADIUS servers and clients, including the FortiGate unit. Typically only user credentials are encrypted. Additional security can be configured through IPsec tunnels.

RADIUS attribute value pairs

RADIUS packets include a set of attribute value pairs (AVP) to identify information about the user, their location and other information. The FortiGate unit sends the following RADIUS attributes.

Table 76: FortiOS supported RADIUS attributes

RADIUS Attribute	Name	Description	AVP type
------------------	------	-------------	----------

Table 76: FortiOS supported RADIUS attributes

1	Acct-Session-ID	Unique number assigned to each start and stop record to make it easy to match them, and to eliminate duplicate records.	44
2	username	Name of the user being authenticated	1
3	NAS-Identifier	Identifier or IP address of the Network Access Server (NAS) that is requesting authentication. In this case, the NAS is the FortiGate unit.	32
4	Framed-IP-Address	Address to be configured for the user.	8
5	Fortinet-VSA	See “Vendor-specific attributes” on page 1203	26
6	Acct-Input-Octets	Number of octets received from the port over the course of this service being provided. Used to charge the user for the amount of traffic they used.	42
7	Acct-Output-Octets	Number of octets sent to the port while delivering this service. Used to charge the user for the amount of traffic they used.	43

[Table 77](#) describes the supported authentication events and the RADIUS attributes that are sent in the RADIUS accounting message.

Table 77: RADIUS attributes sent in RADIUS accounting message

	RADIUS ATTRIBUTE						
AUTHENTICATION METHOD	1	2	3	4	5	6	7
Web	✓	✓	✓		✓		
XAuth of IPsec (without DHCP)	✓	✓	✓		✓		
XAuth of IPsec (with DHCP)	✓	✓	✓	✓	✓		
PPTP/L2TP (in PPP)	✓	✓	✓	✓	✓	✓	✓
SSL-VPN	✓	✓	✓		✓		

Vendor-specific attributes

Vendor specific attributes (VSA) are the method RADIUS servers and client companies use to extend the basic functionality of RADIUS. Some major vendors, such as Microsoft, have published their VSAs, however many do not.

In order to support vendor-specific attributes (VSA), the RADIUS server requires a dictionary to define which VSAs to support. This dictionary is typically supplied by the client or server vendor.

The Fortinet RADIUS vendor ID is 12365.

The FortiGate unit RADIUS VSA dictionary is supplied by Fortinet and is available through the Fortinet Knowledge Base (<http://kb.forticare.com>) or through Technical Support. Fortinet's dictionary for FortiOS 4.0 and up is configured this way:

```
##
Fortinet's VSA's
#
VENDOR fortinet 12356
BEGIN-VENDOR fortinet
ATTRIBUTE Fortinet-Group-Name 1 string
ATTRIBUTE Fortinet-Client-IP-Address 2 ipaddr
ATTRIBUTE Fortinet-Vdom-Name 3 string
ATTRIBUTE Fortinet-Client-IPv6-Address 4 octets
ATTRIBUTE Fortinet-Interface-Name 5 string
ATTRIBUTE Fortinet-Access-Profile 6 string
#
# Integer Translations
#
END-VENDOR Fortinet
```

Note that using the Fortinet-Vdom-Name, users can be tied to a specific VDOM on the FortiGate unit. See the documentation provided with your RADIUS server for configuration details.

Role Based Access Control

In Role Based Access Control (RBAC), network administrators and users have varying levels of access to network resources based on their role, and that role's requirement for access specific resources. For example, a junior accountant does not require access to the sales presentations, or network user account information.

There are three main parts to RBAC: role assignment, role authorization, and transaction authorization. Role assignment is accomplished when someone in an organization is assigned a specific role by a manager or HR. Role authorization is accomplished when a network administrator creates that user's RADIUS account and assigns them to the required groups for that role. Transaction authorization occurs when that user logs on and authenticates before performing a task.

RBAC is enforced when FortiOS network users are remotely authenticated via a RADIUS server. For users to authenticate, an identity-based security policy must be matched. That policy only matches a specific group of users. If VDOMs are enabled, the matched group will be limited to a specific VDOM. Using this method network administrators can separate users into groups that match resources, protocols, or VDOMs. It is even possible to limit users to specific FortiGate units if the RADIUS servers serve multiple FortiOS units.

For more information on identity-based policies, see [“Authentication in security policies” on page 1246](#).

Configuring the FortiGate unit to use a RADIUS server

The information you need to configure the FortiGate unit to use a RADIUS server includes

- the RADIUS server's domain name or IP address
- the RADIUS server's shared secret key.

You can optionally specify the NAS IP or Called Station ID. When configuring the FortiGate to use a RADIUS server, the FortiGate is a Network Access Server (NAS). If the FortiGate interface has multiple IP addresses, or you want the RADIUS requests to come from a different address you can specify it here. Called Station ID applies to carrier networks. However, if the NAS IP is not included in the RADIUS configuration, the IP of the FortiGate unit interface that communicates with the RADIUS server is used instead.

A maximum of 10 remote RADIUS servers can be configured on the FortiGate unit. One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 1224](#).

On the FortiGate unit, the default port for RADIUS traffic is 1812. Some RADIUS servers use port 1645. If this is the case with your server, you can either:

- Re-configure the RADIUS server to use port 1812. See your RADIUS server documentation for more information on this procedure.

or

- Change the FortiGate unit default RADIUS port to 1645 using the CLI:

```
config system global
    set radius_port 1645
end
```

One wildcard admin account can be added to the FortiGate unit when using RADIUS authentication. This uses the wildcard character to allow multiple admin accounts on RADIUS to use a single account on the FortiGate unit. See [“Example — wildcard admin accounts - CLI” on page 1211](#).

To configure the FortiGate unit for RADIUS authentication - web-based manager

- 1 Go to *User > Remote > RADIUS* and select *Create New*.

2 Enter the following information and select OK.

Name	A name to identify the RADIUS server on the FortiGate unit.
Type	Select Query. Selecting a Type of <i>Dynamic Start</i> configures the Dynamic Profile feature. See “ Dynamic profiles and end points ” on page 1325.
Primary Server Name/IP	Enter the domain name (such as fgt.exmaple.com) or the IP address of the RADIUS server.
Primary Server Secret	Enter the server secret key, such as radiusSecret. This can be a maximum of 16 characters long. This must match the secret on the RADIUS primary server.
Secondary Server Name/IP	Optionally enter the domain name (such as fgt.exmaple.com) or the IP address of the secondary RADIUS server.
Secondary Server Secret	Optionally, enter the secondary server secret key, such as radiusSecret2. This can be a maximum of 16 characters long. This must match the secret on the RADIUS secondary server.
Authentication Scheme	If you know the RADIUS server uses a specific authentication protocol, select it from the list. Otherwise select Use Default Authentication Scheme. The Use Default option will usually work.
NAS IP/ Called Station ID	Enter the IP address to be used as an attribute in RADIUS access requests. NAS-IP-Address is RADIUS setting or IP address of FortiGate interface used to talk to RADIUS server, if not configured. Called Station ID is same value as NAS-IP Address but in text format.
Include in every User Group	When enabled this RADIUS server will automatically be included in all user groups. This is useful if all users will be authenticating with the remote RADIUS server.



For MAC OS and iOS devices to authenticate, you must use MS-CHAP-v2 authentication. In the CLI, the command is `set auth-type ms_chap_v2`.

3 Select OK.

To configure the FortiGate unit for RADIUS authentication - CLI example

```
config user radius
edit ourRADIUS
set auth-type auto
set server 10.11.102.100
set secret radiusSecret
end
```

For more information about RADIUS server options, refer to the [FortiGate CLI Reference](#).

Troubleshooting RADIUS

To test the connection to the RADIUS server use the following command:

```
diagnose test authserver radius-direct <server_name or IP> <port  
number> <secret>
```

For the port number, enter -1 to use the default port. Otherwise enter the port number to check.

Additional RADIUS related troubleshooting is located at [“Troubleshooting FSSO” on page 1318](#)

LDAP servers

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, and printers. LDAP consists of a data-representation scheme, a set of defined operations, and a request/response network.

The scale of LDAP servers range from big public servers such as BigFoot and Infospace, to large organizational servers at universities and corporations, to small LDAP servers for workgroups that may be using OpenLDAP. This document focuses on the institutional and workgroup applications of LDAP.

This section includes:

- [Components and topology](#)
- [LDAP directory organization](#)
- [Configuring the FortiGate unit to use an LDAP server](#)
- [Example — wildcard admin accounts - CLI](#)
- [Example of LDAP to allow Dial-in through member-attribute - CLI](#)
- [Troubleshooting LDAP](#)

Components and topology

LDAP organization starts with directories. A directory is a set of objects with similar attributes organized in a logical and hierarchical way. Generally, an LDAP directory tree reflects geographic and organizational boundaries, with the Domain name system (DNS) names to structure the top level of the hierarchy. The common name identifier for most LDAP servers is `cn`, however some servers use other common name identifiers such as `uid`.

When LDAP is configured and a user is required to authenticate the general steps are:

- 1 The FortiGate unit contacts the LDAP server for authentication.
- 2 To authenticate with the FortiGate unit, the user enters a username and password.
- 3 The FortiGate unit sends this username and password to the LDAP server.
- 4 If the LDAP server can authenticate the user, the user is successfully authenticated with the FortiGate unit.
- 5 If the LDAP server cannot authenticate the user, the connection is refused by the FortiGate unit.

Binding

Binding is the step where the LDAP server authenticates the user. If the user is successfully authenticated, binding allows the user access to the LDAP server based on that user's permissions.

The FortiGate unit can be configured to use one of three types of binding:

- anonymous - bind using anonymous user search
- regular - bind using username/password and then search
- simple - bind using a simple password authentication without a search

You can use simple authentication if the user records all fall under one domain name (dn). If the users are under more than one dn, use the anonymous or regular type, which can search the entire LDAP database for the required username.

If your LDAP server requires authentication to perform searches, use the regular type and provide values for username and password.

Supported versions

The FortiGate unit supports LDAP protocol functionality defined in RFC 2251: Lightweight Directory Access Protocol v3, for looking up and validating usernames and passwords. FortiGate LDAP supports all LDAP servers compliant with LDAP v3, including FortiAuthenticator. In addition, FortiGate LDAP supports LDAP over SSL/TLS, which can be configured only in the CLI.

FortiGate LDAP does not support proprietary functionality, such as notification of password expiration, which is available from some LDAP servers. FortiGate LDAP does not supply information to the user about why authentication failed.

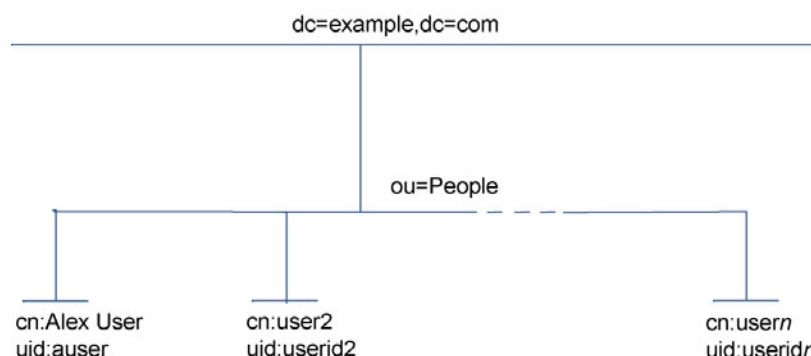
LDAP directory organization

To configure your FortiGate unit to work with an LDAP server, you need to understand the organization of the information on the server.

The top of the hierarchy is the organization itself. Usually this is defined as Domain Component (DC), a DNS domain. If the name contains a dot, such as `example.com`, it is written as two parts separated by a comma: `dc=example,dc=com`.

In this example, Common Name (CN) identifiers reside at the Organization Unit (OU) level, just below DC. The Distinguished Name (DN) is `ou=People,dc=example,dc=com`.

Figure 103: LDAP object hierarchy



In addition to the DN, the FortiGate unit needs an identifier for the individual person. Although the FortiGate unit GUI calls this the Common Name (CN), the identifier you use is not necessarily CN. On some servers, CN is the full name of a person. It might be more convenient to use the same identifier used on the local computer network. In this example, User ID (UID) is used.

Locating your identifier in the hierarchy

You need to determine the levels of the hierarchy from the top to the level that contain the identifier you want to use. This defines the DN that the FortiGate unit uses to search the LDAP database. Frequently used distinguished name elements include:

- uid (user identification)
- pw (password)
- cn (common name)
- ou (organizational unit)
- o (organization)
- c (country)

One way to test this is with a text-based LDAP client program. For example, OpenLDAP includes a client, `ldapsearch`, that you can use for this purpose.

Enter the following at the command line:

```
ldapsearch -x '(objectclass=*)'
```

The output is lengthy, but the information you need is in the first few lines:

```
version: 2
#
# filter: (objectclass=*)
# requesting: ALL

dn: dc=example,dc=com
dc: example
objectClass: top
objectClass: domain

dn: ou=People,dc=example,dc=com
ou: People
objectClass: top
objectClass: organizationalUnit
...
dn: uid=tbrown,ou=People,dc=example,dc=com
uid: tbrown
cn: Tom Brown
```

In the output above, you can see `tbrown` (uid) and `Tom Brown` (cn). Also note the dn is `ou=People, dc=example, dc=com`.

Configuring the FortiGate unit to use an LDAP server

After you determine the common name and distinguished name identifiers and the domain name or IP address of the LDAP server, you can configure the server on the FortiGate unit. The maximum number of remote LDAP servers that can be configured is 10.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 1224](#).

To configure the FortiGate unit for LDAP authentication - web-based manager

- 1 Go to *User > Remote > LDAP* and select *Create New*.
- 2 Enter a name for the LDAP server.
- 3 In *Server Name/IP* enter the server's FQDN or IP address.
- 4 If the server does not use port 389, enter the port number in the *Server Port* field.
- 5 Enter the *Common Name Identifier* (20 characters maximum).
cn is the default, and is used by most LDAP servers.
- 6 In the *Distinguished Name* field, enter the base distinguished name for the server using the correct X.500 or LDAP format.
The FortiGate unit passes this distinguished name unchanged to the server. The maximum number of characters is 512.
If you don't know the distinguished name, leave the field blank and select the Query icon to the right of the field. See the ["Using the Query icon" on page 1211](#).
- 7 In *Bind Type*, select *Regular*.
- 8 In *User DN*, enter the LDAP administrator's distinguished name.
- 9 In *Password*, enter the LDAP administrator's password.
- 10 Select *OK*.



To verify your Distinguished Name field is correct, you can select the Query icon. If your DN field entry is valid, you will see the part of the LDAP database it defines. If your DN field entry is not valid, it will display an error message and return no information.

For detailed information about configuration options for LDAP servers, see the Online Help on your FortiGate unit or the FortiGate CLI Reference.

To configure the FortiGate unit for LDAP authentication - CLI example

```
config user ldap
  edit ourLDAPsrv
    set server 10.11.101.160
    set cnid cn
    set dn cn=users,dc=office,dc=example,dc=com
    set type regular
    set username
      cn=administrator,cn=users,dc=office,dc=example,dc=com
    set password w5AiGVMLkgyPQ
    set password-expiry-warning enable
    set password-renewal enable
  end
```

password-expiry-warning and password-renewal

In SSLVPN, when an LDAP user is connecting to the LDAP server it is possible for them to receive any pending password expiry or renewal warnings. When the password renewal or expiry warning exists, SSLVPN users will see a prompt allowing them to change their password.

`password-expiry-warning` allows FortiOS to detect from the LDAP server when a password is expiring or has expired using server controls or error codes.

`password-renewal` allows FortiOS to perform the online LDAP password renewal operations the LDAP server expects.

On an OpenLDAP server, when a user attempts to logon with an expired password they are allowed to logon on but only to change their password.

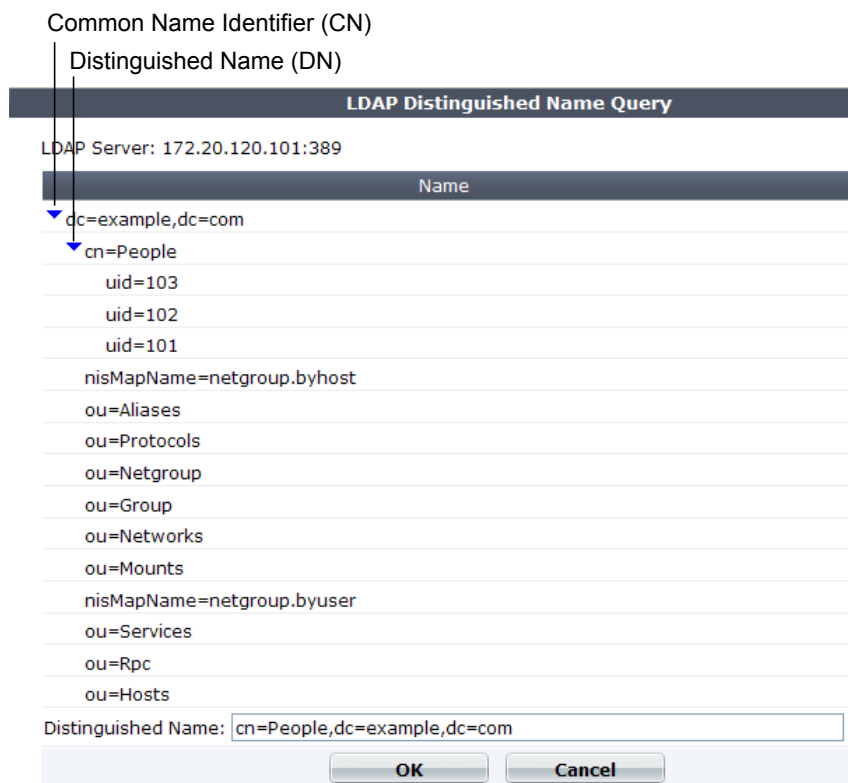
When changing passwords on a Windows AD system, the connection must be SSL-protected.

Using the Query icon

The LDAP Distinguished Name Query list displays the LDAP directory tree for the LDAP server connected to the FortiGate unit. This helps you to determine the appropriate entry for the DN field. To see the distinguished name associated with the Common Name identifier, select the Expand icon next to the CN identifier. Select the DN from the list. The DN you select is displayed in the Distinguished Name field. Select OK and the Distinguished Name you selected will be saved in the Distinguished Name field of the LDAP Server configuration.

To see the users within the LDAP Server user group for the selected Distinguished Name, expand the Distinguished Name in the LDAP Distinguished Name Query tree.

Figure 104: LDAP server Distinguished Name Query tree



Example — wildcard admin accounts - CLI

A wildcard admin account is an administrator account with the wildcard option enabled. This option allows multiple different remote administration accounts to match one local administration account, avoiding the need to set up individual admin accounts on the FortiGate unit. Instead multiple LDAP admin accounts will all be able to use one FortiGate admin account.

The initial benefit of wildcard admin accounts is fast configuration of the FortiGate unit's administration account to work with your LDAP network. The many to one ratio saves on effort, and potential errors.

The ongoing benefit is that as long as the users on the LDAP system belong to that group, and the test admin user settings don't change on the FortiGate unit, no other work is required. This point is important as it can help avoid many system updates or changes that would otherwise require changes to the LDAP administrator account configuration. Even if a user is added to or removed from the LDAP group, no changes are required on the FortiGate unit.

Two potential issues with wildcard admin accounts are that multiple users may be logged on to the same account at the same time. This becomes an issue if they are changing the same information at the same time. The other potential issue is that security is reduced because multiple people have login access for the same account. If each user was assigned their own account, a hijacking of one account would not affect the other users.

Note that wildcard admin configuration also applies to RADIUS. When configuring for RADIUS, configure the RADIUS server, and RADIUS user group instead of LDAP. When using web-based management, wildcard admin is the only type of remote administrator account that does not require you to enter a password on account creation. That password is normally used when the remote authentication server is unavailable during authentication.

In this example, default values are used where possible. If a specific value is not mentioned, it is set to its default value.

Configuring the LDAP server

The important parts of this configuration are the username and group lines. The username is the domain administrator account. The group binding allows only the group with the name GRP to access.



The dn used here is as an example only. On your network use your own domain name.

To configure LDAP server - CLI

```
config user ldap
  edit "ldap_server"
    set server "192.168.201.3"
    set cnid "sAMAccountName"
    set dn "DC=example,DC=com,DC=au"
    set type regular
    set username "CN=Administrator,CN=Users,DC=example,DC=COM"
    set password *
    set group "CN=GRP,OU=training,DC=example,DC=COM"
    set filter ""
  next
end
```

To configure the user group and add the LDAP server - CLI

```
config user group
  edit "ldap_grp"
    set member "ldap"
  config match
```

```
edit 1
    set server-name "ldap"
    set group-name "TRUE"
next
end
next
end
```

Configuring the admin account

The wildcard part of this example is only available in the CLI for admin configuration. When enabled, this allows all LDAP group members to login to the FortiGate unit without the need to create a separate admin account for each user. In effect the members of that group will each be able to login as “test”.

To configure the admin account - CLI

```
config system admin
    edit "test"
        set remote-auth enable
        set accprofile "super_admin"
        set wildcard enable
        set remote-group "ldap_grp"
    next
end
```

For troubleshooting, test that the admin account is operational, and see [“Troubleshooting LDAP” on page 1214](#).

Example of LDAP to allow Dial-in through member-attribute - CLI

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to setup a VPN connection simply based on an attribute that is set to TRUE, instead of based on being part of a specific group.

In AD, the "Allow Dial-In" property is activated in the user properties, and this sets the msNPAllowDialin attribute to "TRUE".

This same procedure can be used for other member attributes, as your system requires.

This example works with FortiOS 4.0 MR2. The filter command was removed in FortiOS 4.0 MR3.

Configuring LDAP member-attribute settings

To accomplish this with a FortiGate unit, the member attribute must be set. Setting member attributes can only be accomplished through the CLI using the `member-attr` keyword - the option is not available through the web-based manager.

Before configuring the FortiGate unit, the AD server must be configured and have the msNPAllowDialin attribute set to "TRUE" for the users in question. If not, those users will not be able to properly authenticate.

The dn used here is as an example only. On your network use your own domain name.

To configure user LDAP member-attribute settings - CLI

```
config user ldap
    edit "ldap_server"
        set server "192.168.201.3"
        set cnid "sAMAccountName"
```

```

set dn "DC=fortinet,DC=com,DC=au"
set type regular
set username "fortigate@example.com"
set password *****
set filter "(&(uid=%u) (msNPAllowDialin=TRUE))"
set member-attr "msNPAllowDialin"
next
end

```

Configuring LDAP group settings

A user group that will use LDAP must be configured. This example adds the member `ldap` to the group which is the LDAP server name that was configured earlier.

To configure LDAP group settings - CLI

```

config user group
edit "ldap_grp"
set member "ldap"
config match
edit 1
set server-name "ldap"
set group-name "TRUE"
next
end
next
end

```

Once these settings are in place, users can authenticate.

Troubleshooting LDAP

The examples in this section use the values from the previous example.

LDAP user test

A quick way to see if the LDAP configuration is correct is to run a diagnose CLI command with LDAP user information. The following command tests with a user called `netAdmin` and a password of `fortinet`. If the configuration is correct the test will be successful.

```

FGT# diag test authserver ldap ldap_server netAdmin fortinet
'ldap_server' is not a valid ldap server name — an LDAP server by that
name has not been configured on the FortiGate unit, check your spelling.
authenticate 'netAdmin' against 'ldap_server' failed! — the user
netAdmin does not exist on ldap_server, check your spelling of both the user and
server and ensure the user has been configured on the FortiGate unit.

```

LDAP authentication debugging

For a more in-depth test, you can use a `diag debug` command. The sample output from a shows more information about the authentication process that may prove useful if there are any problems.

Ensure the "Allow Dial-in" attribute is still set to "TRUE" and run the following CLI command. `fnbamd` is the Fortinet non-blocking authentication daemon.

```

FGT# diag debug enable
FGT# diag debug reset
FGT# diag debug application fnbamd -1

```

```
FGT# diag debug enable
```

The output will look similar to:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='TRUE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Passed group matching
```

If the "Allow Dial-in" attribute is not set but it is expected, the last line of the above output will instead be:

```
fnbamd_auth_poll_ldap-Failed group matching
```

TACACS+ servers

When users connect to their corporate network remotely, they do so through a remote access server. As remote access technology has evolved, the need for security when accessing networks has become increasingly important. This need can be filled using a Terminal Access Controller Access-Control System (TACACS+) server.

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices via one or more centralized servers. TACACS+ allows a client to accept a username and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies the user access to the network.

TACACS+ offers fully encrypted packet bodies, and supports both IP and AppleTalk protocols. TACACS+ uses TCP port 49, which is seen as more reliable than RADIUS's UDP protocol.

There are several different authentication protocols that TACACS+ can use during the authentication process:

Protocol	Definition
ASCII	Machine-independent technique that uses representations of English characters. Requires user to type a username and password that are sent in clear text (unencrypted) and matched with an entry in the user database stored in ASCII format.
PAP	Password Authentication Protocol (PAP) Used to authenticate PPP connections. Transmits passwords and other user information in clear text.
CHAP	Challenge-Handshake Authentication Protocol (CHAP) Provides the same functionality as PAP, but is more secure as it does not send the password and other user information over the network to the security server.
MS-CHAP	MicroSoft Challenge-Handshake Authentication Protocol v1 (MSCHAP) Microsoft-specific version of CHAP.
default	The default protocol configuration, Auto, uses PAP, MS-CHAP, and CHAP, in that order.

Configuring a TACACS+ server on the FortiGate unit

A maximum of 10 remote TACACS+ servers can be configured for authentication.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [“Creating users” on page 1224](#).

To configure the FortiGate unit for TACACS+ authentication - web-based manager

- 1 Go to *User > Remote > TACACS+* and select *Create New*.
- 2 Enter the following information, and select *OK*.

Name	Enter the name of the TACACS+ server.
Server Name/IP	Enter the server domain name or IP address of the TACACS+ server.
Server Key	Enter the key to access the TACACS+ server.
Authentication Type	Select the authentication type to use for the TACACS+ server. <i>Auto</i> tries PAP, MSCHAP, and CHAP (in that order).

To configure the FortiGate unit for TACACS+ authentication - CLI

```
config user tacacs+
  edit <server_name>
    set auth-type {ascii | auto | chap | ms_chap | pap}
    set key <server_key>
    set tacacs+-port <tacacs+_port_num>
    set server <domain>
  end
```

FSSO servers

Novell and Microsoft Windows networks provide user authentication based on directory services: eDirectory for Novell, Active Directory for Windows. Users can log on at any computer in the domain and have access to resources as defined in their user account. The Fortinet Single Sign On (FSSO) agent enables FortiGate units to authenticate these network users for security policy or VPN access without asking them again for their username and password.

When a user logs in to the Windows or Novell domain, the FSSO agent sends the FortiGate unit the user's IP address and the names of the user groups to which the user belongs. The FortiGate unit uses this information to maintain a copy of the domain controller user group database. Because the domain controller authenticates users, the FortiGate unit does not perform authentication. It recognizes group members by their IP address.

In the FortiOS FSSO configuration, you specify the server where the FSSO Collector agent is installed. The Collector agent retrieves the names of the Novell or Active Directory user groups from the domain controllers on the domains, and then the FortiGate unit gets them from the Collector agent. You cannot use these groups directly. You must define FSSO type user groups on your FortiGate unit and then add the Novell or Active Directory user groups to them. The FSSO user groups that you created are used in security policies and VPN configurations to provide access to different services and resources.

FortiAuthenticator servers can replace the Collector agent when FSSO is using polling mode. The benefits of this is that FortiAuthenticator is a stand-alone server that has the necessary FSSO software pre-installed. For more information, see the [FortiAuthenticator Administration Guide](#).

For more information about Directory Services and FSSO, see “FSSO integration with Windows AD or Novell” on page 1283.

RSA ACE (SecurID) servers

SecurID is a two-factor system that uses one-time password (OTP) authentication. It is produced by the company RSA. This system includes portable tokens carried by users, an RSA ACE/Server, and an Agent Host. In our configuration, the FortiGate unit is the Agent Host.

Components

When using SecurID, users carry a small device or “token” that generates and displays a random password. According to RSA, each SecurID authenticator token has a unique 64-bit symmetric key that is combined with a powerful algorithm to generate a new code every 60 seconds. The token is time-synchronized with the SecurID RSA ACE/Server.

The RSA ACE/Server is the management component of the SecurID system. It stores and validates the information about the SecurID tokens allowed on your network. Alternately the server could be an RSA SecurID 130 Appliance.

The Agent Host is the server on your network, in this case it is the FortiGate unit, that intercepts user logon attempts. The Agent Host gathers the user ID and password entered from their SecurID token, and sends that information to the RSA ACE/Server to be validated. If valid, a reply comes back indicating it is a valid logon and the FortiGate unit allows the user access to the network resources specified in the associated security policy.

Configuring the SecurID system

To use SecurID with a FortiGate unit, you need:

- to configure the RSA server and the RADIUS server to work with each other (see RSA server documentation)
- [To configure the RSA SecurID 130 Appliance](#)
- or
- [To configure the FortiGate unit as an Agent Host on the RSA ACE/Server](#)
- [To configure the FortiGate unit to use the RADIUS server](#)
- [To create a SecurID user group and user](#)
- [To configure a security policy with SecurID authentication](#)

The following instructions are based on RSA ACE/Server version 5.1, or RSA SecurID 130 Appliance, and assume that you have successfully completed all the external RSA and RADIUS server configuration steps listed above.

For this example, the RSA server is on the internal network, with an IP address of 192.128.100.100. The FortiGate unit internal interface address is 192.168.100.3, RADIUS shared secret is fortinet123, RADIUS server is at IP address 192.168.100.102.

To configure the RSA SecurID 130 Appliance

- 1 Go to the IMS Console for SecurID and logon.

- 2 Go to *RADIUS > RADIUS Clients*, and select *Add New*.
- 3 Enter the following information to configure your FortiGate as a SecurID Client, and select *Save*.

RADIUS Client Basics	
Client Name	FortiGate
Associated RSA Agent	FortiGate
RADIUS Client Settings	
IP Address	192.168.100.3 The IP address of the FortiGate unit internal interface.
Make / Model	Select Standard Radius
Shared Secret	fortinet123 The RADIUS shared secret.
Accounting	Leave unselected
Client Status	Leave unselected

To configure the FortiGate unit as an Agent Host on the RSA ACE/Server

- 1 On the RSA ACE/Server computer, go to *Start > Programs > RSA ACE/Server*, and then *Database Administration - Host Mode*.
- 2 On the *Agent Host* menu, select *Add Agent Host*.
- 3 Enter and save the following information.

Name	FortiGate
Network Address	192.168.100.3 The IP address of the FortiGate unit.
Secondary Nodes	Optionally enter other IP addresses that resolve to the FortiGate unit.

If needed, refer to the RSA ACE/Server documentation for more information.

To configure the FortiGate unit to use the RADIUS server

- 1 Go to *User > Remote > RADIUS* and select *Create New*.
- 2 Enter the following information, and select *OK*.

Name	RSA
Type	Query
Primary Server Address	192.168.100.102 Optionally select <i>Test</i> to ensure the IP address is correct and the FortiGate can contact the RADIUS server.

Primary Server Secret	fortinet123
Authentication Scheme	Select <i>Use Default Authentication Scheme</i> .

To create a SecurID user group and user

- 1 Go to *User > User Group*, and select *Create New*.
- 2 Enter the following information, and select OK.

Name	RSA_group
Remote Authentication servers	Select the RSA server.

- 3 Go to *User > User > User*, and select *Create New*.
- 4 Enter the following information, and select OK.

User Name	wloman
Match User on RADIUS server	RSA
Add this user to groups	Select RSA_group

To test this configuration, on your FortiGate unit use the CLI command:

```
diag test auth rad RSA auto wloman 111111111
```

The series of 1s is the one time password that your RSA SecurID token generates and you enter.

Using the SecurID user group for authentication

You can use the SecurID user group in several FortiOS features that authenticate by user group including

- [Security policy](#)
- [IPsec VPN XAuth](#)
- [PPTP VPN](#)
- [SSL VPN](#)

The following sections assume the SecurID user group is called `securIDgrp` and has already been configured. Unless otherwise states, default values are used.

Security policy

To use SecurID in a security policy, you must include the SecurID user group in an identity-based security policy. This procedure will create a security policy that allows HTTP, FTP, and POP3 traffic from the `internal` interface to `wan1`. If these interfaces are not available on your FortiGate unit, substitute other similar interfaces.

To configure a security policy with SecurID authentication

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Select `internal` for *Source Interface/Zone*.

- 4 Select `wan1` for *Destination Interface/Zone*.
- 5 For both *Source Address* and *Destination Address* select *all*.
- 6 Select *Enable identity-based Policy* and select *Add*.
- 7 Move `securIDgrp` from the *Available User Groups* to the list on the right.
- 8 Move `HTTP`, `FTP`, and `POP3` from *Available Services* to the list on the right.
- 9 Select *OK*.

You are returned to the security policy creation page, with the information you just entered in the identity-based policy (IBP) table as Rule ID 1.

- 10 Select *OK*.

The SecurID security policy is configured.

To generate usage reports on traffic authenticated with this policy, when you are adding the IBP rule, enable *Log Traffic*.

To either limit traffic or guarantee minimum bandwidth for traffic that uses the SecurID security policy, when you are adding the IBP rule, select traffic shaping and select one of the default shapers from the list such as `guarantee-100kbps`.

To customize any challenge pages or logon pages users will see when authenticating against this security policy, select *Customize Authentication Messages* and select the icon that appears. This takes you to the Edit Message page where you can customize the login challenge page, login failed page, and others.

For more details on configuring security policies, see the [FortiOS Handbook FortiGate Fundamentals chapter](#).

IPsec VPN XAuth

Extended Authentication (XAuth) increases security by requiring additional user authentication information in a separate exchange at the end of the VPN Phase 1 negotiation. If the SecurID user group is used, this extended information will require users to enter their SecurID code. For more on XAuth, see “[Configuring XAuth authentication](#)” on page 1259.

This Phase 1 configuration will be named `securIDxAuth` and it will connect with IP address `10.11.101.155` on the `wan1` interface.

To configure IPsec VPN XAuth with SecurID authentication - web-based manager


- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*.
- 3 Enter `securIDxAuth` for *Name*.
- 4 Select *Dialup User* for *Remote Gateway*.
- 5 Select *Preshared Key* for *Authentication Method*.
- 6 Enter `fortinet` for the *Pre-shared Key*.
- 7 Select *Advanced*.
- 8 For XAuth, select *Enable as Server*.
- 9 Select *AUTO* for *Server Type*.
- 10 Select `securIDgrp` for *User Group*.
- 11 Select *OK*.

PPTP VPN

PPTP VPN is configured in the CLI. In the PPTP configuration (`config vpn pptp`), set `usrgrp` to the SecurID user group.

SSL VPN

In the SecurID user group, select the appropriate web portal for these users. In the security policy for the SSL VPN, include the SecurID user group in the list of selected user groups.



Users and user groups

FortiGate authentication controls system access by user group. By assigning individual users to the appropriate user groups you can control each user's access to network resources. The members of user groups are user accounts, of which there are several types. Local users and peer users are defined on the FortiGate unit. User accounts can also be defined on remote authentication servers.

This section describes how to configure local users and peer users and then how to configure user groups. For information about configuration of authentication servers see [“Authentication servers” on page 1201](#).

This section contains the following topics:

- [Users](#)
- [User groups](#)

Users

A user is a user account consisting of username, password, and in some cases other information, configured on the FortiGate unit or on an external authentication server. Users can access resources that require authentication only if they are members of an allowed user group. There are several different types of user accounts with slightly different methods of authentication:

Table 78: How the FortiGate unit authenticates different types of users

User type	Authentication
Local user, password stored on the FortiGate unit	The username and password must match a user account stored on the FortiGate unit. Authentication by FortiGate security policy.
Local user, password stored on a remote server	The username must match a user account stored on the FortiGate unit and the username and password must match a user account stored on the remote authentication server.
Authentication server user	A FortiGate user group can include user accounts or groups that exist on a remote authentication server.
FSSO user	With Fortinet Single Sign On (FSSO), users on a Microsoft Windows or Novell network can use their network authentication to access resources through the FortiGate unit. Access is controlled through FSSO user groups which contain Windows or Novell user groups as their members.
Peer user with certificate authentication	A peer user is a digital certificate holder that authenticates using a client certificate. No password is required, unless two-factor authentication is enabled.

This section includes:

- [Local users](#)

- [PKI or peer users](#)
- [Two-factor authentication](#)
- [FortiToken](#)
- [Monitoring users](#)

Local users

Local users are defined on the FortiGate unit. To define a local user you need:

- a username
- a password or the name of the authentication server that contains the user account



If the user is authenticated externally, the username on the FortiGate unit must be identical to the username on the authentication server. You may also still be prompted for a password. If the connection to the remote authentication server is interrupted during authentication the local password will be used if it exists to ensure access.

Local users are authenticated through authentication security policies. See

This section includes:

- [Creating users](#)
- [Removing users](#)
- [Removing references to users](#)

Creating users

Before configuring any authentication, except dynamic profiles, you must first create local users. For more about dynamic profiles, see [“Configuring dynamic profile” on page 1332](#).

When creating a new user, there are only two differences between a local and a remote user:

- local users require a password to be configured
- remote users do not require a password, but do require a remote authentication server to be configured

To create a local user - web-based manager

- 1 Go to *User > User* and select *Create New*.
- 2 Enter the username in the *username* field.
 - Select *Password* and type a password. Best practices dictate that the password be at least six characters long.



To authenticate this user using an external authentication server, select the *Match user* option for the appropriate type of server and select the server name. Password is not required. You must configure the remote server access first. See [“Authentication servers” on page 1201](#).

- 3 Optionally select *Enable Two-factor Authentication* to use that option with this user. When enabled, additional options will be displayed. Select one of the following options and configure it as stated.
 - Select *FortiToken*, and choose the FortiToken serial number to associate with this user.
 - Select *Email to* and enter the user's email address to email them the token code.
 - Select *SMS* and enter the Mobile Provider from the list, and enter the user's mobile phone number that will receive the token code in a text message.
- 4 Select *OK*.



The Mobile Provider for SMS must be entered in the CLI using the `config user sms-provider` command before it will be available to select in the web-based manager.

To create a local user - CLI examples

Locally authenticated user

```
config user local
edit user1
set type password
set passwd ljt_pj2gpepfdw
end
```

User authenticated on an LDAP server

```
config user local
edit user2
set type ldap
set ldap_server ourLDAPsrv
end
```

User authenticated on a RADIUS server

```
config user local
edit user3
set type radius
set radius_server ourRADIUSsrv
end
```

User authenticated on a TACACS+ server

```
config user local
edit user4
set type tacacs+
set tacacs+_server ourTACACS+srv
end
```

User authenticated with a FortiToken

```
config user local
edit user5
set type password
set passwd ljt_pj2gpepfdw
set two_factor fortitoken
set fortitoken 182937197
end
```

User authenticated using email

```
config user local
  edit user6
    set type password
    set passwd ljt_pj4h7epfdw
    set two_factor email
    set email-to user6@sample.com
  end
```

User authenticated using SMS text message

```
config user sms-provider
  edit "Sample Mobile Inc"
    set mail-server mail.sample.com
  end

config user local
  edit user7
    set type password
    set passwd 3ww_pjt68dw
    set two_factor sms
    set sms-provider "Sample Mobile Inc"
    set sms-phone 2025551234
  end
```

Removing users

Best practices dictate that when a user account is no longer in use, it be deleted. Removing local and remote users from FortiOS involve the same steps.

If the user account is referenced by any configuration objects those references must be removed before the user can be deleted. See [“Removing references to users” on page 1226](#).

To remove a user from the FortiOS configuration - web-based manager

- 1 Go to *User > User*.
- 2 Select the check box of the user that you want to remove.
- 3 Select *Delete*.
- 4 Select *OK*.

To remove a user from the FortiOS configuration - CLI example

```
config user local
  delete user4444
end
```

Removing references to users

You cannot remove a user that belongs to a user group. Remove the user from the user group first, and then delete the user.

To remove references to a user - web-based manager

- 1 Go to *User > User > User*.
- 2 If the number in the far right column for the selected user contains any number other than zero, select it.

- 3 A more detailed list of object references to this user is displayed. Use its information to find and remove these references to allow you to delete this user.

PKI or peer users

A PKI, or peer user, is a digital certificate holder. A PKI user account on the FortiGate unit contains the information required to determine which CA certificate to use to validate the user's certificate. Peer users can be included in firewall user groups or peer certificate groups used in IPsec VPNs. For more on certificates, see [“Certificates overview” on page 1264](#).

To define a peer user you need:

- a peer username
- the text from the subject field of the user's certificate, or the name of the CA certificate used to validate the user's certificate

Creating a peer user

The configuration page for PKI users in the web-based manager is not available unless there is at least one peer user defined. Follow the CLI-based instructions to create the first peer user. Optionally, you can then logon to the web-based manager to configure additional PKI (peer) users.

To create a peer user for PKI authentication - CLI example

```
config user peer
edit peer1
set subject E=peer1@mail.example.com
set ca CA_Cert_1
end
```



If you create a PKI user in the CLI with no values in `subject` or `ca`, you will not be able to open the user's record in the web-based manager. The CLI will prompt you to add a value in *Subject* (`subject`) or *CA* (`ca`).

To create a peer user for PKI authentication - web-based manager

- 1 Go to *User > PKI* and select *Create New*.
- 2 Enter the user name.
- 3 Fill in at least one of the following fields:

Subject	The text string that appears in the Subject field of the user's certificate.
CA	Select the CA certificate that must be used to authenticate this peer user.

- 4 Optionally you can select Two-factor authentication. See [“Two-factor authentication” on page 1228](#).
- 5 Select *OK*.

There are other configuration settings that can be added or modified for PKI authentication. For example, you can configure the use of an LDAP server to check access rights for client certificates. For information about the detailed PKI configuration settings only available through the CLI, see the [FortiGate CLI Reference](#).

Two-factor authentication

The standard logon requires a username and password. This is one factor authentication—your password is one piece of information you need to know to gain access to the system.

Two factor authentication adds the requirement for another piece of information for you logon. Generally the two factors are something you know (password) and something you have (certificate, token, etc.). This makes it harder for a hacker to steal your logon information. For example if you have a FortiToken device, the hacker would need to both use it and know your password to gain entry to your account.

Two-factor authentication is available on both user and admin accounts. But before you enable two-factor authentication on an administrator account, you need to ensure you have a second administrator account configured to guarantee administrator access to the FortiGate unit if you are unable to authenticate on the main admin account for some reason.



Two-factor authentication does not work with explicit proxies.

The methods of two-factor authentication include:

- [Certificate](#)
- [Email](#)
- [SMS](#)
- [FortiToken](#)

Certificate

You can increase security by requiring both certificate and password authentication for PKI users. Certificates are installed on the user's computer. Requiring a password also protects against unauthorized use of that computer.

Optionally peer users can enter the code from their FortiToken instead of the certificate.

To create a peer user with two-factor authentication - web-based manager

While configuring a peer user (see [“PKI or peer users” on page 1227](#)), select *Require two-factor authentication* and enter a password.

To create a peer user with two-factor authentication - CLI example

```
config user peer
  edit peer1
    set subject E=peer1@mail.example.com
    set ca CA_Cert_1
    set two-factor enable
    set passwd fdktguefheygfe
  end
```

For more information on certificates, see [“Certificates overview” on page 1264](#).

Email

Two-factor email authentication sends a randomly generated six digit numeric code to the specified email address. Enter that code when prompted at logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted.

A benefit is that you do not require mobile service to authenticate. However, a potential issue is if your email server does not deliver the email before the 60 second life of the token expires.

The code will be generated and emailed at the time of logon, so you must have email access at that time to be able to receive the code.

To enable email two-factor authentication - web-based manager

- 1 Go to the email server under *Log&Report->Log Config->Alert e-mail*.
- 2 Enter the SMTP Server and Email from address.
- 3 If applicable, enable Authentication on the SMTP server and enter the SMTP username and password to use.
- 4 Select *Apply*.
- 5 To modify an administrator account, go to *System > Admin > Administrators*. To modify a user account go to **User > User**.
- 6 Select an existing account or select Create New.
- 7 Select *Enable Two-factor Authentication*.
- 8 Select *Email to*.
- 9 Enter the email address.
- 10 Select *OK*.

SMS

SMS two-factor authentication sends the token code in an SMS text message to the mobile device indicated when this user attempts to logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted. Enter this code when prompted at logon to be authenticated.

SMS two-factor authentication has the benefit that you do not require email service before logging on. A potential issue is if the mobile service provider does not send the SMS text message before the 60 second life of the token expires.

Before configuring SMS, you must configure the email server for sending email from the FortiGate unit and one or more SMS providers in the CLI.

To configure the SMTP email address for your FortiGate unit - web-based manager

- 1 Go to the email server under *Log&Report->Log Config->Alert e-mail*.
- 2 Enter the SMTP Server and Email from address.
- 3 If applicable, enable Authentication on the SMTP server and enter the SMTP username and password to use.
- 4 Select *Apply*.

To configure an SMS provider - CLI

```
config user sms-provider
edit <provider_name>
set mail-server <server_email>
next
end
```

To configure SMS two-factor authentication - web-based manager

- 1 To modify an:
 - administrator account, go to *System > Admin > Administrators*, or
 - user account go to **User > User**.
 - 2 Select an existing account or select *Create New*.
 - 3 Select *Enable Two-factor Authentication*.
 - 4 Select SMS.
 - 5 Choose the SMS provider from the drop down list.
 - 6 Enter the phone number of the mobile device that will receive the SMS text messages.
- If you have problems receiving the token codes via SMS messaging, contact your mobile provider to ensure you are using the correct phone number format to receive text messages and that your current mobile plan allows text messages.

FortiToken

FortiToken is a disconnected one-time password (OTP) generator. It is a small physical device with a button that when pressed displays a six digit authentication code. This code is entered with a user's username and password as two-factor authentication. The code displayed changes every 60 seconds, and when not in use the LCD screen is blanked to extend the battery life.

FortiTokens have a small hole in one end. This is intended for a lanyard to be inserted so the device can be worn around the neck, or easily stored with other electronic devices. Do not put the FortiToken on a key ring as the metal ring and other metal objects can damage it. The FortiToken is an electronic device like a cell phone and must be treated with similar care.

Any time information about the FortiToken is transmitted, it is encrypted. When the FortiGate unit receives the code that matches the serial number for a particular FortiToken, it is delivered and stored encrypted. This is in keeping with the Fortinet's commitment to keeping your network highly secured.

FortiTokens can be added to user accounts that are local, IPsec VPN, SSL VPN, and even Administrators. See [“Associating FortiTokens with accounts” on page 1233](#).

A FortiToken can only be associated with one account at a time on a FortiGate unit. However, multiple FortiGate units can have the same FortiToken registered. This is useful for employees who travel between offices.

If a user loses their FortiToken, it can be locked out using the FortiGate so it will not be used to falsely access the network. Later if found, that FortiToken can be unlocked on the FortiGate to allow access once again. See [“FortiToken maintenance” on page 1234](#).

There are three tasks to complete before FortiTokens can be used to authenticate accounts:

- 1 [Adding FortiTokens to the FortiGate](#)
- 2 [Activating a FortiToken on the FortiGate](#)
- 3 [Associating FortiTokens with accounts](#)

The FortiToken authentication process

The steps during FortiToken two-factor authentication are as follows.

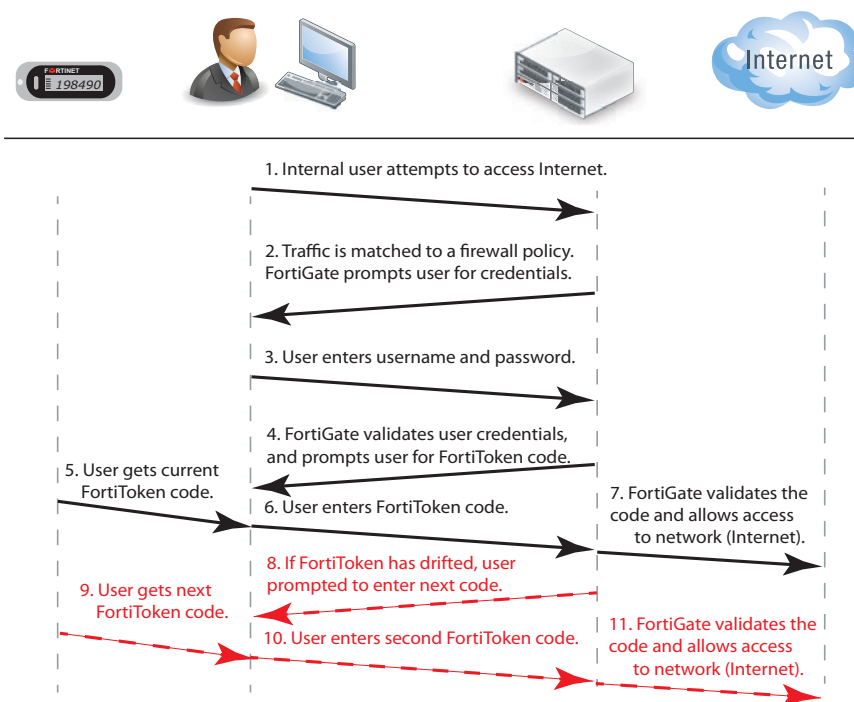
- 1 User attempts to access a network resource.

- 2 FortiGate unit matches the traffic to an authentication security policy, and FortiGate unit prompts the user for username and password.
- 3 User enters their username and password.
- 4 FortiGate unit verifies their information, and if valid prompts the user for the FortiToken code.
- 5 User gets the current code from their FortiToken device.
- 6 User enters current code at the prompt.
- 7 FortiGate unit verifies the FortiToken code, and if valid allows access to the network resources such as the Internet.

The following steps are only if the time on the FortiToken has drifted from the time on the FortiGate unit and needs to be synchronized.

- 8 If time on FortiToken has drifted, FortiGate unit will prompt user to enter a second code to confirm.
- 9 User gets the next code from their FortiToken device
- 10 User enters the second code at the prompt.
- 11 FortiGate unit uses both codes to update its clock to match the FortiToken and then proceeds as in step 7.

Figure 105: FortiToken authentication process



When configured the FortiGate unit accepts the username and password, authenticates them either locally or remotely, and prompts the user for the FortiToken code. The FortiGate then authenticates the FortiToken code. When FortiToken authentication is enabled, the prompt field for entering the FortiToken code is automatically added to the authentication screens.

Even when an Administrator is logging in through a serial or Telnet connection and their account is linked to a FortiToken, that Administrator will be prompted for the token's code at each login.



If you have attempted to add invalid FortiToken serial numbers, there will be no error message. The serial numbers will simply not be added to the list.

Adding FortiTokens to the FortiGate

Before one or more FortiTokens can be used to authenticate logons, they must be added to the FortiGate. The import feature is used to enter many FortiToken serial numbers at one time.

One FortiToken can be added to multiple FortiGate units. This is useful for maintaining two-factor authentication for employees over multiple office locations, such as for employees who travel frequently between offices.

To manually add a FortiToken to the FortiGate - web-based manager

- 1 Go to *User > FortiToken > FortiToken*.
- 2 Select *Create New*.
- 3 Enter one or more serial numbers for the FortiToken or FortiTokens you have.
- 4 Select OK.

To import multiple FortiTokens to the FortiGate - web-based manager

- 1 Go to *User > FortiToken > FortiToken*.
- 2 Select *Import*.
- 3 Browse to the local file location on your local computer.
The file must be a text file with one FortiToken serial number per line.
- 4 Select OK.

To add two FortiTokens to the FortiGate - CLI

```
config user fortitoken
  edit <serial_number>
  next
  edit <serial_number2>
  next
end
```

Activating a FortiToken on the FortiGate

Once one or more FortiTokens have been added to the FortiGate unit, they must be activated before being available to be associated with accounts. The process of activation involves the FortiGate querying FortiGuard servers about the validity of each FortiToken. The serial number and information is encrypted before it is sent for added security.



A FortiGate unit requires a connection to FortiGuard servers to activate a FortiToken.

To activate a FortiToken on the FortiGate unit - web-based manager

- 1 Go to *User > FortiToken*.
 - 2 Select one or more FortiTokens with a status of New.
 - 3 Select *Activate*.
 - 4 Refresh web browser. The status of selected FortiTokens will change to Activated.
- The selected FortiTokens are now available for use with user and admin accounts.

To activate a FortiToken on the FortiGate unit - CLI

```
config user fortitoken
  edit <token_serial_num>
    set status activate
  next
end
```

Associating FortiTokens with accounts

The final step before using the FortiTokens to authenticate logons is associating a FortiToken with an account. The accounts can be local user or administrator accounts.

To add a FortiToken to a local user account - web-based manager

- 1 Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Activated.
- 2 Go to *User > User*, and select *Create New*.
- 3 Enter the username and password for this user account.
- 4 Select *Enable Two-factor Authentication*.
- 5 Select FortiToken, and select the serial number from the list that matches that user's FortiToken.
- 6 Select OK.

To add a FortiToken to a local user account - CLI

```
config user local
  edit <username>
    set type password
    set passwd "myPassword"
    set two-factor fortitoken
    set fortitoken <serial_number>
    set status enable
  next
end
```

To add a FortiToken to an administrator account - web-based manager

- 1 Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Activated.
- 2 Go to *System > Admin > Administrators*, and select an admin account.
This account is assumed to be configured except for two-factor authentication.
- 3 Select *Enable Two-factor Authentication*.
- 4 Select FortiToken, and select the serial number from the list that matches that user's FortiToken.

5 Select OK.

To add a FortiToken to a local user account - CLI

```
config user local
  edit <username>
    set type password
    set passwd "myPassword"
    set two-factor fortitoken
    set fortitoken <serial_number>
    set status enable
  next
end
```

The `fortitoken` keyword will not be visible until `fortitoken` is selected for the two-factor keyword.



Before a new FortiToken can be used, it may need to be synchronized due to clock drift.

FortiToken maintenance

Once FortiTokens are entered into the FortiGate unit, there are only two tasks to maintain them — changing the status,

To change the status of a FortiToken between Activated and Locked - CLI

```
config user fortitoken
  edit <token_serial_num>
    set status lock
  next
end
```

Any user attempting to login using this FortiToken will not be able to authenticate.

To list the drift on all FortiTokens configured on this FortiGate unit - CLI

```
# diag fortitoken drift
FORTITOKEN      DRIFT
```

This command lists the serial number and drift for each FortiToken configured on this FortiGate unit. This command is useful to check if it is necessary to synchronize the FortiGate and any particular FortiTokens.

Monitoring users

To monitor user activity in the web-based manager, go to *Users > Monitor > Firewall*. The list of users who are logged on is displayed with some information about them such as their user group, security policy ID, how long they have been logged on, their IP address, traffic volume, and their authentication method as one of FSSO, NTLM, or firewall (FW-auth).

From this screen you can de-authenticate all users who are logged on. The de-authenticate button is at the top left of this screen.

To see information about banned users go to *User > Monitor > Banned Users*. Displayed information about users who have been banned includes what application triggered the ban (Application Protocol), the reason for the ban (Cause or rule), Created, and when the ban expires.

Filtering the list of users

When there are many users logged on, it can be difficult to locate a specific user or multiple users to analyze. Applying filters to the list allows you to organize the user list to meet your needs, or only display some the users that meet your current requirements.

Select *Column Settings* at the bottom of the screen to adjust columns that are displayed for users, including what order they are displayed in. This can be very helpful in locating information you are looking for.

The *username* column includes a green arrow to the right of the title. Select this arrow to sort the list of users by ordering them in ascending (down arrow) or descending order. This is the only column that allows this.

Each column heading has a grey filter icon. Click on the filter icon to configure a filter for the data displayed in that column. Each column has similar options including a field to enter the filtering information, a checkbox to select the negative of the text in the field, and the options to add more fields, apply the filter, clear all filters, or cancel without saving. To enter multiple terms in the field, separate each of them with a comma. To filter entries that contain a specific prefix, use an * (asterisk).



When removing existing filters, you must select *Apply* for the removal to take place.

For example, to create a filter to display only users with an IP address of 10.11.101.x who authenticated using one of security policies five through eight, and who belong to the user group *Accounting*.

To configure a user monitor filter - web-based manager

- 1 Go to *User > Monitor > Firewall*.
- 2 Select filter icon for IP address.
- 3 Enter 10.11.101.0.
- 4 Select *Add new filter*.
- 5 From the list of fields select *Policy ID* and enter 5-8.
- 6 Select *Add new filter*.
- 7 From the list of fields select *User Group* and enter *Accounting*.
- 8 Select *Apply*.

User groups

A user group is a list of user identities. An identity can be:

- a local user account (username/password) stored on the FortiGate unit
- a local user account with the password stored on a RADIUS, LDAP, or TACACS+ server
- a PKI user account with digital client authentication certificate stored on the FortiGate unit

- a RADIUS, LDAP, or TACACS+ server, optionally specifying particular user groups on that server
- a user group defined on an FSSO server.

Identity-based policies and some types of VPN configurations allow access to specified user groups only. This restricted access enforces Role Based Access Control (RBAC) to your organization's network and its resources. Users must be in a group and that group must be part of the security policy.



You cannot change the type of a group unless the group is empty.

In most cases, the FortiGate unit authenticates users by requesting their username and password. The FortiGate unit checks local user accounts first. If a match is not found, the FortiGate unit checks the RADIUS, LDAP, or TACACS+ servers that belong to the user group. Authentication succeeds when a matching username and password are found. If the user belongs to multiple groups on a server, those groups will be matched as well.



FortiOS does not allow username overlaps between RADIUS, LDAP, or TACACS+ servers.

There are two types of FortiGate user groups: Firewall user groups, and FSSO user groups.

Firewall user groups

Firewall user groups are used locally as part of authentication and can contain any type of user identity except an FSSO group. When a user attempts to access resources controlled by an Identity-Based Policy (IBP), the FortiGate unit requires authentication from that user. If the user authenticates successfully and is a member of one of the permitted groups, the session is allowed to proceed.

This section includes:

- [SSL VPN access](#)
- [IPsec VPN access](#)
- [Configuring a firewall user group](#)
- [User group timeouts](#)
- [Viewing, editing and deleting user groups](#)

SSL VPN access

In any firewall user group, you can enable SSL VPN access and select the web-portal that the users can access. When the user connects to the FortiGate unit via HTTPS on the SSL VPN port (default 10443), the FortiGate unit requests a username and password.

SSL VPN access also requires an SSL VPN security policy (*Action* is *SSL VPN*) with an identity-based rule enabling access for the user group. For more information, see the [FortiOS Handbook SSL VPN chapter](#).

IPsec VPN access

A firewall user group can provide access for dialup users of an IPsec VPN. In this case, the IPsec VPN phase 1 configuration uses the *Accept peer ID in dialup group peer* option. The user's VPN client is configured with the username as peer ID and the password as pre-shared key. The user can connect successfully to the IPsec VPN only if the username is a member of the allowed user group and the password matches the one stored on the FortiGate unit.



A user group cannot be used as a dialup group if any member of the group is authenticated using an external authentication server.

For more information, see the [FortiOS Handbook IPsec VPN chapter](#).

Configuring a firewall user group

A user group can contain:

- local users, whether authenticated by the FortiGate unit or an authentication server
- PKI users
- authentication servers, optionally specifying particular user groups on the server

To create a Firewall user group - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter a name for the user group.
- 3 In *Type*, select *Firewall*.
- 4 From the *Available Users* list, select users and then select the right arrow button to move the names to the *Members* list.

If you select an authentication server as a group member, by default all user accounts on the authentication server are members of this FortiGate user group. Follow steps 5 through 8 if you want to include only specific user groups from the authentication server. Otherwise, select *OK*.
- 5 Select *Add*.
- 6 To add a remote authentication server, select *Add* and select the authentication server from the drop down *Remote Server* list.

The option to add remote servers is available only if at least one remote server has been configured.
- 7 In the *Group Name* field, either select *Any* to match all possible groups, or select *Specify* and enter the group name in the appropriate format for the type of server.

For example, an LDAP server requires LDAP format, such as:
cn=users,dn=office,dn=example,dn=com
- 8 Repeat steps 5 through 7 to add all the authentication server user groups that are required.
- 9 Select *OK*.

To create a firewall user group - CLI example

In this example, the members of `accounting_group` are `User1` and all of the members of `rad_accounting_group` on myRADIUS external RADIUS server.

```
config user group
  edit accounting_group
    set group-type firewall
    set member User1 myRADIUS
  config match
    edit 0
      set server-name myRADIUS
      set group-name rad_accounting_group
    end
  end
end
```



Matching user group names from an external authentication server might not work if the list of group memberships for the user is longer than 8000 bytes. Group names beyond this limit are ignored.

`server_name` is the name of the RADIUS, LDAP, or TACACS+ server, but it must be a member of this group first and must also be a configured remote server on the FortiGate unit.

`group_name` is the name of the group on the RADIUS, LDAP, or TACACS+ server such as "engineering" or "cn=users,dc=test,dc=com".

Before using group matching with TACACS+, you must first enable authentication. For example if you have a configured TACACS+ server called myTACS, use the following CLI commands.

```
config user tacacs+
  edit myTACS
    set authorization enable
  next
end
```

For more information about user group CLI commands, see the [Fortinet CLI Guide](#).

Multiple group enforcement support

Previously, when a user belonged to multiple user groups, this user could only access the group services that were within one group. With multiple group enforcement, a user can access the services within the groups that the user is part of.

For example, `userA` belongs to `user_group1`, `user_group2`, `user_group3`, and `user_group4`; previously `userA` could only access services within one of those four groups, typically the group that matches the first security policy. This can be annoying if HTTP access is in `user_group1`, FTP access is in `user_group2`, and email access is in `user_group3`. Now `userA` can access services within `user_group1`, `user_group2`, `user_group3`, and `user_group4`.

This feature is available only in the CLI and is enabled by default. It applies to RADIUS, LDAP, and TACACS+ servers. The new command for this feature is `auth-multi-group` found in `config user settings` and checks all groups a user belongs to for authentication.

User group timeouts

User groups can have timeout values per group in addition to FortiGate-wide timeouts. There are essentially three different types of timeouts that are configurable for user authentication on the FortiGate unit — idle timeout, hard timeout, and session timeout. These are in addition to any external timeouts such as those associated with RADIUS servers.

If VDOMs are enabled, the global level user setting `authtimeout` is the default all VDOMs inherit. If VDOMs are not enabled, user settings `authtimeout` is the default. The default timeout value is used when the `authtimeout` keyword for a user group is set to zero.

Each type of timeout will be demonstrated using the existing user group `example_group`. Timeout units are minutes. A value of zero indicates the global timeout is used.

Membership in multiple groups

When a user belongs to multiple groups in RADIUS groups, the group auth-timeout values are ignored. Instead the global timeout value is used. The default value is 5 minutes, but it can be set from 1 to 480 minutes.

```
config user setting
    set auth-timeout-type idle-timeout
    set auth-timeout 300
end
```

Idle timeout

The default type of timeout is idle timeout. When a user initiates a session, it starts a timer. As long as data is transferred in this session, the timer continually resets. If data flow stops, the timer is allowed to advance until it reaches its limit. At that time the user has been idle for too long, and the user is forced to re-authenticate before traffic is allowed to continue in that session.

To configure user group authentication idle timeout - CLI

```
config user settings
    set auth-timeout-type idle-timeout
end

config user group
    edit example_group
        set auth-timeout 480
    next
end
```

Hard timeout

Where the idle timeout is reset with traffic, the hard timeout is absolute. From the time the first session a user establishes starts, the hard timeout counter starts. When the timeout is reached, all the sessions for that user must be re-authenticated. This timeout is not affected by any event.

To configure user group authentication hard timeout - CLI

```
config user settings
    set auth-timeout-type hard-timeout
end
config user group
    edit example_group
        set auth-timeout 480
    next
```

```
end
```

Session timeout

The session timeout works much like the hard timeout in that its an absolute timer that can not be affected by events. However, when the timeout is reached existing sessions may continue but new sessions are not allowed until re-authentication takes place. The timeout can be set from 1 to 480 minutes. Setting the timeout value to zero removes the timeout value allowing the user to remain logged on without limit.

To configure a user group authentication new session hard timeout - CLI

```
config user setting
    set auth-timeout-type new-session
end

config user group
    edit example_group
        set authtimeout 30
    next
end
```

FSSO user groups

FSSO user groups are part of FSSO authentication and contain only Windows or Novell network users. No other user types are permitted as members. Information about the Windows or Novell user groups and the logon activities of their members is provided by the Fortinet Single Sign On (FSSO) which is installed on the network domain controllers. You can specify FSSO user groups in identity-based security policies in the same way as you specify firewall user groups. FSSO user groups cannot have SSL VPN or dialup IPsec VPN access.

For information about configuring FSSO user groups, see [“Creating Fortinet Single Sign-On \(FSSO\) user groups” on page 1312](#). For complete information about installing and configuring FSSO, see [“FSSO integration with Windows AD or Novell” on page 1283](#).

Configuring Peer user groups

Peer user groups can only be configured using the CLI. Peers are digital certificate holders defined using the `config user peer` command. The peer groups you define here are used in dialup IPsec VPN configurations that accept RSA certificate authentication from members of a peer certificate group. For more information, see [“Authenticating IPsec VPN users with security certificates” on page 1276](#).

To create a peer group - CLI example

```
config user peergrp
    edit vpn_peergrp1
        set member pki_user1 pki_user2 pki_user3
    end
```

Viewing, editing and deleting user groups

To view the list of FortiGate user groups, go to *User > User Group*.

Editing a user group

When editing a user group in the CLI you must set the type of group this will be — either a firewall group, or a Fortinet Single Sign-On Service group. Once the type of group is set, and members are added you cannot change the group type without removing the members.

In the web-based manager, if you change the type of the group any members will be removed automatically.

To edit a user group - web-based manager

- 1 Go to *User > User Group*.
- 2 Select the check box for the user group that you want to edit.
- 3 Select the *Edit* button.
- 4 Modify the user group as needed.
- 5 Select *OK*.

To edit a user group - CLI example

This example adds user3 to Group1. Note that you must re-specify the full list of users:

```
config user group
  edit Group1
    set group-type firewall
    set member user2 user4 user3
  end
```

Deleting a user group

Before you delete a user group, you must ensure there are no objects referring to, it such as security policies. If there are, you must remove those references before you are able to delete the user group.

To remove a user group - web-based manager

- 1 Go to *User > User Group*.
- 2 Select the check box for the user group that you want to remove.
- 3 Select the *Delete* button.
- 4 Select *OK*.

To remove a user group - CLI example

```
config user group
  delete Group2
end
```




Configuring authenticated access

When you have configured authentication servers, users, and user groups, you are ready to configure security policies and certain types of VPNs to require user authentication.

This section describes:

- [Authentication timeout](#)
- [Password policy](#)
- [Authentication protocols](#)
- [Authentication in security policies](#)
- [VPN authentication](#)

Authentication timeout

An important feature of the security provided by authentication is that it is temporary—a user must re-authenticate after logging out. Also if a user is logged on and authenticated for an extended period of time, it is a good policy to have them re-authenticate at set periods. This ensures a user's session is cannot be spoofed and used maliciously for extended periods of time — re-authentication will cut any spoof attempts short. Shorter timeout values are more secure.

Security authentication timeout

You set the security user authentication timeout to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 480 minutes (8 hours).

To set the security authentication timeout - web-based manager

- 1 Go to *User > User > Authentication*.
- 2 Enter the *Authentication Timeout* value in minutes.
The default authentication timeout is 5 minutes.
- 3 Select *Apply*.

SSL VPN authentication timeout

You set the SSL VPN user authentication timeout (*Idle Timeout*) to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 28 800 seconds. The default timeout is 300 seconds.

To set the SSL VPN authentication timeout - web-based manager

- 1 Go to *VPN > SSL > Config*.
- 2 Enter the *Idle Timeout* value (seconds).
- 3 Select *Apply*.

Password policy

Password authentication is effective only if the password is sufficiently strong and is changed periodically. By default, the FortiGate unit requires only that passwords be at least eight characters in length. You can set a password policy to enforce higher standards for both length and complexity of passwords. Password policies can apply to administrator passwords or IPsec VPN preshared keys.

To set a password policy in the web-based manager, go to *System > Admin > Settings*. In the CLI, use the `config system password-policy` command.

The default minimum password length on the FortiGate unit is eight characters, but up to 32 characters is permitted. Fortinet suggests a minimum length of 14 characters.

Users usually create passwords composed of alphabetic characters and perhaps some numbers. Password policy can require the inclusion of uppercase letters, lowercase letters, numerals or punctuation characters.

Configuring password minimum requirement policy (best practices)

Best practices dictate that passwords include:

- one or more uppercase characters
- one or more lower care characters
- one or more of the numerals
- one or more non alphanumeric characters, such as punctuation marks.

The minimum number of each of these types of characters can be set in both the web-based manager and the CLI.

The following procedures show how to force administrator passwords to contain at least two uppercase, four lower care, two digits, and one non-alphanumeric characters. Leave the minimum length at the default of eight characters.

To change administrator password minimum requirements - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 Select *Enable Password Policy*.
- 3 Select *Must Contain*.
- 4 Enter the following information:

uppercase Letters	2
lower care Letters	4
Numerical Digits	2
Non-alphanumeric Letters	1

- 5 Under *Apply Password Policy to*, select *Admin Password*.
- 6 Select *Apply*.

To change administrator password minimum requirements - CLI

```
config system password-policy
  set status enable
  set apply-to admin-password
  set min-upper-case-letter 2
  set min-lower-case-letter 4
```

```
set min-number 2
set min-non-alphanumeric 1
set change-4-characters enable
next
end
```

The `change-4-characters` option forces new passwords to change a minimum of four characters in the old password. Changing fewer characters results in the new password being rejected. This option is only available in the CLI.

Password best practices

In addition to length and complexity, there are security factors that cannot be enforced in a policy. Guidelines issued to users will encourage proper password habits.

Best practices dictate that password expiration also be enabled. This forces passwords to be changed on a regular basis. You can set the interval in days. The more sensitive the information this account has access to, the shorter the password expiration interval should be. For example 180 days for guest accounts, 90 days for users, and 60 days for administrators.

Avoid:

- real words found in any language dictionary
- numeric sequences, such as “12345”
- sequences of adjacent keyboard characters, such as “qwerty”
- adding numbers on the end of a word, such as “hello39”
- adding characters to the end of the old password, such as “hello39” to “hello3900”
- repeated characters
- personal information, such as your name, birthday, or telephone number.

Maximum logon attempts and blackout period

When you logon and fail to enter the correct password you could be a valid user, or a hacker attempting to gain access. For this reason, best practices dictate to limit the number of failed attempts to logon before a blackout period where you cannot logon.

To set a maximum of five failed authentication attempts before the blackout, using the following CLI command:

```
config user setting
set auth-invalid-max 5
next
end
```

To set the length of the blackout period to five minutes, or 300 seconds, once the maximum number of failed logon attempts has been reached, use the following CLI command:

```
config user setting
set auth-blackout-time 300
next
end
```

Authentication protocols

When user authentication is enabled on a security policy, the authentication challenge is normally issued for any of the four protocols, HTTP, HTTPS, FTP, and Telnet, which are dependent on the connection protocol. By making selections in the Protocol Support list, the user controls which protocols support the authentication challenge. The user must connect with a supported protocol first, so that they can subsequently connect with other protocols.

For example, if you have selected HTTP, FTP, or Telnet, a username and password-based authentication occurs. The FortiGate unit then prompts network users to input their security username and password. If you have selected HTTPS, certificate-based authentication (HTTPS, or HTTP redirected to HTTPS only) occurs.



FTP and Telnet authentication replacement messages cannot be customized. For HTTP and HTTPS replacement messages see [“Authentication replacement messages” on page 1248](#).

For certificate-based authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users. If you do not install certificates on the network user's web browser, the network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate. The network user's web browser may deem the default certificate as invalid.

When you use certificate authentication, if you do not specify any certificate when you create the security policy, the global settings are used. If you specify a certificate, the per-policy setting will overwrite the global setting. For more information about the use of certification authentication see [“Certificate-based authentication” on page 1263](#).

To set the authentication protocols

- 1 Go to *User > User > Authentication*.
- 2 In *Protocol Support*, select the required authentication protocols.
- 3 If using HTTPS protocol support, in *Certificate*, select a Local certificate from the drop-down list.
- 4 Select *Apply*.

Authentication in security policies

Security policies control traffic between FortiGate interfaces, both physical interfaces and VLAN subinterfaces. Without authentication, a security policy enables access from one network to another for all users on the source network. Authentication enables you to allow access only for users who are members of selected user groups. To include authentication in a security policy, you must create an identity-based policy.

The style of the authentication method varies by the authentication protocol. If you have selected HTTP, FTP or Telnet, a username and password-based authentication occurs. The FortiGate unit prompts network users to input their security username and password. If you have selected HTTPS, certificate-based authentication (HTTPS or HTTP redirected to HTTPS only) occurs. You must install customized certificates on the FortiGate unit and on the browsers of network users, which the FortiGate unit matches.



You can configure user authentication for security policies only when *Action* is set to *Accept*. If the policy is set to *Deny*, *IPsec*, or *SSL VPN* the options will be different.

This section includes:

- [Enabling authentication protocols](#)
- [Authentication replacement messages](#)
- [Access to the Internet](#)
- [Configuring authentication security policies](#)
- [Identity-based policy](#)
- [FSSO authentication](#)
- [NTLM authentication](#)
- [Certificate authentication](#)
- [Dynamic profile](#)
- [Restricting number of concurrent user logons](#)

Enabling authentication protocols

Users can authenticate using FTP, HTTP, HTTPS, and Telnet. However, these protocols must be enabled first.

Another authentication option is to redirect any attempts to authenticate using HTTP to a more secure channel that uses HTTPS. This forces users to a more secure connection before entering their user credentials.

To enable support for authentication protocols - web-based manager

- 1 Go to *User > User > Authentication*.
- 2 Select one or more of HTTP, HTTPS, FTP, Telnet, or Redirect HTTP Challenge to a Secure Channel (HTTPS). Only selected protocols will be available for use in authentication.
- 3 Select the *Certificate* to use, for example `Fortinet_Factory`.
- 4 Select *Apply*.

To enable support for authentication protocols - CLI

```
config user setting
    set auth-type ftp http https telnet
    set auth-cert Fortinet_Factory
end
```

Authentication replacement messages

A replacement message is the body of a webpage containing a message about a blocked website message, a file too large message, a disclaimer, or even a login page for authenticating. The user is presented with this message instead of the blocked content.

Authentication replacement messages are the prompts a user sees during the security authentication process such as login page, disclaimer page, and login success or failure pages. These are different from most replacement messages because they are interactive requiring a user to enter information, instead of simply informing the user of some event as other replacement messages do.

Replacement messages have a system-wide default configuration, a per-VDOM configuration, and disclaimers can be customized for multiple security policies within a VDOM.

These replacement messages are used for authentication using HTTP and HTTPS. Authentication replacement messages are HTML messages. You cannot customize the security authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

More information about replacement messages can be found in the `config system replacemsg` section of the [FortiOS CLI Reference](#).

Table 79: List of authentication replacement messages

Replacement message name (CLI name)	Description
Login challenge page (auth-challenge-page)	<p>This HTML page is displayed if security users are required to answer a question to complete authentication. The page displays the question and includes a field in which to type the answer. This feature is supported by RADIUS and uses the generic RADIUS challenge-access auth response. Usually, challenge-access responses contain a Reply-Message attribute that contains a message for the user (for example, "Please enter new PIN"). This message is displayed on the login challenge page. The user enters a response that is sent back to the RADIUS server to be verified.</p> <p>The Login challenge page is most often used with RSA RADIUS server for RSA SecurID authentication. The login challenge appears when the server needs the user to enter a new PIN. You can customize the replacement message to ask the user for a SecurID PIN.</p> <p>This page uses the <code>%%QUESTION%%</code> tag.</p>

Table 79: List of authentication replacement messages

Replacement message name (CLI name)	Description
Disclaimer page (auth-disclaimer-page-1) (auth-disclaimer-page-2) (auth-disclaimer-page-3)	<p>Prompts user to accept the displayed disclaimer when leaving protected network.</p> <p>The web-based manager refers to this as User Authentication Disclaimer, and it is enabled with a security policy that also includes at least one identity-based policy. When a security user attempts to browse a network through the FortiGate unit using HTTP or HTTPS this disclaimer page is displayed.</p> <p>The extra pages seamlessly extend the size of the page from 8 192 characters to 16 384 and 24 576 characters respectively. When configuring the disclaimer page in the web-based manager this is shown by its size being 24 576 characters.</p> <p>See “Disclaimer” on page 1251.</p>
Email token page (auth-email-token-page)	<p>The page prompting a user to enter their email token. See “Email” on page 1228.</p>
FortiToken page (auth-fortitoken-page)	<p>The page prompting a user to enter their FortiToken code. See “FortiToken” on page 1230.</p>
Keepalive page (auth-keepalive-page)	<p>The HTML page displayed with security authentication keepalive is enabled using the following CLI command:</p> <pre>config system global set auth-keepalive enable end</pre> <p>Authentication keepalive keeps authenticated firewall sessions from ending when the authentication timeout ends. In the web-based manager, go to <i>User > Options</i> to set the <i>Authentication Timeout</i>.</p> <p>This page includes %%TIMEOUT%%.</p>
Login failed page (auth-login-failed-page)	<p>The Disclaimer page replacement message does not redirect the user to a redirect URL or the security policy does not include a redirect URL. When a user selects the button on the disclaimer page to decline access through the FortiGate unit, the Declined disclaimer page is displayed.</p>
Login page (auth-login-page)	<p>The authentication HTML page displayed when users who are required to authenticate connect through the FortiGate unit using HTTP or HTTPS.</p> <p>Prompts the user for their username and password to login.</p> <p>This page includes %%USERNAMEID%% and %%PASSWORDID%% tags.</p>
Declined disclaimer page (auth-reject-page)	<p>The page displayed if a user declines the disclaimer page. See “Disclaimer” on page 1251.</p>

Table 79: List of authentication replacement messages

Replacement message name (CLI name)	Description
SMS Token page (auth-sms-token-page)	The page prompting a user to enter their SMS token. See “SMS” on page 1229 .
Success message (auth-success-msg)	The page displayed when a user successfully authenticates. Prompts user to attempt their connection again (as the first was interrupted for authentication).

Access to the Internet

A policy for accessing the Internet is similar to a policy for accessing a specific network, but the destination address is set to *all*. The destination interface is the one that connects to the Internet Service Provider (ISP). For general purpose Internet access, the Service is set to ANY.

Access to HTTP, HTTPS, FTP and Telnet sites may require access to a domain name service. DNS requests do not trigger authentication. You must configure a policy to permit unauthenticated access to the appropriate DNS server, and this policy must **precede** the policy for Internet access. Failure to do this will result in the lack of a DNS connection and a corresponding lack of access to the Internet.

Configuring authentication security policies

To include authentication in a security policy, you must create an identity-based policy. An identity-based policy can authenticate by certificate, FSSO, and NTLM. The two exceptions to this are dynamic profiles and FSSO Agents. See [“Configuring dynamic profile” on page 1332](#), and [“Introduction to FSSO” on page 1283](#).

Before creating an identity-based security policy, you need to configure one or more users and firewall user groups. For more information, see [“Users and user groups” on page 1223](#).

Creating the security policy is the same as a regular security policy except you must select the action specific to your authentication method:

Table 80: Authentication methods allowed for each policy Action

Action	Authentication method	Where authentication is used
ACCEPT	FSSO Agent or identity-based policy — FSSO	See “FSSO integration with Windows AD or Novell” on page 1283 .
	identity-based policy — NTLM	See “NTLM authentication” on page 1255 .
	identity-based policy — Certificates	See “Configuring certificate-based authentication” on page 1275 .
	Dynamic Profile	See “Configuring dynamic profile-based security policies” on page 1339 .
IPSEC	IPsec Phase 1 and 2	See “Configuring authentication of remote IPsec VPN users” on page 1258 .
SSL-VPN	SSL certificates	See “Configuring authentication of SSL VPN users” on page 1258 .

Table 80: Authentication methods allowed for each policy Action

Action	Authentication method	Where authentication is used
DENY	none	none

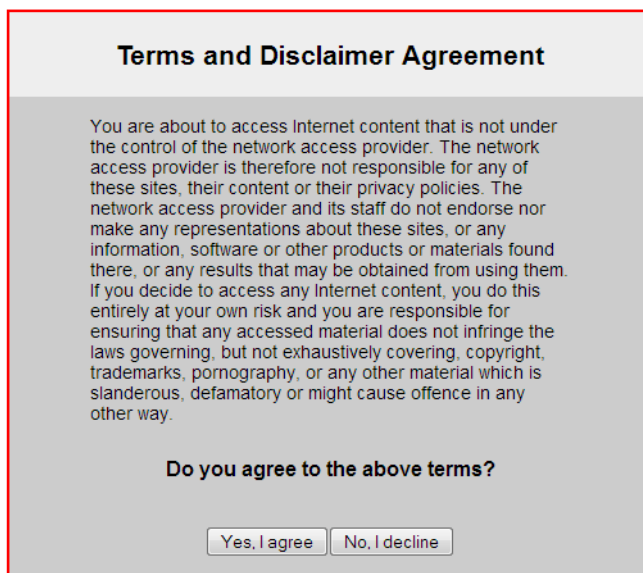
Disclaimer

When configuring any authentication security policy, there is an option to enable a disclaimer. The disclaimer is a replacement message that when enabled, web traffic matching this policy will be presented with the disclaimer that the user must choose to agree or decline.

The default disclaimer contains a warning that any content the user is about to access is the responsibility of the user and not the company or owner of the network. It is presented in [Figure 106](#). You can customize the text and the appearance to as required.

To change the disclaimer

- 1 Go to *System > Config > Replacement messages*.
- 2 Expand Authentication
- 3 Select Disclaimer page to edit.
- 4 Select the existing disclaimer text, and copy it to a separate file.
- 5 Make changes

Figure 106: Default disclaimer message

Terms and Disclaimer Agreement

You are about to access Internet content that is not under the control of the network access provider. The network access provider is therefore not responsible for any of these sites, their content or their privacy policies. The network access provider and its staff do not endorse nor make any representations about these sites, or any information, software or other products or materials found there, or any results that may be obtained from using them. If you decide to access any Internet content, you do this entirely at your own risk and you are responsible for ensuring that any accessed material does not infringe the laws governing, but not exhaustively covering, copyright, trademarks, pornography, or any other material which is slanderous, defamatory or might cause offence in any other way.

Do you agree to the above terms?

Customizing authentication replacement messages

Customizing disclaimers or other authentication replacement messages involves changing the text of the disclaimer message, and possibly the overall appearance of the message.

Disclaimers are useful in many situations. Often companies find it useful to brand the disclaimers with their specific company policy text, logo, and design. One example of this is at an Internet cafe where customers pay for usage and must accept terms of usage before accessing the internet. The cafe benefits from a customized disclaimer that alerts the customer to their online policies. The same is true for other authentication replacement messages such as the login page.

Changing the disclaimer at *System > Config > Replacement messages* is not the same as selecting to customize a disclaimer used in a policy. The *System > Config* location is the default message that all disclaimers inherit. The security policy location is a customized disclaimer that inherits the default format for the disclaimer message, but then can be customized for this policy.



If identity-based policy is enabled, the disclaimer option is not available. Instead Customize Authentication messages is available in the Enable identity-based Policy section. Selecting the Edit icon allows you to customize the listed authentication replacement messages which includes the disclaimer.

To customize the disclaimer for a security policy - web-based manager

- 1 Go to *Policy > Policy*. Either select an existing policy and edit or select *Create New*.
- 2 Enable *Disclaimer*, and select *Customize Disclaimer Message*.
- 3 Select the edit icon to edit the Disclaimer page, and change your text or layout as needed.

Enabling security logging

There are two types of logging that relate to authentication — event logging, and security logging.

When enabled, event logging records system events such as configuration changes, and authentication. To configure event logging, go to *Log&Report > Log Config > Log Setting* and in the *Event Log Settings* section, select the check box beside *Enable*. Select the events you want to log, such as authentication.

When enabled security logging will log UTM and security policy traffic

You must enable logging within a security policy, as well as the options that are applied to a security policy, such as UTM features. Event logs are enabled within the Event Log page,

For more information on logging, see the [FortiOS Log and Reporting chapter](#).

For more information on specific types of log messages, see the [FortiOS Log Message Reference](#).



You need to set the logging severity level to *Notification* when configuring a logging location to record traffic log messages .

To enable logging within an existing security policy - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Expand to reveal the policy list of a policy.
- 3 Select the security policy you want to enable logging on and then select *Edit*.
- 4 To log all general firewall traffic, select the check box beside *Log Allowed Traffic*.
- 5 On the security policy's page, select the check box beside *UTM*.

- 6 Select the protocol option list from the drop-down list beside *Protocol Options*.
By default, the *Protocol Options* check box is selected. You must choose a list from the drop-down list.
- 7 For each row under *UTM*, select the check box beside each of the profiles and/or sensors that you want applied to the policy; then select the profile or sensor from the drop-down list as well.
- 8 To apply other options, such as application lists, repeat step 7 and choose the option from the drop-down list.
- 9 Select OK.

Identity-based policy

An identity-based policy (IBP) performs user authentication in addition to the normal security policy duties. If the user does not authenticate, access to network resources is refused. This enforces Role Based Access Control (RBAC) to your organization's network and resources.

User authentication can occur through any of the following supported protocols, including: HTTP, HTTPS, FTP, and Telnet. The authentication style depends on which of these protocols is included in the selected security services group and which of those enabled protocols the network user applies to trigger the authentication challenge.

For username and password-based authentication (HTTP, FTP, and Telnet) the FortiGate unit prompts network users to enter their username, password, and token code if two-factor authentication is selected for that user account. See [“Two-factor authentication” on page 1228](#). For certificate-based authentication, including HTTPS or HTTP redirected to HTTPS only, see [“Certificate authentication” on page 1256](#).



Enable identity-based Policy and *Enable Dynamic Profile* are mutually exclusive. When one of the two options is selected, the other is hidden. When both are not selected, both are visible again.

Set the *ACCEPT Action*, and select *Enable identity-based Policy*. In the web-based manager you can confirm this by changing *Action* from the default of *ACCEPT* to *DENY*. Note that nearly all of the fields are not available until you switch back to *ACCEPT*.

Set these commands in the CLI to see the other identity-based commands that were hidden before. In the following procedure, this is policy number 7.

```
config firewall policy
edit 7
set action ACCEPT
set identity-based enable
next
end
```

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-policies for the user groups and services. This means unique security policies must be placed **before** an identity-based policy to be effective.

When the identity-based policy has been configured, the option to customize authentication messages is available. This allows you to change the text, style, layout, and graphics of the replacement messages associated with this firewall policy. When enabled, customizing these messages follows the same method as changing the disclaimer. See [“Disclaimer” on page 1251](#).

The types of authentication available in identity-based policies are

- [FSSO authentication](#)
- [NTLM authentication](#)
- [Certificate authentication](#)

Identity-based sub-policies

Once IBP is enabled in a policy, a table appears. Selecting *Add* allows you to configure authentication rules which are added to this table as sub-policies.

Just as with regular security policies, with these identity-based sub-policies traffic is matched from the top of the list of sub-policies down until the criteria is met. If there is no matching policy packets are dropped, even if they have been authenticated. Each sub-policy has its own UTM profile fields, traffic shaping, logging, and so on that take effect when the User Group, Service and Schedule are matched.

The order of these sub-policies is just as important as with regular security policies. For example if a user is a member of two groups, and each group has a separate sub-policy entry, the top one in the list will be matched first.

FSSO authentication

Identity-based security policies are an integral part of any FSSO configuration. These policies are how user information is acquired to sent to the FSSO Collector agent and the AD domain controllers. See [“FSSO authentication” on page 1254](#).

In the following procedure, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fsso_group`, and the FSSO RADIUS server is `fsso_rad_server`.

To configure an FSSO authentication security policy - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information.

Source Interface/Zone	port1
Source Address	company_network
Destination Interface/Zone	port2
Destination Address	all
Action	ACCEPT
Enable NAT	enabled

- 4 Select *Enable identity-based Policy*.
- 5 Select *Add* to add groups of users to this authentication policy.

- 6 Select the `fsso_group`, and the `FSSO_Guest_users` usergroups in the Available User Groups list and move them to the Selected User Groups list.
FSSO_Guest_users is a default user group enabled when FSSO is configured. It allows guest users on the network who do not have FSSO account to still authenticate and have access to network resources. See [“Enabling guest access through FSSO security policies” on page 1315](#).
- 7 Select HTTP, HTTPS, FTP, and Telnet for in the *Available Services* list, and move them to the Selected Services list.
- 8 Select always for the *Schedule*.
- 9 Enable *Log Allowed Traffic*.
- 10 Select UTM, and enable default AntiVirus, IPS, Web Filter, an Email filter.
- 11 Select OK.
 A new line of information will appear in the identity-based policy table. The table lists the ID, user group or groups, the service or services, schedule, UTM, and logging selected for the rule. Use this display to verify your information was entered correctly.
- 12 Select *Fortinet Single Sign-On (FSSO)*.
- 13 Optionally select *Customize Authentication messages* to change the default authentication messages to suit example.com’s company design and policies.
- 14 Select OK.
- 15 Ensure the FSSO authentication policy is at the top of the list so it will be attempted to be matched before any other policy.

An Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user’s identity and credentials, and then if everything is verified the user is allowed access to the Internet.

NTLM authentication

The NT LAN Manager (NTLM) protocol is used when the MS Windows Active Directory (AD) domain controller can not be contacted. NTLM uses web browsers to send and receive authentication information. See [“NTLM” on page 1167](#) and [“NTLM authentication with FSSO” on page 1287](#).

NTLM authentication is enabled when you configure FSSO and enable NTLM in the identity-based policy (IBP). There must be at least one FSSO Collector agent configured on the FortiGate. Any users and user groups associated with the security policy will use NTLM to authenticate without further configuration. However some extra configuration in the CLI may be required for certain cases including guest access, and defining NTLM enabled browsers.



If there are multiple domains, a trust relation must exist between them. This is automatic if they are in a forest. With the trust relation, only one FSSO DC agent needs to be installed. Without the trust relation, FSSO DC agents must be installed on each domain controller.

NTLM guest access - CLI

Guest profile access may be granted to users failing NTLM authentication, such as visitors who have no user credentials on the network. To allow guest users in NTLM, use the following CLI command:

```
config firewall policy
edit 8
    set action accept
    set identity-based enable
    set ntlm enable
    set ntlm-guest enable
next
end
```

NTLM enabled browsers - CLI

User agent strings for NTLM enabled browsers allow the inspection of initial HTTP-User-Agent values, so that non-supported browsers are able to go straight to guest access without needlessly prompting the user for credentials that will fail. `ntlm-guest` must be enabled to use this option.

```
config firewall policy
edit 9
    set action accept
    set identity-based enable
    set ntlm enable
    set ntlm-guest enable
    set ntlm-enabled-browsers <user_agent_string>
next
end
```

`<user_agent_string>` is the name of the browser that is NTLM enabled. Examples of these values include “MSIE”, “Mozilla” (which includes FireFox), and “Opera”.

Value strings can be up to 63 characters in length, and may not contain cross site scripting (XSS) vulnerability characters such as brackets. The FortiGate unit prevents use of these characters to prevent exploit of cross site scripting (XSS) vulnerabilities.

Certificate authentication

Certificates can be used as part of an identity-based policy. A customized certificate must be installed on the FortiGate unit and in the web browser, which the FortiGate unit will attempt to match.

All users being authenticated against the policy are required to have the proper certificate, which must be imported into the FortiGate unit. See [“Certificate-based authentication” on page 1263](#).

To require the user to accept a disclaimer to connect to the destination, select *Enable Disclaimer*. If the user is to be redirected after accepting the disclaimer, enter the URL in the *Redirect URL to* field. You can edit the User Authentication Disclaimer replacement message text in *System > Config > Replacement Messages*.

Certificate redirect authentication

Under *User > User > Authentication* select Redirect HTTP Challenge to a Secure Channel (HTTPS). This forces users to use secure connections to send their authentication information.

The following steps happen during a redirect:

- 1 User tries to access the Internet and the HTTP traffic hits the FortiGate security policy with authentication and HTTPS redirect enabled.
- 2 The FortiGate redirects the user with the HTTPS port and IP address of the interface connected to the user, such as internal.

- 3 User authenticates over the HTTPS connection as with normal authentication.
- 4 On successful authentication, the FortiGate provides access to the Internet as originally requested.

Dynamic profile

Only one security policy can be configured for dynamic profile in a VDOM. Also only one RADIUS server and one dynamic profile group can be configured per VDOM. When one dynamic profile security policy has been configured, the option is not visible when creating other policies. After deleting the dynamic profile security policy, the option is again visible when configuring other security policies.



Enable identity-based Policy and *Enable Dynamic Profile* are mutually exclusive. When one of the two options is selected, the other is hidden. When both are not selected, both are visible again.

By enabling the Dynamic Profile Users Only option, other non-dynamic profile users will not match this policy. This can be useful if you want to enforce all users to be part of the dynamic profile group—in which case you have a deny all profile after this one.

For more information, see [“Configuring dynamic profile-based security policies” on page 1339](#).

Restricting number of concurrent user logons

Some users on your network may often have multiple account sessions open at one time either to the same network resource or accessing to the admin interface on the FortiGate unit.

While there are valid reasons for having multiple concurrent sessions open, hackers also do this to speed up their malicious work. Often a hacker is making multiple attempts to gain access to the internal network or the admin interface of the FortiGate unit, usually from different IP addresses to appear to the FortiGate unit as legitimate users. For this reason, the more concurrent sessions a hacker has open at once, the faster they will achieve their goal.

To help prevent this, you can limit concurrent user sessions to the same IP address. This allows valid users to continue their legitimate work while limiting hackers activity.

To enable administrator concurrent session restrictions - CLI

```
config system global
    set admin-concurrent enable
end
```

VPN authentication

All VPN configurations require users to authenticate. Authentication based on user groups applies to:

- SSL VPNs
- PPTP and L2TP VPNs
- an IPsec VPN that authenticates users using dialup groups
- a dialup IPsec VPN that uses XAUTH authentication (Phase 1)

You must create user accounts and user groups before performing the procedures in this section. If you create a user group for dialup IPsec clients or peers that have unique peer IDs, their user accounts must be stored locally on the FortiGate unit. You cannot authenticate these types of users using a RADIUS or LDAP server.

Configuring authentication of SSL VPN users

The general procedure for authenticating SSL VPN users is:

- 1 Configure user accounts.
- 2 Create one or more user groups for SSL VPN users.
See “Configuring user accounts and user groups for SSL VPN” in the [FortiOS Handbook SSL VPN chapter](#).
- 3 Enable SSL VPN.
- 4 Optionally, set inactivity and authentication timeouts.
- 5 Configure a security policy with SSL VPN action. Add an identity-based rule to allow access for the user groups you created for SSL VPN users.

See “Configuring security policies” in the [FortiOS Handbook SSL VPN chapter](#).

Configuring authentication timeout

By default, the SSL VPN authentication expires after 8 hours (28 800 seconds). You can change it only in the CLI, and the time entered must be in seconds. For example, to change this timeout to one hour, you would enter:

```
config vpn ssl settings
  set auth-timeout 3600
end
```

If you set the authentication timeout (`auth-timeout`) to 0 when you configure the timeout settings, the remote client does not have to re-authenticate unless they log out of the system. To fully take advantage of this setting, the value for `idle-timeout` has to be set to 0 also, so that the client does not time out if the maximum idle time is reached. If the `idle-timeout` is not set to the infinite value, the system will log out if it reaches the limit set, regardless of the `auth-timeout` setting.

Configuring authentication of remote IPsec VPN users

An IPsec VPN on a FortiGate unit can authenticate remote users through a dialup group. The user account name is the peer ID and the password is the pre-shared key.

Authentication through user groups is supported for groups containing only local users. To authenticate users using a RADIUS or LDAP server, you must configure XAUTH settings. See “Configuring XAuth authentication” on page 1259.

To configure user group authentication for dialup IPsec - web-based manager

- 1 Configure the dialup users who are permitted to use this VPN. Create a user group with Type:Firewall and add them to it.

For more information, see “Users and user groups” on page 1223.

- 2 Go to *VPN > IPsec > Auto Key (IKE)*, select *Create Phase 1* and enter the following information.

Name	Name for group of dialup users using the VPN for authentication.
Remote Gateway	List of the types of remote gateways for VPN. Select <i>Dialup User</i> .
Authentication Method	List of authentication methods available for users. Select <i>Preshared Key</i> and enter the preshared key.
Peer Options	Select <i>Accept peer ID in dialup group</i> . Select the user group that is to be allowed access to the VPN. The listed user groups contain only users with passwords on the FortiGate unit.



The *Accept peer ID in dialup group* option does not support authentication of users through an authentication server. The user accounts must exist on the FortiGate unit.

- 3 Select *Advanced* to reveal additional parameters and configure other VPN gateway parameters as needed.
- 4 Select *OK*.

To configure user group authentication for dialup IPsec - CLI example

The `peertype` and `usrgrp` options configure user group-based authentication.

```
config vpn ipsec phase1
  edit office_vpn
    set interface port1
    set type dynamic
    set psksecret yORRAzltNGhzgtV32jend
    set proposal 3des-sha1 aes128-sha1
    set peertype dialup
    set usrgrp Group1
  end
```

Configuring XAuth authentication

Extended Authentication (XAuth) increases security by requiring additional user authentication information in a separate exchange at the end of the VPN Phase 1 negotiation. The FortiGate unit asks the user for a username and password. It then forwards the user's credentials (the password is encrypted) to an external RADIUS or LDAP server for verification.

XAuth can be used in addition to or in place of IPsec phase 1 peer options to provide access security through an LDAP or RADIUS authentication server. You must configure a dialup user group whose members are all externally authenticated.



None of the users in this dialup user group can have their passwords stored on the FortiGate unit.

To configure authentication for a dialup IPsec VPN - web-based manager

- 1 Configure the users who are permitted to use this VPN. Create a user group and add the users to the group.

For more information, see [“Users and user groups” on page 1223](#).

- 2 Go to *VPN > IPsec > Auto Key (IKE)*.
- 3 Select *Create Phase 1* and configure the basic VPN phase1 settings.
Remote Gateway must be *Dialup User*.
- 4 Select *Advanced* to reveal additional parameters and enter the following information.

XAuth	Select <i>Enable as Server</i> .
Server Type	Select <i>PAP</i> , <i>CHAP</i> , or <i>AUTO</i> . Use CHAP whenever possible. Use PAP with all implementations of LDAP and with other authentication servers that do not support CHAP, including some implementations of Microsoft RADIUS. Use AUTO with the Fortinet Remote VPN Client and where the authentication server supports CHAP but the XAuth client does not.
User Group	Select the user group that is to have access to the VPN. The list of user groups does not include any group that has members whose password is stored on the FortiGate unit.

- 5 Select *OK*.

For more information about XAUTH configuration, see the IPsec VPN chapter of this FortiOS Handbook.

To configure authentication for a dialup IPsec VPN - CLI example

The `xauthtype` and `authusrgrp` fields configure XAuth authentication.

```
config vpn ipsec phase1
  edit office_vpn
    set interface port1
    set type dynamic
    set psksecret yORRAz1tNGhzgtV32jend
    set proposal 3des-sha1 aes128-sha1
    set peertype dialup
    set xauthtype pap
    set authusrgrp Group1
  end
```

Some parameters specific to setting up the VPN itself are not shown here. For detailed information about configuring IPsec VPNs, see the [FortiOS Handbook IPsec VPN chapter](#).

Configuring authentication of PPTP VPN users and user groups

Configuration of a PPTP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.

To configure authentication for a PPTP VPN

- 1 Configure the users who are permitted to use this VPN. Create a security user group and add them to it.

For more information, see [“Users and user groups” on page 1223](#).

- 2 Configure the PPTP VPN in the CLI as in this example.

```
config vpn pptp
    set status enable
    set sip 192.168.0.100
    set eip 192.168.0.110
    set usrgrp PPTP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to PPTP clients.

- 3 Configure a security policy. The source interface is the one through which the clients will connect. The source address is the PPTP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.

Configuring authentication of L2TP VPN users/user groups

Configuration of a L2TP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.

To configure authentication for a PPTP VPN

- 1 Configure the users who are permitted to use this VPN. Create a user group and add them to it.


For more information, see [“Users and user groups” on page 1223](#).

- 2 Configure the L2TP VPN in the CLI as in this example.

```
config vpn l2tp
    set status enable
    set sip 192.168.0.100
    set eip 192.168.0.110
    set usrgrp L2TP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to L2TP clients.

- 3 Configure a security policy. The source interface is the one through which the clients will connect. The source address is the L2TP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.



Certificate-based authentication

This section provides an overview of how the FortiGate unit verifies the identities of administrators, SSL VPN users, or IPsec VPN peers using X.509 security certificates.

The following topics are included in this section:

- [What is a security certificate?](#)
- [Certificates overview](#)
- [Managing X.509 certificates](#)
- [Configuring certificate-based authentication](#)
- [Example — Generate a CSR on the FortiGate unit](#)
- [Example — Generate and Import CA certificate with private key pair on OpenSSL](#)
- [Example — Generate an SSL certificate in OpenSSL](#)

What is a security certificate?

A security certificate is a small text file that is part of a third-party generated public key infrastructure (PKI) to help guarantee the identity of both the user logging on and the website they where they are logging in.

A certificate includes identifying information such as the company and location information for the website, as well as the third-party company name, the expiry date of the certificate, and the encrypted public key.

FortiGate units use X.509 certificates to authenticate single sign-on (SSO) for users. The X.509 standard has been in use since before 2000, but has gained popularity with the Internet's increased popularity. X.509 v3 is defined in RFC 5280 and specifies standard formats for public key certificates, certificate revocation lists, and a certification path validation algorithm. The unused earlier X.509 version 1 was defined in RFC 1422.

The main difference between X.509 and PGP certificates is that where in PGP anyone can sign a certificate, for X.509 only a trusted authority can sign certificates. This limits the source of certificates to well known and trustworthy sources. Where PGP is well suited for one-on-one communications, the X.509 infrastructure is intended to be used in many different situations including one-to-many communications. Some common filename extensions for X.509 certificates are listed in [Table 81](#).

Table 81: Common certificate filename extensions

Filetype	Format name	Description
.pem	Privacy Enhanced Mail (PEM)	Base64 encoded DER certificate, that uses "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----"
.cer .crt .der	Security CERTificate	Usually binary DER form, but Base64-encoded certificates are common too.

Table 81: Common certificate filename extensions

.p7b .p7c	PKCS#7 SignedData	Structure without data, just certificates or CRLs. PKCS#7 is a standard for signing or encrypting (officially called "enveloping") data.
.p12	PKCS#12	May contain certificate(s) (public) and private keys (password protected)
.pfx	personal information exchange (PFX)	Older format. Came before PKCS#12. Usually today data is in PKCS#12 format.

Certificates overview

Certificates play a major role in authentication of clients connecting to network services via HTTPS, both for administrators and SSL VPN users. Certificate authentication is optional for IPsec VPN peers.

- [Certificates and protocols](#)
- [IPsec VPNs and certificates](#)
- [Certificate types on the FortiGate unit](#)

Certificates and protocols

There are a number of protocols that are commonly used with certificates including SSL and HTTPS, and other certificate-related protocols.

SSL and HTTPS

The secure HTTP (HTTPS) protocol uses SSL. Certificates are an integral part of SSL. When a web browser connects to the FortiGate unit via HTTPS, a certificate is used to verify the FortiGate unit's identity to the client. Optionally, the FortiGate unit can require the client to authenticate itself in return.

By default, the FortiGate unit uses a self-signed security certificate to authenticate itself to HTTPS clients. When the certificate is offered, the client browser displays two security messages.

- The first message prompts users to accept and optionally install the FortiGate unit's self-signed security certificate. If the user does not accept the certificate, the FortiGate unit refuses the connection. When the user accepts the certificate, the FortiGate login page is displayed, and the credentials entered by the user are encrypted before they are sent to the FortiGate unit. If the user chooses to install the certificate, the prompt is not displayed again.
- Just before the FortiGate login page is displayed, a second message informs users that the FortiGate certificate distinguished name differs from the original request. This message is displayed because the FortiGate unit redirects the connection (away from the distinguished name recorded in the self-signed certificate) and can be ignored.

Optionally, you can install an X.509 server certificate issued by a certificate authority (CA) on the FortiGate unit. You can then configure the FortiGate unit to identify itself using the server certificate instead of the self-signed certificate. For more information, see the [FortiOS Handbook SSL VPN chapter](#), or ["Authenticating SSL VPN users with security certificates"](#) on page 1275.

After successful certificate authentication, communication between the client browser and the FortiGate unit is encrypted using SSL over the HTTPS link.

Certificate-related protocols

There are multiple protocols that are required for handling certificates. These include the Online Certificate Status Protocol (OCSP), Secure Certificate Enrollment Protocol (SCEP), and Server-based Certificate Validation Protocol (SCVP).

Online Certificate Status Protocol

Online Certificate Status Protocol (OCSP) allows the verification of X.509 certificate expiration dates. This is important to prevent hackers from changing the expiry date on an old certificate to a future date.

Normally certificate revocation lists (CRLs) are used, but OCSP is an alternate method available. However a CRL is a public list, and some companies may want to avoid the public exposure of their certificate structure even if it is only invalid certificates.

The OCSP check on the certificate's revocation status is typically carried out over HTTP with a request-response format. The authority responding can reply with a status of good, revoked, or unknown for the certificate in question.

Secure Certificate Enrollment Protocol

Secure Certificate Enrollment Protocol (SCEP) is an automated method of signing up for certificates. Typically this involves generating a request you send directly to the SCEP service, instead of generating a file request that may or may not be signed locally.

Server-based Certificate Validation Protocol

Server-based Certificate Validation Protocol (SCVP) is used to trace a certificate back to a valid root level certificate. This ensures that each step along the path is valid and trustworthy.

IPsec VPNs and certificates

Certificate authentication is a more secure alternative to preshared key (shared secret) authentication for IPsec VPN peers. Unlike administrators or SSL VPN users, IPsec peers use HTTP to connect to the VPN gateway configured on the FortiGate unit. The VPN gateway configuration can require certificate authentication before it permits an IPsec tunnel to be established. See [“Authenticating IPsec VPN users with security certificates” on page 1276](#).

Certificate types on the FortiGate unit

There are different types of certificates available that vary depending on their intended use. FortiOS supports local, remote, CA, and CRL certificates.

Local certificates

Local certificates are issued for a specific server, or web site. Generally they are very specific, and often for an internal enterprise network. For example a personal web site for John Smith at www.example.com (such as <http://www.example.com/home/jsmith>) would have its own local certificate.

These can optionally be just the certificate file, or also include a private key file and PEM passphrase for added security.

Remote certificates

Remote certificates are public certificates without a private key. For dynamic certificate revocation, you need to use an Online Certificate Status Protocol (OCSP) server. The OCSP is configured in the CLI only. Installed Remote (OCSP) certificates are displayed in the Remote Certificates list.

CA root certificates

CA root certificates are similar to local certificates, however they apply to a broader range of addresses or to whole company; they are one step higher up in the organizational chain. Using the local certificate example, a CA root certificate would be issued for all of www.example.com instead of just the smaller single web page.

Certificate revocation list

Certificate revocation list (CRL) is a list of certificates that have been revoked and are no longer usable. This list includes certificates that have expired, been stolen, or otherwise compromised. If your certificate is on this list, it will not be accepted. CRLs are maintained by the CA that issues the certificates and includes the date and time when the next CRL will be issued as well as a sequence number to help ensure you have the most current version of the CRL.

Certificate signing

The trust in a certificate comes from the authority that signs it. For example if VeriSign signs your CA root certificate, it is trusted by everyone. While these certificates are universally accepted, it is cumbersome and expensive to have all certificates on a corporate network signed with this level of trust.

With self-signed certificates nobody, except the other end of your communication, knows who you are and therefore they do not trust you as an authority. However this level is useful for encryption between two points — neither point may care about who signed the certificate, just that it allows both points to communicate. This is very useful for internal networks and communications.

A general rule is that CA signed certificates are accepted and sometimes required, but it is easier to self-sign certificates when you are able.

For more on the methods of certificate signing see [“Generating a certificate signing request” on page 1267](#).

Managing X.509 certificates

Managing security certificates is required due to the number of steps involved in both having a certificate request signed, and then distributing the correct files for use.

You use the FortiGate unit or CA software such as OpenSSL to generate a certificate request. That request is a text file that you send to the CA for verification, or alternately you use CA software to self-validate. Once validated, the certificate file is generated and must be imported to the FortiGate unit before it can be used. These steps are explained in more detail later in this section.

This section provides procedures for generating certificate requests, installing signed server certificates, and importing CA root certificates and CRLs to the FortiGate unit.

For information about how to install root certificates, CRLs, and personal or group certificates on a remote client browser, refer to your browser's documentation.

This section includes:

- [Generating a certificate signing request](#)
- [Generating certificates with CA software](#)
- [Obtaining a signed server certificate from an external CA](#)
- [Installing a CA root certificate and CRL to authenticate remote clients](#)
- [Troubleshooting certificates](#)

Generating a certificate signing request

Whether you create certificates locally with a software application or obtain them from an external certificate service, you will need to generate a certificate signing request (CSR).

When you generate a CSR, a private and public key pair is created for the FortiGate unit. The generated request includes the public key of the FortiGate unit and information such as the FortiGate unit's public static IP address, domain name, or email address. The FortiGate unit's private key remains confidential on the FortiGate unit.

After you submit the request to a CA, the CA will verify the information and register the contact information on a digital certificate that contains a serial number, an expiration date, and the public key of the CA. The CA will then sign the certificate, and you install the certificate on the FortiGate unit.

The Certificate Request Standard is a public key cryptography standard (PKCS) published by RSA, specifically PKCS10 which defines the format for CSRs. This is defined in RFC 2986.

To generate a certificate request in FortiOS - web-based manager

- 1 Go to *System > Certificates > Local Certificates*.
- 2 Select *Generate*.
- 3 In the *Certificate Name* field, enter a unique meaningful name for the certificate request. Typically, this would be the hostname or serial number of the FortiGate unit or the domain of the FortiGate unit such as example.com.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.

- 4 Enter values in the *Subject Information* area to identify the FortiGate unit:
 - If the FortiGate unit has a static IP address, select *Host IP* and enter the public IP address of the FortiGate unit. If the FortiGate unit does not have a public IP address, use an email address (or fully qualified domain name (FQDN) if available) instead.
 - If the FortiGate unit has a static IP address and subscribes to a dynamic DNS service, use a FQDN if available to identify the FortiGate unit. If you select *Domain Name*, enter the FQDN of the FortiGate unit. Do not include the protocol specification (http://) or any port number or path names.



If a domain name is not available and the FortiGate unit subscribes to a dynamic DNS service, an "unable to verify certificate" type message may be displayed in the user's browser whenever the public IP address of the FortiGate unit changes.

- If you select *E-Mail*, enter the email address of the owner of the FortiGate unit.

- 5 Enter values in the *Optional Information* area to further identify the FortiGate unit.

Organization Unit	Name of your department. You can enter a series of OUs up to a maximum of 5. To add or remove an OU, use the plus (+) or minus (-) icon.
Organization	Legal name of your company or organization.
Locality (City)	Name of the city or town where the FortiGate unit is installed.
State/Province	Name of the state or province where the FortiGate unit is installed.
Country	Select the country where the FortiGate unit is installed.
e-mail	Contact email address.

- 6 From the *Key Size* list, select *1024 Bit*, *1536 Bit* or *2048 Bit*. Larger keys are slower to generate but more secure.
- 7 In *Enrollment Method*, you have two methods to choose from. Select *File Based* to generate the certificate request, or *Online SCEP* to obtain a signed SCEP-based certificate automatically over the network. For the SCEP method, enter the URL of the SCEP server from which to retrieve the CA certificate, and the CA server challenge password.
- 8 Select *OK*.
- The request is generated and displayed in the *Local Certificates* list with a status of *PENDING*.
- 9 Select the *Download* button to download the request to the management computer.
- 10 In the *File Download* dialog box, select *Save* and save the Certificate Signing Request on the local file system of the management computer.
- 11 Name the file and save it on the local file system of the management computer.
- The certificate request is ready for the certificate to be generated.

Generating certificates with CA software

CA software allows you to generate unmanaged certificates and CA certificates for managing other certificates locally without using an external CA service. Examples of CA software include *ssl-ca* from OpenSSL (available for Linux, Windows, and Mac) or *gensslcert* from SuSE, MS Windows Server 2000 and 2003 come with a CA as part of their certificate services, and in MS Windows 2008 CA software can be installed as part of the Active Directory installation. See [“Example — Generate and Import CA certificate with private key pair on OpenSSL” on page 1278](#).

The general steps for generating certificates with CA software are

- 1 Install the CA software as a stand-alone root CA.
- 2 Provide identifying information for your self-administered CA.

While following these steps, the methods vary slightly when generating server certificates, CA certificates, and PKI certificates.

Server certificate

- 1 Generate a Certificate Signing Request (CSR) on the FortiGate unit.

- 2 Copy the CSR base-64 encoded text (PKCS10 or PKCS7) into the CA software and generate the certificate.
PKCS10 is the format used to send the certificate request to the signing authority.
PKCS7 is the format the signing authority can use for the newly signed certificate.
- 3 Export the certificate as a X.509 DER encoded binary file with .CER extension
- 4 Upload the certificate file to the FortiGate unit Local Certificates page (type is Certificate).

CA certificate

- 1 Retrieve the CA Certificate from the CA software as a DER encoded file.
- 2 Upload the CA certificate file to the FortiGate unit CA Certificates page at *System > Certificates > CA Certificates*.

PKI certificate

- 1 Generate a Certificate Signing Request (CSR) on the FortiGate unit.
- 2 Copy the CSR base-64 encoded text (PKCS#10 or PKCS#7) into the CA software and generate the certificate.
PKCS10 is the format used to send the certificate request to the signing authority.
PKCS7 is the format the signing authority can use for the newly signed certificate.
- 3 Export the certificate as a X.509 DER encoded binary file with .CER extension.
- 4 Install the certificate in the user's web browser or IPsec VPN client as needed.

Obtaining a signed server certificate from an external CA

To obtain a signed server certificate for a FortiGate unit, you must send a request to a CA that provides digital certificates that adhere to the X.509 standard. The FortiGate unit provides a way for you to generate the request.

To submit the certificate signing request (file-based enrollment)

- 1 Using the web browser on the management computer, browse to the CA web site.
- 2 Follow the CA instructions for a base-64 encoded PKCS#10 certificate request and upload your certificate request.
- 3 Follow the CA instructions to download their root certificate and CRL.

When you receive the signed server certificate from the CA, install the certificate on the FortiGate unit.

To install or import the signed server certificate - web-based manager

- 1 On the FortiGate unit, go to *System > Certificates > Local Certificates*.
- 2 Select *Import*.
- 3 From *Type*, select *Local Certificate*.
- 4 Select *Browse*, browse to the location on the management computer where the certificate was saved, select the certificate, and then select *Open*.
- 5 Select *OK*, and then select *Return*.

Installing a CA root certificate and CRL to authenticate remote clients

When you apply for a signed personal or group certificate to install on remote clients, you can obtain the corresponding root certificate and CRL from the issuing CA. When you receive the signed personal or group certificate, install the signed certificate on the remote client(s) according to the browser documentation. Install the corresponding root certificate (and CRL) from the issuing CA on the FortiGate unit according to the procedures given below.



FortiGate IPsec VPNs do not support CRL lookups.

To install a CA root certificate

- 1 After you download the root certificate of the CA, save the certificate on the management computer. Or, you can use online SCEP to retrieve the certificate.
- 2 On the FortiGate unit, go to *System > Certificates > CA Certificates*.
- 3 Select *Import*.
- 4 Do one of the following:
 - To import using SCEP, select *SCEP*. Enter the URL of the SCEP server from which to retrieve the CA certificate. Optionally, enter identifying information of the CA, such as the filename.
 - To import from a file, select *Local PC*, then select *Browse* and find the location on the management computer where the certificate has been saved. Select the certificate, and then select *Open*.
- 5 Select *OK*, and then select *Return*.

The system assigns a unique name to each CA certificate. The names are numbered consecutively (CA_Cert_1, CA_Cert_2, CA_Cert_3, and so on).

To import a certificate revocation list

A Certificate Revocation List (CRL) is a list of the CA certificate subscribers paired with certificate status information. The list contains the revoked certificates and the reason(s) for revocation. It also records the certificate issue dates and the CAs that issued them.

When configured to support SSL VPNs, the FortiGate unit uses the CRL to ensure that the certificates belonging to the CA and remote peers or clients are valid. You must download the CRL from the CA web site on a regular basis.

- 1 After you download the CRL from the CA web site, save the CRL on the management computer.
- 2 Go to *System > Certificates > CRL*.
- 3 Select *Import*.

- 4 Do one of the following:
 - To import using an HTTP server, select *HTTP* and enter the URL of the HTTP server.
 - To import using an LDAP server, select *LDAP* and select the LDAP server from the list.
 - To import using an SCEP server, select *SCEP* and select the Local Certificate from the list. Enter the URL of the SCEP server from which the CRL can be retrieved.
 - To import from a file, select *Local PC*, then select *Browse* and find the location on the management computer where the CRL has been saved. Select the CRL and then select *Open*.
- 5 Select *OK*, and then select *Return*.

Troubleshooting certificates

There are times when there are problems with certificates — a certificate is seen as expired when its not, or it can't be found. Often the problem is with a third party website, and not FortiOS. However, some problems can be traced back to FortiOS such as DNS or routing issues.

Certificate is reported as expired when it is not

Certificates often are issued for a set period of time such as a day or a month, depending on their intended use. This ensures everyone is using up-to-date certificates. It is also more difficult for hackers to steal and use old certificates.

Reasons a certificate may be reported as expired include:

- It really has expired based on the “best before” date in the certificate
- The FortiGate unit clock is not properly set. If the FortiGate clock is fast, it will see a certificate as expired before the expiry date is really here.
- The requesting server clock is not properly set. A valid example is if your certificate is 2 hours from expiring, a server more than two time zones away would see the certificate as expired. Otherwise, if the server's clock is set wrongly it will also have the same effect.
- The certificate was revoked by the issuer before the expiry date. This may happen if the issuer believes a certificate was either stolen or misused. Its possible it is due to reasons on the issuer's side, such as a sytem change or such. In either case it is best to contact the certificate issuer to determinewhat is happening and why.

A secure connection cannot be completed (Certificate cannot be found)

Everyone who uses a browser has encountered a message such as *This connection is untrusted*. Normally when you try to connect securly to a website, that website will present its valid certificate to prove their identity is valid. When the website's certificate cannot be verified as valid, the message appears stating *This connection is untrusted* or something similar. If you usually connect to this website without problems, this error could mean that someone is trying to impersonate or hijack the website, and best practices dictates you not continue.

Reasons a website's certificate cannot be validated include:

- The website uses an unrecognized self-signed certificate. These are not secure because anyone can sign them. If you accept self-signed certificates you do so at your own risk. Best practices dictate that you must confirm the ID of the website using some other method before you accept the certificate.

- The certificate is valid for a different domain. A certificate is valid for a specific location, domain, or sub-section of a domain such as one certificate for `support.example.com` that is not valid for `marketing.example.com`. If you encounter this problem, contact the webmaster for the website to inform them of the problem.
- There is a DNS or routing problem. If the website's certificate cannot be verified, it will not be accepted. Generally to be verified, your system checks with the third party certificate signing authority to verify the certificate is valid. If you cannot reach that third party due to some DNS or routing error, the certificate will not be verified.
- Firewall is blocking required ports. Ensure that any firewalls between the requesting computer and the website allow the secure traffic through the firewall. Otherwise a hole must be opened to allow it through. This includes ports such as 443 (HTTPS) and 22 (SSH).

Online updates to certificates and CRLs

If you obtained your local or CA certificate using SCEP, you can configure online renewal of the certificate before it expires. Similarly, you can receive online updates to CRLs.

Local certificates

In the `config vpn certificate local` command, you can specify automatic certificate renewal. The relevant fields are:

<code>scep-url <URL_str></code>	The URL of the SCEP server. This can be HTTP or HTTPS.
<code>scep-password <password_str></code>	The password for the SCEP server.
<code>auto-regenerate-days <days_int></code>	How many days before expiry the FortiGate unit requests an updated local certificate. The default is 0, no auto-update.
<code>auto-regenerate-days-warning <days_int></code>	How many days before local certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate local
edit mycert
set scep-url http://scep.example.com/scep
set scep-server-password my_pass_123
set auto-regenerate-days 3
set auto-regenerate-days-warning 2
end
```

CA certificates

In the `config vpn certificate ca` command, you can specify automatic certificate renewal. The relevant fields are:

<code>scep-url <URL_str></code>	The URL of the SCEP server. This can be HTTP or HTTPS.
<code>auto-update-days <days_int></code>	How many days before expiry the FortiGate unit requests an updated CA certificate. The default is 0, no auto-update.
<code>auto-update-days-warning <days_int></code>	How many days before CA certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate ca
  edit mycert
    set scep-url http://scep.example.com/scep
    set auto-update-days 3
    set auto-update-days-warning 2
  end
```

Certificate Revocation Lists

If you obtained your CRL using SCEP, you can configure online updates to the CRL using the `config vpn certificate crl` command. The relevant fields are:

Variable	Description
<code>http-url <http_url></code>	URL of the server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>scep-cert <scep_certificate></code>	Local certificate used for SCEP communication for CRL auto-update.
<code>scep-url <scep_url></code>	URL of the SCEP CA server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>update-interval <seconds></code>	How frequently, in seconds, the FortiGate unit checks for an updated CRL. Enter 0 to update the CRL only when it expires.
<code>update-vdom <update_vdom></code>	VDOM used to communicate with remote SCEP server for CRL auto-update.

In this example, an updated CRL is requested only when it expires.

```
config vpn certificate crl
  edit cert_crl
    set http-url http://scep.example.com/scep
    set scep-cert my-scep-cert
    set scep-url http://scep.ca.example.com/scep
    set update-interval 0
    set update-vdom root
  end
```

Backing up and restoring local certificates

The FortiGate unit provides a way to export and import a server certificate and the FortiGate unit's personal key through the CLI. If required (to restore the FortiGate unit configuration), you can import the exported file through the *System > Certificates > Local Certificates* page of the web-based manager.



As an alternative, you can back up and restore the entire FortiGate configuration through the *System > Maintenance > Backup & Restore* page of the web-based manager. The backup file is created in a FortiGate-proprietary format.

To export a server certificate and private key - CLI

This procedure exports a server (local) certificate and private key together as a password protected PKCS12 file. The export file is created through a customer-supplied TFTP server. Ensure that your TFTP server is running and accessible to the FortiGate unit before you enter the command.

- 1 Connect to the FortiGate unit through the CLI.
- 2 Type the following command:

```
execute vpn certificate local export tftp <cert_name> <exp_filename>  
<tftp_ip>
```

where:

- <cert_name> is the name of the server certificate; typing ? displays a list of installed server certificates.
 - <exp_filename> is a name for the output file.
 - <tftp_ip> is the IP address assigned to the TFTP server host interface.
- 3 Move the output file from the TFTP server location to the management computer for future reference.

To import a server certificate and private key - web-based manager

- 1 Go to *VPN > Certificates > Local Certificates* and select *Import*.
- 2 In *Type*, select *PKCS12 Certificate*.
- 3 Select *Browse*. Browse to the location on the management computer where the exported file has been saved, select the file, and then select *Open*.
- 4 In the *Password* field, type the password needed to upload the exported file.
- 5 Select *OK*, and then select *Return*.

To import separate server certificate and private key files - web-based manager

Use the following procedure to import a server certificate and the associated private key file when the server certificate request and private key were not generated by the FortiGate unit. The two files to import must be available on the management computer.

- 1 Go to *VPN > Certificates > Local Certificates* and select *Import*.
- 2 In *Type*, select *Certificate*.
- 3 Select the *Browse* button beside the *Certificate file* field. Browse to the location on the management computer where the certificate file has been saved, select the file, and then select *Open*.

- 4 Select the *Browse* button beside the *Key file* field. Browse to the location on the management computer where the key file has been saved, select the file, and then select *Open*.
- 5 If required, in the *Password* field, type the associated password, and then select *OK*.
- 6 Select *Return*.

Configuring certificate-based authentication

You can configure certificate-based authentication for FortiGate administrators, SSL VPN users, and IPsec VPN users.

In Microsoft Windows 7, you can use the certificate manager to keep track of all the different certificates on your local computer. To access certificate manager, in Windows 7 press the Windows key, enter “certmgr.msc” at the search prompt, and select the displayed match. Remember that in addition to these system certificates, many applications require you to register certificates with them directly.

To see FortiClient certificates, open the FortiClient Console, and select VPN. The VPN menu has options for My Certificates (local or client) and CA Certificates (root or intermediary certificate authorities). Use Import on those screens to import certificate files from other sources.

Authenticating administrators with security certificates

You can install a certificate on the management computer to support strong authentication for administrators. When a personal certificate is installed on the management computer, the FortiGate unit processes the certificate after the administrator supplies a username and password.

To enable strong administrative authentication:

- Obtain a signed personal certificate for the administrator from a CA and load the signed personal certificate into the web browser on the management computer according to the browser documentation.
- Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [“Installing a CA root certificate and CRL to authenticate remote clients” on page 1270](#)).
- Create a PKI user account for the administrator.
- Add the PKI user account to a firewall user group dedicated to PKI-authenticated administrators.
- In the administrator account configuration, select *PKI* as the account *Type* and select the *User Group* to which the administrator belongs.

Authenticating SSL VPN users with security certificates

While the default self-signed certificates can be used for HTTPS connections, it is preferable to use the X.509 server certificate to avoid the redirection as it can be misinterpreted as possible session hijacking. However, the server certificate method is more complex than self-signed security certificates. Also the warning message is typically displayed for the initial connection, and future connections will not generate these messages.

X.509 certificates can be used to authenticate IPsec VPN peers or clients, or SSL VPN clients. When configured to authenticate a VPN peer or client, the FortiGate unit prompts the VPN peer or client to authenticate itself using the X.509 certificate. The certificate supplied by the VPN peer or client must be verifiable using the root CA certificate installed on the FortiGate unit in order for a VPN tunnel to be established.

To enable certificate authentication for an SSL VPN user group

- 1 Install a signed server certificate on the FortiGate unit and install the corresponding root certificate (and CRL) from the issuing CA on the remote peer or client.
- 2 Obtain a signed group certificate from a CA and load the signed group certificate into the web browser used by each user. Follow the browser documentation to load the certificates.
- 3 Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [“Installing a CA root certificate and CRL to authenticate remote clients” on page 1270](#)).
- 4 Create a PKI user for each SSL VPN user. For each user, specify the text string that appears in the Subject field of the user’s certificate and then select the corresponding CA certificate.
- 5 Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the SSL VPN users who are authenticated by certificate.
- 6 Go to *VPN > SSL > Config*.
- 7 Select *Enable SSL-VPN*.
- 8 Go to *Policy*.
- 9 Select the *Edit* icon in the row that corresponds to the SSL-VPN security policy for traffic generated by holders of the group certificate.
- 10 Select *SSL Client Certificate Restrictive*.
- 11 Select *OK*.

Authenticating IPsec VPN users with security certificates

To require VPN peers to authenticate by means of a certificate, the FortiGate unit must offer a certificate to authenticate itself to the peer.

To enable the FortiGate unit to authenticate itself with a certificate:

- 1 Install a signed server certificate on the FortiGate unit.
See [“To install or import the signed server certificate - web-based manager” on page 1269](#).
- 2 Install the corresponding CA root certificate on the remote peer or client. If the remote peer is a FortiGate unit, see [“To install a CA root certificate” on page 1270](#).
- 3 Install the certificate revocation list (CRL) from the issuing CA on the remote peer or client. If the remote peer is a FortiGate unit, see [“To import a certificate revocation list” on page 1270](#).
- 4 In the VPN phase 1 configuration, set *Authentication Method* to *RSA Signature* and from the *Certificate Name* list select the certificate that you installed in Step 1.

To authenticate a VPN peer using a certificate, you must install a signed server certificate on the peer. Then, on the FortiGate unit, the configuration depends on whether there is only one VPN peer or if this is a dialup VPN that can have multiple peers.

To configure certificate authentication of a single peer

- 1 Install the CA root certificate and CRL.
- 2 Create a PKI user to represent the peer. Specify the text string that appears in the Subject field of the user's certificate and then select the corresponding CA certificate.
- 3 In the VPN phase 1 *Peer Options*, select *Accept this peer certificate only* and select the PKI user that you created.

To configure certificate authentication of multiple peers (dialup VPN)

- 1 Install the corresponding CA root certificate and CRL.
- 2 Create a PKI user for each remote VPN peer. For each user, specify the text string that appears in the Subject field of the user's certificate and then select the corresponding CA certificate.
- 3 Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the PKI users who will use the IPsec VPN.
In the VPN phase 1 *Peer Options*, select *Accept this peer certificate group only* and select the peer group that you created.



FortiGate IPsec VPNs do not support CRL lookups.

Example — Generate a CSR on the FortiGate unit

This example follows all the steps required to create and install a local certificate on the FortiGate unit, without using CA software.

The FortiGate unit is called myFortiGate60, and is located at 10.11.101.101 (a private IP address) and <http://myfortigate.example.com>. Mr. John Smith (john.smith@myfortigate.example.com) is the IT administrator for this FortiGate unit, and the unit belongs to the Sales department located in Greenwich, London, England.

To generate a certificate request on the FortiGate unit - web-based manager

- 1 Go to *System > Certificates > Local Certificates*.
- 2 Select *Generate*.
- 3 In the *Certificate Name* field, enter myFortiGate60.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.

Since the IP address is private, we will use the FQDN instead.

- 4 Select *Domain Name*, and enter <http://myfortigate.example.com>.
- 5 Enter values in the *Optional Information* area to further identify the FortiGate unit.

Organization Unit	Sales.
Organization	Example.com
Locality (City)	Greenwich
State/Province	London

Country	England
e-mail	john.smith@myfortigate.example.com

- 6 From the *Key Size* list, select *2048 Bit* or the most secure option available to you.
- 7 In *Enrollment Method*, select *File Based* to generate the certificate request
- 8 Select *OK*.
The request is generated and displayed in the *Local Certificates* list with a status of *PENDING*.
- 9 Select the *Download* button to download the request to the management computer.
- 10 In the *File Download* dialog box, select *Save* and save the Certificate Signing Request on the local file system of the management computer.
- 11 Name the file and save it on the local file system of the management computer.

Example — Generate and Import CA certificate with private key pair on OpenSSL

This example explains how to generate a certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

Assumptions

Before starting this procedure, ensure that you have downloaded and installed OpenSSL on Windows. One source is <http://www.slproweb.com/products/Win32OpenSSL.html>.

Generating and importing the CA certificate and private key

The two following procedures will generate a CA certificate file and private key file, and then import it to the FortiGate unit as a local certificate.

To generate the private key and certificate

- 1 At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

- 2 Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as *fortinet123*.

```
openssl genrsa -des3 -out fgtpcapriv.key 2048
```

This command generates an RSA DES3 2038-bit encryption key.

- 3 The following command will generate the certificate using the key from the previous step.

```
openssl req -new -x509 -days 3650 -extensions v3_ca -key  
fgtpcapriv.key -out fgtpca.crt
```

This step generates an X509 CA certificate good for 10 years that uses the key generated in the previous step. The certificate filename is *fgtpca.crt*.

You will be prompted to enter information such as PEM Pass Phrase from the previous step, Country Name, State, Organization Name, Organizational Unit (such as department name), Common Name (the FQDN), and Email Address.

To import the certificate to the FortiGate unit - web-based manager

- 1 Go to *System > Certificates > Local Certificates*.
- 2 Select *Import*.
- 3 Select Certificate for Type.
Fields for Certificate file, Key file, and Password are displayed.
- 4 For Certificate file, enter `c:\OpenSSL-Win32\bin\fgtca.crt`.
- 5 For Key file, enter `c:\OpenSSL-Win32\bin\fgtcapriv.key`.
- 6 For Password, enter the PEM Pass Phrase you entered earlier, such as `fortinet123`.
- 7 Select OK.

The Certificate will be added to the list of Local Certificates and be ready for use. It will appear in the list as the filename you uploaded — `fgtca`. You can add comments to this certificate to make it clear where its from and how it is intended to be used. If you download the certificate from FortiOS, it is a .CER file.

It can now be used in “[Authenticating IPsec VPN users with security certificates](#)” on page 1276, and “[Authenticating SSL VPN users with security certificates](#)” on page 1275.

Optionally, you can install the certificate as a CA Certificate. CA certificates are used in HTTPS proxy/inspection. To do this, under CA Certificates select Import. Select Local PC and enter the certificate file `c:\OpenSSL-Win32\bin\fgtca.crt`. Then select OK. This certificate will be displayed in the CA Certificate list under the name `CA_Cert_1`.

Example — Generate an SSL certificate in OpenSSL

This example explains how to generate a CA signed SSL certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

This example includes:

- [Assumptions](#)
- [Generating a CA signed SSL certificate](#)
- [Generating a self-signed SSL certificate](#)
- [Import the SSL certificate into FortiOS](#)

Assumptions

- Before starting this procedure, ensure that you have downloaded and installed OpenSSL on MS Windows. One download source is <http://www.slproweb.com/products/Win32OpenSSL.html>.

Generating a CA signed SSL certificate

This procedure assumes:

- you have already completed “[Example — Generate and Import CA certificate with private key pair on OpenSSL](#)” on page 1278 successfully.

To generate the CA signed SSL certificate

- 1 At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

- 2 Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -des3 -out fgtssl.key 2048
```

This command generates an RSA DES3 2038-bit encryption key.

- 3 Create a certificate signing request for the SSL certificate. This step requires you to enter the information listed in step 3 of the previous example — [“To generate the private key and certificate” on page 1278](#). You can leave the Challenge Password blank.

```
openssl req -new -key fgtssl.key -out fgtssl.csr
```

- 4 Using the CSR from the previous step, you can now create the SSL certificate using the CA certificate that was created in [“Example — Generate and Import CA certificate with private key pair on OpenSSL” on page 1278](#).

```
openssl x509 -req -days 365 -in fgtssl.csr -CA fgtca.crt -  
CAkey fgtpcapriv.key -set_serial 01 -out fgtssl.crt
```

This will generate an X.509 certificate good for 365 days signed by the CA certificate fgtca.crt .

Generating a self-signed SSL certificate

This procedure does not require any existing certificates.

- 1 At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

- 2 Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -des3 -out fgtssl.key 2048  
openssl req -new -key fgtssl.key -out fgtssl.csr  
openssl x509 -req -days 365 -in fgtssl.csr -signkey fgtssl.key  
-out fgtssl.crt
```

These commands:

- generate an RSA 3DES 2048-bit private key,
- generate an SSL certificate signing request, and
- sign the CSR to generate an SSL .CRT certificate file.

Import the SSL certificate into FortiOS

To import the certificate to FortiOS- web-based manager

- 1 Go to *System > Certificates > Local Certificates*.
- 2 Select *Import*.
- 3 Select Certificate for Type.
Fields for Certificate file, Key file, and Password are displayed.
- 4 For Certificate file, enter `c:\OpenSSL-Win32\bin\fgtssl.crt`.
- 5 For Key file, enter `c:\OpenSSL-Win32\bin\fgtssl.key`.
- 6 For Password, enter the PEM Pass Phrase you entered, such as fortinet.
- 7 Select OK.

The SSL certificate you just uploaded can be found under *System > Certificates > Local Certificates* under the name of the file you uploaded — fgtssl.

To confirm the certificate is uploaded properly - CLI

```
config vpn certificate local
edit fgtssl
get
end
```

The get command will display all the certificate's information. If it is not there or the information is not correct, you will need to remove the corrupted certificate (if it is there) and upload it again from your PC.

To use the new SSL certificate - CLI

```
config vpn ssl settings
set servercert fgtssl
end
```

This assigns the fgtssl certificate as the SSL server certificate. For more information see the [FortiOS Handbook SSL VPN chapter](#).



FSSO integration with Windows AD or Novell

This chapter provides information, installation instructions, and troubleshooting for Fortinet Single Sign On (FSSO) agent. Earlier versions of this product were named Fortinet Server Authentication Extension (FSAE).

The following topics are included:

- [Introduction to FSSO](#)
- [FSSO for Windows AD](#)
- [FSSO for Novell eDirectory](#)
- [Configuring FSSO on FortiGate units](#)
- [FortiOS FSSO log messages](#)
- [Testing FSSO](#)
- [Troubleshooting FSSO](#)

Introduction to FSSO

The Fortinet Single Sign On (FSSO) agent connects FortiGate Fortinet security appliances to the corporate authentication servers, such as Microsoft Active Directory and Novell E-Directory, allowing security policies to be defined on the FortiGate unit based on the user information residing on the corporate authentication servers. FSSO, a component installed on the authentication server or a standalone server, provides user authentication information to the FortiGate unit so users can automatically gain access to the permitted resources with a single sign on. Older versions were called Fortinet Server Authentication Extension (FSAE).

On a Microsoft Windows or Novell network, users authenticate with the Active Directory or Novell eDirectory at logon. It would be inconvenient if users then had to enter another username and password for network access through the FortiGate unit. FSSO provides authentication information to the FortiGate unit so that users automatically get access to permitted resources.

There are several mechanisms for passing user authentication information to the FortiGate unit:

- FSSO software installed on a Windows AD network monitors user logons and sends the required information to the FortiGate unit. The FSSO software can obtain this information by polling the AD domain controllers or by using an FSSO agent on each AD domain controller that monitors user logons in real time. Optionally, a FortiGate unit running FortiOS 3.0 MR6 or later can obtain group information directly from AD using Lightweight Directory Access Protocol (LDAP). See [“Using FSSO in a Windows AD environment” on page 1284](#).

- On a Windows AD network, the FSSO software can also serve NT LAN Manager (NTLM) requests coming from client browsers (forwarded by the FortiGate unit) with only one or more Controller agents installed. See [“NTLM authentication with FSSO” on page 1287](#).
- FSSO software installed on a Novell network monitors user logons and sends the required information to the FortiGate unit. The FSSO software can obtain information from the Novell eDirectory using either the Novell API or LDAP. See [“Using FSSO in a Novell eDirectory environment” on page 1290](#).
- A FortiAuthenticator server can act as a replacement for the Collector agent in polling mode in a Windows AD network. FortiAuthenticator can also be configured with internal or external LDAP and RADIUS servers. For more information, see the [FortiAuthenticator Administration Guide](#).

Consult the latest FortiOS and FSSO Release Notes for operating system compatibility information.

Using FSSO in a Windows AD environment

FSSO installed in a Windows AD environment can provide two kinds of services:

- Monitor user logon activity and send the information to FortiGate unit so that the FortiGate unit can support Single user Sign On (SSO).
- Provide NTLM authentication service for requests coming from FortiGate.

SSO is very convenient for users, but may not be supported across all platforms. NTLM is not as convenient, but it enjoys wider support.

FSSO is certified for the Microsoft Windows Server 2003 (32- and 64-bit editions) and is supported on the Microsoft Windows Server 2008 (32- and 64-bit editions) operating systems.

FSSO security

When the different components of FSSO are communicating there are some inherent security features.

FSSO installation requires an account with network admin privileges. The security inherent in these types of accounts helps ensure access to FSSO configurations is not tampered with.

User passwords are never sent between FSSO components. The information that is sent is information to identify a user including the username, group or groups, and IP address.

NTLM uses base-64 encoded packets, and uses a unique randomly generated challenge nonce to avoid sending user information and password between the client and the server. For more information on NTLM, see [“NTLM authentication with FSSO” on page 1287](#).

FSSO Controller agent (CA)

The FSSO Controller agent (CA), or FSSO agent, is a service installed on a Windows computer that has access to both the FortiGate unit and each of the domain controller agents (DC agents).

The CA is responsible for DNS lookups, group verification, workstation checks, and as mentioned FortiGate updates of logon records. The FSSO Collector Agent sends Domain Local Security Group and Global Security Group information to FortiGate units. The CA communicates with the FortiGate over TCP port 8000 and it listens on UDP port 8002 for updates from the DC agents.

The FortiGate unit can have up to five CAs configured for redundancy. If the first on the list is unreachable, the next is attempted, and so on down the list until one is contacted. See [“Configuring FSSO on FortiGate units” on page 1308](#).

All DC agents must point to the correct Collector agent port number and IP address on domains with multiple DCs.

See [“Configuring Collector agent settings” on page 1296](#).

FSSO user logon event monitoring

A FSSO agent installed in a Windows AD environment monitors which users logon to which workstations and pass that information to the FortiGate unit. The FortiGate uses that information to apply its security policies.

When a Windows AD user logs on at a workstation in a monitored domain, FSSO

- detects the logon event and records the workstation name, domain, and user,
- resolves the workstation name to an IP address,
- uses Active Directory to determine which groups the user belongs to,
- sends the user logon information, including IP address and groups list, to the FortiGate unit
- creates one or more log entries on the FortiGate unit for this logon event as appropriate.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. The selection consists of matching the FSSO group or groups the user belongs to with the security policy or policies that match that group. If the user belongs to one of the permitted user groups associated with that policy, the connection is allowed. Otherwise the connection is denied.

With Windows AD, FSSO can use one of two different working modes to monitor user logon activity: DC Agent mode or Polling mode.

Table 82: FSSO DC Agent mode versus Polling mode

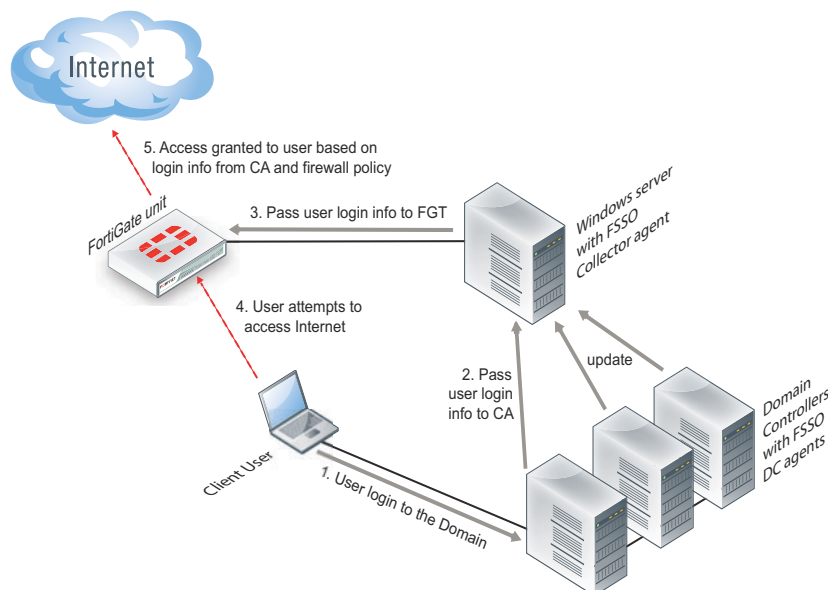
	DC Agent mode	Polling Mode
Installation	Complex — Multiple installations: one agent per DC plus Controller agent, requires a reboot	Easy — only Controller agent installation, no reboot required
Resources	Shares resources with DC system	Has own resources
Network load	Each DC agent requires minimum 64kpbs bandwidth, adding to network load	Increase polling period during busy period to reduce network load
Level of Confidence	Captures all logons	Potential to miss a login if polling period is too great

DC Agent mode

DC Agent mode is the standard mode for FSSO. In DC Agent mode (see [Figure 107](#)), a Fortinet authentication agent is installed on each domain controller. These DC agents monitor user logon events and pass the information to the Collector agent, which stores the information and sends it to the FortiGate unit.

The DC agent installed on the domain controllers is not a service like the Collector agent — it is a DLL file called `dcagent.dll` and is installed in the `Windows\system32` directory. It must be installed on all domain controllers of the domains that are being monitored.

Figure 107: FSSO in DC agent mode



DC Agent mode provides reliable user logon information, however you must install a DC agent on every domain controller. A reboot is needed after the agent is installed. Each installation requires some maintenance as well. For these reasons it may not be possible to use the DC Agent mode.

Each domain controller connection needs a minimum guaranteed 64kpbs bandwidth to ensure proper FSSO functionality. You can optionally configure traffic shapers on the FortiGate unit to ensure this minimum bandwidth is guaranteed for the domain controller connections.


Polling mode

In Polling mode there are two options — NetAPI polling, and Event log polling. Both share the advantages of being transparent and agentless.

NetAPI polling is used to retrieve server logon sessions. This includes the logon event information for the Controller agent. NetAPI runs faster than Event log polling but it may miss some user logon events under heavy system load. It requires a query round trip time of less than 10 seconds.

Event log polling may run a bit slower, but will not miss events, even when the installation site has many users that require authentication. It does not have the 10 second limit or NetAPI polling. Event log polling requires fast network links. Event log polling is required if there are Mac OS users logging into Windows AD.

In Polling mode (see [Figure 108](#)), the Collector agent polls port 445 of each domain controller for user logon information every few seconds and forwards it to the FortiGate unit. There are no DC Agents installed, so the Collector agent polls the domain controllers directly.

Figure 108: FSSO in Polling mode

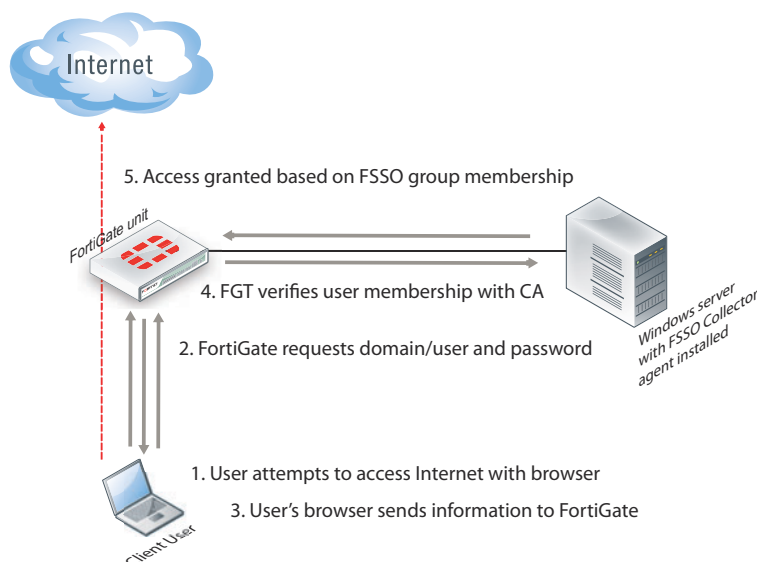
A major benefit of Polling mode is that no FSSO DC Agents are required. If it is not possible to install FSSO DC Agents on your domain controllers, this is the alternate configuration available to you. Polling mode results in a less complex install, and reduces ongoing maintenance. The minimum permissions required in Polling mode are to read the event log or call NetAPI. To install FSSO with minimum permissions, see [“Installing FSSO without using an administrator account” on page 1293](#).

NTLM authentication with FSSO

In a Windows AD network, FSSO can also provide NTLM authentication service to the FortiGate unit. When the user makes a request that requires authentication, the FortiGate unit initiates NTLM negotiation with the client browser. The FortiGate unit does not process the NTLM packets itself. Instead, it forwards all the NTLM packets to the FSSO service to process.

NTLM has the benefit of not requiring an FSSO agent, but it is not transparent to users, and the user’s web browser must support NTLM.

The NTLM protocol protects the user’s password by not sending it over the network. Instead, the server sends the client a random number that the client must encrypt with the hash value of the user’s password. The server compares the result of the client’s encryption with the result of its own encryption. The two will match only if both parties used the same password.

Figure 109: NTLM authentication

If the NTLM authentication with the Windows AD network is successful, and the user belongs to one of the groups permitted in the applicable security policy, the FortiGate unit allows the connection.


Fortinet has tested NTLM authentication with Internet Explorer and Firefox browsers.

NTLM in a multiple domain environment

In a multiple domain environment for NTLM, the important factor is that there is a trust relation between the domains. In a forest, this relation is automatically created. So you can install FSSO agent on one of the domain controllers without worry.

But in case of multiple domains that are not in a forest, you need to create a trust relation between the domains. If you do not want to have a trust relation between your multiple domains, you need to use FSAE 4.0 MR1 and the DC agent needs to be installed once on each domain. Then you can use security policies to configure server access.

In [Figure 110](#), three domains are shown connected to the FSSO Collector agent server. The Client logs on to their local Domain Controller, which then sends the user logon event information to the Collector Agent. When the Client attempts to access the Internet, the FortiGate unit contacts the Collector Agent for the logon information, sees the Client is authenticated, and allows access to the Internet. There are multiple domains each with a domain controller agent (DCagent) that sends logon information to the Collector agent. If the multiple domains have a trust relationship, only one DCagent is required instead of one per domain.

Figure 110: FSSO NTLM with multiple domains not in a forest**Understanding the NTLM authentication process**

- 1 The user attempts to connect to an external (internet) HTTP resource. The client application (browser) on the user's computer issues an unauthenticated request through the FortiGate unit.
- 2 The FortiGate is aware that this client has not authenticated previously, so responds with a 401 Unauthenticated status code, and tells the client which authentication method to reply with in the header: Proxy-Authenticated: NTLM. Then the initial session is dismantled.
- 3 The client application connects again to the FortiGate, and issues a GET-request, with a Proxy-Authorization: NTLM <negotiate string> header. <negotiate-string> is a base64-encoded NTLM Type 1 negotiation packet.
- 4 The FortiGate unit replies with a 401 "proxy auth required" status code, and a Proxy-Authenticate: NTLM <challenge string> (a base 64-encoded NTLM Type 2 challenge packet). In this packet is the challenge nonce, a random number chosen for this negotiation that is used once and prevents replay attacks.



The TCP connection must be kept alive, as all subsequent authentication-related information is tied to the TCP connection. If it is dropped, the authentication process must start again from the beginning.

- 5 The client sends a new GET-request with a header: Proxy-Authenticate: NTLM <authenticate string>, where <authenticate string> is a NTLM Type 3 Authentication packet that contains:
 - username and domain
 - the challenge nonce encoded with the client password (it may contain the challenge nonce twice using different algorithms).

- 6 If the negotiation is successful and the user belongs to one of the groups permitted in the security policy, the connection is allowed. Otherwise, the FortiGate unit denies the authentication by issuing a 401 return code and prompts for a username and password. Unless the TCP connection is broken, no further credentials are sent from the client to the proxy.



If the authentication policy reaches the authentication timeout period, a new NTLM handshake occurs.

Using FSSO in a Novell eDirectory environment

FSSO in a Novell eDirectory environment works similar to the FSSO Polling mode in the Windows AD environment. The eDirectory agent polls the eDirectory servers for user logon information and forwards it to the FortiGate unit.

When a user logs on at a workstation, FSSO:

- detects the logon event by polling the eDirectory server and records the IP address and user ID,
- looks up in the eDirectory which groups this user belongs to,
- sends the IP address and user groups information to the FortiGate unit.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

FSSO is supported on the Novell E-Directory 8.8 operating system

FSSO for Windows AD

This section explains what the FSSO components are for Windows AD, how to install them, and how to configure them.

- [FSSO components for Windows AD](#)
- [Standard versus Advanced mode](#)
- [Installing FSSO for Windows AD](#)
- [Configuring Fortinet Single Sign On with Windows AD](#)

FSSO components for Windows AD

FSSO has two components, or agents, to install on your network:

- the Collector agent must be installed on one or more network computers
- the Domain Controller (DC) agent must be installed on every domain controller if you will use DC Agent mode, but is not required if you use Polling mode.

FSSO is supported on Microsoft Windows Server 2003 (32- and 64-bit editions) and Microsoft Windows Server 2008 (32- and 64-bit editions) operating systems.

FSSO Installation

The FSSO installer first installs the Collector agent. You can then continue with installation of the DC agent, or you can install it later by going to *Start > Programs > Fortinet > Fortinet Server Authentication Extension > Install DC Agent*. The installer will install a DC agent on the domain controllers of all of the trusted domains in your network.



Each domain controller connection needs a minimum guaranteed 64kpbs bandwidth to ensure proper FSSO functionality. Traffic shapers configured on the FortiGate can help guarantee these minimum bandwidths.

You can create a redundant configuration if you install the Collector agent on two or more servers. This provides improved reliability. If the current Collector agent fails, the FortiGate unit will switch to the next Collector agent in its list. The list can have up to five Collector agents.



In Windows 2008 by default, you do not have administrative user rights if you are logged on as a user other than as the built-in administrator, even if you were added to the local Administrators group on the computer.

Ensure you have administrative rights on the servers where you are installing FSSO. Best practices dictate that you install FSSO using the built-in local administrator account. If you are logged on to another account and try to install FSSO, you may see a security alert dialog box requesting your permission to use or install the program.

Best practices dictate that before installing FSSO, you create a dedicated account with administrator privileges and a password that does not expire. Optionally, you can install FSSO without an admin account. See [“Installing FSSO without using an administrator account” on page 1293](#).

Standard versus Advanced mode

Part of installing FSSO for Windows is choosing Standard or Advanced mode. The main difference between Standard and Advanced mode is the naming convention used when referring to username information.

Standard mode uses regular Windows convention: Domain\Username. Advanced mode uses LDAP convention: CN=User, OU=Name, DC=Domain.

If there is no special requirement to use LDAP— best practices dictate you setup FSSO in Standard mode. This mode is easier to setup, and is usually easier to maintain and troubleshoot.

Standard and advanced modes have the same level of functionality with the following exceptions:

- 1 Users have to create Group filters on the Collector agent. This differs from Advanced mode where Group filters are configured from the FortiGate unit. Fortinet strongly encourages users to create filters from CA.
- 2 Advanced mode supports nested or inherited groups. This means that users may be a member of multiple monitored groups. Standard mode does not support nested groups so a user must be a direct member of the group being monitored.

Installing FSSO for Windows AD

To install FSSO, you must obtain the FSSO Setup file from the [Fortinet Support web site](#). Then you follow these two installation procedures on the server that will run the Collector agent. This can be any server or domain controller that is part of your network. These procedures also install the DC Agent on all of the domain controllers in your network.

To install the Collector agent

- 1 Create an account with administrator privileges and a password that does not expire. See Microsoft Advanced Server documentation for help with this task.

To use a non-admin read only account, see [“Installing FSSO without using an administrator account” on page 1293](#).

- 2 logon to the account that you created in Step 1.
- 3 Double-click the `FSSOSetup.exe` file.
The FSSO InstallShield Wizard starts.
- 4 Select *Next*. Optionally, you can change the installation location.
- 5 Select *Next*.



This procedure will install using the currently running account. If you want FSSO to use another existing admin account, change the *username:* field in the InstallShield Wizard to the correct username using the format `DomainName \ UserName`. For example if the account is `jsmith` and the domain is `example_corp` you would enter `example_corp\jsmith` in the *UserName:* field.

- 6 In the *Password* field, enter the password for the account listed in the *username* field. This is the account you are logged onto currently.
- 7 Select *Next*.
- 8 By default, FSSO authenticates users both by monitoring logons and by accepting authentication requests using the NTLM protocol.

If you want to support only NTLM authentication

- Clear the *Monitor user logon events and send the information to Fortinet* check box.
- Select the *Serve NTLM authentication requests coming from FortiGate* check box.

If you do not want to support NTLM authentication

- Clear the *Serve NTLM authentication requests coming from FortiGate* check box.
- Select the *Monitor user logon events and send the information to Fortinet* check box.

You can change these options after installation.

- 9 Select the access method to use for Windows Directory:
 - Select *Standard* to use Windows domain and username credentials.
 - Select *Advanced* if you will set up LDAP access to Windows Directory.

See [“Standard versus Advanced mode” on page 1291](#).

- 10 Select *Next* and then select *Install*.
- 11 For DC Agent mode, when the FSSO InstallShield Wizard completes Collector agent installation, ensure that *Launch DC Agent Install Wizard* is selected and then select *Finish*.



If you see an error such as Service Fortinet Single Sign On agent (service_FSAE) failed to start, there are two possible reasons for this. Verify the user account you selected has sufficient privileges to run the FSSO service. Also verify the computer system you are attempting to install on is a supported operating system and version.

To install the DC Agent

- 1 If you have just installed the Collector agent, the FSSO - Install DC Agent wizard starts automatically. Otherwise, go to *Start > Programs > Fortinet > Fortinet Single Sign On > Install DC Agent*.
- 2 Verify the *Collector agent IP address*.
If the Collector agent computer has multiple network interfaces, ensure that the one that is listed is on your network. The listed *Collector agent listening port* is the default. Only change this if the port is already used by another service.
- 3 Select *Next*.
- 4 Select the domains to monitor and select *Next*.
If any of your required domains are not listed, cancel the wizard and set up the proper trusted relationship with the domain controller. Then run the wizard again by going to *Start > Programs > Fortinet > Fortinet Single Sign On > Install DC Agent*.
- 5 Optionally, select users that you do not want monitored. These users will not be able to authenticate to FortiGate units using FSSO. You can also do this later. See [“Configuring Fortinet Single Sign On with Windows AD” on page 1295](#).
- 6 Select *Next*.
- 7 Optionally, clear the check boxes of domain controllers on which you do not want to install the DC Agent.
- 8 Select the *Working Mode* as DC Agent Mode. While you can select Polling Mode here, in that situation you would not be installing a DC Agent. For more information, see [“DC Agent mode” on page 1285](#) and [“Polling mode” on page 1286](#).
- 9 Select *Next*.
- 10 Select Yes when the wizard requests that you reboot the computer.



If you reinstall the FSSO software on this computer, your FSSO configuration is replaced with default settings.

If you want to create a redundant configuration, repeat the procedure [“To install the Collector agent” on page 1292](#) on at least one other Windows AD server.



When you start to install a second Collector agent, cancel the Install Wizard dialog appears the second time. From the configuration GUI, the monitored domain controller list will show your domain controllers un-selected. Select the ones you wish to monitor with this Collector agent, and select *Apply*.

Before you can use FSSO, you need to configure it on both Windows AD and on the FortiGate units. The next section and [“Configuring FSSO on FortiGate units” on page 1308](#) will help you accomplish these two tasks.

Installing FSSO without using an administrator account

Normally when installing services in Windows, it is best to use the Domain Admin account, as stated earlier. This ensures installation goes smoothly and uninterrupted, and when using the FSSO agent there will be no permissions issues. However, it is possible to install FSSO with a non-admin account in Windows 2003 or 2008 AD.



The following instructions for Windows 2003 are specific to the event log polling mode only. Do not use this procedure with other FSSO configurations.

Windows 2003

There are two methods in Windows 2003 AD for installing FSSO without an admin account — add the non-admin user to the security log list, and use a non-admin account with read-only permissions. A problem with the first method is that full rights (read, write, and clear) are provided to the event log. This can be a problem when audits require limited or no write access to logs. In those situations, the non-admin account with read-only permissions is the solution.

To add the non-admin user account to the Windows 2003 security log list

- 1 Go to *Default Domain Controller Security Settings > Security Settings > User Rights Assignment > Manage auditing and security log*.
- 2 Add the user account to this list.
- 3 Repeat these steps on every domain controller in Windows 2003 AD.
- 4 A reboot is required.

To use a non-admin account with read-only permissions to install FSSO on Windows 2003

The following procedure provides the user account specified with read only access to the Windows 2003 AD Domain Controller Security Event Log which allows FSSO to function.

- 1 Find out the SID of the account you intend to use.
Tools for this can be downloaded for free from <http://technet.microsoft.com/en-us/sysinternals/bb897417>.
- 2 Then create the permission string. For example:
(A;;0x1;;;S-1-5-21-4136056096-764329382-1249792191-1107)
A means Allow,
0x1 means Read, and
S-1-5-21-4136056096-764329382-1249792191-1107 is the SID.
- 3 Then, append it to the registry key
- 4 HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Eventlog\Security\CustomSD.
- 5 Repeat these steps on every domain controller in Windows 2003 AD.
- 6 A reboot is required.

Windows 2008

In Windows 2008 AD, if you do not want to use the Domain Admin account then the user account that starts the FSSO agent needs to be added to the Event Log Readers group.

When the user is added to the Event Log Readers group, that user is now allowed to have read only access to the event log and this is the minimal rights required for FSSO to work.

Updating FSSO with Windows AD

After FSSO is installed on your network, you may want to upgrade to a newer version. The following procedure helps ensure you have a trouble free upgrade. How you update FSSO depends on if you are using polling mode or DCAgent mode.

For polling mode, since there are no DC agents you only need to upgrade the Collector. However in DCAgent mode, each DC Agent must be updated as well.

To update FSSO in DC Agent mode

- 1 Go to the system32 directory on all DC's and rename the `dcagent.dll` file to `dcagent.dll.old`.

This ensures that when the upgrade is pushed to the DC it does not overwrite the old file. If there are any problems this makes it easy to revert to the old version.

- 2 Run the FSSO setup .exe file to update the collector. When this is completed, ignore any reboot message.
- 3 Go to *Programs > Fortinet > Fortinet Single Sign On > Install DC Agent* and push the DC agent out to all servers. All DC's will now need to be rebooted so that the new dll file is loaded.
- 4 After the reboot, go to all DC's and delete the `dcagent.dll.old` files.

Configuring Fortinet Single Sign On with Windows AD

On the FortiGate unit, security policies control access to network resources based on user groups. With Fortinet Single Sign On, this is also true but each FortiGate user group is associated with one or more Windows AD user groups. This is how Windows AD user groups get authenticated in the FortiGate security policy.

Fortinet Single Sign On sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units.

To avoid this problem, you can configure the Fortinet Single Sign On Collector agent to send logon information only for groups named in the FortiGate unit's security policies. See [“Configuring FortiGate group filters” on page 1301](#).

On each server with a Collector agent, you will be

- [Configuring Windows AD server user groups](#)
- [Configuring Collector agent settings](#), including the domain controllers to be monitored
- [Configuring Directory Access settings](#)
- [Configuring the Ignore User List](#)
- [Configuring FortiGate group filters](#) for each FortiGate unit
- [Configuring FSSO ports](#)
- [Configuring alternate user IP address tracking](#)



In some environments where user IP addresses change frequently, it might be necessary to configure the alternate IP address tracking method. For more information, see [“Configuring alternate user IP address tracking” on page 1303](#).

Configuring Windows AD server user groups

FortiGate units control network resource access at the group level. All members of a user group have the same network access as defined in FortiGate security policies.

You can use existing Windows AD user groups for authentication to FortiGate units if you intend that all members within each group have the same network access privileges.

Otherwise, you need to create new user groups for this purpose.



If you change a user's group membership, the change does not take effect until the user logs off and then logs on again.



The FSSO Agent sends only Domain Local Security Group and Global Security Group information to FortiGate units. You cannot use Distribution group types for FortiGate access. No information is sent for empty groups.

Refer to Microsoft documentation for information about creating and managing Windows AD user groups.

Configuring Collector agent settings

You need to configure which domain controllers the Collector agent will use and which domains to monitor for user logons. You can also alter default settings and settings you made during installation. These tasks are accomplished by configuring the FSSO Collector Agent, and selecting either Apply to enable the changes.

At any time to refresh the FSSO Agent settings, select Apply.

To configure the Collector agent

- 1 From the Start menu, select *Programs > FortiNet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.

- 2 Enter the following information and then select *Save&Close*.

Monitoring user logon events	Select to automatically authenticate users as they logon to the Windows domain.
Support NTLM authentication	Select to facilitate logon of users who are connected to a domain that does not have the FSSO DC Agent installed.
Collector Agent Status	Shows RUNNING when Collector agent is active.
Listening ports	You can change FSSO Collector Agent related port numbers if necessary.
FortiGate	TCP port for FortiGate units. Default 8000.
DC Agent	UDP port for DC Agents. Default 8002.
Logging	
Log level	Select the minimum severity level of logged messages.
Log file size limit (MB)	Enter the maximum size for the log file in MB.
View Log	View all Fortinet Single Sign On agent logs.
Log logon events in separate logs	Record user login-related information separately from other logs. The information in this log includes data received from DC agents user logon/logoff information workstation IP change information data sent to FortiGate units
View Logon Events	If <i>Log logon events in separate logs</i> is enabled, you can view user login-related information.
Authentication	
Require authenticated connection from FortiGate	Select to require the FortiGate unit to authenticate before connecting to the Collector agent.
Password	Enter the password that FortiGate units must use to authenticate. The maximum password length is 16 characters. The default password is "fortinetcanada".

Timers		
	Workstation verify interval (minutes)	<p>Enter the interval in minutes at which the Fortinet Single Sign On Collector agent connects to client computers to determine whether the user is still logged on. The default is every 5 minutes. The interval may be increased if your network has too much traffic.</p> <p>Note: This verification process creates security log entries on the client computer.</p> <p>If ports 139 or 445 cannot be opened on your network, set the interval to 0 to prevent checking. See “Configuring FSSO ports” on page 1302.</p>
	Dead entry timeout interval	<p>Enter the interval in minutes after which Fortinet Single Sign On Agent purges information for user logons that it cannot verify. The default is 480 minutes (8 hours).</p> <p>Dead entries usually occur because the computer is unreachable (such as in standby mode or disconnected) but the user has not logged off. A common reason for this is when users forget to logoff before leaving the office for the day.</p> <p>You can also prevent dead entry checking by setting the interval to 0.</p>
	IP address change verify interval	<p>Fortinet Single Sign On Agent periodically checks the IP addresses of logged-in users and updates the FortiGate unit when user IP addresses change. IP address verification prevents users from being locked out if they change IP addresses, as may happen with DHCP assigned addresses.</p> <p>Enter the verification interval in seconds. The default is 60 seconds. You can enter 0 to prevent IP address checking if you use static IP addresses.</p> <p>This does not apply to users authenticated through NTLM.</p>
	Cache user group lookup result	<p>Enable caching.</p> <p>Caching can reduce group lookups and increase performance.</p>
	Cache expire in (minutes)	<p>Fortinet Single Sign On Agent caches group information for logged-in users.</p> <p>Enter the duration in minutes after which the cache entry expires. If you enter 0, the cache never expires.</p> <p>A long cache expire interval may result in more stale user group information. This can be an issue when a user's group information is changed.</p>
	Clear Group Cache	<p>Clear group information of logged-in users.</p> <p>This affects all logged-in users, and may force them to re-login.</p>

Common Tasks		
	Show Service Status	View information about the status of the Collector agent and connected FortiGate units. See “Viewing FSSO component status” on page 1303 .
	Show Monitored DCs	Shows detailed information about connected Domain Controller agents. Use the <i>Select DC to Monitor</i> button to select domain controllers to monitor and choose <i>Working Mode</i> . See “Selecting Domain Controllers and working mode for monitoring” on page 1304 .
	Show Logon Users	View a list of currently logged-in users. Select the column headers to sort the list.
	Select Domains to Monitor	Select this button to remove domains that you do not want to monitor. From the <i>Domain Filter</i> dialog box that displays, clear check boxes for unwanted domains and select <i>OK</i> .
	Set Directory Access Information	See “Configuring Directory Access settings” on page 1299 .
	Set Group Filters	Configure group filtering for each FortiGate unit. See “Configuring FortiGate group filters” on page 1301 .
	Set Ignore User List	Exclude users such as system accounts that do not authenticate to any FortiGate unit. See “Configuring the Ignore User List” on page 1300 .
	Sync Configuration With Other Agents	Copy this Collector agent's Ignore User List and Group Filters to the other Collector agents to synchronize the configuration. You are asked to confirm synchronization for each Collector agent.
	Export Configuration	Export Fortinet Single Sign On Agent configuration to a text file. The file is named <code>saved_config.txt</code> and is saved in the Fortinet Single Sign On Agent program directory.
Save & Close		Save the modified settings and exit.
Apply		Apply changes now.
Default		Change all settings to the default values.
Help		View the online Help.



To view the version and build number information for your FSSO Collector Agent configuration, selecting the Fortinet icon in the upper left corner of the Collector agent Configuration screen and select *About Fortinet Single Sign On Agent configuration*.

Configuring Directory Access settings

The FSSO Collector Agent can access Windows Active Directory in one of two modes:

- **Standard** — the FSSO Collector Agent receives group information from the Collector agent in the *domain\user* format. This option is available on FortiOS 3.0 and later.

- **Advanced** — the FSSO Collector Agent obtains user group information using LDAP. The benefit of this method is that it is possible to nest groups within groups. This option is available on FortiOS 3.0 MR6 and later. The group information is in standard LDAP format.



If you change AD access mode, you must reconfigure your group filters to ensure that the group information is in the correct format.

To configure Directory Access settings

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Set Directory Access Information*.
The *Set Directory Access Information* dialog box opens.
- 3 From the *AD access mode* list, select either *Standard* or *Advanced*.
- 4 If you selected *Advanced* AD access mode, select *Advanced Setting* and configure the following settings and then select *OK*:

AD server address	Enter the address of your network's global catalog server.
AD server port	The default AD server port is 3268. This must match your server port.
BaseDN	Enter the Base distinguished name for the global catalog. This is the point in the tree that will be considered the starting point by default.
username	If the global catalog accepts your Fortinet Single Sign On Agent agent's credentials, you can leave these fields blank. Otherwise, enter credentials for an account that can access the global catalog.
Password	

BaseDN example

An example DN for Training Fortinet, Canada is DN = ou=training, ou=canada, dc=fortinet, dc=com. If you set the *BaseDN* to ou=canada, dc=fortinet, dc=com then when Fortinet Single Sign On Agent is looking up user credentials, it will only search the Canada organizational unit, instead of all the possible countries in the company. Its a short cut to entering less information and faster searches.

However, you may have problems if you narrow the BaseDN too much when you have international employees from the company visiting different offices. If someone from Fortinet Japan is visiting the Canada office in the example above, their account credentials will not be matched because they are in DN = ou=japan, dc=fortinet, dc=com instead of the BaseDN ou=canada, dc=fortinet, dc=com. The easy solution is to change the BaseDN to simply be dc=fortinet, dc=com. Then any search will check all the users in the company.

Configuring the Ignore User List

The Ignore User List excludes users that do not authenticate to any FortiGate unit, such as system accounts. The logons of these users are not reported to FortiGate units. This reduces the amount of required resources on the FortiGate unit especially when logging logon events to memory.

To configure the Ignore User List

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Set Ignore User List*.
The current list of ignored users is displayed. To view ignored usernames, expand each domain.
- 3 Do any of the following:
 - To remove a user from the list, select the check box beside the username and then select *Remove*. The user's login is no longer ignored.
 - To add users to be ignored, select *Add*, select the check box beside each required username, and then select *Add*.
- 4 Select *OK*.

Configuring FortiGate group filters

FortiGate group filters actively control which user logon information is sent to each FortiGate unit. You need to configure the group filter list so that each FortiGate unit receives the correct user logon information for the user groups that are named in its security policies. These group filters help limit the traffic sent to the FortiGate unit, and help limit the logon events logged.

The maximum number of Windows AD user groups allowed on a FortiGate depends on the model. Low end models up to 300A support 256 Windows AD user groups, where mid and high end models support 1024 groups. This is per VDOM if VDOMs are enabled on the FortiGate unit.

You do not need to configure a group filter on the Collector agent if the FortiGate unit retrieves group information from Windows AD using LDAP. In that case, the Collector agent uses the list of groups you selected on the FortiGate unit as its group filter.

The filter list is initially empty. You need to configure filters for your FortiGate units using the Add function. At a minimum, create a default filter that applies to all FortiGate units without a defined filter.



If no filter is defined for a FortiGate unit and there is no default filter, the Collector agent sends all Windows AD group and user logon events to the FortiGate unit. While this normally is not a problem, limiting the amount of data sent to the FortiGate unit improves performance by reducing the amount of memory the unit uses to store the group list and resulting logs.

To configure a FortiGate group filter

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Set Group Filters*.

The FortiGate Filter List opens. It has the following columns:

FortiGate SN	The serial number of the FortiGate unit to which this filter applies.
Description	An optional description of the role of this FortiGate unit.
Monitored Groups	The Windows AD user groups that are relevant to the security policies on this FortiGate unit.
Add	Create a new filter.

Edit	Modify the filter selected in the list.
Remove	Remove the filter selected in the list.
OK	Save the filter list and exit.
Cancel	Cancel changes and exit.

- 3 Select *Add* to create a new filter. If you want to modify an existing filter, select it in the list and then select *Edit*.

- 4 Enter the following information and then select *OK*.

Default filter	Select to create the default filter. The default filter applies to any FortiGate unit that does not have a specific filter defined in the list.
FortiGate Serial Number	Enter the serial number of the FortiGate unit to which this filter applies. This field is not available if Default is selected.
Description	Enter a description of this FortiGate unit's role in your network. For example, you could list the resources accessed through this unit. This field is not available if Default is selected.
Monitor the following groups	The Collector agent sends to the FortiGate unit the user logon information for the Windows AD user groups in this list. Edit this list using the Add, Advanced and Remove buttons.
Add	In the preceding single-line field, enter the Windows AD domain name and user group name, and then select Add. If you don't know the exact name, use the Advanced button instead. The format of the entry depends on the AD access mode (see "Configuring Directory Access settings" on page 1299): Standard: Domain\Group Advanced: cn=group, ou=corp, dc=domain
Advanced	Select <i>Advanced</i> , select the user groups from the list, and then select <i>Add</i> .
Remove	Remove the user groups selected in the monitor list.

Configuring FSSO ports

For FSSO to function properly a small number of TCP and UDP ports must be open through all firewalls on the network. There ports listed in this section assume the default FSSO ports are used.

TCP ports for FSSO agent with client computers

Windows AD records when users log on but not when they log off. For best performance, Fortinet Single Sign On Agent monitors when users log off. To do this, Fortinet Single Sign On Agent needs read-only access to each client computer's registry over TCP port 139 or 445. Open at least one of these ports — ensure it is not blocked by firewalls.

If it is not feasible or acceptable to open TCP port 139 or 445, you can turn off Fortinet Single Sign On Agent logoff detection. To do this, set the Collector agent workstation verify interval to 0. The FSSO Collector Agent assumes that the logged on computer remains logged on for the duration of the Collector agent dead entry timeout interval — by default this is eight hours.

Configuring ports on the Collector agent computer

On the computer where you install the Collector agent, you must make sure that the firewall does not block the listening ports for the FortiGate unit and the DC Agent. By default, these are TCP port 8000 and UDP port 8002. For more information about setting these ports, see [“Configuring Collector agent settings” on page 1296](#).

Configuring alternate user IP address tracking

In environments where user IP addresses change frequently, you can configure Fortinet Single Sign On Agent to use an alternate method to track user IP address changes. Using this method, Fortinet Single Sign On Agent responds more quickly to user IP address changes because it directly queries workstation IP addresses to match users and IP addresses.

This feature requires FSAE version 3.5.27 or later, Fortinet Single Sign On Agent any version, and FortiOS 3.0 MR7 or later.

To configure alternate user IP address tracking

- 1 On the computer where the Collector agent is installed, go to *Start > Run*.
- 2 Enter `regedit` or `regedt32` and select *OK*.
The Registry Editor opens.
- 3 Find the registry key
`HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\collectoragent`.
- 4 Set the `supportFSAEauth` value (dword) to `00000001`.
- 5 Close the Registry Editor.
- 6 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 7 Select *Apply*.
The Fortinet Single Sign On Agent service restarts with the updated registry settings.

Viewing FSSO component status

It is important to know the status of both your Collector agents and DC agents.

Viewing Collector agent status

Use the *Show Service Status* to view your Collector agent information in the Status window. The Status window displays:

- the version of the software
- the status of the service
- the number of connected FortiGate units
- connected FortiGate information such as serial number, IP address, and time connected

To view Collector agent status

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Show Service Status*.

The Fortinet Single Sign On Collector agent Status window opens.

Optionally select *Get NTLM statistics* in the Status window to display NTLM information such as number of messages received, processed, failed, in the queue.

Viewing DC agent status

Use the *Show Monitored DCs* to view the status of DC agents.

To view domain controller agent status

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Show Monitored DCs*.

For each DC Agent, the following information is displayed:

- IP address
- number of logon events received
- the last logon event
- when last logon was received.

To change which DC agents are monitored or change the working mode for logon event monitoring, select *Select DC to Monitor*.

Selecting Domain Controllers and working mode for monitoring

You can change which DC agents are monitored or change the working mode for logon event monitoring between DC agent mode and polling mode.

When polling mode is selected, it will poll port 445 of the domain controller every few seconds to see who is logged on.

- 1 From the Start menu select *Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent*.
- 2 In the *Common Tasks* section, select *Show Service Status*.
- 3 Select *Select DC to Monitor*.

Working Mode	
	DC Agent mode — a Domain Controller agent monitors user logon events and passes the information to the Collector agent. This provides reliable user logon information, however you must install a DC agent on every domain controller in the domain.
	Polling mode — the Collector agent polls each domain controller for user logon information. Under heavy system load this might provide information less reliably. However installing a DC agent on each domain controller is not required in this mode.

FSSO for Novell eDirectory

The components you need to install depend on whether you are installing FSSO on Windows AD or Novell eDirectory. FSSO supports the Novell E-Directory 8.8 operating system.

This section includes:

- [FSSO components for Novell eDirectory](#)
- [Installing FSSO for Novell](#)

FSSO components for Novell eDirectory

For a Novell network, there is only one FSSO component to install — the eDirectory agent. In some cases, you also need to install the Novell Client.

Installing FSSO for Novell

To install FSSO, you must obtain the FSAE_Setup_eDirectory file from the [Fortinet Support web site](#). Perform the following installation procedure on the computer that will run the eDirectory agent. This can be any server or domain controller that is part of your network.

This section includes:

- [Configuring the eDirectory agent](#)
- [Adding an eDirectory server](#)
- [Configuring a group filter](#)

To install the FSSO eDirectory agent

- 1 Create an account with administrator privileges and a password that does not expire. See Novell documentation for more information.
- 2 logon to the account that you created in Step 1.
- 3 Double-click the FSSO_Setup_edirectory.exe file.
The Fortinet eDirectory Agent InstallShield Wizard starts.
- 4 Optionally, fill in the *username* and *Organization* fields.
- 5 Select the *Anyone who uses this computer (all users)* option.
- 6 Select *Next*.
- 7 Optionally, enter any of the following information:

eDirectory Server		
	Server Address	Enter the IP address of the eDirectory server.
	Use secure connection (SSL)	Select to connect to the eDirectory server using SSL security.
	Search Base DN	Enter the base Distinguished Name for the user search.
eDirectory Authentication		
	username	Enter a username that has access to the eDirectory, using LDAP format.
	User password	Enter the password.

- 8 Select *Next*.
- 9 Select *Install*.

Configuring Fortinet Single Sign On with Novell networks

You need to configure the eDirectory agent for it to communicate with eDirectory servers. You may have provided some of this information during installation.

This section includes:

- [Configuring the eDirectory agent](#)

- [Adding an eDirectory server](#)
- [Configuring a group filter](#)

Configuring the eDirectory agent

You need to configure the eDirectory agent for it to communicate with eDirectory servers.

To configure the eDirectory agent

- 1 From the Start menu select *Programs > Fortinet > eDirectory Agent > eDirectory Config Utility*.
- 2 The eDirectory Agent Configuration Utility dialog opens. Enter the following information and select **OK**.

eDirectory Authentication	
username	Enter a username that has access to the eDirectory, using LDAP format.
User password	Enter the password.
Listening port	Enter the TCP port on which Fortinet Single Sign On Agent listens for connections from FortiGate units. The default is 8000. You can change the port if necessary.
Refresh interval	Enter the interval in seconds between polls of the eDirectory server to check for new logons. The default is 30 seconds.
FortiGate Connection Authentication	
Require authenticated connection from FortiGate	Select to require the FortiGate unit to authenticate before connecting to the eDirectory Agent.
Password	Enter the password that FortiGate units must use to authenticate. The maximum password length is 16 characters. The default password is "FortinetCanada".
User logon info search method	Select how the eDirectory agent accesses user logon information: <i>LDAP</i> or <i>Native</i> (Novell API). LDAP is the default. If you select <i>Native</i> , you must also have the Novell Client installed on the PC.
Logging	
Log level	Select <i>Debug</i> , <i>Info</i> , <i>Warning</i> or <i>Error</i> as the minimum severity level of message to log or select <i>None</i> to disable logging.
Log file size limit (MB)	Enter the maximum size for the log file in MB.
View Log	View the current log file.
Dump Session	List the currently logged-on users in the log file. This can be useful for troubleshooting.

eDirectory server list	If you specified an eDirectory server during installation, it appears in this list.
Add	Add an eDirectory server. See .
Delete	Delete the selected eDirectory server.
Edit	Modify the settings for the selected server.
Group Filter	Select the user groups whose user logons will be reported to the FortiGate unit. This is used only if user groups are not selected on the FortiGate unit.

Adding an eDirectory server

Once the eDirectory agent is configured, you add one or more eDirectory servers.

To add an eDirectory server

- 1 In the eDirectory Agent Configuration Utility dialog box (see the preceding procedure, “Configuring the eDirectory agent”), select *Add*.
- 2 The eDirectory Setup dialog box opens. Enter the following information and select OK:

eDirectory Server Address	Enter the IP address of the eDirectory server.
Port	If the eDirectory server does not use the default port 389, clear the Default check box and enter port number.
Use default credential	Select to use the credentials specified in the eDirectory Configuration Utility. See . Otherwise, leave the check box clear and enter a username and Password below.
username	Enter a username that has access to the eDirectory, using LDAP format.
User password	Enter the password.
Use secure connection (SSL)	Select to connect to the eDirectory server using SSL security.
Search Base DN	Enter the base Distinguished Name for the user search.

Configuring a group filter

The eDirectory agent sends user logon information to the FortiGate unit for all user groups unless you either configure an LDAP server entry for the eDirectory on the FortiGate unit and select the groups that you want to monitor configure the group filter on the eDirectory agent.

If both the FortiGate LDAP configuration and the eDirectory agent group filter are present, the FortiGate user group selections are used.

To configure the group filter

- 1 From the Start menu select *Programs > Fortinet > eDirectory Agent > eDirectory Config Utility*.
- 2 Select *Group Filter*.

- 3 Do one of the following:
 - Enter group names, then select *Add*.
 - Select *Advanced*, select groups, and then select *Add*.
- 4 Select *OK*.

Configuring FSSO on FortiGate units

To configure your FortiGate unit to operate with either a Windows AD or a Novell eDirectory FSSO install, you

- Configure LDAP access to the Novell eDirectory or Windows AD global catalog. Skip this step if you are using FSSO Standard mode. See [“Configuring LDAP server access” on page 1308](#).
- Specify the Collector agent or Novell eDirectory agent that will provide user logon information. See [“Specifying your Collector agents or Novell eDirectory agents” on page 1310](#).
- Add Active Directory user groups to FortiGate user groups. See [“Selecting Windows user groups \(LDAP only\)” on page 1311](#).
- Create security policies for FSSO-authenticated groups. See [“Creating Fortinet Single Sign-On \(FSSO\) user groups” on page 1312](#) and [“Creating security policies” on page 1312](#).
- Optionally, specify a guest protection profile to allow guest access. See [“Enabling guest access through FSSO security policies” on page 1315](#).

Configuring LDAP server access

LDAP access is required if your network has a Novell eDirectory agent or a Collector agent using Windows Advanced AD access mode. If you are using FSSO Standard mode, go to [“Specifying your Collector agents or Novell eDirectory agents” on page 1310](#).

The LDAP configuration on the FortiGate unit not only provides access to the LDAP server, it sets up the retrieval of Windows AD user groups for you to select in FSSO. The LDAP Server configuration (in *User > Remote > LDAP*) includes a function to preview the LDAP server’s response to your distinguished name query. If you already know the appropriate Distinguished Name (DN) and User DN settings, you may be able to skip some of the following steps.

- 1 Go to *User > Remote > LDAP* and select *Create New*.
- 2 Select the *Query distinguished name* button to the right of the *Distinguished Name* field.

A new window opens.
- 3 If more than one name is listed, you might need to explore each name following the steps below to determine which one is relevant to your needs.
- 4 Copy the name string to the *Distinguished Name* field and select *OK*.

This closes the window and copies the name string to the *Distinguished Name* field of the LDAP Server configuration.
- 5 Set *Bind Type* to *Regular*.
- 6 In the *User DN* field, enter the administrative account name that you created for FSSO.

For example, if the account is FSSO_Admin, enter “cn=FSSO_Admin,cn=users”.

- 7 Make sure that the *User DN* entry ends with a comma and append the string from the *Distinguished Name* field to the end of it.

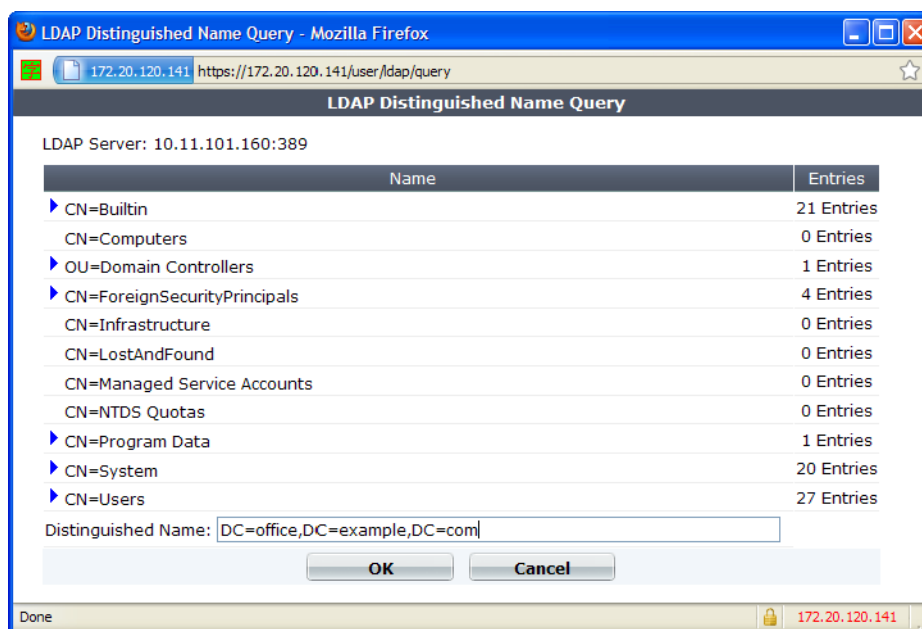
Example: cn=FSSO_Admin,cn=users,dc=office,dc=example,dc=com

- 8 Enter the administrative account password in the *Password* field.

- 9 Select the *Query distinguished name* button again.

The LDAP Distinguished Name Query window opens:

Figure 111: Authenticated DN query



You can expand any of the DNs that contain entries. When you select an expandable DN, the *Distinguished Name* field is updated. Look for the DN that contains the users or groups whose logon you want to monitor.

- 10 Select the DN that you want to monitor and then select **OK**.

This closes the window and updates the *Distinguished Name* field of the LDAP Server configuration with the selected Domain Name Identifier (DNI).

- 11 Check the following fields and select **OK**:

Name	Enter a name to identify the LDAP server.
Common Name Identifier	The default common name identifier is <code>cn</code> . This is correct for most LDAP servers. However some servers use other identifiers such as <code>uid</code> .
Secure Connection	Do not select. The Collector agent does not support secure connection.

To configure LDAP for FSSO - CLI example

```
config user ldap
edit "ADserver"
set server "10.11.101.160"
set cnid "cn"
set dn "cn=users,dc=office,dc=example,dc=com"
```

```

set type regular
set username
    "cn=administrator,cn=users,dc=office,dc=example,dc=com"
set password set_a_secure_password
next
end

```

Specifying your Collector agents or Novell eDirectory agents

You need to configure the FortiGate unit to access at least one Collector agent or Novell eDirectory agent. You can specify up to five servers on which you have installed a Collector or eDirectory agent. The FortiGate unit accesses these servers in the order that they appear in the list. If a server becomes unavailable, the next one in the list is tried.

To specify Collector agents - web-based manager

- 1 Go to *User > FSSO > FSSO Agent* and select *Create New*.
- 2 Enter a Name for the Windows AD server. This name appears in the list of Windows AD servers when you create user groups.
- 3 Enter the following information for each of up to five collector agents and select *OK*:

FSSO Agent IP/Name	Enter the IP address or the name of the server where this agent is installed. Maximum name length is 63 characters.
Port	Enter the TCP port used for FSSO. You must enter the port number for the server. This must be the same as the FortiGate listening port specified in the Novell eDirectory or Collector agent configuration. TCP port 8000 is used by default. See “Configuring Collector agent settings” on page 1296 .
Password	Enter the password for the Collector agent or eDirectory agent. For the Collector agent, this is required only if you configured the agent to require authenticated access.
LDAP Server	For Novell eDirectory, enable. For Windows AD, enable if the Collector agent is configured to use Advanced AD access mode. Select the LDAP server you configured previously. See “Configuring LDAP server access” on page 1308 .

To specify the FSSO Collector agent - CLI

```

config user fsso
edit WinGroups
set ldap-server ADserver
set password ENC
    G7GQV7NEqilCM9jKmVmJJFVvhQ2+wtNEe9T0iYA5Sa+EQT2J8zhOrbkJFD
    r0RmY3c4LaoXdsoBczAldONmcGfthTxxwGsigzGpbJdC7lspFlQYtj
set server 10.11.101.160
set port 8000
end

```

Selecting Windows user groups (LDAP only)

If the Collector agent uses Advanced AD access mode, the FortiGate unit obtains user group information using LDAP. You need to select the Windows user groups that you want to monitor. These user group names are then available to add to FSSO user groups.

To select Windows user groups - web-based manager

- 1

Go to *User > FSSO > FSSO Agent*.
The list of FSSO agent servers is displayed.
- 2

Select the *Edit Users/Groups* icon.
The FortiGate unit performs an LDAP query and displays the result.
- 3

Select the check boxes of the user groups that you want to monitor and then select *OK*.
You can also use the *Add User/Group* icon to select a group by entering its distinguished name.

Viewing information imported from the Windows AD server

You can view the domain and group information that the FortiGate unit receives from the AD Server. Go to *User > FSSO > FSSO Agent*. The display differs for Standard and Advanced AD access mode.

Figure 112: List of groups from Active Directory server (Standard AD access mode)

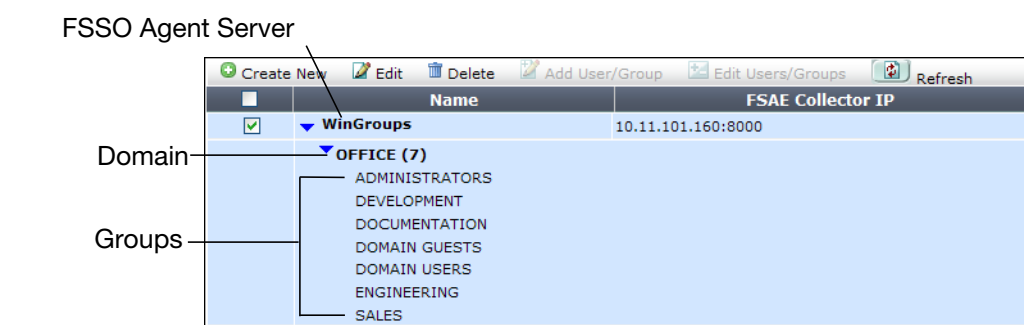
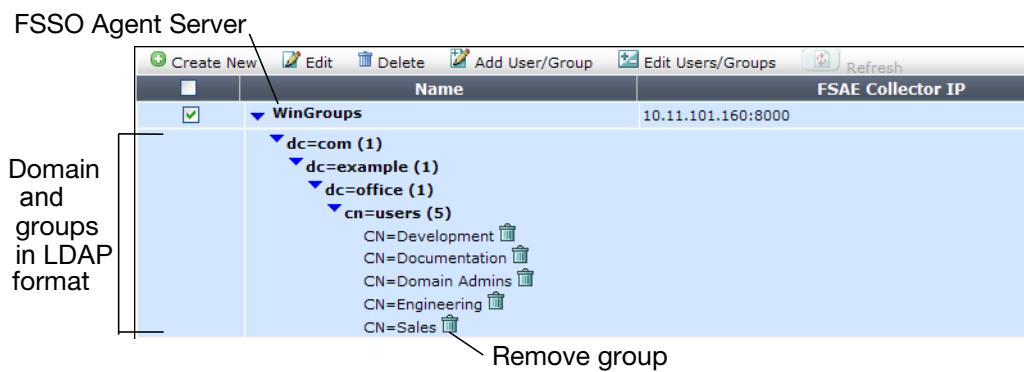


Figure 113: List of monitored groups (Advanced AD access mode)



Create New	Add a new FSSO agent server.
Name	
Server	The name defined for the FSSO agent server.

	Domain	Domain name imported from the FSSO agent server.
	Groups	The group names imported from the FSSO agent server.
	FSSO Collector IP	The IP address of the FSSO agent on the server
	Delete icon	Delete this server definition.
	Edit icon	Edit this server definition.
	Refresh icon	Get user group information from the FSSO agent server.
	Add User/Group	Add a user or group to the list. You must know the distinguished name for the user or group. This is available for Windows AD in Advanced AD access mode only.
	Edit Users/Groups	Select users and groups to add to the list. See “Selecting Windows user groups (LDAP only)” on page 1311 . This is available in Advanced AD access mode only.

Creating Fortinet Single Sign-On (FSSO) user groups

You cannot use Windows or Novell groups directly in FortiGate security policies. You must create FortiGate user groups of the FSSO type and add Windows or Novell groups to them.

To create a user group for FSSO authentication - web-based manager

- 1 Go to *User > User Group*.
- 2 Select *Create New*.
The New User Group dialog box opens.
- 3 In the *Name* box, enter a name for the group, *FSSO_Internet_users* for example.
- 4 In *Type*, select *Fortinet Single Sign-On (FSSO)*.
- 5 From the *Available Members* list, select the required FSSO groups.
Using the CTRL or SHIFT keys, you can select multiple groups.
- 6 Select the green right arrow button to move the selected groups to the *Members* list.
- 7 Select *OK*.

To create the FSSO_Internet-users user group - CLI

```
config user group
edit FSSO_Internet_users
set group-type fsso-service
set member
    CN=Engineering,cn=users,dc=office,dc=example,dc=com
    CN=Sales,cn=users,dc=office,dc=example,dc=com
end
```

Creating security policies

Policies that require FSSO authentication are very similar to other security policies. Using identity-based policies, you can configure access that depends on the FSSO user group. This allows each FSSO user group to have its own level of access to its own group of services

In this situation, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fsso_group`, and the FSSO RADIUS server is `fsso_rad_server`.

To configure an FSSO authentication security policy - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information.

Source Interface/Zone	port1
Source Address	company_network
Destination Interface/Zone	port2
Destination Address	all
Action	ACCEPT
Enable NAT	enabled

- 4 Select *Enable identity-based Policy*.
- 5 Select *Add* to add groups of users to this authentication policy.
- 6 Select the `fsso_group`, and the `FSSO_Guest_users` usergroups in the *Available User Groups* list and move them to the *Selected User Groups* list.

`FSSO_Guest_users` is a default user group enabled when FSSO is configured. It allows guest users on the network who do not have FSSO account to still authenticate and have access to network resources. See [“Enabling guest access through FSSO security policies” on page 1315](#).
- 7 Select HTTP, HTTPS, FTP, and Telnet for in the *Available Services* list, and move them to the *Selected Services* list.
- 8 Select always for the *Schedule*.
- 9 Enable *Log Allowed Traffic*.

Logging FSSO logon events helps troubleshoot any FSSO related issues.
- 10 Select UTM, and enable default AntiVirus, IPS, Web Filter, an Email filter.
- 11 Select OK.

A new line of information will appear in the identity-based policy table. The table lists the ID, user group or groups, the service or services, schedule, UTM, and logging selected for the rule. Use this display to verify your information was entered correctly.
- 12 Select *Fortinet Single Sign-On (FSSO)*.
- 13 Optionally select *Customize Authentication messages* to change the default authentication messages to suit example.com’s company design and policies.
- 14 Select OK.

- 15 Ensure the FSSO authentication policy is at the top of the list so it will be attempted to be matched before any other policy.

To create a security policy for FSSO authentication - CLI

```
config firewall policy
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr company_network
    set dstaddr all
    set action accept
    set identity-based enable
    set nat enable
    config identity-based-policy
      edit 1
        set schedule any
        set groups company_network FSSO_guest_users
        set service HTTP HTTPS FTP TELNET
      end
    end
  end
```

Here is an example of how this FSSO authentication policy is used. Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user's identity and credentials, and then if everything is verified the user is allowed access to the Internet.

Users belonging to multiple groups

Before FSSO 4.0 MR3, if a user belonged to multiple user groups, the first security policy to match any group that user belonged too was the only security policy applied. If that specific group did not have access to this protocol or resource where another group did, the user was still denied access. For example, `test_user` belongs to `group1` and `group2`. There are two FSSO authentication policies — one matches `group1` to authenticate FTP traffic and one matches `group2` to authenticate email traffic. The `group1` policy is at the top of the list of policies. If `test_user` wants to access an email server, the first policy encountered for a group `test_user` belongs to is the `group1` policy which does not allow email access and `test_user` is denied access. This is despite the next policy allowing access to email. If the order was reversed in this case, the traffic would be matched and the user's traffic would be allowed through the firewall. However if the policy order was reversed, FTP traffic would not be matched.

As of FSSO 4.0 MR3, if a user belongs to multiple groups multiple then attempts to match the group are attempted if applicable. Using the above example, when the attempt to match the `group1` policy is made and fails, the next policy with a group that `test_user` is a member of is attempted. In this case, the next policy is matched and access is granted to the email server.

When configuring this example the only difference between the policies is the services that are listed and the FSSO user group name.

Authenticating through multiple groups allows administrators to assign groups for specific services, and users who are members of each group have access to those services. For example there could be an FTP group, an email group, and a Telnet group.

Resolve usernames Using FSSO Agent

When configuring a security policy, there is an option to enable the feature *Resolve usernames Using FSSO Agent* without configuring an identity-based policy.

If an FSSO server and user group are configured, but no identity-based policy is enabled then logon events do not generate a log entry — essentially the logon event goes unnoticed by the FortiGate unit.

If this option is enabled then when there is no identity-based policy configured, the logon event will still be logged on the traffic log along with the username. The username will be resolved from the FSSO information.

Enabling guest access through FSSO security policies

You can enable guest users to access FSSO security policies. Guests are users who are unknown to the Windows AD or Novell network and servers that do not logon to a Windows AD domain.

To enable guest access in your FSSO security policy, add an identity-based policy assigned to the built-in user group `FSSO_Guest_Users`. Specify the services, schedule and protection profile that apply to guest users — typically guests receive reduced access to a reduced set of services. See [“Creating security policies” on page 1312](#).

FortiOS FSSO log messages

There are two types of FortiOS log messages — firewall and event. FSSO related log messages are generated from authentication events. These include user logon and log off events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues. For more information on firewall logging, see [“Enabling security logging” on page 1252](#). For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

This section includes:

- [Enabling authentication event logging](#)
- [Viewing FSSO log messages](#)

Enabling authentication event logging

For the FortiGate unit to log events, that specific type of event must be enabled under logging.

When VDOMs are enabled certain options may not be available, such as CPU and memory usage events. You can enable event logs only when you are logged on to a VDOM; you cannot enable event logs globally.

To ensure you log all the events need, set the minimum log level to Notification or Information. Firewall logging requires Notification as a minimum. The closer to Debug level, the more information will be logged. While this extra information is useful, you must

To enable event logging

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 In the *Event Log Settings* section, select the check box beside *Enable*.

- 3 You must select the following events:

System activity event	All system-related events, such as ping server failure and gateway status.
Admin event	All administration events, such as user logins, resets, and configuration updates.
Firewall authentication event	All firewall-related events, such as user authentication.
CPU & memory usage (every 5 min)	Real-time CPU and memory events only, at 5-minute intervals.

Optionally you can enable any or all of the other logging event options.

- 4 Select *Apply*.
5 Select *OK*.

Viewing FSSO log messages

If you have configured logging to a FortiAnalyzer unit, FortiGate unit's local disk or system memory, you can view log messages from either the web-based manager or CLI. The following procedures explain how to view log messages from the Log Access menu in the web-based manager, and how to view log messages from within the CLI.

Figure 114 on page 1317 shows the event log filtered for authentication log entries, and the display columns customized.

Table 83 on page 1317 lists the log message IDs and descriptions for FSSO related log messages. Use this table to help identify FSSO log messages on the FortiGate unit. Note that for the log messages to be displayed, the Minimum Log Level must be set to the Severity or lower for the log entry to be logged.



When viewing log messages in the Raw format in *Memory*, the ten-digit log ID number is used; however, when viewing the same log messages, in Raw format, in *Disk*, the five-digit log ID number is used (except for traffic logs which have only one-digit log IDs). This five-digit log identification number is used because of log size reduction that occurred in FortiOS 4.0 MR1.

To view log messages - web-based manager

- Go to *Log&Report > Log & Archive Access*.
- Select the log menu that you want to view log messages in.
For example, the attack log messages in *Log&Report > Log & Archive Access > Traffic Log*.
- Within the page, use any one of the following to view each log message:
 - Download Raw Log* – downloads the log file to your PC. *Column Settings* – customize what columns display on the page.
 - Filter Settings* – filter the information within the page.
 - Detailed Information* – display the log table on the right side of the page, at the bottom (default), or hide the log table.

For more on viewing logs, see the [FortiOS Handbook Log and Reporting chapter](#).

- 4 To display current log messages on the page, select *Refresh*.

To view log messages - CLI

- 1 Enter the following to configure how the log messages will be displayed, as well as what log messages you want to display:

```
execute log filter category <category_number>
execute log filter start-line <line_number>
execute log filter view-lines <lines_per_view>
```
- 2 Enter the following to display the logs messages within the CLI:

```
execute log display
```
- 3 Log messages appear and stop when the maximum number of view-lines is reached.

Figure 114: Authentication log messages

#	Date	Time	Level	Sub Type	ID	Action	Message	User Interface
1	2011-05-02	11:51:47	notice	auth	43011	authentication	User from 172.20.120.230 was timed out	test(172.20.120.230)
2	2011-05-02	11:43:31	notice	auth	43008	authentication	User test succeeded in authentication	HTTPS(172.20.120.230)
3	2011-05-02	11:43:23	notice	auth	43009	authentication	User techdoc failed in authentication	HTTPS(172.20.120.230)

Date	2011-05-02	Time	11:51:47
Level	notice	Sub Type	auth
ID	43011	Virtual Domain	root
Src	172.20.120.230	Dst	N/A
User	test	Group	testee
Policy ID	4	User Interface	test(172.20.120.230)
Action	authentication	Status	timed_out
Reason	Authentication timed out	Message	User from 172.20.120.230 was timed out

Table 83: List of FSSO related log messages

Message ID	Severity	Description
43008	Notification	Authentication was successful
43009	Notification	Authentication session failed
43010	Warning	Authentication locked out
43011	Notification	Authentication timed out
43012	Notification	FSSO authentication was successful
43013	Notification	FSSO authentication failed
43014	Notification	FSSO user logged on
43015	Notification	FSSO user logged off
43016	Notification	NTLM authentication was successful
43017	Notification	NTLM authentication failed

For more information on logging, see the [FortiOS Handbook Log and Reporting chapter](#).

Testing FSSO

Once FSSO is configured, you can easily test to ensure your configuration is working as expected. For additional FSSO testing, see “[Troubleshooting FSSO](#)” on page 1318.

- 1 Logon to one of the stations on the FSSO domain, and access an Internet resource.
- 2 Connect to the CLI of the FortiGate unit, and if possible log the output.

- 3 Enter the following command:

```
diagnose debug authd fsso list
```

- 4 Check the output. If FSSO is functioning properly you will see something similar to the following:

```
----FSSO logons----
IP: 192.168.1.230  User: ADMINISTRATOR  Groups: VLAD-AD/DOMAIN
        USERS
IP: 192.168.1.240  User: ADMINISTRATOR  Groups: VLAD-AD/DOMAIN
        USERS
Total number of users logged on: 2
----end of FSSO logons----
```

The exact information will vary based on your installation.

- 5 Check the FortiGate event log, for FSSO-auth action or other FSSO related events with FSSO information in the message field. For a list of FSSO log message IDs, see [Table 83 on page 1317](#).
- 6 To check server connectivity, run the following commands from the CLI:

```
FGT# diagnose debug enable
FGT# diagnose debug authd fsso server-status
FGT# Server Name          Connection Status
-----
SBS-2003                  connected
```

Troubleshooting FSSO

When installing, configuring, and working with FSSO some problems are quite common. A selection of these problems follows including explanations and solutions.

Some common Windows AD problems include:

- [General troubleshooting tips for FSSO](#)
- [User status “Not Verified” on the Collector agent](#)
- [After initial configuration, there is no connection to the Collector agent](#)
- [FortiGate performance is slow on a large network with many users](#)
- [Users from the Windows AD network are not able to access the network](#)
- [Users on a particular computer \(IP address\) can not access the network](#)
- [Guest users do not have access to network](#)
- [Can’t find the DCagent service](#)
- [User logon events not received by FSSO Collector agent](#)
- [User list from Windows AD is empty](#)
- [Mac OS X users can’t access external resources after waking from sleep mode](#)

General troubleshooting tips for FSSO

The following tips are useful in many FSSO troubleshooting situations.

- To help locate the problem, configure a sniffer policy to capture FSSO logon messages along with other information.

If FSSO is in use the log messages captured by a sniffer policy will include a user name if the IP address in the log message corresponds to the IP address of a user who has been authenticated with FSSO.

- Ensure all firewalls are allowing the FSSO required ports through.
FSSO has a number of required ports that must be allowed through all firewalls or connections will fail. These include: ports 139, 389 (LDAP), 445, 636 (LDAP), 8000, and 8002.
- Ensure the Collector agent has at least 64kbps bandwidth to the FortiGate unit.
If not the Collector agent does not have this amount of bandwidth, information FSSO information may not reach the FortiGate unit resulting in outages. The best solution is to configure traffic shaping between the FortiGate unit and the Collector agent to ensure that minimum bandwidth is always available.

User status “Not Verified” on the Collector agent

When selecting “Show logon Users” in the Collector agent, some users may have their status set as “Not Verified”.

The Collector agent receives logon events for users from the DC agents, but Windows does not generate log out events. As such, the Collector agent needs to verify that the user is still logged on by checking the registry on that host.

If the Collector agent cannot connect to the host on ports 139 and 445 to perform this check, the host status is set to “Not Verified” and a log entry will be added to the Collector agent logs:

```
"01/01/2010 01:23:45 [ 1884] name_ip_match: failed to connect to  
workstation: <Workstation Name> (192.168.1.1) "
```

Solution

There are a few things that can cause the Collector agent not to be able to connect to the user's work station. Below is a list of the most common causes:

- Most commonly, a host firewall on the user's workstation or a router on the network prevents remote access on ports 139 and/or 445. Try opening the ports on the host firewall.
- If the remote registry service is not running on the user's workstation, the Collector agent will not be able to connect to the registry remotely. Make sure the remote registry service is running.
- This problem may also be caused by [a known MS upgrade issue](#).

Using Regedit.exe, edit

```
"HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\  
Control\SecurePipeServers", set permissions for winreg and allow Local  
Service with R and W permissions.
```

After initial configuration, there is no connection to the Collector agent

The Collector Agent has been configured but now cannot be contacted. This may be a regular connectivity problem. The section *Troubleshooting Connectivity* in the [FortiOS Handbook Troubleshooting chapter](#) will help locate and identify any network problems. Other solutions specific to FSSO are listed here.

Solution

If there are no network problems that can be identified, try the following solutions.

- The Windows AD network must be configured before configuring the FortiGate unit. This includes the domain controller agents, and Collector agents.
- Ensure the DC agents point to the correct collector agent port and IP address.

- Ensure that TCP port 8000, and UDP port 8002 are not blocked.
- FSSO is very dependent on DNS, ensure the forward DNS zone has no stale records and after adding it to the domain if the DNS entry is not in the zone add it.
- An error in the DNI field on the FortiGate unit will prevent connections. Select the browse button next to the field to confirm it can connect correctly to the Windows AD server and return information. See
- If the secure check box is selected, ensure that LDAP v3 is being used since earlier LDAP does not support secure TLS connections.
- Ensure that the default LDAP ports are not being blocked on the network. These ports include port 389, and port 636. If you change the default ports, ensure both the FortiGate unit and the Windows AD server are using the same port numbers and that those ports are allowed through all firewalls on your network.
- If you are using FSSO in polling mode, ensure that port 445 is not blocked by firewalls.

Collector Agent service freezing and shutting down

FSSO problem.

Solution

-

FortiGate performance is slow on a large network with many users

FSSO sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units. Logon tracking is logged to memory, and may reduce performance in extreme situations.

To avoid this problem, you can configure the Collector agent to send logon information only for groups named in the FortiGate unit's security policies. Also you can configure the Ignore User list on the FortiGate unit to avoid tracking unnecessary logons.

Also logging to memory can consume large amounts of FortiGate system memory. To lessen the memory used, change the logging from the default level of Information to a less frequent level such as Error or Warning. This results in less information being logged and frees system memory to improve overall FortiGate system performance. However, if you are trying to troubleshoot a problem one of the first things to do is to change the logging severity to Information or possibly even Debug to provide you with additional information while solving your problem.

Solution

Add users to the Ignore User list. This a best practice for admin accounts whose logon information will not be sent to the FortiGate unit. This is useful for automated accounts that may logon many times. Examples of accounts in this category include:

- IIS services
- AV
- other system accounts

For more information on configuring the Ignore users list, see [“Configuring the Ignore User List” on page 1300](#).

Users from the Windows AD network are not able to access the network

If nobody can access the network and your network has only one Collector agent, when it goes offline no users will have access. However if only some users can not access the network, it is likely that user group changes were made recently that are causing the problems.

Solutions

- If there is only one Collector agent, configure additional Collector agents in the domain to act as backups. They will provide the redundancy required if the original collector goes offline. Remember to add them to the FortiGate FSSO entry under *User > FSSO > FSSO Agent* on the web-based manager or `config user fssso` in the CLI. If the server and port for the new agent are not in the list, it will not be contacted.
- Ensure the Collector agent has at least 64kbps bandwidth to the FortiGate unit. If not, information FSSO information may not reach the FortiGate unit resulting in outages. The best solution is to configure traffic shaping between the FortiGate unit and the Collector agent to ensure that minimum bandwidth is always available.
- If some users can not connect, verify their Windows AD records to find groups in common, and investigate the state of those groups focusing on any recent changes. It may be a group or permission change is the reason.
- There may be a problem with the user list. See [“User list from Windows AD is empty” on page 1322](#).

Users on a particular computer (IP address) can not access the network

Windows AD Domain Controller agent gets the username and workstation where the logon attempt is coming from. If there are two computers with the same IP address and the same user trying to logon, it is possible for the authentication system to become confused and believe that the user on computer_1 is actually trying to access computer_2.

Windows AD does not track when a user logs out. It is possible that a user logs out on one computer, and immediately logs onto a second computer while the system still believes the user is logged on the original computer. While this is allowed, information that is intended for the session on one computer may mistakenly end up going to the other computer instead. The result would look similar to a hijacked session.

Solutions

- Ensure each computer has separate IP addresses.
- Encourage users to logout on one machine before logging onto another machine.
- If multiple users have the same username, change the usernames to be unique.
- Shorten timeout timer to flush inactive sessions after a shorter time.

Guest users do not have access to network

A group of guest users was created, but they don't have access.

Solution

The group of the guest users was not included in a policy, so they do not fall under the guest account. To give them access, associate their group with a security policy.

Additionally, there is a default group called `FSSO_Guest_Users`. Ensure that group is part of an identity-based security policy to allow traffic.

Can't find the DCagent service

The DCagent service can't be found in the list of regular windows services. This is because it has no associated Windows service.

Instead DCagent is really `dcagent.dll` and is located in the `Windows\system32` folder. This DLL file is loaded when windows boots up and it intercepts all logon events processed by the domain controller to send these events to the Collector agent (CA).

Solution

To verify that the DCagent is installed properly

- 1 Check that `DCagent.dll` exists in `%windir%\system32` folder.
- 2 Check that the registry key exists:
`[HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\dcagent]`
If both exist, the DCagent is properly installed.

User logon events not received by FSSO Collector agent

When a warning dialog is present on the screen on the Collector agent computer, the Collector agent will not receive any logon events. Once the dialog has been closed normal operation will resume.

If polling mode is enabled, it is possible the polling interval is too large. Use a shorter polling interval to ensure the collector agent is capturing all logon events.

If NetAPI polling mode is enabled, consider switching to Event log polling as it provides better accuracy.

User list from Windows AD is empty

FSSO server is configured. I have received a list of windows AD groups. However, a user list is empty.

Solution

There could be 2 problems:

In most cases, the FortiGate receives login information, but can't translate the Windows AD group into the protection profile. Make sure that all the required Windows AD groups are included in the FortiGate user groups and that all FortiGate user groups are included into the authentication security policy.

There may be a problem with AD FSSO service running on the Windows AD server.

To ensure the problem is on windows side

- 1 Go to *Log&Report > Log Config*.
- 2 Enable firewall authentication event logging and debug level logging on the FortiGate.
- 3 Ask one or more users to log in into windows.
- 4 Check the FortiGate logs for the logon event from the Windows AD server.

If there is no new logon event entry in the logs, the problem is with Windows side. Use MS Windows AD documentation to troubleshoot the problem.

Mac OS X users can't access external resources after waking from sleep mode

When client computers running Mac OS X (10.6.X and higher) wake up from sleep mode, the user must authenticate again to be able to access external resources. If the user does not re-authenticate, the user will maintain access to internal web sites, but will be unable to access any external resources.

This issue is caused by Mac OS X not providing sufficient information to the FSAE. This results in the FortiGate blocking access to the user because they cannot be authenticated.

Solution

The security settings on client computer(s) must be configured to require that a username and password be entered when exiting sleep mode or screen saver. With this feature enabled in Mac OS X, the FortiGate will receive the authentication information it requires to authenticate the user and allow them access.

Note that if the user reverts their settings to disable the password requirement, this will cause the issue to reappear.



Dynamic profiles and end points

This section explains how to set up the FortiOS dynamic profile and end point features to identify users and communication sessions. This section also describes configuring FortiOS Carrier HTTP header options, and end points.

This section describes:

- [Overview](#)
- [Configuring dynamic profile](#)
- [Configuring dynamic profile-based security policies](#)
- [Configuring end points](#)
- [Timeout options](#)
- [Log settings](#)
- [Troubleshooting dynamic profiles](#)

Overview

Network administrators can add customer identifying information and profile group names to their RADIUS server accounts. In response to a user connecting to the company network, if the RADIUS server successfully authenticates that user it sends a RADIUS Start record to FortiOS. In real time, FortiOS can extract identifying information and profile group names from these RADIUS Start records and match the identifying information with the customer communication session. FortiOS can then dynamically select and apply the profile group named in the RADIUS Start record to the communication session.

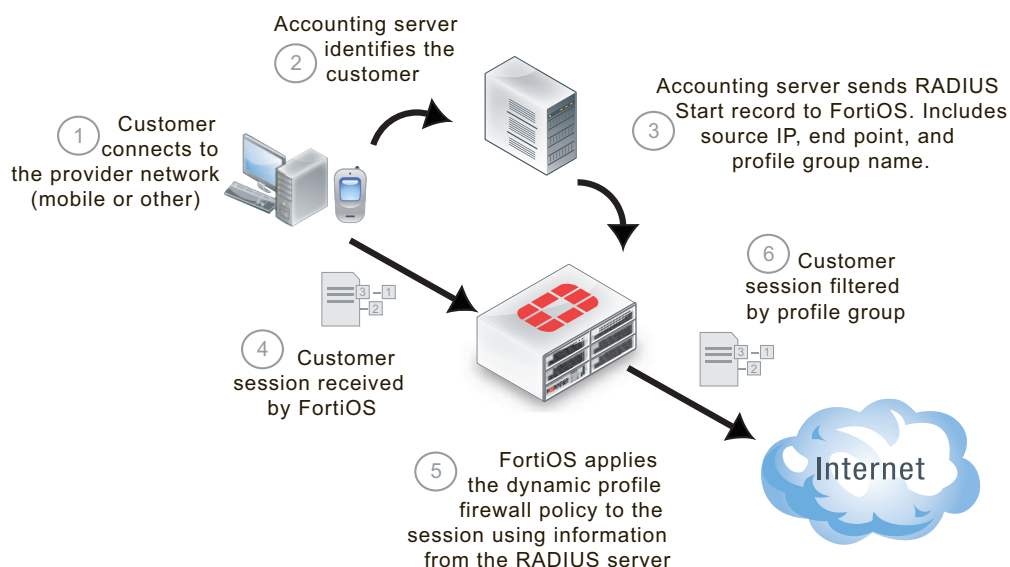
The application of parental-type controls to user sessions at a school is a useful example of the dynamic profile and multiple profile groups with different levels of access. Administrators can create profile groups that provide different levels of parental controls. Then these different levels of controls can be applied to communication sessions; the level of parental control depending on the profile group name added to the user account in the accounting system. This would allow teachers and student accounts to have different parental control settings. It would also allow teachers and students to logon to the school network from different locations, with different IP addresses, and still have the same restrictions for their level of account. It would also allow for difference between younger and older student controls, as appropriate.

Another example is on a mobile phone network with a mobile service provider. The mobile phone may access the network from many physical or network locations. With FortiOS Carrier dynamic profiles enabled on the provider network, that user will be identified as the same user and given the same privileges no matter where they logon. This makes it easier for mobile providers to track customers and provide extended services where warranted, or block access to some resources as may also be warranted. It also allows carriers to treat customers from other carriers who are using their network differently according to any agreements that might exist between carriers or even countries.



While FortiOS Carrier dynamic profiles and end points are very similar to those objects defined in FortiOS, in FortiOS Carrier there are some differences. These differences will be noted in the documentation. For definitions of carrier related terms, see the [FortiOS Handbook Carrier chapter](#).

Figure 115: Information flow between customers, the accounting system, and FortiOS



When to use FSSO or dynamic profiles

FSSO and dynamic profiles have much in common. Both use RADIUS servers for authentication and accounting. Both are single sign on (SSO) solutions. For these and other reasons its valid to wonder why use one method over the other.

There are three main points FSSO and dynamic profiles differ on.

When to use FSSO

- 1 Your large network has a complex Windows AD authentication server already configured.** FSSO is designed to work with an established Windows AD or Novell eDirectory network. Its agent software installations sit on the various servers to help send the required information to the FortiGate unit for authentication. Dynamic profiles does not require this extra software to be installed, but its performance may suffer as a result on large networks.
- 2 You have multiple authentication groups that require different levels of access.** FSSO can be configured to work with 512 groups or more (depending on your FortiGate model). Dynamic profiles allows one group per VDOM.

- 3 Dynamic profiles requires no user or user group configuration. For Dynamic profile, configure the RADIUS server, and security policy. FSSO requires more configuration, but as a result allows more flexibility as well.

End points

The term for customer identifying information is an **end point**. An end point can be any information the service provider uses to identify a customer and the device that the customer is using to connect to the network. For example, if the customer is using a mobile phone, the end point could be the phone's MSISDN number. The end point information must be included in the RADIUS Start record and must be available in the customer communication session (for example, in the HTTP header).

This is not the Endpoint Control feature referred to in other parts of FortiOS — that is a different feature for application control. End points in this chapter are used with the dynamic profile feature and RADIUS records to identify users.



In most cases, FortiOS Carrier can find the end point and IP address in customer communication sessions. An important exception is WAP traffic. Because WAP traffic may have the source IP address changed from the customer's IP address to the IP address of the WAP server, extra configuration may be required to extract the end point and source IP address from WAP traffic. See [FortiOS Handbook Carrier chapter](#) for information on configuring FortiOS Carrier for WAP traffic.

Without the end point, customers can only be identified by the IP address of the device that they are using. Because IP addresses may not be permanent or multiple users may be behind a NAT device and sharing a single IP address, the additional end point information is a more reliable and accurate way to identify individual customers. See [“Configuring end points” on page 1341](#).

Dynamic profiles and security policies

For FortiOS to use the dynamic profile, you must add a security policy that includes the feature. Then all traffic that matches that profile will be authenticated by the configured RADIUS server.



Dynamic profile and identity-based policy are mutually exclusive options. When one is enabled, the other is hidden.

You can configure a maximum of one dynamic profile security policy per VDOM.

The general steps to configure a dynamic profile security policy are:

- 1 Create a new UTM profile group with a name that matches the dynamic profile group name that will be in the RADIUS attribute.
- 2 Create a new security policy and set the required fields (source, destination, schedule, service and action) as normal.
- 3 Enable *Dynamic Profile*, and select the dynamic profile from the list.
- 4 Enable *Dynamic Profile Users Only*.
- 5 Select other optional fields as required, and select OK when done.

Dynamic profile UTM profile groups

UTM profile groups allow you to create a group of UTM services and associate them together for a specific purpose. This group is then applied to one or more security policies. The benefit of this method is that you can have multiple configurations for AntiVirus, IPS, Application control and such, but only select the versions or services you require. To make changes in the future you change that IPS configuration for example and all UTM profile groups that use that IPS configuration will be updated as well.

When creating a dynamic profile UTM profile group, the group name must match the RADIUS server profile group name that was specified in the RADIUS server configuration. This group name is used to match traffic.

To create a dynamic profile UTM profile group - web-based manager

- 1 Go to *UTM > Profile Group*.
If this options is not visible, go to *System > Admin > Settings* and enable firewall related displays.
- 2 Select *Create New*.
- 3 Enter the *Name*. For example, `my_dyn_prof_group`. The name has to be the same as the dynamic profile group name used when configuring the RADIUS server for dynamic profiles. See [“RADIUS server configuration for dynamic profiles” on page 1333](#).
- 4 Select default for the various UTM services. All entries have at least a default entry, and may have additional entries if they have been configured. For additional UTM configuration details, see the *FortiOS Handbook UTM chapter*.
- 5 Select the web and email protocols to apply dynamic profile to. This can include any or all of HTTP, HTTPS, FTP, IMAP, POP3, SMTP, IM, and NNTP. Best practices dictate that if you do not use some of these protocols, do not configure those ones. For all protocols selected, enable monitor to ensure statistics are collected for the Dashboard.
- 6 Select *OK*.

To create a dynamic profile UTM profile group - CLI

```
config firewall profile-group
edit my_dyn_prof_group
set av_profile default
set profile_protocol_options default
set webfilter-profile default
set spamfilter-profile default
set ips-sensor default
next
end
```

Dynamic security policies

The following procedures use a dynamic profile group called `my_dyn_prof_group` to create a security policy to allow traffic between `wan1` and `internal` for all traffic at any time of day. It allows a group of services called `allowed_protocols` that includes protocols for web traffic, email, and other useful business applications.

Even though you have configured a UTM profile with AV, webfilter, spam, and IPS you must also configure UTM in the security policy. The UTM profile is used to match the dynamic profile group name on the RADIUS server. Configuring UTM within the security policy ensures that level UTM will be applied to any users who fail the dynamic profile authentication check. Otherwise those users would be left without any UTM, which can leave a gap in your security depending how those users are handled.

When this policy has been created, place it at the top of the security policy list to ensure it will be matched first. Otherwise it is possible another policy will match the users but without Authentication resulting in an unsecure user gaining access. Ensure there are policies lower in the list to catch any traffic that needs access through the FortiGate unit without requiring Authentication, as well as the final implicit deny policy at the bottom. You may have to enable *Display Implicit security Policies* under *System > Admin > Settings* to see it.

To create a dynamic profile security policy - web-based manager

- 1 Go to *Policy*.
- 2 Select *Create New*.
- 3 Enter the following information.

Source Interface/Zone	wan1
Source Address	all
Destination Interface/Zone	internal
Destination Address	all
Action	ACCEPT

- 4 Select *Enable Dynamic Profile*.
- 5 Select *my_dyn_prof_group* for Profile.
This is the UTM profile group created earlier.
- 6 Select UTM, and enable default configuration for AntiVirus, IPS, Web Filter, Email Filter, and Protocol Options.
- 7 Add the comment "Dynamic Profile policy" to help identify this security policy in a list of policies.
- 8 Select *OK*.

To create a dynamic profile security policy - CLI

```
config firewall policy
edit 0
    set srcintf wan1
    set srcaddr any
    set dstintf internal
    set dstaddr any
    set action accept
    set dynamic-profile enable
    set dynamic-profile-access http imap pop3 smtp
    set dynamic-profile-group my_dyn_prof_group
    set utm-status enable
    set av-profile default
    set webfilter-profile default
    set spamfilter-profile default
```

```
set ips-sensor default
set profile-protocol-options default
set comments "Dynamic Profile policy"
next
end
```

Accounting system RADIUS configuration

You can configure FortiOS dynamic profiles to work with most RADIUS-based accounting systems. In most cases, you only need to do the following to your RADIUS accounting system before you can use dynamic profiles.

- Add a profile group name field to customer accounts on the RADIUS server so that the name is added to the RADIUS Start record sent by the accounting system to the FortiOS unit. Profile group names do not need to be added for all users, only to the accounts of customers who will use the dynamic profile feature on the FortiOS unit. If a profile group is not found in a RADIUS Start record regular RADIUS authentication will be used instead of dynamic profile if the user is configured for that type of authentication. See [“RADIUS authentication with a FortiGate unit” on page 1202](#).
- Configure your accounting system to send RADIUS Start records to the FortiOS unit. You can send the RADIUS Start records to any FortiOS network interface. If your FortiOS unit is operating with virtual domains (VDOMs) enabled, the RADIUS Start records must be sent to a network interface in the management VDOM.

User context list

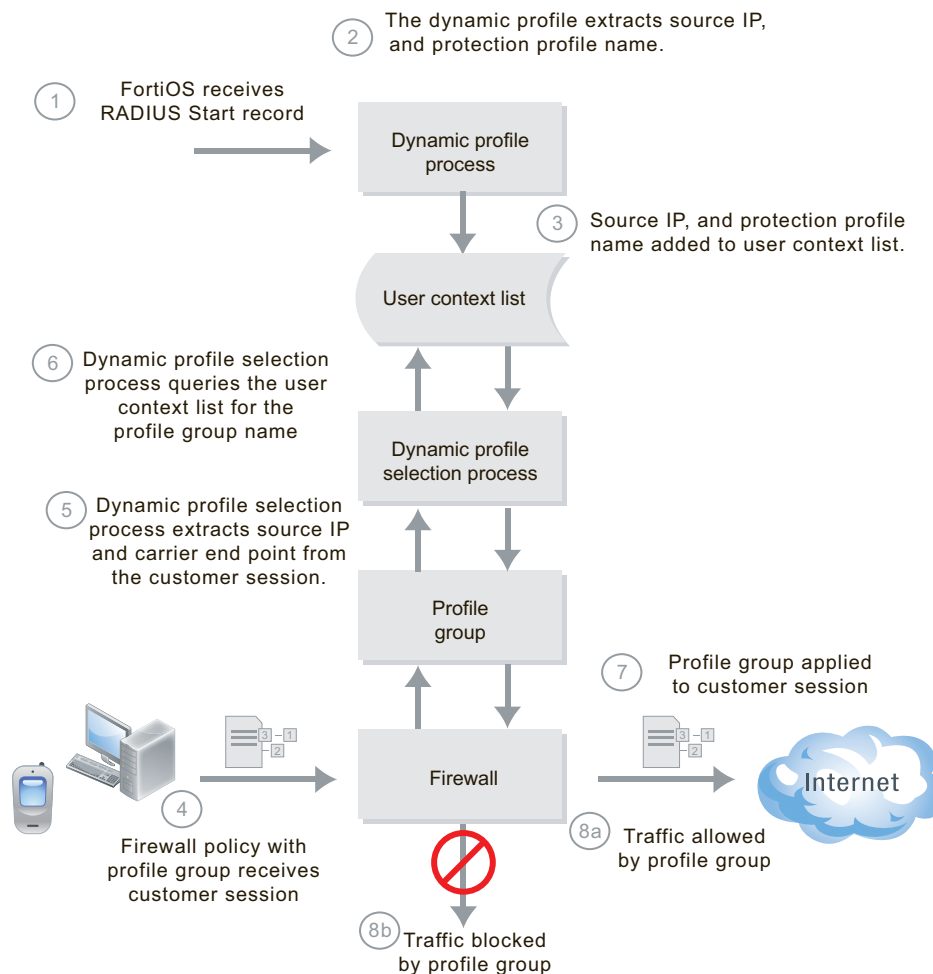
FortiOS maintains a dynamic **user context list**—a list of current end points, IP addresses, and the profile group name received in RADIUS Start records. FortiOS uses timeouts to make sure that the list contains only current information, removing entries that are no longer needed (see [“Timeout options” on page 1349](#)).

FortiOS can also remove entries from the user context list if the accounting system sends a RADIUS Stop record when a customer finishes a communication session. When FortiOS receives a RADIUS Stop record, the end point in the record is removed from the user context list. The RADIUS Stop records are optional, but they make sure FortiOS maintains an accurate user context list.

Figure 116: FortiOS dynamic profile information flow



You can use the IP Filter list to block access through FortiOS for end points. To use this feature, add end points to the list and select block traffic. FortiOS uses the user context list to look up an end point in the IP filter list to find the source IP address that the end point is using. Then at the IP level, FortiOS blocks all sessions from the source IP address.



Accepting sessions only from dynamic profile users

Extracting an end point from the content of a communication session creates extra processing overhead for the FortiOS unit. This extra overhead is acceptable for communication sessions where you want to apply dynamic profiles because of the ability to examine the content of the communication session.

However, communication sessions that do not have an end point in the user context list also create the same extra processing overhead only to be dropped when no match is found. In some cases you can reduce the amount of processing overhead by adding specific source and destination addresses to a dynamic profile security policy so that the policy matches fewer sessions. However, this may or may not work depending on factors such as your network design and security policy.

A better solution is to select *Dynamic Profile Users Only* in the security policy. If this option is selected, the dynamic profile policy only accepts sessions with source addresses that are in the user context list. Sessions with source addresses that are not in the user context list do not match the policy. For sessions that don't match the policy, the FortiOS unit continues searching down the policy list for a match.

You can add policies below the dynamic profile policy to apply various FortiOS features to non-dynamic profile sessions. For example:

- To block all non-dynamic profile sessions, make the next policy in the list a deny policy that matches all traffic. You could select *Log Violation Traffic* in this policy to log all non-dynamic profile sessions.
- To use traffic shaping to reduce the bandwidth available for non-dynamic profile sessions, apply a restrictive traffic shaper to the policy below the dynamic profile policy.

Dynamic Profile Users Only option

The *Dynamic Profile Users Only* option also allows you to differentiate between dynamic profile and regular sessions without including a profile group in the dynamic profile policy. You can use this property to apply a profile group only to non-dynamic profile traffic by adding a profile group to the next policy in the list. All dynamic profile sessions would use the dynamic profile policy and all non-dynamic profile sessions would use the next policy.

One example use of this configuration would be if you can assume that your dynamic profile users are not a security risk and you want to give them the benefit of enhanced performance of firewall sessions that do not apply a profile group.



With *Dynamic Profile User Only* selected, the FortiOS unit does not wait for the *User Context Creation Timeout* to see if a matching entry is added to the user context list. If there is a delay in receiving the RADIUS record and adding entries to the user context list, it is possible that sessions may not be matched with the dynamic profile security policy when as expected. If users experience this problem you may need to improve the performance of your RADIUS server, network, or FortiOS Carrier unit. For more information about the *User Context Creation Timeout*, see [“Timeout options” on page 1349](#).

Configuring dynamic profile

For dynamic profile to work, your RADIUS accounting system needs to send all the user information to the FortiOS unit which will then match the user session to a dynamic profile security policy. The users do not need to be configured on the FortiOS unit as they are being authenticated by the RADIUS server. Once the RADIUS server authenticates the user, it's just a matter of the security policy allowing the type of traffic in the session—just like an identity-based security policy after the user has authenticated.

You can configure one dynamic profile RADIUS server and one dynamic profile security policy per VDOM on your FortiOS unit. If you provide services to multiple networks by using VDOMs, this allows you to provide each network with their own dynamic profile configuration. Alternately you can configure a separate profile group in each VDOM's RADIUS server configuration to provide different levels of service to different user groups.

To enable dynamic profiles, the following configuration is required:

- ensure dynamic profile is visible in web-based manager
- configure the single RADIUS server for dynamic profiles under *User > Remote > RADIUS*
- optionally configure UTM Profile Group
- configure dynamic profile security policy

This section includes:

- [Make dynamic profiles visible](#)
- [RADIUS server configuration for dynamic profiles](#)

Make dynamic profiles visible

Dynamic profiles are not required by everyone, and the extra information displayed in the web-based manager may be confusing. For this reason, dynamic profiles can optionally be hidden.

If you want to configure dynamic profiles, you must first ensure dynamic profiles are not hidden in FortiOS.

To make dynamic profiles visible - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 Enable *Dynamic Profile Support on GUI*.
- 3 Select *Apply*.

To make dynamic profiles visible in GUI - CLI

```
config global
  config system global
    set gui-dynamic-profile-display enable
end
```

RADIUS server configuration for dynamic profiles

A RADIUS server must be configured on FortiOS for dynamic profile to work. If it is not, then the FortiOS unit does not listen for the RADIUS start record it needs for notification that the user was authenticated.

Restrictions and requirements

There are no restrictions on which brand of RADIUS servers can be used with FortiOS. The only requirements are that your RADIUS server supports the configuration FortiOS requires. Read this configuration section to ensure your RADIUS server can be properly configured using the documentation for your RADIUS server. This documentation does not explain how to configure your RADIUS server.

Only one RADIUS server can be configured for dynamic profiles per VDOM. Once one server is configured, the option to create more in the web-based manager is grayed-out and in the CLI attempting to enable the feature returns an error.

Profile group name

The profile group name is a means to identify which users on the RADIUS server are valid for the FortiOS unit and which are not. For example some users may have internal only access or be for automated processes that are internal network only. These users do not need to access resources outside their subnet, and will not use the FortiOS unit firewall. Using this method, FortiOS enforces Role Based Access Control (RBAC).

Also the profile group name allows multiple FortiOS units to share the Same RADIUS server by assigning each FortiOS unit a different group name and possibly even use a different RADIUS attribute as well. This can be useful if you only have one RADIUS server on your networks but have multiple FortiOS units protecting different network segments—for security reasons different users may have different levels of access to different parts of the network.

This example uses a RADIUS server called `myRADIUS.example.com`. This server uses “Login-IP Host” for both endpoint attribute and endpoint blocking. It uses “Vendor-Specific” for profile along with a value of “forti_dyn_prof”. This means when dynamic profile is trying to match the profile group, the group name it is looking to match will be in the *Vendor-Specific* RADIUS attribute and the value will be `forti_dyn_prof`. You may use different attributes for your configuration as is required.

The RADIUS server configuration will use the default settings where possible. This procedure assumes there is no existing dynamic profile server already configured.

To configure a dynamic profile RADIUS server - web-based manager

- 1 Go to *User > Remote > Radius*.
- 2 Select *Create New*.
- 3 For *Name*, enter `myRADIUS.example.com`. This is the name of the RADIUS server on the network.
- 4 For *Type*, select *Dynamic Start*. This enables dynamic profiles.

5 Enter the following information, and select OK:

Endpoint Attribute	Select the RADIUS attribute “Login-IP Host” from drop-down list. This attribute will be used to uniquely identify the user and enable extraction of endpoint information. Some RADIUS attributes are only for use with FortiOS Carrier and carrier networks. These attributes relate to mobile devices such as *-Station-Id.
Profile Attribute	Select the RADIUS attribute “Vendor-Specific” from drop-down list. This attribute will be used to uniquely identify the user and enable extraction of security profile information. RBAC is enabled with Vendor-Specific attributes. See “Role Based Access Control” on page 1204 . Some RADIUS attributes are only for use with FortiOS Carrier and carrier networks. These attributes relate to mobile devices such as *-Station-Id.
Endpoint Blocking Attribute	Select the RADIUS attribute “Login-IP Host” from drop-down list. This attribute will be used to uniquely identify the user and enable extraction of endpoint blocking information. Some RADIUS attributes are only for use with FortiOS Carrier and carrier networks. These attributes relate to mobile devices such as *-Station-Id.
Send RADIUS Response	Select this option. FortiOS will send RADIUS responses after receiving RADIUS Start and Stop records. This setting may be required by your accounting system.
Validate RADIUS Secret	Select to use RADIUS shared secret for responses and validating requests for dynamic profile.
Flush IP sessions on STOP	Select to terminate all firewall TCP and UDP sessions associated with an IP after receiving a RADIUS STOP message for a dynamic profile. If left open it can result in unwanted traffic from an earlier user being assigned to a new user if the same IP address was assigned.
Profile Key	For this example, enter “forti_dyn_prof”, short for (fortios dynamic profile). Contains the profile group name to be matched in Profile Attribute. This is the profile group from the RADIUS server that the user belongs to. The profile group is how RADIUS users can be divided into different groups and access levels.
Log All Events	Select to log all possible dynamic profile related events. To select individual events, go to the CLI command <code>dp-log-dyn_flags</code> and select individual events to log. See “Advanced dynamic profile RADIUS server configuration - CLI” on page 1336 .



Dynamic profiles always use the RADIUS `framed-ip-addr` field to get the IP address associated with the end point.

Advanced dynamic profile RADIUS server configuration - CLI

In the CLI, the following keywords are specific to configuring a RADIUS server for dynamic profile. The CLI keywords provide more in-depth options than found in the web-based manager. In the following example the RADIUS server uses `secret123` for the server secret, and the server is located at the IP address of `10.21.101.10`.

The dynamic profile RADIUS server keywords are found under `config user radius`.

To repeat the example above:

```
config vdom
edit root
config user radius
edit myRADIUS.example.com
  set dynamic_profile enable
  set auth-type auto
  set server myRADIUS.example.com
  set secret secret123
  set source-ip 10.21.101.10
  set use-management-vdom enable
  set dp-context-timeout 7200
  set dp-log-dyn-flags accounting-event accounting-stop-missed
    context-missing endpoint-block profile-missing protocol-
    error radiusd-other
  set dp-profile-attribute Vendor-Specific
  set dp-profile-attribute-key "forti_dyn_prof"
  set dp-radius-response enable
next
end
```

`nas-ip` defines the station ID, commonly used with FortiOS Carrier, and is not used in this example.

`use-management-vdom` ensures that all requests are sent over the management VDOM instead of the current VDOM. When multiple VDOMs are enabled, the management VDOM always has access to servers such as FortiGuard which may help ensure connections to RADIUS servers as well. If each VDOM sends its own requests, changes to that VDOM may unknowingly cause problems with the sending of RADIUS requests.

`dp-context-timeout` defines how long before the user context is cleared from the table. The default is eight hours (28 800 seconds) which is generally OK. However, in this example we are using two hours (7 200 seconds). This is more secure, but may require clients to reconnect if they maintain an open connection for periods longer than two hours.

`dp-log-dyn-flags` allows you to set which dynamic profile related events you want to log. In the example, all flags were selected. This customizing allows you to ensure all the necessary information is logged while also allowing you to remove any unused events from the log to save space and time. Available log events include `accounting-event`, `accounting-stop-missed`, `ontext-missing`, `endpoint-block`, `none`, `profile-missing`, `protocol-error`, and `radiusd-other`. See [“Log settings” on page 1350](#).

`dp-profile-attribute` and `dp-profile-attribute-key` work together. Essentially the first field tells FortiOS which RADIUS attribute will contain the name of the profile group (such as Fortinet-Group-Name in the Vendor Specific Attribute), and the second field has the group name, or profile group name, to look for.

It is likely that different users will have different strings in the selected RADIUS attribute, but FortiOS is only concerned with the one string that identifies its users. This is the same name as the UTM profile that is applied to the dynamic profile in the security policy. Role based access control is accomplished this way. See [“Dynamic profiles and security policies” on page 1327](#).

`dp-radius-response` tells FortiOS to send a response when the system receives a RADIUS start message. Some RADIUS servers require this to be enabled. It allows the RADIUS server to ensure the Start Record was received, which will prevent retries and such which will in turn improve system performance overall.

For additional dynamic profile configuration, see the [FortiOS CLI reference](#).

Carrier endpoints

When you are configuring the dynamic profile RADIUS server in the CLI, you have the option of configuring carrier endpoints as well. If you set `endpoint-translation enable`, this enables the endpoint series of CLI commands. See [“Configuring end points” on page 1341](#).

<code>dynamic_profile</code>	Enable dynamic profile feature. Default is disable. Note that identity-based policy and dynamic profile are mutual exclusive — when one is enabled, the other is hidden. When disabled, all related keywords (dp-*) are hidden.
<code>dp-carrier-endpoint-attribute</code>	Enter the RADIUS Attribute used to hold End Point name. By default this value is Calling-Station-Id.
<code>dp-carrier-endpoint-block-attribute</code>	Select “Login-IP Host”. This is the RADIUS Attribute used to hold the endpoint to block. The list of available attributes to use is extensive. By default this value is Calling-Station-Id. Some RADIUS attributes are only for use with FortiOS Carrier and carrier networks. These attributes relate to mobile devices such as *-Station-Id.
<code>dp-context-timeout</code>	Timeout value in seconds for user context table entries (0 = infinite). Default is 28 800 seconds. See “Timeout options” on page 1349 .
<code>dp-flush-ip-session</code>	Enable to flush user IP sessions on RADIUS accounting stop. By default this is set to disabled.
<code>dp-hold-time</code>	Enter the time in seconds to hold in proxy connection state to receive RADIUS START. Default time is 5 seconds.

dp-log-dyn_flags	<p>Select one or more dynamic profile events to log. Events available include:</p> <ul style="list-style-type: none"> • accounting-event • accounting-stop-missed • context-missing • endpoint-block • none • profile-missing • protocol-error • radiusd-other <p>By default, all event types are enabled except none.</p>
dp-log-period	Enter the minimum time period in seconds to use for event logs. Default is 0.
dp-mem-percent	Enter the maximum percentage of system memory to use for context tables. The default for this is set to 4%.
dp-profile-attribute	<p>Select "Vendor-Specific". This is the RADIUS Attribute used to hold the security profile group name. By default this is set to Class.</p> <p>To extract a profile group name from the RADIUS Start record, set this field to the name of the RADIUS attribute that contains the profile group name.</p> <p>Some RADIUS attributes are only for use with FortiOS Carrier and carrier networks. These attributes relate to mobile devices such as *-Station-Id.</p>
dp-profile-attribute-key	<p>Enter a string in this field if the <i>Profile Attribute</i> always contains the same text string directly before the profile group name. For example, if the <i>Profile Attribute</i> always includes the string <code>profile_name=</code> before the profile group name (for example, <code>profile_name=<profile_name_str></code>), set the <i>Profile Key</i> to <code>profile_name</code>. FortiOS uses the string in the <i>Profile Key</i> to extract the profile name from the complete <i>Profile Attribute</i> string.</p>
dp-radius-response	Enable to send RADIUS response packets. By default this is set to disable.
dp-radius-server-port	If required, change the UDP port number used by the RADIUS accounting server for sending RADIUS records. The default is 1813. FortiOS Carrier listens for RADIUS Start and Stop records on this port.

dp-secret	Enter the RADIUS shared secret for responses and validating requests for dynamic profile by the RADIUS accounting server.
dp-validate-request-secret	Enable to validate RADIUS request shared secret. Select if you want FortiOS Carrier to verify that the RADIUS secret matches the RADIUS secret in the RADIUS Start or End record. You can use the RADIUS secret to verify that the RADIUS record is valid.

Configuring dynamic profile-based security policies

Only one security policy can be configured for dynamic profile in a VDOM. Also only one RADIUS server and one dynamic profile group can be configured per VDOM. When one dynamic profile security policy has been configured, the option is not visible when creating other policies. After deleting the dynamic profile security policy, the option is again visible when configuring other security policies.

By enabling the *Dynamic Profile Users Only* option, other non-dynamic profile users will not match this policy. This can be useful if you want to enforce all users to be part of the dynamic profile group—in which case you have a deny all profile after this one.

If you want to allow all users access to the resources this security policy allows, do not enable the *Dynamic Profile Users Only* option.

To create dynamic profile-based security policies - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the Source, Destination, Schedule, Service and Action information as you would for any security policy.
- 4 Select *Enable Dynamic Profile*.
If this option does not appear, another security policy has dynamic profile enabled already.
- 5 Select *Profile Group*, and choose a configured group from the list.
The dynamic profile Profile Group uses the UTM Profile Group list, located under *UTM > Profile Group > Profile Group*.
- 6 Optionally select *Dynamic Profile Users Only*. See [“Accepting sessions only from dynamic profile users” on page 1331](#) and [“Dynamic Profile Users Only option” on page 1332](#).
- 7 Select UTM, Traffic Shaping, Endpoint Security, and other settings as you would for any other security policy.
- 8 Select OK.

If the RADIUS server and users have been configured, you can now start authenticating users with dynamic profile.
chapter of FortiOS Handbook



In FortiOS Carrier, while creating or editing the identity based policy the service list includes MMS as an option.

Configuration concepts

Between identity-based policies, dynamic profiles, and traffic shaping there are many possible ways to configure your security policies to work with dynamic profiles. A few common ideas are presented here.

While many security options such as AV, webfiltering, and endpoint control are not mentioned here they are perfectly valid for use with dynamic profile security policies and are critical to maintaining network security.

These concepts include:

- Only allow dynamic profile users
- Only allow dynamic profile users, but traffic shape protocols
- Allow multiple dynamic profile groups
- Promote dynamic profile users
- Schedule-based policies

Only allow dynamic profile users

In this security policy configuration, there are only two policies needed. The first one will allow all dynamic profile authenticated users, and the second will deny everyone else.

This can be a problem for guest users visiting the office. It may also prove to be too restrictive in other ways as well.

Only allow dynamic profile users, but traffic shape protocols

This configuration is the same as the last one with one change. Different protocols (HTTPS, HTTP, telnet, ftp, etc.) will have different bandwidths allotted to each of them, and each will need its own security policy and traffic shaper to implement this. Each policy will be an IBP allow policy.

This is useful to ensure that important protocols or applications have the needed bandwidth to function properly. An example is VoIP applications that require a minimum level of bandwidth to guarantee a level of service.

Allow multiple dynamic profile groups

This configuration is the same as the last one, but uses multiple VDOMs to allow one profile group per VDOM. This configuration can be useful if you provide services to multiple networks with multiple VDOMs, or if you have one network with multiple levels of user access and want to differentiate between different dynamic profile users.

Promote dynamic profile users

It is possible that non-dynamic profile users are allowed on your network, but they will not have the same amount of resources available to them. This configuration is the same as the last one, but for the traffic shaping protocols there is no identity-based policy involved.

It is possible to have the users still authenticate if you want, by using IBPs and certificates. This would maintain a higher level of security on your network by requiring all users to authenticate in some form or another.

Schedule-based policies

Up to now, no mention has been made of scheduling. It is possible to have a policy only active at certain times of the day, or certain days of the week. This would allow a school to configure “during class” and “after school” schedules that may vary. Or at a company, different policies might be used over the lunch hour, for example relaxing company Internet browsing policies.

Schedules can be used by any security policy, but remember when that policy is offline you likely want a different policy in place allowing some traffic through. If you deny all traffic at times, you may run into problems with blocking even administrator access and be forced to use the console interface to access the FortiOS unit.

Configuring end points

An end point contains enough information about a user to be unique. This includes the IP address, profile, and group the user belongs to at a minimum. The end point information must be included in the RADIUS Start record and must be available in the customer communication session (for example, in the HTTP header).

Note that while the feature name is carrier end point, the feature is common to both FortiOS and FortiOS Carrier except where noted in the text.

This section includes:

- [Configuring end points - CLI](#)
- [Controlling MMS service access based on a user's end point - FortiCarrier Only](#)
- [Blocking access to the network based on end points - FortiOS Carrier only](#)
- [Extracting carrier end points for notifications - FortiOS Carrier only](#)

Configuring end points - CLI

Dynamic profile end points are configured per RADIUS user that will be authentication using the dynamic profile. End points are only configurable in the CLI.

The following CLI keywords are specific to configuring end points. The end point keywords are found under `config user radius`, when `set dynamic-profile` is enabled.

endpoint-translation	<p>Enable/disable endpoint-translation configure. Default is disabled.</p> <p>When disabled, all carrier endpoint translation related keywords (ep-*) are hidden.</p> <p>The example dynamic profile in Figure expects the end point to be in the Calling-Station-ID attribute.</p> <p>For details about RADIUS attributes see RFC 2138 and RFC 2866.</p>
ep-carrier-endpoint-convert-hex	<p>Enable to enable converting end point to and from HEX string.</p>
ep-carrier-endpoint-header	<p>Enable to allow custom information to be added to HTTP headers. This option is commonly used on carrier networks.</p> <p>Default is x-up-calling-line-id.</p>

ep-carrier-endpoint-header-suppress	Select Enable to prevent use of HTTP headers
ep-carrier-endpoint-prefix	Enable to prefix end point values with additional information
ep-ip-header-suppress	Enable/disable HTTP header suppression.
ep-missing-header-fallback	Specify either the policy-profile or the session-ip address to act as a backup when the extracted header is not present. Default is policy-profile.
ep-profile-query-type	Specify the type of dynamic profile query as one of <ul style="list-style-type: none"> • extract-carrier-endpoint • extract-ip • session-ip Use the IP address when the full carrier endpoint information is not available. The default is session-ip.

Controlling MMS service access based on a user's end point - FortiCarrier Only

You can control access to MMS services for users according to their end point by configuring *end point filtering* (also called end point blocking). End point filtering can filter MM1, MM3, MM4, and MM7 messages according to the end points in the *From* or *To* addresses of the messages. For a definition of end points, see [“End points” on page 1327](#).

You configure end point filtering by creating an end point filter list containing end point patterns. A end point pattern can match one end point or can use wildcards or regular expressions to match multiple end points.

For each pattern, you select the action that FortiOS Carrier takes on a message when the pattern matches a end point in the message. Actions include blocking the message, exempting the message from mass MMS scanning and exempting the message from all scanning. You can also intercept the message and archive the message to a FortiAnalyzer unit.

To apply an end point filter list, you need to add the list to the *MMS Scanning > Carrier End Point Block* section of an MMS protection profile.

Configuring end point filtering - FortiOS Carrier only

To apply end point filtering- web-based manager

- 1 Go to *UTM > Carrier > Carrier End point Filter Lists*.
- 2 Add or edit an end point filter list.
- 3 Add or edit end point patterns in the list. See [“Configuring end point filter lists” on page 1343](#).
- 4 Go to *UTM > Carrier > MMS Profile* and add or edit an MMS protection profile.
- 5 Expand *MMS Scanning* and select *Carrier End Point Block*.
- 6 Select the MMS protocols to apply the end point filter list to—one or more of MM1, MM3, MM4 and MM7.
- 7 Select the end point filter list to apply.

- 8 Add the MMS protection profile to a protection profile.
- 9 Add the protection profile to the dynamic profile security policy that accepts the MMS messages that you want to filter.

Configuring end point filter lists

To configure an end point filter list - web-based manager

- 1 Go to *UTM > Carrier > Carrier End Point Filter Lists*.
- 2 Select Create New.
- 3 Enter the following information and select OK.

Name	Name of the end point filter list. You select this name in an MMS protection profile.
Comments	Optional description of the end point filter list.
Check/Uncheck All	Select the check box to enable all end point patterns in the MMS filter list. Clear the check box to disable all entries on the MMS filter list. You can also select or clear individual check boxes to enable or disable individual end point patterns.
Pattern	The pattern that FortiOS Carrier uses to match with end points. The pattern can be a single end point or consist of wildcards or Perl regular expressions that will match more than one end point. See “Regular expression vs. wildcard match pattern” on page 1344 .
Action	Select the action taken by FortiOS Carrier for messages from a carrier end point that matches the end point pattern: None - No action is taken. Block - MMS messages from the end point are not delivered and FortiOS Carrier records a log message. Exempt from mass MMS - MMS messages from the end point are delivered and are exempt from mass MMS filtering. Mass MMS filtering is configured in MMS protection profiles and is also called MMS Bulk Email Filtering and includes MMS message flood protection and MMS duplicate message detection. Exempt from all scanning - MMS messages from the end point are delivered and are exempt from all MMS protection profile scanning. MMS messages are not subject to protection profile filtering, just MMS protection profile filtering.
Content Archive	MMS messages from the end point are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
Intercept	MMS messages from the end point are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.

Pattern Type	The pattern type: <i>Wildcard</i> , <i>Regular Expression</i> , or <i>Single end point</i> . See “Using wildcards and Perl regular expressions” in the FortiOS Handbook UTM chapter .
Enable	Select to enable this end point filter pattern.

Regular expression vs. wildcard match pattern

A wildcard character is a special character that represents one or more other characters. The most commonly used wildcard characters are the asterisk (*), which typically represents zero or more characters in a string of characters, and the question mark (?), which typically represents any one character.

In Perl regular expressions, the ‘.’ character refers to any single character. It is similar to the ‘?’ character in wildcard match pattern. As a result:

- fortinet.com not only matches fortinet.com but also fortinetacom, fortinetbcom, fortinetccom, and so on.



To add a question mark (?) character to a regular expression from the FortiGate CLI, enter Ctrl+V followed by ?. To add a single backslash character (\) to a regular expression from the CLI you must add precede it with another backslash character. For example, fortinet\\.com.

To match a special character such as ‘.’ and ‘*’ use the escape character ‘\’. For example:

- To match fortinet.com, the regular expression will be: fortinet\\.com

In Perl regular expressions, ‘*’ means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- forti*.com matches fortiiii.com but does not match fortinet.com

To match any character 0 or more times, use ‘.*’ where ‘.’ means any character and the ‘*’ means 0 or more times. For example, the wildcard match pattern forti*.com will therefore be fort.*\\.com.

Blocking access to the network based on end points - FortiOS Carrier only

You can use end point IP filtering to block traffic from source IP addresses associated with end points. You can also configure FortiOS Carrier to record log messages whenever end point IP filtering blocks traffic. End point IP filtering blocks traffic at the IP level, before the traffic is accepted by a security policy..

To configure end point IP filtering, go to *UTM > Carrier > IP Filter* and add end points to the IP filter list. For each end point you can enable or disable both blocking traffic and logging blocked traffic.



You cannot add end point patterns to the end point IP filter list. You must enter complete and specific end points that are valid for your network.



The only action available is block. You cannot use end point IP filtering to exempt end points from IP filtering or to content archive or quarantine communication sessions.

FortiOS Carrier looks in the current user context list for the end points in the IP filter list and extracts the source IP addresses for these end points. Then any communication session with a source IP address that matches one of these IP addresses is blocked at the IP level, before the communication session is accepted by a security policy.

FortiOS Carrier dynamically updates the list of IP addresses to block as the user context list changes. Only these updated IP addresses are blocked by end point IP filtering.

Carrier Endpoint Filter Lists

A carrier end point filter list contains carrier end point patterns. A pattern can match one carrier end point or can use wildcards or regular expressions to match multiple carrier end points. For each pattern, you select the action that the unit takes on a message when the pattern matches a carrier end point in the message. Actions include blocking the message, exempting the message from MMS scanning, and exempting the message from all scanning. You can also configure the pattern to intercept the message and content archive the message to a FortiAnalyzer unit.

Viewing and defining an end point IP filter list

To view the end point IP filter list, go to *UTM > Carrier > IP Filter*.

You define a end point IP filter list by adding IP addresses to it. From this list, you can select *Create New* to add a new IP address or select the *Edit* icon beside an entry that you want to change.

There is only one IP filter list, and each entry can be blocked or allowed. The single list prevents configuration issues by applying all IP filters to all MMS protection profiles that use IP filtering.

Once the IP filters are configured, it is applied to all Carrier Endpoints as their traffic hits the security policies.

Carrier end point IP filter list		
	Create New	Add a end point to the end point IP filter list.
	Edit	Edit the IP filter list entry. You can change the end point, enable or disable blocking, and enable or disable logging blocked traffic. If multiple entries are selected, Edit is not available.
	Delete	Delete one or more selected end points from the list.
	Enable	Select an end point entry, and select enable to enable blocking for that entry. The Enable icon will be green with a checkmark.
	Disable	Select an end point entry, and select disable to disable blocking for that entry. This effectively turns off this entry. The Enable Icon will be grey with an X.
	Remove All Entries	Remove all entries from the end point IP filter list.

	Check / uncheck all entries on this page	Select the check box to enable all entries on the end point IP filter list. FortiOS Carrier will then block all communication sessions from source IP addresses associated with these end points. Clear the check box to disable all entries on the IP filter list. Disabling all entries disables end point IP filtering. You can also select or clear individual check boxes to enable or disable individual end points.
Create new entry		
	Adding or modifying a carrier end point pattern	Select <i>Create New</i> to add a new end point IP address or select <i>Edit</i> to change an existing address. The end point and logging selections will appear in the end point filter list.
	Carrier End Point	Enter the end point. You must enter a single end point and not a end point pattern.
	Log Blocked Traffic	Select to record a log message when carrier end point IP filtering blocks a communication session.
	Block Traffic	Select to block traffic from the carrier end point. <i>Block Traffic</i> is selected by default.

Example

You can use IP filtering on Carrier Endpoints when a handheld on your carrier network is known to be sending spam or malware. You can easily use IP filtering to block the traffic from that unit, log it, or both. This allows you to escalate your response, and even monitor afterwards as well.

In this example the hand held IP address is 10.11.101.99 and since they are sending malware this IP address will be both logged and blocked.

To create an IP filter list entry to block an address - web-based manager

- 1 Go to *UTM > Carrier > IP Filter*.
- 2 Select *Create New*.
- 3 Enter *10.11.101.199* for the *Carrier End Point*.
- 4 Select *Log Blocked Traffic* and *Block Traffic*.
- 5 Select *OK*.

Adding additional blocked IP addresses

In the future, additional IP filters can be added both to the Carrier End Point filtering list, and to the firewall address list. To add additional filters to the filtering list, repeat the procedure *To create an IP filter list entry to block an address* for each new additional IP address to be filtered.

To add additional firewall addresses to be matched in the security policy, enter each additional new IP address as a separate firewall address, and in the security policy for the address select multiple and then select all the IP_Filter_xx addresses. Using a logical naming convention for additional IP filters such as IP_Filter_02, IP_Filter_xx, etc. will help with configuration.

Extracting carrier end points for notifications - FortiOS Carrier only

The sender's carrier end point is used to provide logging and reporting details to the mobile operator and to identify the sender of infected content.

When MMS messages are transmitted, the *From* field may or may not contain the sender's address. When the address is not included, the sender information will not be present in the logs and the FortiOS Carrier unit will not be able to notify the user if the message is blocked unless the sender's address is made available elsewhere in the request. One reason for this is if multiple users are behind the same NAT device.

Beyond logging, the sender address is also important for billing, end point control, and applying security policies. It is also important for differing levels of URL access.

FortiOS Carrier can extract the sender's address from an extended HTTP header field in the HTTP request. This field must be added to the HTTP request before it is received by FortiOS Carrier. If this field is present, it will be used instead of the sender's address in the MMS message for logging and notification. If this header field is present when a message is retrieved, it will be used instead of the *To* address in the message. If this header field is not present the content of the *To* header field is used instead.

Alternatively, FortiOS Carrier can extract the sender's address from a cookie. The cookie is sent as part of the HTTP header.

You can configure MMS address translation to extract the sender's carrier end point so that it can be added to log and notification messages. You can configure MMS address translation settings to extract carrier end points from HTTP header fields or from cookies. You can also configure MMS address translation to add an end point prefix to the extracted carrier end points.

To configure MMS address translation, go to *Firewall > MMS Profile*. Select *Create New* or select the *Edit* icon beside an existing profile. Expand *MMS Address Translation*. Complete the fields as described in the following table and select *OK*.

Configuring MMS address translation - FortiOS Carrier only

MMS address translation changes the address from using the one embedded in the MMS message to using the additional HTTP Header Field (if present) or a cookie to get the address.

This applies to MM1 and MM7 messages — messages sent to or from handsets, and messages sent to or from content providers. These are the only message types that use the HTTP headers that enable this feature.

To configure MMS address translation - web-based manager

- 1 Go to *UTM > Carrier > MMS Profile*.
- 2 Select *Create New*.
- 3 Expand *MMS Address Translation*.
- 4 Select settings for MM1 and MM7 as required.

5 Select *OK*.

Sender Address Source	Select to extract the sender's address from the <i>HTTP Header Field</i> or a <i>Cookie</i> . You must also specify the identifier that contains the carrier end point.
Sender Address Identifier	<p>Enter the sender address identifier that includes the carrier end point. The default identifier is <code>x-up-calling-line-id</code>.</p> <p>If the <i>Sender Address Source</i> is <i>HTTP Header Field</i>, the address and its identifier in the HTTP request header takes the format:</p> <p><code><Sender Address Identifier>: <MSISDN_value></code></p> <p>If the <i>Sender Address Source</i> is <i>Cookie</i>, the address and its identifier in the HTTP request header's <i>Cookie</i> field takes the format of attribute-value pairs:</p> <p><code>Cookie: id=<cookie-id>;</code> <code><Sender Address Identifier>=<MSISDN Value></code></p>
Convert Sender Address From / To Hex	Select to convert the sender address from ASCII to hexadecimal or from hexadecimal to ASCII. This is required by some applications.
Add End Point Prefix for Logging / Notification	
Enable	Select to enable adding the country code to the extracted carrier end point, such as the MSISDN, for logging and notification purposes. You can limit the number length for the test numbers used for internal monitoring without a country code.
Prefix	Enter a carrier end point prefix that will be added to all carrier end points. Use the prefix to add extra information to the carrier end point in the log entry.
Minimum Length	Enter the minimum length of the number. If this and Maximum Length are set to zero (0), length is not limited.
Maximum Length	Enter the maximum length of the number. If this and Minimum Length are set to zero (0), length is not limited.

HTTP header field example

For this example, in FortiOS Carrier we are concerned about MMS traffic between content providers — MM7 traffic only. The default `x-up-calling-line-id` will be used in the HTTP header along with a country code of 9811. The *Sender Address* does need converting from hex. The prefix will be added to *Logging/Notification* using the MSISDN for the prefix.

To configure MMS address translation using HTTP header field

- 1 Go to *UTM > Carrier > MMS Profile*.
- 2 Select *Create New*.
- 3 Enter `MMS_addr_http_header` for the *Profilename*.
- 4 Expand *MMS Address Translation*.
- 5 Under MM7, select *HTTP Header Field* for *Sender Address Source*.
- 6 Under MM7, enter `x-up-calling-line-id` for *Sender Address Identifier*.
- 7 Under MM7, select *Convert Sender Address From / To HEX*.

8 Select **Enable** for *Add End Point Prefix for Logging / Notification*.

9 Enter **9811** for *Prefix*.

10 Select **OK**.

If the Sender Address Source is HTTP Header Field, the address and its identifier in the HTTP request header takes the format:

```
<Sender Address Identifier>: <MSISDN_value>
```

Where the <MSISDN_value> is the carrier end point. For example, the HTTP header might contain:

```
x-up-calling-line-id: 9811301234
```

where `x-up-calling-line-id` would be the Sender Address Identifier, and `9811301234` would be the MSISDN.

Cookie example

If you want the address to persist, then use a cookie. Keep in mind that cookies are a less secure method than the HTTP header field option because of their persistence.

For this example we are concerned about traffic to and from handsets and will only be using MM1. A non-standard field for Sender Address Identifier will be used: `x-up-calling-cookie`, and a country code of `467`. The *Sender Address* does not need converting from hex. The prefix will be added to *Logging/Notification* using the MSISDN for the prefix.

To configure MMS address translation using cookies - web-based manager

1 Go to *UTM > Carrier > MMS Profile*.

2 Select **Create New**.

3 Enter `MMS_addr_cookie` for the *Profilename*.

4 Expand *MMS Address Translation*.

5 Under MM1, select **Cookie** for *Sender Address Source*.

6 Under MM1, enter `x-up-calling-cookie` for *Sender Address Identifier*.

7 Select **Enable** for *Add End Point Prefix for Logging / Notification*.

8 Enter **467** for *Prefix*.

9 Select **OK**.

For MM1 messages, a cookie can now be referenced for the Sender's address. Any messages that trigger logging or notification that use this address translation will include the `467` prefix for added identification.

A sample HTTP request header resulting from this configuration would be:

```
Cookie: id=0123jfa; x-up-calling-cookie=467301297
```

where `0123jfa` is the cookie id, `x-up-calling-cookie` is the Sender Address Identifier, and `467301297` is the MSISDN.

Timeout options

Dynamic profile timeouts control how long FortiOS keeps entries in the user context list. Entries are removed when a RADIUS stop record is received by FortiOS, but otherwise entries remain in the context list until they timeout.

Usually you would not want entries staying in the user context list if they are not being used. A smaller list is easier and more efficient for FortiOS to manage. As well, because user context information can change, a smaller list means incorrect or out-of-date information is more likely to be removed.

However, some situations may benefit from keeping the entries in the user context list for a longer period of time. This may reduce network load and provide other benefits, but the security risk must be kept in mind as well.

User Context Entry Timeout	<p>Enter the number of seconds that a user context entry can remain in the list without FortiOS Carrier receiving a communication session from the end point. If a user context entry is not being looked up, then the user must no longer be connected to the network.</p> <p>This timeout is only required if FortiOS Carrier does not receive the RADIUS Stop record. However, even if the accounting system does send RADIUS Stop records, this timeout must be set in case FortiOS Carrier misses one.</p> <p>The default user context entry timeout is 28800 seconds (8 hours). You can keep this timeout relatively high because it is not usually a problem to have a long list. But a timeout is usually required because FortiOS Carrier normally removes entries that are no longer used.</p> <p>You might want to reduce this timeout if the accounting server does not send RADIUS Stop records. Also, if customer IP addresses change often, you might want to set this timeout lower so that out-of-date entries are removed from the list.</p> <p>Avoid entering a setting that is too low. FortiOS Carrier may remove user context entries for users who are still connected.</p> <p>Set the timeout to 0 if you do not want FortiOS Carrier to remove entries from the list except in response to RADIUS Stop messages.</p>
User Context Creation Timeout	<p>If FortiOS Carrier receives a communication session and can't find a corresponding end point and IP address in the user context list, the system waits for the <i>User Context Creation Timeout</i>. If a match is not found after this timeout, FortiOS Carrier applies the profile group in the security policy to the communication session.</p> <p>The default user context creation timeout is 5 seconds. You might want to increase this timeout if the default profile group, instead of the dynamic profile, is being applied to users. This could be happening if there is a delay before FortiOS Carrier receives the RADIUS Start record from the accounting server.</p> <p>If you set this timeout to 0, FortiOS Carrier blocks communication sessions that do not have a matching entry in the user context list.</p>

Log settings

You can use Log settings to configure FortiOS Carrier to record event log messages for dynamic profile events. You can also set a log message period to group log messages.

When using a dynamic profile, RADIUS attributes are included in the log information including framed IP address, radius username, and profile group.

Log Message Period	Enter the time in seconds to group event log messages for dynamic profile events. For example, if the log message period is 30 seconds, FortiOS Carrier generates groups of single-event log messages every 30 seconds instead of generating event log messages continuously. The grouped log messages generated each period contain a count of how many events of that type occurred. If you set this period to 0, FortiOS Carrier generates all event log messages in real time.
Protocol Errors	Select to have FortiOS generate event log messages if RADIUS protocol errors occur. One example could be a RADIUS record containing a RADIUS secret that does not match the one added to the dynamic profile.
Missing Profile Errors	Select to have FortiOS generate an event log message whenever FortiOS Carrier cannot find a profile group name in a RADIUS start message that matches the name of a profile group added to FortiOS Carrier.
Missing Context Errors	Select to have FortiOS generate an event log message whenever a user context creation timeout expires indicating that FortiOS Carrier was not able to match a communication session because a matching entry was not found in the user context list.
Missed Accounting 'Stop' Events	Select to have FortiOS generate an event log message whenever a user context entry timeout expires indicating that FortiOS Carrier removed an entry from the user context list without receiving a RADIUS Stop message.
Accounting Events	Select to have FortiOS generate an event log message when FortiOS Carrier does not find the expected information in a RADIUS record. This may happen, for example, if a RADIUS record contains more than the expected number of addresses.
Other Log Messages	Select to have FortiOS generate event log messages for other events. The event is described in the log message. For example, a log message may be generated if the memory limit for the user context list is reached and the oldest entries in the table have been dropped.

Carrier end point filters and blocking

Carrier end points are available on FortiOS units as well as FortiOS Carrier units. However, only FortiOS Carrier units can handle MMS traffic. In the following section, any reference to MMS messages is FortiOS Carrier only material.

This section includes:

- [Controlling MMS service access based on a user's end point - FortiCarrier Only](#)
- [Blocking access to the network based on end points - FortiOS Carrier only](#)
- [Extracting carrier end points for notifications - FortiOS Carrier only](#)

Controlling access to MMS services based on a user's carrier end point

You can control access to MMS services for users according to their carrier end point by configuring *carrier end point filtering* (also called carrier end point blocking). Carrier end point filtering can filter MM1, MM3, MM4, and MM7 messages according to the carrier end points in the *From* or *To* addresses of the messages.

For a definition of carrier end points, see “End points” on page 1327.

You configure carrier end point filtering by creating a carrier end point filter list containing carrier end point patterns. A carrier end point pattern can match one carrier end point or can use wildcards or regular expressions to match multiple carrier end points.

For each pattern, you select the action that FortiOS Carrier takes on a message when the pattern matches a carrier end point in the message. Actions include blocking the message, exempting the message from mass MMS scanning and exempting the message from all scanning. You can also intercept the message and archive the message to a FortiAnalyzer unit.

To apply a carrier end point filter list, you need to add the list to the *MMS Scanning > Carrier End Point Block* section of an MMS protection profile.

Filtering MMS flows from selected carrier end points

You may want to block specific MMS protocols or block all of them depending on the situation. For example if you wanted to block traffic from selected handsets, you only need to select MM1 traffic. The benefit of this is a more specific match allows you to create individual security policies for each situation on your network. Extending the example you would have one policy for MM1 handset traffic, another policy for Internet MM3 traffic, another policy for Content Provider MM7 traffic, and another policy for inter-provider MM4 traffic.

Dividing traffic into smaller groups enables focused logging and reporting, while providing a more secure and customized firewall solution for your specific needs. When something changes, its easier to update a specific smaller security policy than to update a single large complex rambling policy.

Configuring carrier end point filtering

To apply carrier end point filtering- web-based manager

- 1 Go to *UTM > Carrier End Point > Carrier End point Filter Lists*.
- 2 Add or edit a carrier end point filter list.
- 3 Add or edit carrier end point patterns in the list.
- 4 Go to *UTM > Carrier > MMS Profile* and add or edit an MMS protection profile.
- 5 Expand *MMS Scanning* and select *Carrier End Point Block*.
- 6 Select the MMS protocols to apply the carrier end point filter list to (MM1, MM3, MM4 and MM7).
- 7 Select the carrier end point filter list to apply.
- 8 Add the MMS protection profile to a protection profile.
- 9 Add the protection profile to a security policy that accepts the MMS messages that you want to filter.

To configure a carrier end point filter list - web-based manager

To configure a carrier end point filter list go to *UTM > Carrier > Carrier End Point Filter Lists*.

Name	Name of the carrier end point filter list. You select this name in an MMS protection profile.	
Comments	Optional description of the carrier end point filter list.	
Check/Uncheck All	<p>Select the check box to enable all carrier end point patterns in the MMS filter list.</p> <p>Clear the check box to disable all entries on the MMS filter list.</p> <p>You can also select or clear individual check boxes to enable or disable individual carrier end point patterns.</p>	
Pattern	<p>The pattern that FortiOS Carrier uses to match with carrier end points. The pattern can be a single carrier end point or consist of wildcards or Perl regular expressions that will match more than one carrier end point. See “<i>Using wildcards and Perl regular expressions</i>” in the FortiOS Handbook UTM chapter.</p>	
Action	<p>Select the action taken by FortiOS Carrier for messages from a carrier end point that matches the carrier end point pattern:</p> <p>None - No action is taken.</p> <p>Block - MMS messages from the carrier end point are not delivered and FortiOS Carrier records a log message.</p> <p>Exempt from mass MMS - MMS messages from the carrier end point are delivered and are exempt from mass MMS filtering. Mass MMS filtering is configured in MMS protection profiles and is also called MMS Bulk Email Filtering and includes MMS message flood protection and MMS duplicate message detection.</p> <p>Exempt from all scanning - MMS messages from the carrier end point are delivered and are exempt from all MMS protection profile scanning.</p> <p>MMS messages are not subject to protection profile filtering, just MMS protection profile filtering.</p>	
	Content Archive	MMS messages from the carrier end point are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
	Intercept	MMS messages from the carrier end point are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.
Pattern Type	<p>The pattern type: <i>Wildcard</i>, <i>Regular Expression</i>, or <i>Single Carrier End Point</i>. See “<i>Using wildcards and Perl regular expressions</i>” in the FortiOS Handbook UTM chapter.</p>	
Enable	Select to enable this carrier end point filter pattern.	

Blocking network access for IP addresses based on carrier end points

You can use carrier end point IP filtering to block traffic from source IP addresses associated with carrier end points. You can also configure FortiOS Carrier to record log messages whenever carrier end point IP filtering blocks traffic. Carrier end point IP filtering blocks traffic at the IP level, before the traffic is accepted by a security policy.

For a definition of carrier end points, see [“End points” on page 1327](#).

To configure carrier end point IP filtering, go to *UTM > Carrier > IP Filter* and add carrier end points to the IP filter list. For each carrier end point you can enable or disable both blocking traffic and logging blocked traffic.



You cannot add carrier end point patterns to the carrier end point IP filter list. You must enter complete and specific carrier end points that are valid for your network.



The only action available is block. You cannot use carrier end point IP filtering to exempt carrier end points from IP filtering or to content archive or quarantine communication sessions.

FortiOS Carrier looks in the current user context list for the carrier end points in the IP filter list and extracts the source IP addresses for these carrier end points. Then any communication session with a source IP address that matches one of these IP addresses is blocked at the IP level, before the communication session is accepted by a security policy.

FortiOS Carrier dynamically updates the list of IP addresses to block as the user context list changes. Only these updated IP addresses are blocked by carrier end point IP filtering.

Configuring end point IP filtering includes:

- [Viewing and defining a carrier end point IP filter list](#)
- [Adding additional blocked IP addresses](#)

Viewing and defining a carrier end point IP filter list

To view the carrier end point IP filter list, go to *UTM > Carrier > IP Filter*.

You define a carrier end point IP filter list by adding IP addresses to it. From this list, you can select *Create New* to add a new IP address or select the *Edit* icon beside an entry that you want to change.

There is only one IP filter list, and each entry can be blocked or allowed. The single list prevents configuration issues by applying all IP filters to all MMS protection profiles that use IP filtering.

Once the IP filters are configured, it is applied to all Carrier Endpoints as their traffic hits the security policies.

Carrier end point IP filter list		
	Create New	Add a carrier end point to the carrier end point IP filter list.
	Edit	<p>Edit the IP filter list entry. You can change the carrier end point, enable or disable blocking, and enable or disable logging blocked traffic.</p> <p>If multiple entries are selected, Edit is not available.</p>

	Delete	Delete one or more selected carrier end points from the list.
	Enable	Select an end point entry, and select enable to enable blocking for that entry. The Enable icon will be green with a checkmark.
	Disable	Select an end point entry, and select disable to disable blocking for that entry. This effectively turns off this entry. The Enable Icon will be grey with an X.
	Remove All Entries	Remove all entries from the carrier end point IP filter list.
	Check / uncheck all entries on this page	Select the check box to enable all entries on the carrier end point IP filter list. FortiOS Carrier will then block all communication sessions from source IP addresses associated with these carrier end points. Clear the check box to disable all entries on the IP filter list. Disabling all entries disables carrier end point IP filtering. You can also select or clear individual check boxes to enable or disable individual carrier end points.
Create new entry		
	Adding or modifying a carrier end point pattern	Select <i>Create New</i> to add a new carrier end point IP address or select <i>Edit</i> to change an existing address. The carrier end point and logging selections will appear in the carrier end point filter list.
	Carrier End Point	Enter the carrier end point. You must enter a single carrier end point and not a carrier end point pattern.
	Log Blocked Traffic	Select to record a log message when carrier end point IP filtering blocks a communication session.
	Block Traffic	Select to block traffic from the carrier end point. <i>Block Traffic</i> is selected by default.

Example

You can use IP filtering on Carrier Endpoints when a handheld on your carrier network is known to be sending spam or malware. You can easily use IP filtering to block the traffic from that unit, log it, or both. This allows you to escalate your response, and even monitor afterwards as well.

In this example the hand held device IP address is 10.11.101.99 and since they are sending malware this IP address will be both logged and blocked. The following general steps are required.

- 1 Create an IP filter list entry to block an address
- 2 Create an MMS protection profile to use this IP filter list
- 3 Create a firewall address for the blocked IP address, and create incoming and outgoing security policies that use the firewall address and MMS profile.



Firewall configuration is not explained here. See the Fundamentals Handbook chapter.

To create an IP filter list entry to block an address - web-based manager

- 1 Go to *UTM > Carrier > IP Filter*.
- 2 Select *Create New*.
- 3 Enter `10.11.101.199` for the *Carrier End Point*.
- 4 Select *Log Blocked Traffic* and *Block Traffic*.
- 5 Select *OK*.

To create an MMS protection profile to use the IP filter

- 1 Go to *UTM > Carrier > MMS profile*.
- 2 Select *Create New*.
- 3 Enter `ipFilteringProfile` as the *Profilename*.
- 4 Expand *MMS Scanning*.
- 5 Select all MMS protocols for *Carrier End Point Block*.
- 6 Select `my IP filter list` from the *Option* menu for *Carrier End Point Block*.
- 7 Select any other settings as required for this MMS protection profile.
- 8 Select *OK*.

Adding additional blocked IP addresses

In the future, additional IP filters can be added both to the Carrier End Point filtering list, and to the firewall address list. To add additional filters to the filtering list repeat general step 1 for each new additional IP address to be filtered. To add additional firewall addresses to be matched in the security policy, enter each additional new IP address as a separate firewall address, and in the security policy for the address select multiple and then select all the `IP_Filter_xx` addresses. Using a logical naming convention for additional IP filters such as `IP_Filter_02`, `IP_Filter_xx`, etc. will help with configuration.

Troubleshooting dynamic profiles

This section provides you with some commands and methods for troubleshooting dynamic profile issues.

This section includes:

- [General dynamic profile troubleshooting](#)
- [Dynamic profile related diag commands](#)

General dynamic profile troubleshooting

When you are troubleshooting, work your way down the list. The first solution is more general, and each solution becomes more specific as you move down the list.

- If dynamic profiles are not displayed anywhere in the web-based manager they must un-hidden. Go to *System > Admin > Settings* and select *Dynamic Profile Support on GUI*. If VDOMs are enabled, this is at the global level.
- Dynamic server option does not appear in RADIUS server configuration. This means there is already one dynamic RADIUS server configured. To configure another you must first locate and remove the original configuration.

- Dynamic profile option is not displayed in security policies. There are two possible reasons for this. The first is that there is already another security policy with dynamic profile enabled—this must be removed before you can configure another dynamic security policy. Second, when identity-based policy is selected in a security policy dynamic profile is not displayed. For dynamic profile option to reappear, un-select identity-based profile.
- Valid users not being authenticated. It is possible that it is taking too long for the FortiOS unit to receive the RADIUS start record, and the security policy is failing from lacking that information.

Dynamic profile related diag commands

You can use the following FortiOS diagnose commands to debug communication between the RADIUS server and FortiOS:

```
diag test application radiusd 2
```

Clears the user context list. This is similar to having all users log off.

```
diag test application radiusd 3
```

Shows the user context list (user, profile, ip)

```
diag test application radiusd 5
```

Displays RADIUS statistics such as the number of RADIUS Start and Stop packets received, the number of packet errors and so on.

```
diag dynamic-profile query ip 10.0.0.1
```

A filter that displays the same kind of information as `diagnose test application radiusd 3`.

```
diag dynamic-profile query profile-usage <vdom_name>
```

Displays a summary of profile usage for the named VDOM.

```
diag debug application radiusd 3
```


Displays carrier end points and their associated profile group names. An example entry for carrier end point 5551231234 and profile group name `profile_name` could be:

```
<000.000000> [49]: received from radiusd -- msgId=41,
profile=profile_name, endpoint=5551231234 [49]: delayed setup
for profile profile_name)
```

In the following example, the command `diagnose test application radiusd 3` displays three entries in the user context list. In this example, the carrier end points are listed under endpoint and they are all email addresses. This example output uses example IP addresses and domain names.

```
index,"time left (hh:mm:ss)",ip,endpoint,profile,rc,"Default
Profile?","Blacklist?
1,07:08:07,"192.168.23.7","33example@example.com","PackageVAS2",1,
No,No
2,07:23:32,"192.168.33.112","45example@example.com","PackageVAS2",
```

```
1, No, No
3, 07:28:32, "172.20.123.71", "332example@example.com", "PackageVAS1",
1, No, No
```



Monitoring authenticated users

This section describes how to view lists of currently logged-in firewall and VPN users. It also describes how to disconnect users.

The following topics are included in this section:

- [Monitoring firewall users](#)
- [Monitoring SSL VPN users](#)
- [Monitoring IPsec VPN users](#)

Monitoring firewall users

To monitor firewall users, go to *User > Monitor > Firewall*>

Figure 117: Firewall users listed in monitor

[Column Settings] [Clear All Filters] [De-authenticate All Users]							
User Name	User Group	Policy ID	Duration	IP Address	Traffic Volume	Method	
user3	Group1	2	0 day(s) 0 hour(s) 4 minute(s)	10.11.101.20	35 KB	FW-auth	
user4	Group1	2	0 day(s) 3 hour(s) 4 minute(s)	10.11.101.101	421 KB	FW-auth	

You can de-authenticate a user by selecting the Delete icon for that entry.

You can filter the list of displayed users either by selecting the funnel icon for one of the column titles or selecting *Filter Settings*.

Select *Column Settings* to add or remove columns to the display, or rearrange the order of the columns displayed.

Optionally, you can select *De-authenticate all users*. Best practices dictate that this only be used in extreme cases since all users will momentarily lose their network resource connections.

Monitoring SSL VPN users

You can monitor web-mode and tunnel-mode SSL VPN users by username and IP address.

To monitor SSL VPN users, go to *VPN > Monitor > SSL-VPN Monitor*. To disconnect a user, select the user and then select the *Delete* icon.

Figure 118: Monitoring SSL VPN users

Delete					
<input type="checkbox"/>	No.	User	Source IP	Begin Time	Description
<input type="checkbox"/>	1	user2	172.20.120.51	Wed Mar 17 13:17:32 2010	
<input checked="" type="checkbox"/>		Subsession			Tunnel IP:10.0.0.1

The first line, listing the username and IP address, is present for a user with either a web-mode or tunnel-mode connection. The Subsession line is present only if the user has a tunnel mode connection. The *Description* column displays the virtual IP address assigned to the user's tunnel-mode connection.

For more information about SSL VPN, see the [FortiOS Handbook SSL VPN chapter](#).

To monitor SSL VPN users - CLI

To list all of the SSL VPN sessions and their index numbers:

```
execute vpn sslvpn list
```

The output looks like this:

SSL-VPN Login Users:

Index	User	Auth	Type	Timeout	From	HTTPS in/out
0	user1	1		256	172.20.120.51	0/0

SSL-VPN sessions:

Index	User	Source IP	Tunnel/Dest IP
0	user2	172.20.120.51	10.0.0.1

You can use the Index value in the following commands to disconnect user sessions:

To disconnect a tunnel-mode user

```
execute vpn sslvpn del-tunnel <index>
```

To disconnect a web-mode user

```
execute vpn sslvpn del-web <index>
```

You can also disconnect multiple users:

To disconnect all tunnel-mode SSL VPN users in this VDOM

```
execute vpn ssl del-all tunnel
```

To disconnect all SSL VPN users in this VDOM

```
execute vpn ssl del-all
```

Monitoring IPsec VPN users


To monitor IPsec VPN tunnels in the web-based manager, go to *VPN > Monitor > IPsec Monitor*. usernames are available only for users who authenticate with XAuth.

You can close a tunnel by selecting its *Bring Down* link in the *Status* column.

Figure 119: Monitoring dialup VPN users

Type Dialup							
Name	Remote Gateway	Timeout	Status	Incoming Data	Outgoing Data	Username	
dialup1_0	172.20.120.51	1116	Bring Down	79233170 B	171639314 B	user2	

For more information, see the [FortiOS Handbook IPsec chapter](#).



Examples and Troubleshooting


This chapter provides an example of a FortiGate unit providing authenticated access to the Internet for both Windows network users and local users.

The following topics are included in this section:

- [Firewall authentication example](#)
- [LDAP Dial-in using member-attribute](#)
- [Dynamic Profile example](#)
- [Troubleshooting](#)

Firewall authentication example

Figure 120: Example configuration



Overview

In this example, there is a Windows network connected to Port 2 on the FortiGate unit and another LAN, Network_1, connected to Port 3.

All Windows network users authenticate when they logon to their network. Members of the Engineering and Sales groups can access the Internet without entering their authentication credentials again. The example assumes that the Fortinet Single Sign On (FSSO) has already been installed and configured on the domain controller.

LAN users who belong to the Internet_users group can access the Internet after entering their username and password to authenticate. This example shows only two users, User1 is authenticated by a password stored on the FortiGate unit, User2 is authenticated on an external authentication server. Both of these users are referred to as local users because the user account is created on the FortiGate unit.

Creating a locally-authenticated user account

User1 is authenticated by a password stored on the FortiGate unit. It is very simple to create this type of account.

To create a local user - web-based manager

- 1 Go to *User > User* and select *Create New*.
- 2 Enter the following information: username, Password.

username	User1
Password	hardtoguess

- 3 Select *OK*.

To create a local user - CLI

```
config user local
  edit user1
    set type password
    set passwd hardtoguess
  end
```

Creating a RADIUS-authenticated user account

To authenticate users using an external authentication server, you must first configure the FortiGate unit to access the server.

To configure the remote authentication server - web-based manager

- 1 Go to *User > Remote > RADIUS* and select *Create New*.
- 2 Enter the following information and select *OK*:

Name	OurRADIUSsrv
Primary Server Name/IP	10.11.101.15
Primary Server Secret	OurSecret
Authentication Scheme	Select <i>Use Default Authentication Scheme</i> .

To configure the remote authentication server - CLI

```
config user radius
  edit OurRADIUSsrv
    set server 10.11.102.15
    set secret OurSecret
    set auth-type auto
  end
```

Creation of the user account is similar to the locally-authenticated account, except that you specify the RADIUS authentication server instead of the user's password.

To configure a remote user - web-based manager

- 1 Go to *User > User* and select *Create New*.
- 2 Enter the following information and select *OK*:

username	User2
RADIUS	Select <i>Match user on RADIUS server</i> and then select OurRADIUSsrv from the list.

To configure a remote user - CLI

```
config user local
  edit User2
    set name User2
    set type radius
    set radius-server OurRADIUSsrv
  end
```

Creating user groups

There are two user groups: an FSSO user group for FSSO users and a firewall user group for other users. It is not possible to combine these two types of users in the same user group.

Creating the FSSO user group

For this example, assume that FSSO has already been set up on the Windows network and that it uses Advanced mode, meaning that it uses LDAP to access user group information. You need to

- configure LDAP access to the Windows AD global catalog
- specify the collector agent that sends user logon information to the FortiGate unit
- select Windows user groups to monitor
- select and add the Engineering and Sales groups to an FSSO user group

To configure LDAP for FSSO - web-based manager

- 1 Go to *User > Remote > LDAP* and select *Create New*.
- 2 Enter the following information:

Name	ADserver
Server Name / IP	10.11.101.160
Distinguished Name	dc=office,dc=example,dc=com
Bind Type	Regular
User DN	cn=FSSO_Admin,cn=users,dc=office,dc=example,dc=com
Password	set_a_secure_password

Leave other fields at their default values.

- 3 Select *OK*.

To configure LDAP for FSSO - CLI

```

config user ldap
  edit "ADserver"
    set server "10.11.101.160"
    set dn "cn=users,dc=office,dc=example,dc=com"
    set type regular
    set username
      "cn=administrator,cn=users,dc=office,dc=example,dc=com"
    set password set_a_secure_password
  next
end

```

To specify the collector agent for FSSO - web-based manager

- 1 Go to *User > FSSO* and select *Create New*.
- 2 Enter the following information and select *OK*:

Name	WinGroups
Enter on one line	
Collector IP/Name	10.11.101.160
Port	8000
Password	fortinet_canada
LDAP Server	ADserver

To specify the collector agent for FSSO - CLI

```

config user fsso
  edit "WinGroups"
    set ldap-server "ADserver"
    set password ENC
      G7GQV7NEqilCM9jKmVmJJFVvhQ2+wtNEe9T0iYA5Sa+EQT2J8zhOrbkJFD
      r0RmY3c4LaoXdsoBczAldONmcGfthTxxwGsigzGpbJdC7lspFlQYtj
    set server "10.11.101.160"
  end

```

To select Windows user groups to monitor - web-based manager

- 1 Go to *User > FSSO > FSSO Agent*.
- 2 Expand *WinGroups*, then select the *Edit Users/Groups* icon.
- 3 Select the Engineering and Sales groups and then select *OK*.

To create the FSSO_Internet-users user group - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter the group name, *FSSO_Internet_users*.
- 3 Select *Fortinet Single Sign-On (FSSO)*.
- 4 In the *Available Members* list, select the Engineering and Sales groups and then select the right arrow button to move them to the *Members* list.
- 5 Select *OK*.

To create the FSSO_Internet-users user group - CLI

```

config user group
  edit FSSO_Internet_users
    set group-type directory-service
    set member
      CN=Engineering,cn=users,dc=office,dc=example,dc=com
      CN=Sales,cn=users,dc=office,dc=example,dc=com
  end

```

Creating the Firewall user group

The non-FSSO users need a user group too. In this example, only two users are shown, but additional members can be added easily.

To create the firewall user group - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter the following information and then select *OK*:

Name	Internet_users
Type	Firewall
Members	User1, User2

To create the firewall user group - CLI

```

config user group
  edit Internet_users
    set group-type firewall
    set member User1 User2
  end

```

Defining policy addresses

Go to *Firewall Objects > Address* and create the following addresses:

Address Name	Internal_net
Type	Subnet / IP Range
Subnet / IP Range	10.11.102.0/24
Interface	Port 3

Address Name	Windows_net
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0/24
Interface	Port 2

Creating security policies

Two security policies are needed: one for firewall group who connect through port3 and one for FSSO group who connect through port2.

To create a security policy for FSSO authentication - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.

- 2 Enter the following information:

Source interface	Port2
Source address	Windows_net
Destination interface	Port1
Destination address	all
Action	ACCEPT
NAT	Enable

- 3 Select *Enable identity-based Policy* and then select *Add*.

In the *New Authentication Rule* window, enter the following information, and then select OK:

User Group	FSSO_Internet_users
Service	ANY
Schedule	always
UTM	Optionally, enable UTM options.

- 4 Select OK.

To create a security policy for FSSO authentication - CLI

```
config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr Windows_net
    set dstaddr all
    set action accept
    set identity-based enable
    set nat enable
    config identity-based-policy
      edit 1
        set schedule always
        set groups FSSO_Internet_users
        set service ANY
      end
    end
  end
```

To create a security policy for local user authentication - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.
- 2 Enter the following information:

Source interface	Port3
Source address	Internal_net
Destination interface	Port1
Destination address	all

Action	ACCEPT
NAT	Enable

- 3 Select *Enable identity-based Policy* and then select *Add*.

In the *New Authentication Rule* window, enter the following information, and then select OK:

User Group	Internet_users
Service	ANY
Schedule	always
UTM	Optionally, enable UTM options.

- 4 Select OK.

To create a security policy for local user authentication - CLI

```
config firewall policy
  edit 0
    set srcintf port3
    set dstintf port1
    set srcaddr internal_net
    set dstaddr all
    set action accept
    set identity-based enable
    set nat enable
    config identity-based-policy
      edit 1
        set schedule always
        set groups Internet_users
        set service ANY
      end
    end
  end
```

LDAP Dial-in using member-attribute

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to setup a VPN connection simply based on an attribute that is set to TRUE, instead of based on the group they are in like normal. In AD the "Allow Dialin" property is activated in the user properties, and this sets the `msNPAllowDialin` attribute to "TRUE".

This same procedure can be used for other member attributes, as your system requires.

To accomplish this with a FortiGate unit, member-attribute must be set. This can only be accomplished through the CLI - the option is not available through the web-based manager.

Before configuring the FortiGate unit, ensure the AD server has the `msNPAllowDialin` attribute set to "TRUE" for the users in question. If not, those users will not be able to authenticate.

To configure user LDAP member-attribute settings - CLI

```
config user ldap
  edit "ldap_server"
    set server "192.168.201.3"
```

```

set cnid "sAMAccountName"
set dn "DC=fortilabanz,DC=com,DC=au"
set type regular
set username "fortigate@sample.com"
set password *****
set filter "(&(uid=%u)(msNPAllowDialin=TRUE))"
set member-attr "msNPAllowDialin"
next
end

```

To configure LDAP group settings - CLI

```

config user group
edit "ldap_grp"
set member "ldap"
config match
edit 1
set server-name "ldap"
set group-name "TRUE"
next
end
next
end

```

Once these settings are in place, users that are a member of the `ldap` user group will be able to authenticate.

To ensure your settings are correct, here is the sample output from a `diag debug` command that shows the authentication process.

When the "Allow Dial-in" attribute is set to "TRUE" the following will likely be in the output:

```

get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='TRUE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Passed group matching

```

If the attribute is not set but it is expected, the following will likely be in the output:

```

get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='FALSE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Failed group matching

```

The only difference between these two outputs is the last line which is either passed or failed based on if the member-attribute is set to the expected value or not.

Dynamic Profile example

A common dynamic profile topology involves a medium sized company network of users connecting to the Internet through the FortiGate unit, and authenticating with a RADIUS server. Dynamic profile authentication was selected because it is fast and relatively easy to configure.

This section includes:


- [Assumptions](#)
- [Topology](#)
- [General configuration](#)
- [Configuring RADIUS](#)
- [Configuring FortiGate regular and dynamic profile security policies](#)
- [Testing](#)

Assumptions

- VDOMs are not enabled
- The admin super_admin administrator account will be used for all FortiGate unit configuration.
- Any other devices on the network do not affect the topology of this example, and therefore are not included.
- Anywhere settings are not described, they are assumed to be default values.
- third-party RADIUS server is installed on a server box.
- BGP is used for any dynamic routing.
- Authentication event logging under Log&Report has been configured.

Topology

Example.com has an office with 20 users on the internal network. These users need access to the Internet to do their jobs. The office network is protected by a FortiGate-60C unit with access to the Internet through the wan1 interface, the user network on the internal interface, and all the servers are on the DMZ interface. This includes an Ubuntu Linux server running FreeRADIUS. For this example only two users will be configured — Pat Lee with an account name plee, or plee@example.com, and Kelly Green with an account name kgreen, or kgreen@example.com.

Figure 121: Dynamic Profile topology

General configuration

- 1 [Configuring RADIUS with users, user group, and FortiGate information.](#)
- 2 [Configuring FortiGate interfaces](#)
- 3 [Configuring dynamic profile RADIUS server on FortiGate](#)
- 4 [Configuring FortiGate regular and dynamic profile security policies](#)

Configuring RADIUS

Configuring RADIUS includes configuring the RADIUS server such as FreeRADIUS, a radius client on user's computers, and configuring users in the system. For this example the two users will be Pat Lee, and Kelly Green. They belong to a group called `exampledotcom_employees`. When it is all configured, the RADIUS daemon needs to started.

The users have a RADIUS client installed on their PCs that allows them to authenticate through the RADIUS server.

FreeRADIUS can be found on the freeradius.org website. For any problems installing FreeRADIUS, see the FreeRADIUS documentation.

Configuring FortiGate interfaces

Before configuring the dynamic profile security policy, configure FortiGate interfaces. This includes defining DHCP servers for the dmz and internal networks as these networks typically use DHCP. However, wan1 is assigned a static IP address by the ISP and does not need a DHCP server.

Table 84: FortiGate interfaces used in this example

Interface	Subnet	Act as DHCP Server	Devices
wan1	172.20.120.141	No	Internet Service Provider
dmz	10.11.101.100	Yes: x.x.x.110-.250	Servers, including RADIUS server
internal	10.11.102.100	Yes: x.x.x.110-.250	Internal user network

To configure FortiGate interfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select wan1 to edit.
- 3 Enter the following information and select OK.

Alias	Internet
Addressing Mode	Manual
IP/Netmask	172.20.120.141/255.255.255.0
Administrative Access	HTTPS, SSH
Comments	Internet
Administrative Status	Up

- 4 Select dmz to edit.
- 5 Enter the following information and select OK.

Alias	Servers
Addressing Mode	Manual
IP/Netmask	10.11.101.100/255.255.255.0
Administrative Access	HTTPS, SSH, PING, SNMP
Comments	Servers
Administrative Status	Up

- 6 Select internal to edit.
- 7 Enter the following information and select OK.

Alias	Internal network
Addressing Mode	Manual
IP/Netmask	10.11.102.100/255.255.255.0
Administrative Access	HTTPS, SSH, PING
Comments	Internal network
Administrative Status	Up

To configure DHCP servers - web-based manager

- 1 Go to *System > Network > DHCP Server*.
- 2 Select Create New.
- 3 Enter the following information and select OK.

Interface Name	dmz
Mode	Server

Enable	Select
Type	Regular
IP	10.11.101.110 - 10.11.101.250
Network mask	255.255.255.0
Default Gateway	10.11.101.100
DNS Service	Use System DNS Setting

- 4 Select Create New.
- 5 Enter the following information and select OK.

Interface Name	internal
Mode	Server
Enable	Select
Type	Regular
IP	10.11.102.110 - 10.11.102.250
Network mask	255.255.255.0
Default Gateway	10.11.102.100
DNS Service	Use System DNS Setting

Configuring dynamic profile RADIUS server on FortiGate

The FortiGate unit needs the RADIUS server information to be able to properly connect to it. Only one entry can be configured. If you cannot create one, ensure there is not already one configured.

The *Profile Key* entry must match the group name for users on the RADIUS server to be validated. Enable all logging if possible as it will assist with troubleshooting.

To configure the dynamic profile RADIUS server

- 1 Go to *User > Remote > RADIUS*, and *Create New*.
- 2 Enter the following information, and select OK.

Name	dynamic_profile_server
Type	Dynamic Start
Endpoint Attribute	Vendor-Specific
Profile Attribute	Vendor-Specific
Endpoint Blocking Attribute	Vendor-Specific
Send RADIUS Response	enable
Validate Radius Secret	enable
Profile Key	my_dyn_prof_group
Log All Events	enable

Configuring FortiGate regular and dynamic profile security policies

With the RADIUS server and FortiGate interfaces configured, security policies can be configured. This includes both dynamic profile and regular policies, as well as addresses and address groups. All policies require NAT to be enabled.

Table 85: security policies needed for dynamic profile

Seq. No.	From -> To	Type	Schedule	Description
1	internal -> wan1	dynamic profile	business hours	Authenticate outgoing user traffic.
2	internal -> wan1	regular	always	Allow essential network services and VoIP.
3	dmz -> wan1	regular	always	Allow servers to access Internet.
4	internal -> dmz	regular	always	Allow users to access servers.
5	any -> any	deny	always	Implicit deny all traffic that hasn't been matched



The dynamic profile policy must be placed at the top of the policy list so it is matched first. The only exception to this is if you have a policy to deny access to a list of banned users. In this case, that policy must go at the top so the dynamic profile does not mistakenly match a banned user or IP address.

This section includes:

- [Schedules, address groups, and services groups](#)
- [Configuring regular security policies](#)
- [Configuring dynamic profile security policy](#)

Schedules, address groups, and services groups

This section lists the lists that need to be configured before security policies are created. Creating these lists is straight forward, so the essential information has been provided here but not step by step instructions. For more information on firewall related details, see

Schedules

Only one schedule needs to be configured — `business_hours`. This is a fairly standard Monday to Friday 8am to 5pm schedule, or whatever days and hours covers standard work hours at the company.

Address groups

The following address groups need to be configured before the security policies.

Address group Name	Interface	Address range included
internal_network	internal	10.11.102.110 to 10.11.102.250
company_servers	dmz	10.11.101.110 to 10.11.101.250

Service groups

The following service groups need to be configured before the security policies. Note that the services listed are suggestions and may include more or less as required.

Service group Name	Interface	Description of services to be included
essential_network_services	internal	Any network protocols required for normal network operation such as DNS, NTP, BGP.
essential_server_services	dmz	All the protocols required by the company servers such as BGP, HTTP, HTTPS, FTP, IMAP, POP3, SMTP, IKE, SQL, MYSQL, NTP, TRACEROUTE, SOCKs, and SNMP.
dynamic_profile_services	internal	Any protocols required by users HTTP, HTTP, FTP,

The following security policy configurations are basic and only include logging, and default AV and IPS.

Configuring regular security policies

Regular security policies allow or deny all non-dynamic profile traffic. This is essential as there are network services—such as DNS, NTP, and FortiGuard—that require access to the Internet.

To configure regular security policies - web-based manager

- 1 Go to *Policy > Policy*, and select Create New.
- 2 Enter the following information, and select OK.

Source Interface/Zone	Internal
Source Address	internal_network
Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	essential_network_services
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	enable
UTM	enable
Enable Antivirus	enable Default
Enable IPS	enable Default
Enable VoIP	enable Default
Comments	Essential network services

- 3 Select Create New, enter the following information, and select OK.

Source Interface/Zone	dmz
Source Address	company_servers
Destination Interface/Zone	wan1

Destination Address	all
Schedule	always
Service	essential_server_services
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	enable
UTM	enable
Enable Antivirus	enable Default
Enable IPS	enable Default
Comments	Company servers accessing the Internet

- 4 Select Create New, enter the following information, and select OK.

Source Interface/Zone	Internal
Source Address	internal_network
Destination Interface/Zone	dmz
Destination Address	company_servers
Schedule	always
Service	all
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	enable
UTM	enable
Enable Antivirus	enable Default
Enable IPS	enable Default
Comments	Access company servers

Configuring dynamic profile security policy

A dynamic profile security policy has less configuration than a regular security policy. Before configuring the policy, you must configure a UTM Profile Group. This group name must match the RADIUS user group name.

If *UTM > Profile Group* is not visible, to go *System > Admin > Settings* and display it on the GUI.

To configure UTM profile group

- 1 Go to *UTM > Profile Group > Profile Group*, and select *Create New*.
- 2 Enter the name `exampledotdcom_employees`.
- 3 Enable default Antivirus, IPS, Web Filter, Email Filter, and Protocol Options.
- 4 Select OK.

To configure dynamic profile security policy

- 1 Go to *Policy > Policy*, and select *Create New*.
- 2 Enter the following information, and select OK.

Source Interface/Zone	Internal
Source Address	internal_network
Destination Interface/Zone	wan1
Destination Address	all
Schedule	business_hours
Service	all
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	enable
Enable Dynamic Profile	Enable
Profile Group	exampledotcom_employees
Dynamic Profile Users Only	enable
UTM	enable
Enable Antivirus	enable Default
Enable IPS	enable Default
Enable Web Filter	enable Default
Enable Email Filter	enable Default
Comments	Access company servers

Testing


Once configured, a user only needs to logon to their PC using their RADIUS account. After that when they attempt to access an Internet website, the FortiGate unit will use their session information to get their RADIUS information. Once the user is verified, they are allowed access to the website.

To test the configuration perform the following steps.

- 1 Have user pleee logon to their PC, and try to access an Internet website.
- 2 The FortiGate unit will contact the RADUS server for user pleee's information.
- 3 Once confirmed, pleee will have access to the website.

Each step generates log entries that enable you to verify that each step was successful.

If a step is unsuccessful, confirm your configuration is correct and see [“Troubleshooting dynamic profiles” on page 1356](#).

Figure 122: Dynamic profile test

Troubleshooting

In the web-based manager, a good tool for troubleshooting is the packet counter column on the security policy page (*Policy > Policy*). This column displays the number of packets that have passed through this security policy. Its value when you are troubleshooting is that when you are testing your configuration (end to end connectivity, user authentication, policy use) watching the packet count for an increase confirms any other methods you may be using for troubleshooting. It provides the key of which policy is allowing the traffic, useful information if you expect a user to require authentication and it never happens. For more information about authentication security policies, see [“Authentication in security policies” on page 1246](#).

This section addresses how to get more information from the CLI about users and user authentication attempts to help troubleshoot failed authentication attempts.

```
diag firewall iprope authuser
```

Shows the IP of where your computer is connected from. This is useful to confirm authorization and VPN settings.

```
diag firewall iprope resetauth
```

Clear all authorized users from the current list. Useful to force users to re-authenticate after system or group changes. However, this command may easily result in many users having to re-authenticate, so use carefully.

```
diag firewall auth list
```

List all the authorized users on this system.

```
diag debug disable
```

```
diag application auth 1
```

```
diag debug enable
```

See the stream of authentication system messages displayed enter the following commands. Best practices dictate the use of a terminal program that logs output so you can save the system messages to analyze later.

For more information on troubleshooting specific features, go to that section of this document. Most sections have troubleshooting information at the end of the section. In addition to that information, see the [FortiOS Handbook Troubleshooting chapter](#) for general troubleshooting information.



Chapter 8 IPsec VPNs

This FortiOS Handbook chapter contains the following sections:

[IPsec VPN concepts](#) explains the basic concepts that you need to understand about virtual private networks (VPNs).

[IPsec VPN Overview](#) provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

[IPsec VPN in the web-based manager](#) describes the IPsec VPN menu of the web-based manager interface.

[Gateway-to-gateway configurations](#) explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN. In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks.

[Hub-and-spoke configurations](#) describes how to set up hub-and-spoke IPsec VPNs. In a hub-and-spoke configuration, connections to a number of remote peers and/or clients radiate from a single, central FortiGate hub.

[Dynamic DNS configuration](#) describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a dynamic IP address and a domain name.

[FortiClient dialup-client configurations](#) guides you through configuring a FortiClient dialup-client IPsec VPN. In a FortiClient dialup-client configuration, the FortiGate unit acts as a dialup server and VPN client functionality is provided by the FortiClient Endpoint Security application installed on a remote host.

[FortiGate dialup-client configurations](#) explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit with a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.

[Supporting IKE Mode config clients](#) explains how to set up a FortiGate unit as either an IKE Mode Config server or client. IKE Mode Config is an alternative to DHCP over IPsec.

[Internet-browsing configuration](#) explains how to support secure web browsing performed by dialup VPN clients, and hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.

[Redundant VPN configurations](#) discusses the options for supporting redundant and partially redundant tunnels in an IPsec VPN configuration. A FortiGate unit can be configured to support redundant tunnels to the same remote peer if the FortiGate unit has more than one interface to the Internet.

[Transparent mode VPNs](#) describes two FortiGate units that create a VPN tunnel between two separate private networks transparently. In transparent mode, all FortiGate unit interfaces except the management interface are invisible at the network layer.

[Manual-key configurations](#) explains how to manually define cryptographic keys to establish an IPsec VPN tunnel. If one VPN peer uses specific authentication and encryption keys to establish a tunnel, both VPN peers must use the same encryption and authentication algorithms and keys.

[IPv6 IPsec VPNs](#) describes FortiGate unit VPN capabilities for networks based on IPv6 addressing. This includes IPv4-over-IPv6 and IPv6-over-IPv4 tunnelling configurations. IPv6 IPsec VPNs are available in FortiOS 3.0 MR5 and later.

[L2TP and IPsec \(Microsoft VPN\)](#) explains how to support Microsoft Windows native VPN clients.

[GRE over IPsec \(Cisco VPN\) configurations](#) explains how to interoperate with Cisco VPNs that use Generic Routing Encapsulation (GRE) protocol with IPsec.

[Protecting OSPF with IPsec](#) provides an example of protecting OSPF links with IPsec.

[Auto Key phase 1 parameters](#) provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The basic phase 1 parameters identify the remote peer or clients and support authentication through preshared keys or digital certificates. You can increase VPN connection security further using methods such as extended authentication (XAuth).

[Phase 2 parameters](#) provides detailed step-by-step procedures for configuring an IPsec VPN tunnel. During phase 2, the specific IPsec security associations needed to implement security services are selected and a tunnel is established.

[Defining VPN security policies](#) explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN tunnel, and how to define a security encryption policy. Security policies control all IP traffic passing between a source address and a destination address.

[Hardware offloading and acceleration](#) explains how to make use of FortiASIC network processor IPsec accelerated processing capabilities.

[Monitoring and troubleshooting](#) provides VPN monitoring and testing procedures



IPsec VPN concepts

Virtual Private Network (VPN) technology enables remote users to connect to private computer networks to gain access to their resources in a secure way. For example, an employee traveling or working from home can use a VPN to securely access the office network through the Internet.

Instead of remotely logging on to a private network using an unencrypted and unsecure Internet connection, the use of a VPN ensures that unauthorized parties cannot access the office network and cannot intercept any of the information that is exchanged between the employee and the office. It is also common to use a VPN to connect the private networks of two or more offices.

Fortinet offers VPN capabilities in the FortiGate Unified Threat Management (UTM) appliance and in the FortiClient Endpoint Security suite of applications. A FortiGate unit can be installed on a private network, and FortiClient software can be installed on the user's computer. It is also possible to use a FortiGate unit to connect to the private network instead of using FortiClient software.


This chapter discusses VPN terms and concepts including:

- [VPN tunnels](#)
- [VPN gateways](#)
- [Clients, servers, and peers](#)
- [Encryption](#)
- [Authentication](#)
- [Phase 1 and Phase 2 settings](#)
- [Security Association](#)

VPN tunnels

The data path between a user's computer and a private network through a VPN is referred to as a tunnel. Like a physical tunnel, the data path is accessible only at both ends. In the telecommuting scenario, the tunnel runs between the FortiClient application on the user's PC, or a FortiGate unit or other network device and the FortiGate unit on the office private network.

Encapsulation makes this possible. IPsec packets pass from one end of the tunnel to the other and contain data packets that are exchanged between the local user and the remote private network. Encryption of the data packets ensures that any third-party who intercepts the IPsec packets can not access the data.

Figure 123: Encoded data going through a VPN tunnel

You can create a VPN tunnel between:

- a PC equipped with the FortiClient application and a FortiGate unit
- two FortiGate units
- third-party VPN software and a FortiGate unit

Third-party VPN software is not covered in this document. Refer to the [Fortinet Knowledge Base](#) for more information on this topic.


VPN gateways

A gateway is a router that connects the local network to other networks. The default gateway setting in your computer's TCP/IP properties specifies the gateway for your local network.

A VPN gateway functions as one end of a VPN tunnel. It receives incoming IPsec packets, decrypts the encapsulated data packets and passes the data packets to the local network. Also, it encrypts data packets destined for the other end of the VPN tunnel, encapsulates them, and sends the IPsec packets to the other VPN gateway. The VPN gateway is a FortiGate unit because the private network behind it is protected, ensuring the security of the unencrypted VPN data. The gateway can also be FortiClient software running on a PC since the unencrypted data is secure on the PC.

The IP address of a VPN gateway is usually the IP address of the network interface that connects to the Internet. Optionally, you can define a secondary IP address for the interface and use that address as the local VPN gateway address. The benefit of doing this is that your existing setup is not affected by the VPN settings.

The following diagram shows a VPN connection between two private networks with FortiGate units acting as the VPN gateways. This configuration is commonly referred to as Gateway-to-Gateway IPsec VPN.

Figure 124: VPN tunnel between two private networks

Although the IPsec traffic may actually pass through many Internet routers, you can visualize the VPN tunnel as a simple secure connection between the two FortiGate units.


Users on the two private networks do not need to be aware of the VPN tunnel. The applications on their computers generate packets with the appropriate source and destination addresses, as they normally do. The FortiGate units manage all the details of encrypting, encapsulating and sending the packets to the remote VPN gateway.

The data is encapsulated in IPsec packets only in the VPN tunnel between the two VPN gateways. Between the user's computer and the gateway, the data is on the secure private network and it is in regular IP packets.

For example User1 on the Site A network, at IP address 10.10.1.7, sends packets with destination IP address 192.168.10.8, the address of User2 on the Site B network. The Site A FortiGate unit is configured to send packets with destinations on the 192.168.10.0 network through the VPN, encrypted and encapsulated. Similarly, the Site B FortiGate unit is configured to send packets with destinations on the 10.10.1.0 network through the VPN tunnel to the Site A VPN gateway.

In the site-to-site, or gateway-to-gateway VPN shown in [Figure 124](#), the FortiGate units have static (fixed) IP addresses and either unit can initiate communication.

You can also create a VPN tunnel between an individual PC running FortiClient and a FortiGate unit, as shown below. This is commonly referred to as Client-to-Gateway IPsec VPN.

Figure 125: VPN tunnel between a FortiClient PC and a FortiGate unit

On the PC, the FortiClient application acts as the local VPN gateway. Packets destined for the office network are encrypted, encapsulated into IPsec packets, and sent through the VPN tunnel to the FortiGate unit. Packets for other destinations are routed to the Internet as usual. IPsec packets arriving through the tunnel are decrypted to recover the original IP packets.

Clients, servers, and peers

A FortiGate unit in a VPN can have one of the following roles:

- **server** — responds to a request to establish a VPN tunnel.
- **client** — contacts a remote VPN gateway and requests a VPN tunnel.
- **peer** — brings up a VPN tunnel or responds to a request to do so.

The site-to-site VPN shown in [Figure 124](#) is a peer-to-peer relationship. Either FortiGate unit VPN gateway can establish the tunnel and initiate communications. The FortiClient-to-FortiGate VPN shown in [Figure 125](#) is a client-server relationship. The FortiGate unit establishes a tunnel when the FortiClient PC requests one.



A FortiGate unit cannot be a VPN server if it has a dynamically-assigned IP address. VPN clients need to be configured with a static IP address for the server.

A FortiGate unit acts as a server only when the remote VPN gateway has a dynamic IP address or is a client-only device or application, such as FortiClient.

As a VPN server, a FortiGate unit can also offer automatic configuration for FortiClient PCs. The user needs to know only the IP address of the FortiGate VPN server and a valid user name/password. FortiClient downloads the VPN configuration settings from the FortiGate VPN server. For information about configuring a FortiGate unit as a VPN server, see the [FortiClient Administration Guide](#).

Encryption

Encryption mathematically transforms data to appear as meaningless random numbers. The original data is called plaintext and the encrypted data is called ciphertext. The opposite process, called decryption, performs the inverse operation to recover the original plaintext from the ciphertext.

The process by which the plaintext is transformed to ciphertext and back again is called an algorithm. All algorithms use a small piece of information, a key, in the arithmetic process of converted plaintext to ciphertext, or vice-versa. IPsec uses symmetrical algorithms, in which the same key is used to both encrypt and decrypt the data.

The security of an encryption algorithm is determined by the length of the key that it uses. FortiGate IPsec VPNs offer the following encryption algorithms, in descending order of security:

AES256	A 128-bit block algorithm that uses a 256-bit key.
AES192	A 128-bit block algorithm that uses a 192-bit key.
AES128	A 128-bit block algorithm that uses a 128-bit key.
3DES	Triple-DES, in which plain text is DES-encrypted three times by three keys.
DES	Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key



The default encryption algorithms provided on FortiGate units make recovery of encrypted data almost impossible without the proper encryption keys.

There is a human factor in the security of encryption. The key must be kept secret, known only to the sender and receiver of the messages. Also, the key must not be something that unauthorized parties might easily guess, such as the sender's name, birthday or simple sequence such as 123456.

Authentication

In addition to protecting data through encryption, a VPN must ensure that only authorized users can access the private network. You must use either a preshared key on both VPN gateways or RSA X.509 security certificates. The examples in this guide use only preshared key authentication. Refer to the [Fortinet Knowledge Base](#) for articles on RSA X.509 security certificates.

Preshared keys

A preshared key contains at least six random alphanumeric characters. Users of the VPN must obtain the preshared key from the person who manages the VPN server and add the preshared key to their VPN client configuration.

Although it looks like a password, the preshared key, also known as a shared secret, is never sent by either gateway. The preshared key is used in the calculations at each end that generate the encryption keys. As soon as the VPN peers attempt to exchange encrypted data, preshared keys that do not match will cause the process to fail.

Additional authentication

To increase security, you can require additional means of authentication from users:

- an identifier, called a peer ID or a local ID
- extended authentication (XAUTH) which imposes an additional user name/password requirement

A Local ID is an alphanumeric value assigned in the Phase 1 configuration. The Local ID of a peer is called a Peer ID.

Phase 1 and Phase 2 settings

A VPN tunnel is established in two phases: Phase 1 and Phase 2. Several parameters determine how this is done. Except for IP addresses, the settings simply need to match at both VPN gateways. There are defaults that are appropriate for most cases.



FortiClient distinguishes between Phase 1 and Phase 2 only in the VPN Advanced settings and uses different terms. Phase 1 is called the IKE Policy. Phase 2 is called the IPsec Policy.

Phase 1

In Phase 1, the two VPN gateways exchange information about the encryption algorithms that they support and then establish a temporary secure connection to exchange authentication information.

When you configure your FortiGate unit or FortiClient application, you must specify the following settings for Phase 1:

Remote Gateway	The remote VPN gateway's address. FortiGate units also have the option of operating only as a server by selecting the "Dialup User" option.
Preshared key	This must be the same at both ends. It is used to encrypt phase 1 authentication information.
Local interface	The network interface that connects to the other VPN gateway. This applies on a FortiGate unit only.

All other Phase 1 settings have default values. These settings mainly configure the types of encryption to be used. The default settings on FortiGate units and in the FortiClient application are compatible. The examples in this guide use these defaults.

For more detailed information about Phase 1 settings, see the ["Auto Key phase 1 parameters" on page 1407](#).

Phase 2

Similar to the Phase 1 process, the two VPN gateways exchange information about the encryption algorithms that they support for Phase 2. You may choose different encryption for Phase 1 and Phase 2. If both gateways have at least one encryption algorithm in common, a VPN tunnel can be established. Keep in mind that more algorithms each phase does not share with the other gateway, the longer negotiations will take. In extreme cases this may cause timeouts during negotiations.

To configure default Phase 2 settings on a FortiGate unit, you need only select the name of the corresponding Phase 1 configuration. In FortiClient, no action is required to enable default Phase 2 settings.

For more detailed information about Phase 2 settings, see [“Phase 2 parameters” on page 1425](#).

Security Association

The establishment of a Security Association (SA) is the successful outcome of Phase 1 negotiations. Each peer maintains a database of information about VPN connections. The information in each SA can include cryptographic algorithms and keys, keylife, and the current packet sequence number. This information is kept synchronized as the VPN operates. Each SA has a Security Parameter Index (SPI) that is provided to the remote peer at the time the SA is established. Subsequent IPsec packets from the peer always reference the relevant SPI. It is possible for peers to have multiple VPNs active simultaneously, and correspondingly multiple SPIs.



IPsec VPN Overview

This section provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

The following topics are included in this section:

- [Types of VPNs](#)
- [Planning your VPN](#)
- [General preparation steps](#)
- [How to use this guide to configure an IPsec VPN](#)

VPN configurations interact with the firewall component of the FortiGate unit. There must be a security policy in place to permit traffic to pass between the private network and the VPN tunnel.

Security policies for VPNs specify:

- the FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optionally, a schedule that restricts when the VPN can operate
- optionally, the services (types of data) that can be sent

When the first packet of data that meets all of the conditions of the security policy arrives at the FortiGate unit, a VPN tunnel may be initiated and the encryption or decryption of data is performed automatically afterward. For more information, see [“Defining VPN security policies” on page 1431](#).

Types of VPNs

FortiGate unit VPNs can be policy-based or route-based. There is little difference between the two types. In both cases, you specify phase 1 and phase 2 settings. However there is a difference in implementation. A route-based VPN creates a virtual IPsec network interface that applies encryption or decryption as needed to any traffic that it carries. That is why route-based VPNs are also known as interface-based VPNs. A policy-based VPN is implemented through a special security policy that applies the encryption you specified in the phase 1 and phase 2 settings.

Route-based VPNs

For a route-based VPN, you create two security policies between the virtual IPsec interface and the interface that connects to the private network. In one policy the virtual interface is the source. In the other policy the virtual interface is the destination. The Action for both policies is Accept. This creates bidirectional policies that ensure traffic will flow in both directions over the VPN.

Policy-based VPNs

For a policy-based VPN, one security policy enables communication in both directions. You must select IPSEC as the Action and then select the VPN tunnel you defined in the phase 1 settings. You can then enable inbound and outbound traffic as needed within that policy, or create multiple policies of this type to handle different types of traffic differently. For example HTTPS traffic may not require the same level of scanning as FTP traffic.

Comparing policy-based or route-based VPNs

For both VPN types you create phase 1 and phase 2 configurations. Both types are handled in the stateful inspection security layer, assuming there is no IPS or AV. For more information on the security layers, see [“Life of a Packet” on page 699](#).

The main difference is in the security policy.

You create a policy-based VPN by defining an IPSEC security policy between two network interfaces and associating it with the VPN tunnel (phase 1) configuration.

You create a route-based VPN by enabling IPsec interface mode in the VPN phase 1 configuration. This creates a virtual IPsec interface. You then define a regular ACCEPT security policy to permit traffic to flow between the virtual IPsec interface and another network interface. And lastly, configure a static route to allow traffic over the VPN.

Where possible, you should create route-based VPNs. Generally, route-based VPNs are more flexible and easier to configure than policy-based VPNs — by default they are treated as interfaces. However, these two VPN types have different requirements that limit where they can be used.

Table 86: Comparison of policy-based and route-based VPNs

Features	Policy-based	Route-based
• Both NAT and transparent modes available	• Yes	• NAT mode only
• L2TP-over-IPsec supported	• Yes	• No
• GRE-over-IPsec supported	• No	• Yes
• security policy requirements	• Requires a security policy with IPSEC action that specifies the VPN tunnel	• Requires only a simple security policy with ACCEPT action
• Number of policies per VPN	• One policy controls connections in both directions	• A separate policy is required for connections in each direction

Planning your VPN

It is a good idea to plan the VPN configuration ahead of time. This will save time later and help you configure your VPN correctly.

All VPN configurations comprise a number of required and optional parameters. Before you begin, you need to determine:

- where does the IP traffic originate and where does it need to be delivered
- which hosts, servers, or networks to include in the VPN

- which VPN devices to include in the configuration
- through which interfaces the VPN devices communicate
- through which interfaces do private networks access the VPN gateways

Once you have this information, you can select a VPN topology that meets the requirements of your situation.

Network topologies

The topology of your network will determine how remote peers and clients connect to the VPN and how VPN traffic is routed. You can read about various network topologies and find the high-level procedures needed to configure IPsec VPNs in one of these sections.

Table 87: VPN network topologies and brief descriptions

Topology	Description
Gateway-to-gateway configurations	Standard one-to-one VPN between two FortiGate units. See “Gateway-to-gateway configurations” on page 1437 .
Hub-and-spoke configurations	One central FortiGate unit has multiple VPNs to other remote FortiGate units. See “Hub-and-spoke configurations” on page 1453 .
Dynamic DNS configuration	One end of the VPN tunnel has a changing IP address and the other end must go to a dynamic DNS server for the current IP address before establishing a tunnel. See “Dynamic DNS configuration” on page 1469 .
FortiClient dialup-client configurations	Typically remote FortiClient dialup-clients use dynamic IP addresses through NAT devices. The FortiGate unit acts as a dialup server allowing dialup VPN connections from multiple sources. See “FortiClient dialup-client configurations” on page 1483 .
FortiGate dialup-client configurations	Similar to FortiClient dialup-client configurations but with more gateway-to-gateway settings such as unique user authentication for multiple users on a single VPN tunnel. See “FortiGate dialup-client configurations” on page 1501 .
Internet-browsing configuration	Secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. See “Internet-browsing configuration” on page 1515 .
Redundant VPN configurations	Options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches. See “Redundant VPN configurations” on page 1519 .
Transparent mode VPNs	In transparent mode, the FortiGate acts as a bridge with all incoming traffic being broadcast back out on all other interfaces. Routing and NAT must be performed on external routers. See “Transparent mode VPNs” on page 1543 .
Manual-key configurations	Manually define cryptographic keys to establish an IPsec VPN, either policy-based or route-based. See “Manual-key configurations” on page 1551 .
L2TP and IPsec (Microsoft VPN)	Configure VPN for Microsoft Windows dialup clients using the built in L2TP software. Users do not have to install any See “L2TP and IPsec (Microsoft VPN)” on page 1567 .

These sections contain high-level configuration guidelines with cross-references to detailed configuration procedures. If you need more detail to complete a step, select the cross-reference in the step to drill-down to more detail. Return to the original procedure to complete the procedure. For a general overview of how to configure a VPN, see [“General preparation steps”](#) below.

General preparation steps

A VPN configuration defines relationships between the VPN devices and the private hosts, servers, or networks making up the VPN. Configuring a VPN involves gathering and recording the following information. You will need this information to configure the VPN.

- **The private IP addresses of participating hosts, servers, and/or networks.** These IP addresses represent the source addresses of traffic that is permitted to pass through the VPN. A IP source address can be an individual IP address, an address range, or a subnet address.
- **The public IP addresses of the VPN end-point interfaces.** The VPN devices establish tunnels with each other through these interfaces.
- **The private IP addresses associated with the VPN-device interfaces to the private networks.** Computers on the private networks behind the VPN gateways will connect to their VPN gateways through these interfaces.

How to use this guide to configure an IPsec VPN

This guide uses a task-based approach to provide all of the procedures needed to create different types of VPN configurations. Follow the step-by-step configuration procedures in this guide to set up the VPN.

The following configuration procedures are common to all IPsec VPNs:

- 1 Define the phase 1 parameters that the FortiGate unit needs to authenticate remote peers or clients and establish a secure a connection. See [“Auto Key phase 1 parameters”](#) on page 1407.
- 2 Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with a remote peer or dialup client. See [“Phase 2 parameters”](#) on page 1425.
- 3 Specify the source and destination addresses of IP packets that are to be transported through the VPN tunnel. See [“Defining policy addresses”](#) on page 1431.
- 4 Create an IPsec security policy to define the scope of permitted services between the IP source and destination addresses. See [“Defining VPN security policies”](#) on page 1432.



These steps assume you configure the FortiGate unit to generate unique IPsec encryption and authentication keys automatically. In situations where a remote VPN peer or client requires a specific IPsec encryption and authentication key, you must configure the FortiGate unit to use manual keys instead of performing Steps 1 and 2. For more information, see [“Manual-key configurations”](#) on page 1551.



IPsec VPN in the web-based manager

The IPsec VPN menu in FortiOS provides settings to configure an IPsec VPN. IPsec VPNs that are configured by using the general procedure:

- 1 Define phase 1 parameters to authenticate remote peers and clients for a secure connection. See [“Phase 1 configuration” on page 1394](#).
- 2 Define phase 2 parameters to create a VPN tunnel with a remote peer or dialup client. See [“Phase 2 configuration” on page 1399](#).



With these steps, your FortiGate unit will automatically generate unique IPsec encryption and authentication keys. If a remote VPN peer or client requires a specific IPsec encryption or authentication key, you must configure your FortiGate unit to use manual keys instead. See [“Manual Key” on page 1403](#).

- 3 Create a security policy to permit communication between your private network and the VPN. Policy-based VPNs have an action of IPSEC, where for interface-based VPNs the security policy action is ACCEPT. See [“Defining VPN security policies” on page 1431](#).

The FortiGate unit implements the Encapsulated Security Payload (ESP) protocol. Internet Key Exchange (IKE) is performed automatically based on pre-shared keys or X.509 digital certificates. As an option, you can specify manual keys. Interface mode, supported in NAT mode only, creates a virtual interface for the local end of a VPN tunnel.

This topic contains the following:

- [Auto Key \(IKE\)](#)
- [Manual Key](#)
- [Concentrator](#)

Auto Key (IKE)

You can configure VPN peers (or a FortiGate dialup server and a VPN client) to generate unique Internet Key Exchange (IKE) keys automatically during the IPsec phase 1 and phase 2 exchanges.

When you define phase 2 parameters, you can choose any set of phase 1 parameters to set up a secure connection for the tunnel and authenticate the remote peer. Auto Key configuration applies to both tunnel-mode and interface-mode VPNs.

To configure VPN peers go to *VPN > IPsec > Auto Key (IKE)*.

Auto Key (IKE) page	
Create Phase 1	Creates a new phase 1 tunnel configuration. For more information, see “Phase 1 configuration” on page 1394 .
Create Phase 2	Creates a new phase 2 configuration. For more information, see “Phase 2 configuration” on page 1399 .
Create FortiClient VPN	Creates a new FortiClient VPN. For more information, see “FortiClient VPN” on page 1402 .



If you want to control how the IKE negotiation process controls traffic when there is no traffic, as well as the length of time the FortiGate unit waits for negotiations to occur, use the `negotiation-timeout` and `auto-negotiation` commands in the CLI.

Phase 1 configuration

The basic phase 1 settings associate IPsec phase 1 parameters with a remote gateway, if a pre-shared key or digital certificate will be used, and if a special identifier will be used to identify the remote VPN peer or client.

New Phase 1 page	
Name	<p>Type a name for the phase 1 definition. The maximum name length is 15 characters for an interface mode VPN, 35 characters for a policy-based VPN. If <i>Remote Gateway</i> is <i>Dialup User</i>, the maximum name length is further reduced depending on the number of dialup tunnels that can be established: by 2 for up to 9 tunnels, by 3 for up to 99 tunnels, 4 for up to 999 tunnels, and so on.</p> <p>For a tunnel mode VPN, the name normally reflects where the remote connection originates. For a route-based tunnel, the FortiGate unit also uses the name for the virtual IPsec interface that it creates automatically.</p>
Remote Gateway	<p>Select the category of the remote connection:</p> <ul style="list-style-type: none"> • <i>Static IP Address</i> — If the remote peer has a static IP address. • <i>Dialup User</i> — If one or more FortiClient or FortiGate dialup clients with dynamic IP addresses will connect to the FortiGate unit. • <i>Dynamic DNS</i> — If a remote peer that has a domain name and subscribes to a dynamic DNS service will connect to the FortiGate unit.
IP Address	If you selected <i>Static IP Address</i> , enter the IP address of the remote peer.
Dynamic DNS	If you selected <i>Dynamic DNS</i> , enter the domain name of the remote peer.

Local Interface	<p>This option is available in NAT mode only. Select the name of the interface through which remote peers or dialup clients connect to the FortiGate unit.</p> <p>By default, the local VPN gateway IP address is the IP address of the interface that you selected. Optionally, you can specify a unique IP address for the VPN gateway in the <i>Advanced</i> settings.</p>
Mode	<p>Select <i>Main (ID Protection)</i> or <i>Aggressive</i>:</p> <ul style="list-style-type: none"> • <i>Main mode</i> — the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information. • <i>Aggressive mode</i> — the phase 1 parameters are exchanged in single message with authentication information that is not encrypted. <p>When the remote VPN peer has a dynamic IP address and is authenticated by a pre-shared key, you must select <i>Aggressive</i> mode if there is more than one dialup phase1 configuration for the interface IP address.</p> <p>When the remote VPN peer has a dynamic IP address and is authenticated by a certificate, you must select <i>Aggressive</i> mode if there is more than one phase 1 configuration for the interface IP address and these phase 1 configurations use different proposals.</p>
Authentication Method	Select <i>Preshared Key</i> or <i>RSA Signature</i> .
Pre-shared Key	If you selected <i>Pre-shared Key</i> , enter the pre-shared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. You must define the same key at the remote peer or client. The key must contain at least 6 printable characters. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.
Certificate Name	If you selected <i>RSA Signature</i> , select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. For information about obtaining and loading the required server certificate, see the FortiOS User Authentication guide .
Peer Options	Peer options are available to authenticate VPN peers or clients, depending on the <i>Remote Gateway</i> and <i>Authentication Method</i> settings.
Accept any peer ID	<p>Accept the local ID of any remote VPN peer or client. The FortiGate unit does not check identifiers (local IDs). You can set <i>Mode</i> to <i>Aggressive</i> or <i>Main</i>.</p> <p>You can use this option with <i>RSA Signature</i> authentication. But, for highest security, configure a PKI user/group for the peer and set <i>Peer Options</i> to <i>Accept this peer certificate only</i>.</p>

Accept this peer ID	<p>This option is available when <i>Aggressive Mode</i> is enabled. Enter the identifier that is used to authenticate the remote peer. This identifier must match the Local ID that the remote peer's administrator has configured.</p> <p>If the remote peer is a FortiGate unit, the identifier is specified in the <i>Local ID</i> field of the Advanced phase 1 configuration.</p> <p>If the remote peer is a FortiClient user, the identifier is specified in the <i>Local ID</i> field, accessed by selecting <i>Config</i> in the <i>Policy</i> section of the VPN connection's <i>Advanced Settings</i>.</p>
Accept peer ID in dialup group	<p>Authenticate multiple FortiGate or FortiClient dialup clients that use unique identifiers and unique pre-shared keys (or unique pre-shared keys only) through the same VPN tunnel.</p> <p>You must create a dialup user group for authentication purposes. Select the group from the list next to the <i>Accept peer ID in dialup group</i> option.</p> <p>You must set <i>Mode</i> to <i>Aggressive</i> when the dialup clients use unique identifiers and unique pre-shared keys. If the dialup clients use unique pre-shared keys only, you can set <i>Mode</i> to <i>Main</i> if there is only one dialup phase 1 configuration for this interface IP address.</p>
Advanced	<p>Defines advanced phase 1 parameters. For more information, see Phase 1 advanced configuration settings.</p>

Phase 1 advanced configuration settings

You use the advanced parameters to select the encryption and authentication algorithms that the FortiGate unit uses to generate keys for the IKE exchange. You can also select these advanced settings to ensure the smooth operation of phase 1 negotiations.

To configure Phase 1 settings, go to *VPN > Auto Key (IKE)* and select *Create Phase 1*.

Advanced section of the New Phase 1 page	
Enable IPsec Interface Mode	<p>This is available in NAT mode only.</p> <p>Create a virtual interface for the local end of the VPN tunnel. Select this option to create a route-based VPN, clear it to create a policy-based VPN.</p>
IKE Version	<p>Select the version of IKE to use. This is available only if <i>IPsec Interface Mode</i> is enabled. For more information about IKE v2, refer to RFC 4306.</p> <p>IKE v2 is not available if <i>Mode</i> is <i>Aggressive</i>.</p> <p>When <i>IKE Version</i> is 2, <i>Mode</i> and <i>XAUTH</i> are not available.</p>
IPv6 Version	<p>Select if you want to use IPv6 addresses for the remote gateway and interface IP addresses. This is available only when <i>Enable IPsec Interface Mode</i> is selected and IPv6 Support is enabled in the administrative settings (<i>System > Admin > Settings</i>).</p>

Local Gateway IP	<p>If you selected <i>Enable IPsec Interface Mode</i>, specify an IP address for the local end of the VPN tunnel. Select one of the following:</p> <ul style="list-style-type: none"> • <i>Main Interface IP</i> — The FortiGate unit obtains the IP address of the interface from the network interface settings. • <i>Specify</i> — Enter a secondary address of the interface selected in the phase 1 <i>Local Interface</i> field. For more information, see “Local Interface” on page 1395. <p>You cannot configure Interface mode in a transparent mode VDOM.</p>
P1 Proposal	<p>Select the encryption and authentication algorithms used to generate keys for protecting negotiations and add encryption and authentication algorithms as required.</p> <p>You need to select a minimum of one and a maximum of three combinations. The remote peer or client must be configured to use at least one of the proposals that you define.</p>
	<p>Select one of the following symmetric-key encryption algorithms:</p> <ul style="list-style-type: none"> • <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key. • <i>3DES</i> — Triple-DES, in which plain text is encrypted three times by three keys. • <i>AES128</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 128-bit key. • <i>AES192</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 192-bit key. • <i>AES256</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 256-bit key.
	<p>Select either of the following authentication message digests to check the authenticity of messages during phase 1 negotiations:</p> <ul style="list-style-type: none"> • <i>MD5</i> — Message Digest 5, the hash algorithm developed by RSA Data Security. • <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest. • <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest. <p>To specify a third combination, use the <i>Add</i> button beside the fields for the second combination.</p>
DH Group	<p>Select one or more Diffie-Hellman groups from DH group 1, 2, 5 and 14. At least one of the <i>DH Group</i> settings on the remote peer or client must match one the selections on the FortiGate unit. Failure to match one or more DH groups will result in failed negotiations.</p>
Keylife	<p>Enter the time (in seconds) that must pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172 800 seconds.</p>

Local ID	<p>If the FortiGate unit will act as a VPN client and you are using peer IDs for authentication purposes, enter the identifier that the FortiGate unit will supply to the VPN server during the phase 1 exchange.</p> <p>If the FortiGate unit will act as a VPN client, and you are using security certificates for authentication, select the distinguished name (DN) of the local server certificate that the FortiGate unit will use for authentication purposes.</p> <p>If the FortiGate unit is a dialup client and will not be sharing a tunnel with other dialup clients (that is, the tunnel will be dedicated to this FortiGate dialup client), set <i>Mode</i> to <i>Aggressive</i>. Note that this Local ID value must match the peer ID value given for the remote VPN peer's Peer Options.</p>
XAuth	<p>This option supports the authentication of dialup clients. It is available for IKE v1 only.</p> <ul style="list-style-type: none"> • <i>Disable</i> — Select if you do not use XAuth. • <i>Enable as Client</i> — If the FortiGate unit is a dialup client, enter the user name and password that the FortiGate unit will need to authenticate itself to the remote XAuth server. • <i>Enable as Server</i> — This is available only if <i>Remote Gateway</i> is set to <i>Dialup User</i>. Dialup clients authenticate as members of a dialup user group. You must first create a user group for the dialup clients that need access to the network behind the FortiGate unit. <p>You must also configure the FortiGate unit to forward authentication requests to an external RADIUS or LDAP authentication server.</p> <p>Select a <i>Server Type</i> setting to determine the type of encryption method to use between the FortiGate unit, the XAuth client and the external authentication server, and then select the user group from the User Group list.</p>
Username	Enter the user name that is used for authentication.
Password	Enter the password that is used for authentication.
NAT Traversal	Select the check box if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared) to connect reliably.

Keepalive Frequency	If you enabled <i>NAT-traversal</i> , enter a keepalive frequency setting.
Dead Peer Detection	<p>Select this check box to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. You can use this option to receive notification whenever a tunnel goes up or down, or to keep the tunnel connection open when no traffic is being generated inside the tunnel. For example, in scenarios where a dialup client or dynamic DNS peer connects from an IP address that changes periodically, traffic may be suspended while the IP address changes.</p> <p>With <i>Dead Peer Detection</i> selected, you can use the <code>config vpn ipsec phase1</code> (tunnel mode) or <code>config vpn ipsec phase1-interface</code> (interface mode) CLI command to optionally specify a retry count and a retry interval.</p>

Phase 2 configuration

After IPsec phase 1 negotiations end successfully, you begin phase 2. You configure the phase 2 parameters to define the algorithms that the FortiGate unit may use to encrypt and transfer data for the remainder of the session. During phase 2, you select specific IPsec security associations needed to implement security services and establish a tunnel.

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration that specifies the remote end point of the VPN tunnel. In most cases, you need to configure only basic phase 2 settings.

To configure Phase 2 settings, go to *VPN > Auto Key (IKE)* and select *Create Phase 2*.

New Phase 2 page	
Name	Type a name to identify the phase 2 configuration.
Phase 1	Select the phase 1 tunnel configuration. For more information on configuring phase 1, see “Phase 1 configuration” on page 1394 . The phase 1 configuration describes how remote VPN peers or clients will be authenticated on this tunnel, and how the connection to the remote peer or client will be secured.
Advanced	Define advanced phase 2 parameters. For more information, see “Phase 2 advanced configuration settings” on page 1399 .

Phase 2 advanced configuration settings

In phase 2, the FortiGate unit and the VPN peer or client exchange keys again to establish a secure communication channel between them. You select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). These are called P2 Proposal parameters. The keys are generated automatically using a Diffie-Hellman algorithm.

You can use a number of additional advanced phase 2 settings to enhance the operation of the tunnel.

Advanced section of New Phase 2 page	
P2 Proposal	<p>Select the encryption and authentication algorithms that will be proposed to the remote VPN peer. You can specify up to three proposals. To establish a VPN connection, at least one of the proposals that you specify must match configuration on the remote peer.</p> <p>Initially there are two proposals. <i>Add</i> and <i>Delete</i> icons are next to the second <i>Authentication</i> field.</p> <p>It is invalid to set both <i>Encryption</i> and <i>Authentication</i> to NULL.</p>
Encryption	<p>Select one of the following symmetric-key algorithms:</p> <ul style="list-style-type: none"> • <i>NULL</i> — Do not use an encryption algorithm. • <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key. • <i>3DES</i> — Triple-DES, in which plain text is encrypted three times by three keys. • <i>AES128</i> — a 128-bit block Cipher Block Chaining (CBC) algorithm that uses a 128-bit key. • <i>AES192</i> — a 128-bit block CBC algorithm that uses a 192-bit key. • <i>AES256</i> — a 128-bit block CBC algorithm that uses a 256-bit key.
Authentication	<p>Select one of the following message digests to check the authenticity of messages during an encrypted session:</p> <ul style="list-style-type: none"> • <i>NULL</i> — Do not use a message digest. • <i>MD5</i> — Message Digest 5, the hash algorithm developed by RSA Data Security. • <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest. • <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest. • <i>SHA384</i> — Secure Hash Algorithm 2, which produces a 384-bit message digest. • <i>SHA512</i> — Secure Hash Algorithm 2, which produces a 512-bit message digest.
Enable replay detection	Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.
Enable perfect forward secrecy (PFS)	Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.
DH Group	Select one Diffie-Hellman group (1, 2, 5 or 14). This must match the DH Group that the remote peer or dialup client uses.
Keylife	Select the method for determining when the phase 2 key expires: <i>Seconds</i> , <i>KBytes</i> , or <i>Both</i> . If you select <i>Both</i> , the key expires when either the time has passed or the number of KB have been processed.

Autokey Keep Alive	Select the check box if you want the tunnel to remain active when no data is being processed.
DHCP-IPSec	<p>Provide IP addresses dynamically to VPN clients. This is available for phase 2 configurations associated with a dialup phase 1 configuration. You also need configure a DHCP server or relay on the private network interface. You must configure the DHCP parameters separately.</p> <p>If you configure the DHCP server to assign IP addresses based on RADIUS user group attributes, you must also set the Phase 1 <i>Peer Options</i> to <i>Accept peer ID in dialup group</i> and select the appropriate user group. See “Phase 1 configuration” on page 1394.</p> <p>If the FortiGate unit acts as a dialup server and you manually assigned FortiClient dialup clients VIP addresses that match the network behind the dialup server, selecting the check box will cause the FortiGate unit to act as a proxy for the dialup clients.</p>
Quick Mode Selector	<p>Specify the source and destination IP addresses to be used as selectors for IKE negotiations. If the FortiGate unit is a dialup server, keep the default value of 0.0.0.0/0 unless you need to circumvent problems caused by ambiguous IP addresses between one or more of the private networks making up the VPN. You can specify a single host IP address, an IP address range, or a network address. You may optionally specify source and destination port numbers and a protocol number.</p> <p>If you are editing an existing phase 2 configuration, the <i>Source address</i> and <i>Destination address</i> fields are unavailable if the tunnel has been configured to use firewall addresses as selectors. This option exists only in the CLI.</p>
Source address	<p>If the FortiGate unit is a dialup server, enter the source IP address that corresponds to the local senders or network behind the local VPN peer (for example, 172.16.5.0/24 or 172.16.5.0/255.255.255.0 for a subnet, or 172.16.5.1/32 or 172.16.5.1/255.255.255.255 for a server or host, or 192.168.10.[80-100] or 192.168.10.80-192.168.10.100 for an address range). A value of 0.0.0.0/0 means all IP addresses behind the local VPN peer.</p> <p>If the FortiGate unit is a dialup client, source address must refer to the private network behind the FortiGate dialup client.</p>
Source port	Enter the port number that the local VPN peer uses to transport traffic related to the specified service (protocol number). The range is from 0 to 65535. To specify all ports, type 0.
Destination address	Enter the destination IP address that corresponds to the recipients or network behind the remote VPN peer (for example, 192.168.20.0/24 for a subnet, or 172.16.5.1/32 for a server or host, or 192.168.10.[80-100] for an address range). A value of 0.0.0.0/0 means all IP addresses behind the remote VPN peer.
Destination port	Enter the port number that the remote VPN peer uses to transport traffic related to the specified service (protocol number). To specify all ports, enter 0.
Protocol	Enter the IP protocol number of the service. To specify all services, enter 0.

FortiClient VPN

Use the FortiClient VPN configuration settings when configuring an IPsec VPN for FortiClient. When configuring a FortiClient VPN connection, the settings for phase 1 and phase 2 settings are automatically configured by the FortiGate unit. They are set to:

- Remote Gateway — Dialup User
- Mode — Aggressive
- IPSec Interface Mode — Enabled
- Default settings for P1 and P2 Proposal
- XAUTH Enable as Server (Auto)
- IKE mode-config will be enabled
- Peer Option — "Accept any peer ID"

The remainder of the settings use the current FortiGate defaults. Note that FortiClient settings need to match these FortiGate defaults.

If you need to configure advanced settings for the FortiClient VPN, select *Edit* on the Auto Key (IKE) page (Go to *VPN > IPsec > Auto Key (IKE)*) and configure the peer options or advanced options.

New FortiClient VPN page	
Name	Enter a name for the FortiClient VPN.
Local Outgoing Interface	Select the local outgoing interface for the VPN.
Authentication Method	Select the type of authentication used when logging in to the VPN.
Preshared Key	If <i>Pre-shared Key</i> was selected in <i>Authentication Method</i> , enter the pre-shared key in the field provided.
User Group	Select a user group. You can also create a user group from the drop-down list by selecting <i>Create New</i> .
Address Range Start IP	Enter the start IP address for the DHCP address range for the client.
Address Range End IP	Enter the end IP address for the address range.
Subnet Mask	Enter the subnet mask.
Enable IPv4 Split Tunnel	Enabled by default, this option enables the FortiClient user to use the VPN to access internal resources while other Internet access is not sent over the VPN, alleviating potential traffic bottlenecks in the VPN connection. Disable this option to have all traffic sent through the VPN tunnel.
DNS Server	Select which DNS server to use for this VPN: <ul style="list-style-type: none"> • <i>Use System DNS</i> — Use the same DNS servers as the FortiGate unit. These are configured at <i>System > Interface > DNS</i>. This is the default option. • <i>Specify</i> — Specify the IP address of a different DNS server.

Manual Key



Use manual keys only if it is unavoidable. There are potential difficulties in keeping keys confidential and in propagating changed keys to remote VPN peers securely.

If required, you can manually define cryptographic keys for establishing an IPsec VPN tunnel. You would define manual keys in situations where:

- you require prior knowledge of the encryption or authentication key (that is, one of the VPN peers requires a specific IPsec encryption or authentication key).
- you need to disable encryption and authentication.

In both cases, you do not specify IPsec phase 1 and phase 2 parameters; you define manual keys by going to *VPN > IPsec > Manual Key* instead.

Manual key configuration settings



If you are not familiar with the security policies, SAs, selectors, and SA databases for your particular installation, do not attempt these procedures without qualified assistance.

If one of the VPN devices is manually keyed, the other VPN device must also be manually keyed with the identical authentication and encryption keys. In addition, it is essential that both VPN devices be configured with complementary Security Parameter Index (SPI) settings. The administrators of the devices need to cooperate to achieve this.

Each SPI identifies a Security Association (SA). The value is placed in ESP datagrams to link the datagrams to the SA. When an ESP datagram is received, the recipient refers to the SPI to determine which SA applies to the datagram. You must manually specify an SPI for each SA. There is an SA for each direction, so for each VPN you must specify two SPIs, a local SPI and a remote SPI, to cover bidirectional communications between two VPN devices.

To add a manual key, go to *VPN > IPsec > Manual Key* and select *Create New*.

New Manual Key page

Name	Type a name for the VPN tunnel. The maximum name length is 15 characters for an interface mode VPN, 35 characters for a policy-based VPN.
Local SPI	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles outbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the Remote SPI value in the manual key configuration at the remote peer.
Remote SPI	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles inbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the Local SPI value in the manual key configuration at the remote peer.
Remote Gateway	Enter the IP address of the public interface to the remote peer. The address identifies the recipient of ESP datagrams.

Local Interface	This option is available in NAT mode only. Select the name of the interface to which the IPsec tunnel will be bound. The FortiGate unit obtains the IP address of the interface from the network interface settings.
Encryption Algorithm	<p>Select one of the following symmetric-key encryption algorithms:</p> <ul style="list-style-type: none"> • <i>NULL</i> — Do not use an encryption algorithm. • <i>DES</i> — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key. • <i>3DES</i> — Triple-DES, where plain text is encrypted three times by three keys. • <i>AES128</i> — a 128-bit block Cipher Block Chaining algorithm that uses a 128-bit key. • <i>AES192</i> — a 128-bit block Cipher Block Chaining) algorithm that uses a 192-bit key. • <i>AES256</i> — a 128-bit block Cipher Block Chaining algorithm that uses a 256-bit key. <p>Note: The algorithms for encryption and authentication cannot both be NULL.</p>
Authentication Algorithm	<p>Select one of the following message digests:</p> <ul style="list-style-type: none"> • <i>NULL</i> — Do not use a message digest. • <i>MD5</i> — Message Digest 5 algorithm, which produces a 128-bit message digest. • <i>SHA1</i> — Secure Hash Algorithm 1, which produces a 160-bit message digest. • <i>SHA256</i> — Secure Hash Algorithm 2, which produces a 256-bit message digest. • <i>SHA384</i> — Secure Hash Algorithm 2, which produces a 384-bit message digest. • <i>SHA512</i> — Secure Has Algorithm 2, which produces a 512-bit message digest. <p>Note: The Algorithms for encryption and authentication cannot both be NULL.</p>
IPsec Interface Mode	<p>Create a virtual interface for the local end of the VPN tunnel. Select this check box to create a route-based VPN, clear it to create a policy-based VPN.</p> <p>This is available only in NAT mode.</p>

Concentrator

In a hub-and-spoke configuration, policy-based VPN connections to a number of remote peers radiate from a single, central FortiGate unit. Site-to-site connections between the remote peers do not exist; however, you can establish VPN tunnels between any two of the remote peers through the FortiGate unit's "hub".

In a hub-and-spoke network, all VPN tunnels terminate at the hub. The peers that connect to the hub are known as "spokes". The hub functions as a concentrator on the network, managing all VPN connections between the spokes. VPN traffic passes from one tunnel to the other through the hub.

You define a concentrator to include spokes in the hub-and-spoke configuration. You create the concentrator in *VPN > IPsec > Concentrator* and select *Create New*. A concentrator configuration specifies which spokes to include in an IPsec hub-and-spoke configuration.

New VPN Concentrator page	
Concentrator Name	Type a name for the concentrator.
Available Tunnels	A list of defined IPsec VPN tunnels. Select a tunnel from the list and then select the right arrow.
Members	A list of tunnels that are members of the concentrator. To remove a tunnel from the concentrator, select the tunnel and select the left arrow.

IPsec Monitor

You can use the IPsec Monitor to view activity on IPsec VPN tunnels and start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels, including tunnel mode and route-based (interface mode) tunnels.

To view the IPsec monitor, go to *VPN > Monitor > IPsec Monitor*.

For dialup VPNs, the list provides status information about the VPN tunnels established by dialup clients, and their IP addresses.

For static IP or dynamic DNS VPNs, the list provides status and IP addressing information about VPN tunnels, active or not, to remote peers that have static IP addresses or domain names. You can also start and stop individual tunnels from the list.



Auto Key phase 1 parameters

This chapter provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The phase 1 parameters identify the remote peer or clients and support authentication through preshared keys or digital certificates. You can increase access security further using peer identifiers, certificate distinguished names, group names, or the FortiGate extended authentication (XAuth) option for authentication purposes.

For more information on phase 1 parameters in the web-based manager, see [“Phase 1 configuration” on page 1394](#).



The information and procedures in this section do not apply to VPN peers that perform negotiations using manual keys. Refer to [“Manual-key configurations” on page 1551](#) instead.

The following topics are included in this section:

- [Overview](#)
- [Defining the tunnel ends](#)
- [Choosing main mode or aggressive mode](#)
- [Authenticating the FortiGate unit](#)
- [Authenticating remote peers and clients](#)
- [Defining IKE negotiation parameters](#)
- [Using XAuth authentication](#)

Overview

To configure IPsec phase 1 settings, go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*. IPsec phase 1 settings define:

- the remote and local ends of the IPsec tunnel
- if phase 1 parameters are exchanged in multiple rounds with encrypted authentication information (main mode) or in a single message with authentication information that is not encrypted (aggressive mode)
- if a preshared key or digital certificates will be used to authenticate the FortiGate unit to the VPN peer or dialup client
- if the VPN peer or dialup client is required to authenticate to the FortiGate unit. A remote peer or dialup client can authenticate by peer ID or, if the FortiGate unit authenticates by certificate, it can authenticate by peer certificate.
- the IKE negotiation proposals for encryption and authentication
- optional XAuth authentication, which requires the remote user to enter a user name and password. A FortiGate VPN server can act as an XAuth server to authenticate dialup users. A FortiGate unit that is a dialup client can also be configured as an XAuth client to authenticate itself to the VPN server.

For all the phase 1 web-based manager fields, see “[Phase 1 configuration](#)” on [page 1394](#).



If you want to control how the IKE negotiation process controls traffic when there is no traffic, as well as the length of time the unit waits for negotiations to occur, use the `negotiation-timeout` and `auto-negotiation` commands in the CLI.

Defining the tunnel ends

To begin defining the phase 1 configuration, go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*. Enter a descriptive name for the VPN tunnel. This is particularly important if you will create several tunnels.

The phase 1 configuration mainly defines the ends of the IPsec tunnel. The remote end is the remote gateway with which the FortiGate unit exchanges IPsec packets. The local end is the FortiGate interface that sends and receives IPsec packets.

The remote gateway can be:

- a static IP address
- a domain name with a dynamic IP address
- a dialup client

A statically addressed remote gateway is the simplest to configure. You specify the IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer has a domain name and subscribes to a dynamic DNS service, you need to specify only the domain name. The FortiGate unit performs a DNS query to determine the appropriate IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer is a dialup client, only the dialup client can bring up the tunnel. The IP address of the client is not known until it connects to the FortiGate unit. This configuration is a typical way to provide a VPN for client PCs running VPN client software such as the FortiClient Endpoint Security application.

The local end of the VPN tunnel, the Local Interface, is the FortiGate interface that sends and receives the IPsec packets. This is usually the public interface of the FortiGate unit that is connected to the Internet. Packets from this interface pass to the private network through a security policy.



The local Interface for a Phase 1 cannot not be a loopback interface. By design, the IPSec tunnel will not be established if this happens.

By default, the local VPN gateway is the IP address of the selected Local Interface. If you are configuring an interface mode VPN, you can optionally use a secondary IP address of the Local Interface as the local gateway.

Choosing main mode or aggressive mode

The FortiGate unit and the remote peer or dialup client exchange phase 1 parameters in either Main mode or Aggressive mode. This choice does not apply if you use IKE version 2, which is available only for route-based configurations.

- In Main mode, the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information
- In Aggressive mode, the phase 1 parameters are exchanged in single message with authentication information that is not encrypted.

Although Main mode is more secure, you must select Aggressive mode if there is more than one dialup phase 1 configuration for the interface IP address, and the remote VPN peer or client is authenticated using an identifier local ID). Descriptions of the peer options in this guide indicate whether Main or Aggressive mode is required.

Choosing the IKE version

If you create a route-based VPN, you have the option of selecting IKE version 2. Otherwise, IKE version 1 is used.

IKEv2, defined in RFC 4306, simplifies the negotiation process that creates the security association (SA).

If you select IKEv2:

- There is no choice in Phase 1 of Aggressive or Main mode.
- FortiOS does not support Peer Options or Local ID.
- Extended Authentication (XAUTH) is not available.
- You can select only one DH Group.

Authenticating the FortiGate unit

The FortiGate unit can authenticate itself to remote peers or dialup clients using either a pre-shared key or an RSA Signature (certificate).

Authenticating the FortiGate unit with digital certificates

To authenticate the FortiGate unit using digital certificates, you must have the required certificates installed on the remote peer and on the FortiGate unit. The signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer. If you use certificates to authenticate the FortiGate unit, you can also require the remote peers or dialup clients to authenticate using certificates.

For more information about obtaining and installing certificates, see the [FortiOS User Authentication guide](#).

To authenticate the FortiGate unit using digital certificates

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Create a new phase 1 configuration or edit an existing phase 1 configuration.
- 3 Include appropriate entries as follows:

Name	Enter a name that reflects the origination of the remote connection. For interface mode, the name can be up to 15 characters long.
Remote Gateway	Select the nature of the remote connection. Each option changes the available fields you must configure. For more information, see “Defining the tunnel ends” on page 1408 .

Local Interface	Select the interface that is the local end of the IPsec tunnel. For more information, see “Defining the tunnel ends” on page 1408 . This interface cannot be a loopback interface.
Mode	Select amode. It is easier to use aggressive mode. <ul style="list-style-type: none"> In Main mode, parameters are exchanged in multiple encrypted rounds. In Aggressive mode, parameters are exchanged in a single unencrypted message. Aggressive mode must be used when the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID). For more information, see “Choosing main mode or aggressive mode” on page 1408 .
Authentication Method	Select <i>RSA Signature</i> .
Certificate Name	Select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. You must obtain and load the required server certificate before this selection. See the FortiOS User Authentication guide . If you have not loaded any certificates, use the certificate named <code>Fortinet_Factory</code> .
Peer Options	Peer options define the authentication requirements for remote peers or dialup clients. They are not for your FortiGate unit itself. See “Authenticating remote peers and clients” on page 1412 .
Advanced	You can use the default settings for most phase 1 configurations. Changes are required only if your network requires them. These settings includes IKE version, DNS server, P1 proposal encryption and authentication settings, and XAuth settings. See “Defining IKE negotiation parameters” on page 1417 .

- 4 If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters in the Advanced section. See [“Using the FortiGate unit as an XAuth server” on page 1422](#).
- 5 Select *OK*.

Authenticating the FortiGate unit with a pre-shared key

The simplest way to authenticate a FortiGate unit to its remote peers or dialup clients is by means of a pre-shared key. This is less secure than using certificates, especially if it is used alone, without requiring peer IDs or extended authentication (XAuth). Also, you need to have a secure way to distribute the pre-shared key to the peers.

If you use pre-shared key authentication alone, all remote peers and dialup clients must be configured with the same pre-shared key. Optionally, you can configure remote peers and dialup clients with unique pre-shared keys. On the FortiGate unit, these are configured in user accounts, not in the phase_1 settings. For more information, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1415](#).

The pre-shared key must contain at least 6 printable characters and best practices dictate that it be known only to network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.

If you authenticate the FortiGate unit using a pre-shared key, you can require remote peers or dialup clients to authenticate using peer IDs, but not client certificates.

To authenticate the FortiGate unit with a pre-shared key

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Create a new phase 1 configuration or edit an existing phase 1 configuration.
- 3 Include appropriate entries as follows:

Name	Enter a name that reflects the origination of the remote connection.
Remote Gateway	Select the nature of the remote connection. For more information, see “Defining the tunnel ends” on page 1408 .
Local Interface	Select the interface that is the local end of the IPsec tunnel. For more information, see “Defining the tunnel ends” on page 1408 .
Mode	<p>Select Main or Aggressive mode.</p> <ul style="list-style-type: none"> In Main mode, the phase 1 parameters are exchanged in multiple rounds with encrypted authentication information. In Aggressive mode, the phase 1 parameters are exchanged in single message with authentication information that is not encrypted. <p>When the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID), you must select Aggressive mode if there is more than one dialup phase 1 configuration for the interface IP address.</p> <p>For more information, see “Choosing main mode or aggressive mode” on page 1408.</p>
Authentication Method	Select <i>Pre-shared Key</i> .
Pre-shared Key	Enter the preshared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during phase 1 negotiations. You must define the same value at the remote peer or client. The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.

Peer options	Peer options define the authentication requirements for remote peers or dialup clients, not for the FortiGate unit itself. You can require the use of peer IDs, but not client certificates. For more information, see “Authenticating remote peers and clients” on page 1412 .
Advanced	You can retain the default settings unless changes are needed to meet your specific requirements. See “Defining IKE negotiation parameters” on page 1417 .

- 4 If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters. See [“Using the FortiGate unit as an XAuth server” on page 1422](#).
- 5 Select **OK**.

Authenticating remote peers and clients

Certificates or pre-shared keys restrict who can access the VPN tunnel, but they do not identify or authenticate the remote peers or dialup clients. You have the following options for authentication:

Table 88: Methods of authenticating remote VPN peers

Certificates or Pre-shared key	Local ID	User account pre-shared keys	Reference
Certificates			See “Enabling VPN access for specific certificate holders” on page 1412
Either	X		See “Enabling VPN access by peer identifier” on page 1414
Pre-shared key		X	See “Enabling VPN access with user accounts and pre-shared keys” on page 1415
Pre-shared key	X	X	See “Enabling VPN access with user accounts and pre-shared keys” on page 1415

For authentication of users of the remote peer or dialup client device, see [“Using XAuth authentication” on page 1422](#).

Enabling VPN access for specific certificate holders

When a VPN peer or dialup client is configured to authenticate using digital certificates, it sends the DN of its certificate to the FortiGate unit. This DN can be used to allow VPN access for the certificate holder. That is, a FortiGate unit can be configured to deny connections to all remote peers and dialup clients except the one having the specified DN.

Before you begin

The following procedures assume that you already have an existing phase 1 configuration (see [“Authenticating the FortiGate unit with digital certificates” on page 1409](#)). Follow the procedures below to add certificate-based authentication parameters to the existing configuration.

Before you begin, you must obtain the certificate DN of the remote peer or dialup client. If you are using the FortiClient application as a dialup client, refer to [FortiClient online Help](#) for information about how to view the certificate DN. To view the certificate DN of a FortiGate unit, see [“To view server certificate information and obtain the local DN” on page 1413](#).

Use the `config user peer` CLI command to load the DN value into the FortiGate configuration. For example, if a remote VPN peer uses server certificates issued by your own organization, you would enter information similar to the following:

```
config user peer
  edit DN_FG1000
    set cn 192.168.2.160
    set cn-type ipv4
  end
```

The value that you specify to identify the entry (for example, DN_FG1000) is displayed in the Accept this peer certificate only list in the IPsec phase 1 configuration when you return to the web-based manager.

If the remote VPN peer has a CA-issued certificate to support a higher level of credibility, you would enter information similar to the following:

```
config user peer
  edit CA_FG1000
    set ca CA_Cert_1
    set subject FG1000_at_site1
  end
```

The value that you specify to identify the entry (for example, CA_FG1000) is displayed in the Accept this peer certificate only list in the IPsec phase 1 configuration when you return to the web-based manager. For more information about these CLI commands, see the “user” chapter of the [FortiGate CLI Reference](#).

A group of certificate holders can be created based on existing user accounts for dialup clients. To create the user accounts for dialup clients, see the “User” chapter of the [FortiGate Administration Guide](#). To create the certificate group afterward, use the `config user peergrp` CLI command. See the “user” chapter of the [FortiGate CLI Reference](#).

To view server certificate information and obtain the local DN

- 1 Go to *System > Certificates > Local Certificates*.
- 2 Note the CN value in the *Subject* field (for example, CN = 172.16.10.125, CN = info@fortinet.com, or CN = www.example.com).

To view CA root certificate information and obtain the CA certificate name

- 1 Go to *System > Certificates > CA Certificates*.
- 2 Note the value in the *Name* column (for example, CA_Cert_1).

Configuring certificate authentication for a VPN

With peer certificates loaded, peer users and peer groups defined, you can configure your VPN to authenticate users by certificate.

To enable access for a specific certificate holder or a group of certificate holders

- 1 At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list of defined configurations, select the phase 1 configuration and edit it.
- 3 From the *Authentication Method* list, select *RSA Signature*.

- 4 From the *Certificate Name* list, select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client
- 5 Under *Peer Options*, select one of these options:
 - To accept a specific certificate holder, select *Accept this peer certificate only* and select the name of the certificate that belongs to the remote peer or dialup client. The certificate DN must be added to the FortiGate configuration through CLI commands before it can be selected here. See “Before you begin” on page 1412.
 - To accept dialup clients who are members of a certificate group, select *Accept this peer certificate group only* and select the name of the group. The group must be added to the FortiGate configuration through CLI commands before it can be selected here. See “Before you begin” on page 1412.
- 6 If you want the FortiGate VPN server to supply the DN of a local server certificate for authentication purposes, select *Advanced* and then from the *Local ID* list, select the DN of the certificate that the FortiGate VPN server is to use.
- 7 Select *OK*.

Enabling VPN access by peer identifier

Whether you use certificates or pre-shared keys to authenticate the FortiGate unit, you can require that remote peers or clients have a particular peer ID. This adds another piece of information that is required to gain access to the VPN. More than one FortiGate/FortiClient dialup client may connect through the same VPN tunnel when the dialup clients share a preshared key and assume the same identifier.

A peer ID, also called local ID, can be up to 63 characters long containing standard regular expression characters. Local ID is set in phase1 Aggressive Mode configuration.

You cannot require a peer ID for a remote peer or client that uses a pre-shared key and has a static IP address.

To authenticate remote peers or dialup clients using one peer ID

- 1 At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select a phase 1 configuration and edit its parameters.
- 3 Select *Aggressive* mode in any of the following cases:
 - the FortiGate VPN server authenticates a FortiGate dialup client that uses a dedicated tunnel
 - a FortiGate unit has a dynamic IP address and subscribes to a dynamic DNS service
 - FortiGate/FortiClient dialup clients sharing the same preshared key and local ID connect through the same VPN tunnel
- 4 Select *Accept this peer ID* and type the identifier into the corresponding field.
- 5 Select *OK*.

To assign an identifier (local ID) to a FortiGate unit

Use this procedure to assign a peer ID to a FortiGate unit that acts as a remote peer or dialup client.

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select a phase 1 configuration and edit its parameters.
- 3 Select *Advanced*.
- 4 In the *Local ID* field, type the identifier that the FortiGate unit will use to identify itself.

- 5 Set *Mode* to *Aggressive* if any of the following conditions apply:
 - The FortiGate unit is a dialup client that will use a unique ID to connect to a FortiGate dialup server through a dedicated tunnel.
 - The FortiGate unit has a dynamic IP address, subscribes to a dynamic DNS service, and will use a unique ID to connect to the remote VPN peer through a dedicated tunnel.
 - The FortiGate unit is a dialup client that shares the specified ID with multiple dialup clients to connect to a FortiGate dialup server through the same tunnel.
- 6 Select *OK*.

To configure the FortiClient application

Follow this procedure to add a peer ID to an existing FortiClient configuration:

- 1 Start the FortiClient application.
- 2 Go to *VPN > Connections*, select the existing configuration.
- 3 Select *Advanced > Edit > Advanced*.
- 4 Under *Policy*, select *Config*.
- 5 In the *Local ID* field, type the identifier that will be shared by all dialup clients. This value must match the *Accept this peer ID* value that you specified previously in the phase 1 gateway configuration on the FortiGate unit.
- 6 Select *OK* to close all dialog boxes.
- 7 Configure all dialup clients the same way using the same preshared key and local ID.

Enabling VPN access with user accounts and pre-shared keys

You can permit access only to remote peers or dialup clients that have pre-shared keys and/or peer IDs configured in user accounts on the FortiGate unit.

If you want two VPN peers (or a FortiGate unit and a dialup client) to accept reciprocal connections based on peer IDs, you must enable the exchange of their identifiers when you define the phase 1 parameters.

The following procedures assume that you already have an existing phase 1 configuration (see [“Authenticating the FortiGate unit with digital certificates” on page 1409](#)). Follow the procedures below to add ID checking to the existing configuration.

Before you begin, you must obtain the identifier (local ID) of the remote peer or dialup client. If you are using the FortiClient Endpoint Security application as a dialup client, refer to the [Authenticating FortiClient Dialup Clients Technical Note](#) to view or assign an identifier. To assign an identifier to a FortiGate dialup client or a FortiGate unit that has a dynamic IP address and subscribes to a dynamic DNS service, see [“To assign an identifier \(local ID\) to a FortiGate unit” on page 1414](#).

If required, a dialup user group can be created from existing user accounts for dialup clients. To create the user accounts and user groups, see the [User Authentication](#) chapter of The Handbook.

The following procedure supports FortiGate/FortiClient dialup clients that use unique preshared keys and/or peer IDs. The client must have an account on the FortiGate unit and be a member of the dialup user group.

The dialup user group must be added to the FortiGate configuration before it can be selected. For more information, see the [User Authentication](#) chapter of The Handbook.

The FortiGate dialup server compares the local ID that you specify at each dialup client to the FortiGate user-account user name. The dialup-client preshared key is compared to a FortiGate user-account password.

To authenticate dialup clients using unique preshared keys and/or peer IDs

- 1 At the FortiGate VPN server, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select the *Edit* icon of a phase 1 configuration to edit its parameters.
- 3 If the clients have unique peer IDs, set *Mode* to *Aggressive*.
- 4 Clear the *Pre-shared Key* field.
The user account password will be used as the preshared key.
- 5 Select *Accept peer ID in dialup group* and then select the group name from the list of user groups.
- 6 Select *OK*.

Follow this procedure to add a unique pre-shared key and unique peer ID to an existing FortiClient configuration.

To configure FortiClient - pre-shared key and peer ID

- 1 Start the FortiClient Endpoint Security application.
- 2 Go to *VPN > Connections*, select the existing configuration.
- 3 Select *Advanced > Edit*.
- 4 In the *Preshared Key* field, type the FortiGate password that belongs to the dialup client (for example, 1234546).
The user account password will be used as the preshared key.
- 5 Select *Advanced*.
- 6 Under *Policy*, select *Config*.
- 7 In the *Local ID* field, type the FortiGate user name that you assigned previously to the dialup client (for example, FortiClient1).
- 8 Select *OK* to close all dialog boxes.

Configure all FortiClient dialup clients this way using unique preshared keys and local IDs.

Follow this procedure to add a unique pre-shared key to an existing FortiClient configuration.

To configure FortiClient - preshared key only

- 1 Start the FortiClient Endpoint Security application.
- 2 Go to *VPN > Connections*, select the existing configuration
- 3 Select *Advanced > Edit*.
- 4 In the *Preshared Key* field, type the user name, followed by a “+” sign, followed by the password that you specified previously in the user account settings on the FortiGate unit (for example, FC2+1FG6LK)
- 5 Select *OK* to close all dialog boxes.

Configure all the FortiClient dialup clients this way using their unique peer ID and pre-shared key values.

Defining IKE negotiation parameters

In phase 1, the two peers exchange keys to establish a secure communication channel between them. As part of the phase 1 process, the two peers authenticate each other and negotiate a way to encrypt further communications for the duration of the session. For more information see [“Authenticating remote peers and clients” on page 1412](#). The P1 Proposal parameters select the encryption and authentication algorithms that are used to generate keys for protecting negotiations.

The IKE negotiation parameters determine:

- which encryption algorithms may be applied for converting messages into a form that only the intended recipient can read
- which authentication hash may be used for creating a keyed hash from a preshared or private key
- which Diffie-Hellman group (DH Group) will be used to generate a secret session key

Phase 1 negotiations (in main mode or aggressive mode) begin as soon as a remote VPN peer or client attempts to establish a connection with the FortiGate unit. Initially, the remote peer or dialup client sends the FortiGate unit a list of potential cryptographic parameters along with a session ID. The FortiGate unit compares those parameters to its own list of advanced phase 1 parameters and responds with its choice of matching parameters to use for authenticating and encrypting packets. The two peers handle the exchange of encryption keys between them, and authenticate the exchange through a preshared key or a digital signature.

Generating keys to authenticate an exchange

The FortiGate unit supports the generation of secret session keys automatically using a Diffie-Hellman algorithm. These algorithms are defined in RFC 2409. The *Keylife* setting in the *P1 Proposal* area determines the amount of time before the phase 1 key expires. Phase 1 negotiations are rekeyed automatically when there is an active security association. See [“Dead peer detection” on page 1421](#).



You can enable or disable automatic rekeying between IKE peers through the `phase1-rekey` attribute of the `config system global` CLI command. For more information, see the “system” chapter of the [FortiGate CLI Reference](#).



When in FIPS-CC mode, the FortiGate unit requires DH key exchange to use values at least 3072 bits long. However most browsers need the key size set to 1024. You can set the minimum size of the DH keys in the CLI.

```
config system global
    set dh-params 3072
end
```


When you use a preshared key (shared secret) to set up two-party authentication, the remote VPN peer or client and the FortiGate unit must both be configured with the same preshared key. Each party uses a session key derived from the Diffie-Hellman exchange to create an authentication key, which is used to sign a known combination of inputs using an authentication algorithm (such as HMAC-MD5, HMAC-SHA-1, or HMAC-SHA-256). Hash-based Message Authentication Code (HMAC) is a method for calculating an authentication code using a hash function plus a secret key, and is defined in RFC 2104. Each party signs a different combination of inputs and the other party verifies that the same result can be computed.



SHA-256, SHA-384 and SHA-512 are not accelerated by some FortiASIC processors (including FortiASIC network processors and security processors). As a result, using SHA-256, SHA-384 and SHA-512 may reduce the performance of the FortiGate unit more significantly than SHA-1 which is accelerated by all FortiASIC processors.

When you use preshared keys to authenticate VPN peers or clients, you must distribute matching information to all VPN peers and/or clients whenever the preshared key changes.

As an alternative, the remote peer or dialup client and FortiGate unit can exchange digital signatures to validate each other's identity with respect to their public keys. In this case, the required digital certificates must be installed on the remote peer and on the FortiGate unit. By exchanging certificate DNs, the signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer.

The following procedure assumes that you already have a phase 1 definition that describes how remote VPN peers and clients will be authenticated when they attempt to connect to a local FortiGate unit. For information about the Local ID and XAuth options, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1415](#) and [“Using the FortiGate unit as an XAuth server” on page 1422](#). Follow this procedure to add IKE negotiation parameters to the existing definition.

Defining IKE negotiation parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select the *Edit* button to edit the phase 1 parameters for a particular remote gateway.

3 Select *Advanced* and include appropriate entries as follows:

P1 Proposal	<p>Select the encryption and authentication algorithms that will be used to generate keys for protecting negotiations.</p> <p>Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.</p> <p>You can select any of the following symmetric-key algorithms:</p> <ul style="list-style-type: none"> • DES-Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key. • 3DES-Triple-DES, in which plain text is encrypted three times by three keys. • AES128-A 128-bit block algorithm that uses a 128-bit key. • AES192-A 128-bit block algorithm that uses a 192-bit key. • AES256-A 128-bit block algorithm that uses a 256-bit key. <p>You can select one of the following message digests to check the authenticity of messages during phase 1 negotiations:</p> <ul style="list-style-type: none"> • MD5-Message Digest 5, the hash algorithm developed by RSA Data Security. • SHA1-Secure Hash Algorithm 1, which produces a 160-bit message digest. • SHA-256 Secure Hash Algorithm 256, which produces a 256-bit message digest • SHA-384 Secure Hash Algorithm 384, which produces a 384-bit message digest • SHA-512 Secure Hash Algorithm 512, which produces a 512-bit message digest <p>To specify a third combination, use the add button beside the fields for the second combination.</p> <p>SHA-256, SHA-384 and SHA-512 are not accelerated by some FortiASIC processors (including FortiASIC network processors and security processors). As a result, using SHA-256, SHA-384 and SHA-512 may reduce the performance of the FortiGate unit more significantly than SHA-1 which is accelerated by all FortiASIC processors.</p>
--------------------	--

DH Group	<p>Select one or more Diffie-Hellman groups from DH group 1, 2, and 5. When using aggressive mode, DH groups cannot be negotiated.</p> <p>If both VPN peers (or a VPN server and its client) have static IP addresses and use aggressive mode, select a single DH group. The setting on the FortiGate unit must be identical to the setting on the remote peer or dialup client.</p> <p>When the remote VPN peer or client has a dynamic IP address and uses aggressive mode, select up to three DH groups on the FortiGate unit and one DH group on the remote peer or dialup client. The setting on the remote peer or dialup client must be identical to one of the selections on the FortiGate unit.</p> <p>If the VPN peer or client employs main mode, you can select multiple DH groups. At least one of the settings on the remote peer or dialup client must be identical to the selections on the FortiGate unit.</p>
Keylife	Type the amount of time (in seconds) that will be allowed to pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172800 seconds.
Nat-traversal	Enable this option if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared). When in doubt, enable NAT-traversal. See “NAT traversal” on page 1420 .
Keepalive Frequency	If you enabled NAT traversal, enter a keepalive frequency setting. The value represents an interval from 0 to 900 seconds where the connection will be maintained with no activity. For additional security this value must be as low as possible. See “NAT keepalive frequency” on page 1421 .
Dead Peer Detection	Enable this option to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. This feature minimizes the traffic required to check if a VPN peer is available or unavailable (dead). See “Dead peer detection” on page 1421 .

4 Select OK.

NAT traversal

Network Address Translation (NAT) is a way to convert private IP addresses to publicly routable Internet addresses and vice versa. When an IP packet passes through a NAT device, the source or destination address in the IP header is modified. FortiGate units support NAT version 1 (encapsulate on port 500 with non-IKE marker), version 3 (encapsulate on port 4500 with non-ESP marker), and compatible versions.

NAT cannot be performed on IPsec packets in ESP tunnel mode because the packets do not contain a port number. As a result, the packets cannot be demultiplexed. To work around this, the FortiGate unit provides a way to protect IPsec packet headers from NAT modifications. When the Nat-traversal option is enabled, outbound encrypted packets are wrapped inside a UDP IP header that contains a port number. This extra encapsulation allows NAT devices to change the port number without modifying the IPsec packet directly.

To provide the extra layer of encapsulation on IPsec packets, the Nat-traversal option must be enabled whenever a NAT device exists between two FortiGate VPN peers or a FortiGate unit and a dialup client such as FortiClient. On the receiving end, the FortiGate unit or FortiClient removes the extra layer of encapsulation before decrypting the packet.

NAT keepalive frequency

When a NAT device performs network address translation on a flow of packets, the NAT device determines how long the new address will remain valid if the flow of traffic stops (for example, the connected VPN peer may be idle). The device may reclaim and reuse a NAT address when a connection remains idle for too long.

To work around this, when you enable NAT traversal specify how often the FortiGate unit sends periodic keepalive packets through the NAT device in order to ensure that the NAT address mapping does not change during the lifetime of a session. To be effective, the keepalive interval must be smaller than the session lifetime value used by the NAT device.

The keepalive packet is a 138-byte ISAKMP exchange.

Dead peer detection

Sometimes, due to routing issues or other difficulties, the communication link between a FortiGate unit and a VPN peer or client may go down—packets could be lost if the connection is left to time out on its own. The FortiGate unit provides a mechanism called Dead Peer Detection (DPD), sometimes referred to as gateway detection or ping server, to prevent this situation and reestablish IKE negotiations automatically before a connection times out: the active phase 1 security associations are caught and renegotiated (rekeyed) before the phase 1 encryption key expires.

By default, DPD sends probe messages every five seconds by default (see `dpd-retryinterval` in the [FortiGate CLI Reference](#)). If you are experiencing high network traffic, you can experiment with increasing the ping interval. However longer intervals will require more traffic to detect dead peers which will result in more traffic.

In the web-based manager, the Dead Peer Detection option can be enabled when you define advanced phase 1 options. The `config vpn ipsec phase1` CLI command supports additional options for specifying a retry count and a retry interval.

For more information about these commands and the related `config router gwdetect` CLI command, see the [FortiGate CLI Reference](#).

For example, enter the following CLI commands to configure dead peer detection on the existing IPsec Phase1 configuration called `test` to use 15 second intervals and to wait for 3 missed attempts before declaring the peer dead and taking action.

```
config vpn ipsec phase1
  edit test
    set dpd enable
    set dpd-retryinterval 15
    set dpd-retrycount 3
  next
end
```

Using XAuth authentication

Extended authentication (XAuth) increases security by requiring the remote dialup client user to authenticate in a separate exchange at the end of phase 1. XAuth draws on existing FortiGate user group definitions and uses established authentication mechanisms such as PAP, CHAP, RADIUS and LDAP to authenticate dialup clients. You can configure a FortiGate unit to function either as an XAuth server or an XAuth client.



If the server or client is attempting a connection using XAuth and the other end is not using XAuth, the failed connection attempts that are logged will not specify XAuth as the reason.

This section includes:

- [Using the FortiGate unit as an XAuth server](#)
- [Using the FortiGate unit as an XAuth client](#)

Using the FortiGate unit as an XAuth server

A FortiGate unit can act as an XAuth server for dialup clients. When the phase 1 negotiation completes, the FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.

If the user records on the RADIUS server have suitably configured Framed-IP-Address fields, you can assign client virtual IP addresses by XAuth instead of from a DHCP address range. See [“Assigning VIPs by RADIUS user group” on page 1487](#).

The authentication protocol to use for XAuth depends on the capabilities of the authentication server and the XAuth client:

- Select PAP whenever possible.
- You must select PAP for all implementations of LDAP and some implementations of Microsoft RADIUS.
- Select AUTO when the authentication server supports CHAP but the XAuth client does not. The FortiGate unit will use PAP to communicate with the XAuth client and CHAP to communicate with the authentication server.

To authenticate a dialup user group using XAuth settings

Before you begin, create user accounts and user groups to identify the dialup clients that need to access the network behind the FortiGate dialup server. If password protection will be provided through an external RADIUS or LDAP server, you must configure the FortiGate dialup server to forward authentication requests to the authentication server. For information about these topics, see the [FortiGate User Authentication Guide](#).

- 1 At the FortiGate dialup server, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select the *Edit* icon of a phase 1 configuration to edit its parameters for a particular remote gateway.
- 3 Select *Advanced*.
- 4 Under *XAuth*, select *Enable as Server*.

- 5 The *Server Type* setting determines the type of encryption method to use between the XAuth client, the FortiGate unit and the authentication server. Select one of the following options:
 - *PAP*—Password Authentication Protocol.
 - *CHAP*— Challenge-Handshake Authentication Protocol.
 - *AUTO*— Use PAP between the XAuth client and the FortiGate unit, and CHAP between the FortiGate unit and the authentication server.
- 6 From the *User Group* list, select the user group that needs to access the private network behind the FortiGate unit. The group must be added to the FortiGate configuration before it can be selected here.
- 7 Select *OK*.

Using the FortiGate unit as an XAuth client

If the FortiGate unit acts as a dialup client, the remote peer, acting as an XAuth server, might require a user name and password. You can configure the FortiGate unit as an XAuth client, with its own user name and password, which it provides when challenged.

To configure the FortiGate dialup client as an XAuth client

- 1 At the FortiGate dialup client, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 In the list, select a phase 1 configuration and select *Edit*.
- 3 Select *Advanced*.
- 4 Under *XAuth*, select *Enable as Client*.
- 5 In the *Username* field, type the FortiGate PAP, CHAP, RADIUS, or LDAP user name that the FortiGate XAuth server will compare to its records when the FortiGate XAuth client attempts to connect.
- 6 In the *Password* field, type the password to associate with the user name.
- 7 Select *OK*.



Phase 2 parameters

This section describes the phase 2 parameters that are required to establish communication through a VPN.

The following topics are included in this section:

- [Basic phase 2 settings](#)
- [Advanced phase 2 settings](#)
- [Configure the phase 2 parameters](#)

Basic phase 2 settings

After IPsec VPN phase 1 negotiations complete successfully, phase 2 negotiation begins. Phase 2 parameters define the algorithms that the FortiGate unit can use to encrypt and transfer data for the remainder of the session. The basic phase 2 settings associate IPsec phase 2 parameters with a phase 1 configuration.

When defining phase 2 parameters, you can choose any set of phase 1 parameters to set up a secure connection and authenticate the remote peer.

For more information on phase 2 settings in the web-based manager, see [“Phase 2 configuration” on page 1399](#)

The information and procedures in this section do not apply to VPN peers that perform negotiations using manual keys. Refer to [“Manual-key configurations” on page 1551](#) instead.

Advanced phase 2 settings

The following additional advanced phase 2 settings are available to enhance the operation of the tunnel:

- [P2 Proposals](#)
- [Replay detection](#)
- [Perfect forward secrecy \(PFS\)](#)
- [Keylife](#)
- [Quick mode selectors](#)

Figure 126: Advanced phase 2 settings

Advanced...

P2 Proposal

1- Encryption: 3DES Authentication: SHA1

2- Encryption: AES128 Authentication: SHA1

☐ Enable replay detection

☐ Enable perfect forward secrecy (PFS)

DH Group: 1 2 5 14

Keylife: Seconds 1800 (Seconds) 4608000 (KBytes)

Autokey Keep Alive ☐ Enable

Quick Mode Selector

Source address	Specify	0.0.0.0/0
	Select	-----Address-----
Source port		0
Destination address	Specify	0.0.0.0/0
	Select	-----Address-----
Destination port		0
Protocol		0

P2 Proposals

In phase 2, the VPN peer or client and the FortiGate unit exchange keys again to establish a secure communication channel. The P2 Proposal parameters select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). The keys are generated automatically using a Diffie-Hellman algorithm.

Replay detection

IPsec tunnels can be vulnerable to replay attacks. Replay detection enables the FortiGate unit to check all IPsec packets to see if they have been received before. If any encrypted packets arrive out of order, the FortiGate unit discards them.

Perfect forward secrecy (PFS)

By default, phase 2 keys are derived from the session key created in phase 1. Perfect forward secrecy forces a new Diffie-Hellman exchange when the tunnel starts and whenever the phase 2 keylife expires, causing a new key to be generated each time. This exchange ensures that the keys created in phase 2 are unrelated to the phase 1 keys or any other keys generated automatically in phase 2.

Keylife

The Keylife setting sets a limit on the length of time that a phase 2 key can be used. The default units are seconds. Alternatively, you can set a limit on the number of kilobytes (KB) of processed data, or both. If you select both, the key expires when either the time has passed or the number of KB have been processed. When the phase 2 key expires, a new key is generated without interrupting service.

Auto-negotiate

By default, the phase 2 security association (SA) is not negotiated until a peer attempts to send data. The triggering packet and some subsequent packets are dropped until the SA is established. Applications normally resend this data, so there is no loss, but there might be a noticeable delay in response to the user.

Automatically establishing the SA can also be important for a dialup peer. This ensures that the VPN tunnel is available for peers at the server end to initiate traffic to the dialup peer. Otherwise, the VPN tunnel does not exist until the dialup peer initiates traffic.

When enabled, auto-negotiate initiates the phase 2 SA negotiation automatically, repeating every five seconds until the SA is established.

The auto-negotiate feature is available only through the Command Line Interface (CLI). Use the following commands to enable it.

```
config vpn ipsec phase2
  edit <phase2_name>
    set auto-negotiate enable
  end
```

If the tunnel goes down, the auto-negotiate feature will attempt to re-establish it. However, the Autokey Keep Alive feature is a better method to ensure your VPN remains up.

Autokey Keep Alive

The phase 2 SA has a fixed duration. If there is traffic on the VPN as the SA nears expiry, a new SA is negotiated and the VPN switches to the new SA without interruption. If there is no traffic, the SA expires and the VPN tunnel goes down. A new SA will not be generated until there is traffic.

The Autokey Keep Alive option ensures that a new SA is negotiated even if there is no traffic so that the VPN tunnel stays up.

DHCP-IPsec

Select this option if the FortiGate unit assigns VIP addresses to FortiClient dialup clients through a DHCP server or relay. This option is available only if the Remote Gateway in the phase 1 configuration is set to Dialup User and it works only on policy-based VPNs.

With the DHCP-IPsec option, the FortiGate dialup server acts as a proxy for FortiClient dialup clients that have VIP addresses on the subnet of the private network behind the FortiGate unit. In this case, the FortiGate dialup server acts as a proxy on the local private network for the FortiClient dialup client. When a host on the network behind the dialup server issues an ARP request that corresponds to the device MAC address of the FortiClient host (when a remote server sends an ARP to the local FortiClient dialup client), the FortiGate unit answers the ARP request on behalf of the FortiClient host and forwards the associated traffic to the FortiClient host through the tunnel.

This feature prevents the VIP address assigned to the FortiClient dialup client from causing possible arp broadcast problems—the normal and VIP addresses can confuse some network switches by two addresses having the same MAC address.

Quick mode selectors

Quick Mode selectors determine which IP addresses can perform IKE negotiations to establish a tunnel. By only allowing authorized IP addresses access to the VPN tunnel, the network is more secure.

The default settings are as broad as possible: any IP address or configured address group, using any protocol, on any port. This enables configurations in which multiple subnets at each end of the tunnel can communicate, limited only by the security policies at each end.

When configuring Quick Mode selector *Source Address* and *Destination address*, valid options include IPv4 and IPv6 single addresses, IPv4 firewall address or group name, IPv4 range, IPv6 range, IPv4 subnet, or IPv6 subnet. For more information on IPv6 IPsec VPN, see [“Overview of IPv6 IPsec support” on page 1555](#).

There are some configurations that require specific selectors:

- the VPN peer is a third-party device that uses specific phase2 selectors
- the FortiGate unit connects as a dialup client to another FortiGate unit, in which case you must specify a source IP address, IP address range or subnet

With FortiOS VPNs, your network has multiple layers of security, with quick mode selectors being an important line of defence.

- Routes guide traffic from one IP address to another.
- Phase 1 and phase 2 connection settings ensure there is a valid remote end point for the VPN tunnel that agrees on the encryption and parameters.
- Quick mode selectors allow IKE negotiations only for allowed peers.
- Security policies control which IP addresses can connect to the VPN.
- Security policies also control what protocols are allowed over the VPN along with any bandwidth limiting.

Configure the phase 2 parameters

Follow this procedure to create an IPsec phase 2 definition.



If you are creating a hub-and-spoke configuration or an Internet-browsing configuration, you may have already started defining some of the required phase 2 parameters. If so, edit the existing definition to complete the configuration.

Specifying the phase 2 parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*.
- 3 Include appropriate entries as follows:

Name	Enter a name to identify the phase 2 configuration.
Phase 1	Select the phase 1 configuration that describes how remote peers or dialup clients will be authenticated on this tunnel, and how the connection to the remote peer or dialup client will be secured.


- 4 Select *Advanced*.
- 5 Include appropriate entries as follows:

P2 Proposal	<p>Select the encryption and authentication algorithms that will be used to change data into encrypted code.</p> <p>Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.</p> <p>It is invalid to set both <i>Encryption</i> and <i>Authentication</i> to null.</p>
--------------------	---

Encryption	<p>You can select any of the following symmetric-key algorithms:</p> <p>NULL — Do not use an encryption algorithm.</p> <p>DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</p> <p>3DES — Triple-DES, in which plain text is encrypted three times by three keys.</p> <p>AES128 — A 128-bit block algorithm that uses a 128-bit key.</p> <p>AES192 — A 128-bit block algorithm that uses a 192-bit key.</p> <p>AES256 — A 128-bit block algorithm that uses a 256-bit key.</p>
Authentication	<p>You can select either of the following message digests to check the authenticity of messages during an encrypted session:</p> <ul style="list-style-type: none"> • NULL — Do not use a message digest. • MD5 — Message Digest 5, the hash algorithm developed by RSA Data Security. • SHA1 — Secure Hash Algorithm 1, which produces a 160-bit message digest. <p>To specify one combination only, set the <i>Encryption</i> and <i>Authentication</i> options of the second combination to NULL. To specify a third combination, use the <i>Add</i> button beside the fields for the second combination.</p>
Enable replay detection	Optionally enable or disable replay detection. Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.
Enable perfect forward secrecy (PFS)	Enable or disable PFS. Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.
DH Group	Select one Diffie-Hellman group (1, 2, 5, or 14). The remote peer or dialup client must be configured to use the same group.
Keylife	Select the method for determining when the phase 2 key expires: <i>Seconds</i> , <i>KBytes</i> , or <i>Both</i> . If you select <i>Both</i> , the key expires when either the time has passed or the number of KB have been processed. The range is from 120 to 172800 seconds, or from 5120 to 2147483648 KB.
Autokey Keep Alive	Enable the option if you want the tunnel to remain active when no data is being processed.
DHCP-IPsec	<p>Select <i>Enable</i> if the FortiGate unit acts as a dialup server and FortiGate DHCP server or relay will be used to assign VIP addresses to FortiClient dialup clients. The DHCP server or relay parameters must be configured separately.</p> <p>If the FortiGate unit acts as a dialup server and the FortiClient dialup client VIP addresses match the network behind the dialup server, select <i>Enable</i> to cause the FortiGate unit to act as a proxy for the dialup clients.</p> <p>This is available only for phase 2 configurations associated with a dialup phase 1 configuration. It works only on policy-based VPNs.</p>

Quick Mode Selector	<p>Optionally specify the source and destination IP addresses to be used as selectors for IKE negotiations. If the FortiGate unit is a dialup server, keep the default value 0.0.0.0/0 unless you need to circumvent problems caused by ambiguous IP addresses between one or more of the private networks making up the VPN. You can specify a single host IP address, an IP address range, or a network address. You may optionally specify source and destination port numbers and/or a protocol number.</p> <p>If you are editing an existing phase 2 configuration, the <i>Source address</i> and <i>Destination address</i> fields are unavailable if the tunnel has been configured to use firewall addresses as selectors. This option exists only in the CLI. See the <code>dst-addr-type</code>, <code>dst-name</code>, <code>src-addr-type</code> and <code>src-name</code> keywords for the <code>vpn ipsec phase2</code> command in the FortiGate CLI Reference.</p>
Source address	<p>If the FortiGate unit is a dialup server, type the source IP address that corresponds to the local sender(s) or network behind the local VPN peer (for example, 172.16.5.0/24 or 172.16.5.0/255.255.255.0 for a subnet, or 172.16.5.1/32 or 172.16.5.1/255.255.255.255 for a server or host, or 192.168.10.[80-100] or 192.168.10.80-192.168.10.100 for an address range). A value of 0.0.0.0/0 means all IP addresses behind the local VPN peer.</p> <p>If the FortiGate unit is a dialup client, source address must refer to the private network behind the FortiGate dialup client.</p>
Source port	Type the port number that the local VPN peer uses to transport traffic related to the specified service (protocol number). The range is 0 to 65535. To specify all ports, type 0.
Destination address	Type the destination IP address that corresponds to the recipient(s) or network behind the remote VPN peer (for example, 192.168.20.0/24 for a subnet, or 172.16.5.1/32 for a server or host, or 192.168.10.[80-100] for an address range). A value of 0.0.0.0/0 means all IP addresses behind the remote VPN peer.
Destination port	Type the port number that the remote VPN peer uses to transport traffic related to the specified service (protocol number). The range is 0 to 65535. To specify all ports, type 0.
Protocol	Type the IP protocol number of the service. The range is 1 to 255. To specify all services, type 0.

6 Select OK.



Defining VPN security policies

This section explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN, and how to define appropriate security policies.

The following topics are included in this section:


- [Defining policy addresses](#)
- [Defining VPN security policies](#)

Defining policy addresses

A VPN tunnel has two end points. These end points may be VPN peers such as two FortiGate gateways. Encrypted packets are transmitted between the end points. At each end of the VPN tunnel, a VPN peer intercepts encrypted packets, decrypts the packets, and forwards the decrypted IP packets to the intended destination.

You need to define firewall addresses for the private networks behind each peer. You will use these addresses as the source or destination address depending on the security policy.

Figure 127: Example topology for the following policies



In general:

- In a gateway-to-gateway, hub-and-spoke, dynamic DNS, redundant-tunnel, or transparent configuration, you need to define a policy address for the private IP address of the network behind the remote VPN peer (for example, 192.168.10.0/255.255.255.0 or 192.168.10.0/24).
- In a peer-to-peer configuration, you need to define a policy address for the private IP address of a server or host behind the remote VPN peer (for example, 172.16.5.1/255.255.255.255 or 172.16.5.1/32 or 172.16.5.1).

For a FortiGate dialup server in a dialup-client or Internet-browsing configuration:

- If you are not using VIP addresses, or if the FortiGate dialup server assigns VIP addresses to FortiClient dialup clients through FortiGate DHCP relay, select the predefined destination address “all” in the security policy to refer to the dialup clients.
- If you assign VIP addresses to FortiClient dialup clients manually, you need to define a policy address for the VIP address assigned to the dialup client (for example, 10.254.254.1/32), or a subnet address from which the VIP addresses are assigned (for example, 10.254.254.0/24 or 10.254.254.0/255.255.255.0).
- For a FortiGate dialup client in a dialup-client or Internet-browsing configuration, you need to define a policy address for the private IP address of a host, server, or network behind the FortiGate dialup server.

To define a security IP address

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.
- 2 In the *Address Name* field, type a descriptive name that represents the network, server(s), or host(s).
- 3 In *Type*, select *Subnet / IP Range*.
- 4 In the *Subnet/IP Range* field, type the corresponding IP address and subnet mask.
For a subnet you could use the format 172.16.5.0/24 or its equivalent 172.16.5.0/255.255.255.0. For a server or host it would likely be 172.16.5.1/32. Alternately you can use an IP address range such as 192.168.10.[80-100] or 192.168.10.80-192.168.10.100.
- 5 Select *OK*.

Defining VPN security policies

Security policies allow IP traffic to pass between interfaces on a FortiGate unit. You can limit communication to particular traffic by specifying source address and destination addresses. Then only traffic from those addresses will be allowed.

Policy-based and route-based VPNs require different security policies.

- A policy-based VPN requires an IPsec security policy. You specify the interface to the private network, the interface to the remote peer and the VPN tunnel. A single policy can enable traffic inbound, outbound, or in both directions.
- A route-based VPN requires an Accept security policy for each direction. As source and destination interfaces, you specify the interface to the private network and the virtual IPsec interface (phase 1 configuration) of the VPN. The IPsec interface is the destination interface for the outbound policy and the source interface for the inbound policy. One security policy must be configured for each direction of each VPN interface.

There are examples of security policies for both policy-based and route-based VPNs throughout this guide. See [“Route-based or policy-based VPN”](#) on page 1471.



If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPSec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

Defining an IPsec security policy for a policy-based VPN

An IPsec security policy enables the transmission and reception of encrypted packets, specifies the permitted direction of VPN traffic, and selects the VPN tunnel. In most cases, a single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel.

Allow outbound and allow inbound

In addition to these operations, security policies specify which IP addresses can initiate a tunnel. When the *Allow outbound* option is selected, traffic from the local private network initiates the tunnel. When the *Allow inbound* option is selected, traffic from a dialup client or computers on the remote network initiates the tunnel. Both can be enabled at the same time for bi-directional initiation of the tunnel.

Outbound and inbound NAT

When a FortiGate unit operates in NAT mode, you can also enable inbound or outbound NAT. Outbound NAT may be performed on outbound encrypted packets, or on IP packets before they are sent through the tunnel. Inbound NAT is performed on IP packets emerging from the tunnel. By default, these options are not selected in security policies.

- When used in conjunction with the `natip` CLI attribute (see the “config firewall” chapter of the [FortiGate CLI Reference](#)), outbound NAT enables you to change the source addresses of IP packets before they go into the tunnel. This feature is often used to resolve ambiguous routing when two or more of the private networks making up a VPN have the same or overlapping IP addresses. .

When inbound NAT is enabled, inbound encrypted packets are intercepted and decrypted, and the source IP addresses of the decrypted packets are translated into the IP address of the FortiGate interface to the local private network before they are routed to the private network. If the computers on the local private network can communicate only with devices on the local private network (that is, the FortiGate interface to the private network is not the default gateway) and the remote client (or remote private network) does not have an IP address in the same network address space as the local private network, enable inbound NAT.

Source and destination addresses

Most security policies control outbound IP traffic. A VPN outbound policy usually has a source address originating on the private network behind the local FortiGate unit, and a destination address belonging to a dialup VPN client or a network behind the remote VPN peer. The source address that you choose for the security policy identifies from where outbound cleartext IP packets may originate, and also defines the local IP address or addresses that a remote server or client will be allowed to access through the VPN tunnel. The destination address that you choose identifies where IP packets must be forwarded after they are decrypted at the far end of the tunnel, and determines the IP address or addresses that the local network will be able to access at the far end of the tunnel.

Enabling other policy features

You can fine-tune a policy for services such as HTTP, FTP, and POP3; enable logging, traffic shaping, antivirus protection, web filtering, email filtering, file transfer, and email services throughout the VPN; and optionally allow connections according to a predefined schedule.



As an option, differentiated services (diffserv or DSCP) can be enabled in the security policy through CLI commands. For more information on this feature, see [FortiOS Handbook Traffic Shaping chapter](#) or the “firewall” chapter of the [FortiGate CLI Reference](#).

When a remote server or client attempts to connect to the private network behind a FortiGate gateway, the security policy intercepts the connection attempt and starts the VPN tunnel. The FortiGate unit uses the remote gateway specified in its phase 1 tunnel configuration to reply to the remote peer. When the remote peer receives a reply, it checks its own security policy, including the tunnel configuration, to determine which communications are permitted. As long as one or more services are allowed through the VPN tunnel, the two peers begin to negotiate the tunnel. To follow this negotiation in the web-based manager, go to *VPN > Monitor > IPsec Monitor*. There you will find a list of the VPN tunnels, their status, and the data flow both incoming and outgoing.

Before you begin

Before you define the IPsec policy, you must:

- Define the IP source and destination addresses. See “[Defining policy addresses](#)” on [page 1431](#).
- Specify the phase 1 authentication parameters. See “[Auto Key phase 1 parameters](#)” on [page 1407](#).
- Specify the phase 2 parameters. See “[Phase 2 parameters](#)” on [page 1425](#).

To define an IPsec security policy

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Include the settings as follows:

Source Interface/Zone	Select the local interface to the internal (private) network.
Source Address Name	Select the name that corresponds to the local network, server(s), or host(s) from which IP packets may originate.
Destination Interface/Zone	Select the local interface to the external (public) network.
Destination Address Name	Select the name that corresponds to the remote network, server(s), or host(s) to which IP packets may be delivered.
Schedule	Keep the default setting (always) unless changes are needed to meet specific requirements.
Service	Keep the default setting (ANY) unless changes are needed to meet your specific requirements.
Action	Select <i>IPSEC</i> .

VPN Tunnel	Select the name of the phase 1 tunnel configuration to which this policy will apply.
Allow Inbound	Select if traffic from the remote network will be allowed to initiate the tunnel.
Allow Outbound	Select if traffic from the local network will be allowed to initiate the tunnel.
Inbound NAT	Select if you want to translate the source IP addresses of inbound decrypted packets into the IP address of the FortiGate interface to the local private network.
Outbound NAT	Select in combination with a <code>natip</code> CLI value to translate the source addresses of outbound cleartext packets into the IP address that you specify. Do not select Outbound NAT unless you specify a <code>natip</code> value through the CLI. When a <code>natip</code> value is specified, the source addresses of outbound IP packets are replaced before the packets are sent through the tunnel. For more information, see the “firewall” chapter of the FortiGate CLI Reference .

- 3 You may enable UTM features, and/or event logging, or select advanced settings to authenticate a user group, or shape traffic. For more information, see the [Firewall](#) chapter of [The Handbook](#).
- 4 Select OK.
- 5 Place the policy in the policy list above any other policies having similar source and destination addresses.

Defining multiple IPsec policies for the same tunnel

You must define at least one IPsec policy for each VPN tunnel. If the same remote server or client requires access to more than one network behind a local FortiGate unit, the FortiGate unit must be configured with an IPsec policy for each network. Multiple policies may be required to configure redundant connections to a remote destination or control access to different services at different times.

To ensure a secure connection, the FortiGate unit must evaluate IPSEC policies before ACCEPT and DENY security policies. Because the FortiGate unit reads policies starting at the top of the list, you must move all IPsec policies to the top of the list. When you define multiple IPsec policies for the same tunnel, you must reorder the IPsec policies that apply to the tunnel so that specific constraints can be evaluated before general constraints.



Adding multiple IPsec policies for the same VPN tunnel can cause conflicts if the policies specify similar source and destination addresses but have different settings for the same service. When policies overlap in this manner, the system may apply the wrong IPsec policy or the tunnel may fail.

For example, if you create two equivalent IPsec policies for two different tunnels, it does not matter which one comes first in the list of IPsec policies—the system will select the correct policy based on the specified source and destination addresses. If you create two different IPsec policies for the same tunnel (that is, the two policies treat traffic differently depending on the nature of the connection request), you might have to reorder the IPsec

policies to ensure that the system selects the correct IPsec policy. Reordering is especially important when the source and destination addresses in both policies are similar (for example, if one policy specifies a subset of the IP addresses in another policy). In this case, place the IPsec policy having the most specific constraints at the top of the list so that it can be evaluated first.

Defining security policies for a route-based VPN

When you define a route-based VPN, you create a virtual IPsec interface on the physical interface that connects to the remote peer. You create ordinary Accept security policies to enable traffic between the IPsec interface and the interface that connects to the private network. This makes configuration simpler than for policy-based VPNs, which require IPsec security policies.


To define security policies for a route-based VPN

- 1 Define an ACCEPT security policy to permit communications between the local private network and the private network behind the remote peer. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select the address name that you defined for the private network behind this FortiGate unit.
Destination Interface/Zone	Select the IPsec Interface you configured.
Destination Address Name	Select the address name that you defined for the private network behind the remote peer.
Action	Select ACCEPT.
NAT	Disable.

- 2 To permit the remote client to initiate communication, you need to define a security policy for communication in that direction. Enter these settings in particular:

Source Interface/Zone	Select the IPsec Interface you configured.
Source Address Name	Select the address name that you defined for the private network behind the remote peer.
Destination Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address Name	Select the address name that you defined for the private network behind this FortiGate unit.
Action	Select <i>ACCEPT</i> .
NAT	Disable.



Gateway-to-gateway configurations

This section explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN.


The following topics are included in this section:

- [Configuration overview](#)
- [General configuration steps](#)
- [Configuring the two VPN peers](#)
- [How to work with overlapping subnets](#)
- [Testing](#)

Configuration overview

In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks. All traffic between the two networks is encrypted and protected by FortiGate security policies.


Figure 128: Example gateway-to-gateway configuration




In some cases, computers on the private network behind one VPN peer may (by coincidence) have IP addresses that are already used by computers on the network behind the other VPN peer. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent. To resolve issues related to ambiguous routing, see [“How to work with overlapping subnets” on page 1444](#).

In other cases, computers on the private network behind one VPN peer may obtain IP addresses from a local DHCP server. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and/or IP-address overlap issues may arise. For a discussion of the related issues, see [“FortiGate dialup-client configurations” on page 1501](#).

You can set up a fully meshed or partially meshed configuration (see [Figure 129](#) and [Figure 130](#)).

Figure 129: Fully meshed configuration

In a fully meshed network, all VPN peers are connected to each other, with one hop between peers. This topology is the most fault-tolerant: if one peer goes down, the rest of the network is not affected. This topology is difficult to scale because it requires connections between all peers. In addition, unnecessary communication can occur between peers. Best practices dictates a hub-and-spoke configuration instead (see [“Hub-and-spoke configurations” on page 1453](#)).

Figure 130: Partially meshed configuration

A partially meshed network is similar to a fully meshed network, but instead of having tunnels between all peers, tunnels are only configured between peers that communicate with each other regularly.

General configuration steps

The FortiGate units at both ends of the tunnel must be operating in NAT mode and have static public IP addresses.

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec phase 1 parameters to establish a secure connection and authenticate that VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the IPsec security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

To support these functions, the following general configuration steps must be performed by both FortiGate units:

- Define the phase 1 parameters that the FortiGate unit needs to authenticate the remote peer and establish a secure connection.
- Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with the remote peer.
- Create security policies to control the permitted services and permitted direction of traffic between the IP source and destination addresses.

Configuring the two VPN peers

Configure the VPN peers as follows. Each step is required, but these are general steps. For more detailed information on each step follow the cross references. See [“Auto Key phase 1 parameters” on page 1407](#).



All steps are required. Cross references point to required information that is repeated. No steps are optional.

Configuring Phase 1 and Phase 2 for both peers

This procedure applies to both peers. Repeat the procedure on each FortiGate unit, using the correct IP address for each. You may wish to vary the Phase 1 names but this is optional. Otherwise all steps are the same for each peer.

The phase 1 configuration defines the parameters that FortiGate_1 will use to authenticate FortiGate_2 and establish a secure connection. For the purposes of this example, a preshared key will be used to authenticate FortiGate_2. The same preshared key must be specified at both FortiGate units.

Before you define the phase 1 parameters, you need to:

- Reserve a name for the remote gateway.
- Obtain the IP address of the public interface to the remote peer.
- Reserve a unique value for the preshared key.

The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must have a minimum of 16 randomly chosen alphanumeric characters.

At the local FortiGate unit, define the phase 1 configuration needed to establish a secure connection with the remote peer. See [“Phase 1 configuration” on page 1394](#).

To create phase 1 to establish a secure connection with the remote peer

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*.
- 3 Enter the following information, and select *OK*.

Name	Enter <code>peer_1</code> . A name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Enter 172.20.0.2 when configuring FortiGate_1. Enter 172.18.0.2 when configuring FortiGate_2. The IP address of the remote peer public interface.
Local Interface	Select <i>wan1</i> . The FortiGate unit's public interface. This interface cannot be a loopback interface.
Enable IPsec Interface Mode	Select <i>Advanced</i> to see this setting. Enable <i>IPsec Interface Mode</i> to have the FortiGate unit create a virtual IPsec interface for a route-based VPN. Disable this option to create a policy-based VPN. For more information, see “Comparing policy-based or route-based VPNs” on page 1390 . After you select <i>OK</i> to create the phase 1 configuration, you cannot change this setting.

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration and specify the remote end point of the VPN tunnel. Before you define the phase 2 parameters, you need to reserve a name for the tunnel. See [“Phase 2 configuration” on page 1399](#).

To configure phase 2 settings

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*.
- 3 Enter the following information, and select *OK*.

Name	Enter <code>peer_1_p2</code> . A name to identify this phase 2 configuration.
Phase 1	Select <code>peer_1</code> . The name of the phase 1 configuration.

Creating security policies

Security policies control all IP traffic passing between a source address and a destination address.

An IPsec security policy is needed to allow the transmission of encrypted packets, specify the permitted direction of VPN traffic, and select the VPN tunnel that will be subject to the policy. A single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel.

Before you define security policies, you must first specify the IP source and destination addresses. In a gateway-to-gateway configuration:

- The IP source address corresponds to the private network behind the local FortiGate unit.
- The IP destination address refers to the private network behind the remote VPN peer.

When you are creating security policies, choose one of either route-based or policy-based methods and follow it for both VPN peers. DO NOT configure both route-based and policy-based policies on the same FortiGate unit for the same VPN tunnel.

The configuration of FortiGate_2 is similar to that of FortiGate_1. You must:

- Define the phase 1 parameters that FortiGate_2 needs to authenticate FortiGate_1 and establish a secure connection.
- Define the phase 2 parameters that FortiGate_2 needs to create a VPN tunnel with FortiGate_1.
- Create the security policy and define the scope of permitted services between the IP source and destination addresses.

When creating security policies it is good practice to include a comment describing what the policy does.

When creating security policies you need to be

- [Creating firewall addresses](#)
- [Creating route-based VPN security policies](#)
- [Configuring a default route for VPN interface](#)

or

- [Creating firewall addresses](#)
- [Creating policy-based VPN security policy](#)

Creating firewall addresses

Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks.

To define the IP address of the network behind FortiGate_1

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter <code>Finance_network</code> A meaningful address name that describes the network behind FortiGate_1.
Subnet/IP Range	Enter <code>10.21.101.0/24</code> . The IP address of the private network behind FortiGate_1.

To specify the address of the network behind FortiGate_2

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter <i>HR_network</i> .
Subnet/IP Range	Enter the IP address of the private network behind FortiGate_2 (for example, <i>10.31.101.0/24</i>).

Creating route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses.

To create route-based VPN security policies

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*.

Source Interface/Zone	Select <i>internal</i> The interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <i>Finance_network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2. The address name that for the private network behind this FortiGate unit.
Destination Interface/Zone	Select <i>peer_1</i> . The VPN Tunnel (IPsec Interface) you configured earlier.
Destination Address Name	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_network</i> when configuring FortiGate_2. The address name that you defined for the private network behind the remote peer.
Action	Select <i>ACCEPT</i> .
NAT	Disable.
Comment	Allow Internal to remote VPN network traffic

- 4 Configure any additional features such as UTM or traffic shaping you may want. (optional).
- 5 Select *Create New* to create another policy for the other direction.

- 6 Enter the following information, and select **OK**.

Source Interface/Zone	Select <i>peer_1</i> . The VPN Tunnel (IPsec Interface) you configured.
Source Address Name	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_Network</i> when configuring FortiGate_2. The address name defined for the private network behind the remote peer.
Destination Interface/Zone	Select <i>internal</i> The interface that connects to the private network behind this FortiGate unit.
Destination Address Name	Select <i>Finance_Network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2. The address name defined for the private network behind this FortiGate unit.
Action	Select <i>ACCEPT</i> .
NAT	Disable.
Comment	Allow remote VPN network traffic to Internal.

- 7 Configure any additional features such as UTM or traffic shaping you may want. (optional).

Configuring a default route for VPN interface

All network traffic must have a static route to direct its traffic to the proper destination. Without a route, traffic will not flow even if the security policies are configured properly.

You may need to create a static route entry for both directions of VPN traffic if your security policies allow bi-directional tunnel initiation.

To configure the route for a route-based VPN

- 1 On FortiGate_2, go to *Router > Static > Static Route*.
- 2 Select *Create New*, enter the following information, and then select **OK**:

Destination IP / Mask	10.21.101.0/24
Device	FGT2_to_FGT1_Tunnel
Gateway	Leave as default: 0.0.0.0.
Distance	Leave this at its default. If there are other routes on this FortiGate unit, you may need to set the distance on this route so the VPN traffic will use it as the default route. However, this normally happens by default because this route is typically a better match than the generic default route.

Creating policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

Source Interface/Zone	Select <i>internal</i> . The interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <i>Finance_network</i> when configuring FortiGate_1. Select <i>HR_network</i> when configuring FortiGate_2. The address name defined for the private network behind this FortiGate unit.
Destination Interface/Zone	Select <i>wan1</i> . The FortiGate unit's public interface.
Destination Address Name	Select <i>HR_network</i> when configuring FortiGate_1. Select <i>Finance_network</i> when configuring FortiGate_2. The address name that you defined in Step for the private network behind the remote peer.
Action	Select <i>IPSEC</i> .
VPN Tunnel	Select <i>peer_1</i> . The name of the phase 1 configuration that you created earlier. Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel. Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.
Comment	Bidirectional policy-based VPN policy




Place VPN policies in the policy list above any other policies having similar source and destination addresses.

How to work with overlapping subnets

A site-to-site VPN configuration sometimes has the problem that the private subnet addresses at each end are the same. You can resolve this problem by remapping the private addresses using virtual IP addresses (VIP).

VIPs allow computers on those overlapping private subnets to each have another set of IP addresses that can be used without confusion. The FortiGate unit maps the VIP addresses to the original addresses. This means if PC1 starts a session with PC2 at 10.31.101.10, FortiGate_2 directs that session to 10.11.101.10 — the actual IP address of PC2. [Figure 131](#) shows this — Finance network VIP is 10.21.101.0/24 and the HR network is 10.31.101.0/24.

Figure 131: Overlapped subnets example

Solution for route-based VPN

You need to:

- Configure IPsec Phase 1 and Phase 2 as you usually would for a route-based VPN. In this example, the resulting IPsec interface is named `FGT1_to_FGT2`.
- Configure virtual IP (VIP) mapping:
 - the 10.21.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate_1
 - the 10.31.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate_2
- Configure an outgoing security policy with ordinary source NAT on both FortiGates.
- Configure an incoming security policy with the VIP as the destination on both FortiGates.
- Configure a route to the remote private network over the IPsec interface on both FortiGates.

To configure VIP mapping on both FortiGates

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Name	Enter a name, for example, <code>my_vip</code> .
External Interface	Select <code>FGT1_to_FGT2</code> . The IPsec interface.
Type	Static NAT
External IP Address/Range	For the external IP address field enter: <ul style="list-style-type: none"> • 10.21.101.1 when configuring FortiGate_1, or • 10.31.101.1 when configuring FortiGate_2.

Mapped IP Address/Range	For the Mapped IP Address enter 10.11.101.1. For the Range enter 10.11.101.254.
Port Forwarding	Disable

To configure the outbound security policy on both FortiGates

Repeat this procedure on both FortiGate_1 and FortiGate_2.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*:

Source Interface/Zone	Select Port 1.
Source Address Name	Select all.
Destination Interface/Zone	Select FGT1_to_FGT2. The IPsec interface.
Destination Address Name	Select all.
Schedule	As required.
Service	As required.
Action	ACCEPT
NAT	Enable

To configure the inbound security policy on both FortiGates

Repeat this procedure on both FortiGate_1 and FortiGate_2.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and then select *OK*:

Source Interface/Zone	Select FGT1_to_FGT2.
Source Address Name	Select all.
Destination Interface/Zone	Select Port 1.
Destination Address Name	Select my-vip.
Schedule	As required.
Service	As required.
Action	ACCEPT
NAT	Disable

To configure the static route for both FortiGates

Repeat this procedure on both FortiGate_1 and FortiGate_2.

- 1 Go to *Router > Static > Static Route*.
- 2 Select *Create New*, enter the following information, and then select *OK*:

Destination IP / Mask	Enter 10.31.101.0/24 when configuring FortiGate_1 Enter 10.21.101.0/24 when configuring FortiGate_2
Device	Select FGT1_to_FGT2.
Gateway	Leave as default: 0.0.0.0.
Distance	Leave at default. If you have advanced routing on your network, you may have to change this value.



If you do not add a static route for the VPN tunnel, no traffic will be sent or received.

Solution for policy-based VPN

As with the route-based solution, users contact hosts at the other end of the VPN using an alternate subnet address. PC1 communicates with PC2 using IP address 10.31.101.10, and PC2 communicates with PC1 using IP address 10.21.101.10.

In this solution however, outbound NAT is used to translate the source address of packets from the 10.11.101.0/24 network to the alternate subnet address that hosts at the other end of the VPN use to reply. Inbound packets from the remote end have their destination addresses translated back to the 10.11.101.0/24 network.

For example, PC1 uses the destination address 10.31.101.10 to contact PC2. Outbound NAT on FortiGate_1 translates the PC1 source address to 10.21.101.10. At the FortiGate_2 end of the tunnel, the outbound NAT configuration translates the destination address to the actual PC2 address of 10.11.101.10. Similarly, PC2 replies to PC1 using destination address 10.21.101.10, with the PC2 source address translated to 10.31.101.10. PC1 and PC2 can communicate over the VPN even though they both have the same IP address.

You need to:

- Configure IPsec Phase 1 as you usually would for a policy-based VPN.
- Configure IPsec Phase 2 with the `use-natip disable` CLI option.
- Define a firewall address for the local private network, 10.11.101.0/24.
- Define a firewall address for the remote private network:
 - define a firewall address for 10.31.101.0/24 on FortiGate_1
 - define a firewall address for 10.21.101.0/24 on FortiGate_2
- Configure an outgoing IPsec security policy with outbound NAT to map 10.11.101.0/24 source addresses:
 - to the 10.21.101.0/24 network on FortiGate_1
 - to the 10.31.101.0/24 network on FortiGate_2

To configure IPsec Phase 2 - CLI

```

config vpn ipsec phase2
edit "FGT1_FGT2_p2"
set keepalive enable
set pfs enable
set phase1name FGT1_to_FGT2
set proposal 3des-sha1 3des-md5
set replay enable
set use-natip disable
end

```

In this example, your phase 1 definition is named FGT1_to_FGT2. `use-natip` is set to `disable`, so you can specify the source selector using the `src-addr-type`, `src-start-ip / src-end-ip` or `src-subnet` keywords. This example leaves these keywords at their default values, which specify the subnet 0.0.0.0/0.

The `pfs` keyword ensures that perfect forward secrecy (PFS) is used. This ensures that each Phase 2 key created is unrelated to any other keys in use.

To define the local private network firewall address

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New* and enter the following information:

Address Name	Enter <code>vpn-local</code> . A meaningful name for the local private network
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	Any

To define the remote private network firewall address

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*:

Address Name	Enter <code>vpn-remote</code> . A meaningful name for the remote private network.
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.0 255.255.255.0 on FortiGate_1 10.21.101.0 255.255.255.0 on FortiGate_2
Interface	Any

To configure the IPsec security policy

In the CLI on FortiGate_1, enter the commands:

```
config firewall policy
edit 1
set srcintf "port1"
set dstintf "port2"
set srcaddr "vpn-local"
set dstaddr "vpn-remote"
set action ipsec
set schedule "always"
set service "ANY"
set inbound enable
set outbound enable
set vpntunnel "FGT1_to_FGT2"
set natoutbound enable
set natip 10.31.101.0 255.255.255.0
end
```

Optionally, you can set everything except `natip` in the web-based manager and then use the CLI to set `natip`.

Enter the same commands on FortiGate_2, but set `natip` be `10.21.101.0 255.255.255.0`.

Testing

The best testing is to look at the packets both as the VPN tunnel is negotiated, and when the tunnel is up.

To determine what the other end of the VPN tunnel is proposing

- 1 Start a terminal program such as `puTTY` and set it to log all output.
When necessary refer to the logs to locate information when output is verbose.
- 2 Logon to the FortiGate unit using a `super_admin` account.
- 3 Enter the following CLI commands.
- 4 Display all the possible IKE error types and the number of times they have occurred:

```
diag vpn ike errors
```
- 5 Check for existing debug sessions:

```
diag debug info
```

If a debug session is running, to halt it enter:

```
diag debug disable
```
- 6 Confirm your proposal settings:

```
diag vpn ike config list
```
- 7 If your proposal settings do not match what you expect, make a change to it and save it to force an update in memory. If that fixes the problem, stop here.
- 8 List the current vpn filter:

```
diag vpn ike filter
```

- 9** If all fields are set to any, there are no filters set and all VPN ike packets will be displayed in the debug output. If your system has only a few VPNs, skip setting the filter.

If your system has many VPN connections this will result in very verbose output and make it very difficult to locate the correct connection attempt.

- 10** Set the VPN filter to display only information from the destination IP address for example 10.10.10.10:

```
diag vpn ike log-filter dst-addr4 10.10.10.10
```

To add more filter options, enter them one per line as above. Other filter options are displayed in [Table 89](#).

Table 89: Filter options for diag vpn ike filter

clear	erase the current filter
dst-addr6	the IPv6 destination address range to filter by
dst-port	the destination port range to filter by
interface	interface that IKE connection is negotiated over
list	display the current filter
name	the phase1 name to filter by
negate	negate the specified filter parameter
src-addr4	the IPv4 source address range to filter by
src-addr6	the IPv6 source address range to filter by
src-port	the source port range to filter by
vd	index of virtual domain. 0 matches all

- 11** Start debugging:

```
diag debug app ike 255
diag debug enable
```

- 12** Have the remote end attempt a VPN connection.



If the remote end attempts the connection they become the initiator. This situation makes it easier to debug VPN tunnels because then you have the remote information and all of your local information. by initiate the connection, you will not see the other end's information.

- 13** If possible go to the web-based manager on your FortiGate unit, go to the VPN monitor and try to bring the tunnel up.

- 14** Stop the debug output:

```
diag debug disable
```

- 15** Go back through the output to determine what proposal information the initiator is using, and how it is different from your VPN P1 proposal settings.


Things to look for in the debug output of attempted VPN connections are shown below.

Table 90: Important terms to look for in VPN debug output

initiator	Starts the VPN attempt, in the above procedure that is the remote end
responder	Answers the initiator's request

Table 90: Important terms to look for in VPN debug output

local ID	In aggressive mode, this is not encrypted
error no SA proposal chosen	There was no proposal match — there was no encryption-authentication pair in common, usually occurs after a long list of proposal attempts
R U THERE and R U THERE ack	dead peer detection (dpd), also known as dead gateway detection — after three failed attempts to contact the remote end it will be declared dead, no farther attempts will be made to contact it
negotiation result	lists the proposal settings that were agreed on
SA_life_soft and SA_life_hard	negotiating a new key, and the key life
R U THERE	If you see this, it means Phase 1 was successful
tunnel up	the negotiation was successful, the VPN tunnel is operational



Hub-and-spoke configurations


This section describes how to set up hub-and-spoke IPsec VPNs. The following topics are included in this section:

- [Configuration overview](#)
- [Configure the hub](#)
- [Configure the spokes](#)
- [Dynamic spokes configuration example](#)

Configuration overview

In a hub-and-spoke configuration, VPN connections radiate from a central FortiGate unit (the hub) to a number of remote peers (the spokes). Traffic can pass between private networks behind the hub and private networks behind the remote peers. Traffic can also pass between remote peer private networks through the hub.

Figure 132: Example hub-and-spoke configuration



The actual implementation varies in complexity depending on

- whether the spokes are statically or dynamically addressed
- the addressing scheme of the protected subnets
- how peers are authenticated

This guide discusses the issues involved in configuring a hub-and-spoke VPN and provides some basic configuration examples.

Hub-and-spoke infrastructure requirements

- The FortiGate hub must be operating in NAT mode and have a static public IP address.
- Spokes may have static IP addresses, dynamic IP addresses (see [“FortiGate dialup-client configurations” on page 1501](#)), or static domain names and dynamic IP addresses (see [“Dynamic DNS configuration” on page 1469](#)).

Spoke gateway addressing

The public IP address of the spoke is the VPN remote gateway as seen from the hub. Statically addressed spokes each require a separate VPN phase 1 configuration on the hub. When there are many spokes, this becomes rather cumbersome.

Using dynamic addressing for spokes simplifies the VPN configuration because then the hub requires only a single phase 1 configuration with “dialup user” as the remote gateway. You can use this configuration even if the remote peers have static IP addresses. A remote peer can establish a VPN connection regardless of its IP address if its traffic selectors match and it can authenticate to the hub. See [“Dynamic spokes configuration example” on page 1463](#) for an example of this configuration.

Protected networks addressing


The addresses of the protected networks are needed to configure destination selectors and sometimes for security policies and static routes. The larger the number of spokes, the more addresses there are to manage. You can

- assign spoke subnets as part of a larger subnet, usually on a new network or
- create address groups that contain all of the needed addresses

Using aggregated subnets

If you are creating a new network, where subnet IP addresses are not already assigned, you can simplify the VPN configuration by assigning spoke subnets that are part of a large subnet.

Figure 133: Aggregated subnets



All spokes use the large subnet address, 10.1.0.0/16 for example, as

- the IPsec destination selector
- the destination of the security policy from the private subnet to the VPN (required for policy-based VPN, optional for route-based VPN)
- the destination of the static route to the VPN (route-based)

Each spoke uses the address of its own protected subnet as the IPsec source selector and as the source address in its VPN security policy. The remote gateway is the public IP address of the hub FortiGate unit.

Using an address group

If you want to create a hub-and-spoke VPN between existing private networks, the subnet addressing usually does not fit the aggregated subnet model discussed earlier. All of the spokes and the hub will need to include the addresses of all the protected networks in their configuration.

On FortiGate units, you can define a named firewall address for each of the remote protected networks and add these addresses to a firewall address group. For a policy-based VPN, you can then use this address group as the destination of the VPN security policy.

For a route-based VPN, the destination of the VPN security policy can be set to All. You need to specify appropriate routes for each of the remote subnets.

Authentication

Authentication is by a common preshared key or by certificates. For simplicity, the examples in this chapter assume that all spokes use the same preshared key.

Configure the hub

At the FortiGate unit that acts as the hub, you need to

- configure the VPN to each spoke
- configure communication between spokes

You configure communication between spokes differently for a policy-based VPN than for a route-based VPN. For a policy-based VPN, you configure a VPN concentrator. For a route-based VPN, you must either define security policies or group the IPsec interfaces into a zone

Define the hub-spoke VPNs

Perform these steps at the FortiGate unit that will act as the hub. Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

To configure the VPN hub

- 1 At the hub, define the phase 1 configuration for each spoke. See [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings in particular:

Name	Enter a name to identify the VPN in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	<p>The remote gateway is the other end of the VPN tunnel. There are three options:</p> <p>Static IP Address — Enter the spoke’s public <i>IP Address</i>. You will need to create a phase 1 configuration for each spoke. Either the hub or the spoke can establish the VPN connection.</p> <p>Dialup User — No additional information is needed. The hub accepts connections from peers with appropriate encryption and authentication settings. Only one phase 1 configuration is needed for multiple dialup spokes. Only the spoke can establish the VPN tunnel.</p> <p>Dynamic DNS — If the spoke subscribes to a dynamic DNS service, enter the spoke’s <i>Dynamic DNS</i> domain name. Either the hub or the spoke can establish the VPN connection. For more information, see “Dynamic DNS configuration” on page 1469.</p>
Local Interface	<p>Select the FortiGate interface that connects to the remote gateway. This is usually the FortiGate unit’s public interface.</p> <p>This interface cannot be a loopback interface.</p>
Enable IPsec Interface Mode	<p>You must select Advanced to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN. For more information, see “Comparing policy-based or route-based VPNs” on page 1390.</p> <p>After you select OK to create the phase 1 configuration, you cannot change this setting.</p>

- 2 Define the phase 2 parameters needed to create a VPN tunnel with each spoke. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Name	Enter a name to identify this spoke phase 2 configuration.
Phase 1	Select the name of the phase 1 configuration that you defined for this spoke.

Define the hub-spoke security policies

- 1 Define a name for the address of the private network behind the hub. For more information, see [“Defining policy addresses” on page 1431](#).
- 2 Define names for the addresses or address ranges of the private networks behind the spokes. For more information, see [“Defining policy addresses” on page 1431](#).
- 3 Define the VPN concentrator. See [“To define the VPN concentrator” on page 1458](#).

- 4 Define security policies to permit communication between the hub and the spokes. For more information, see [“Defining VPN security policies” on page 1432](#).

Route-based VPN security policies

Define ACCEPT security policies to permit communications between the hub and the spoke. You need one policy for each direction. Enter these settings in particular:

Source Interface/Zone	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
Source Address Name	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
Destination Interface/Zone	Select the hub’s interface to the internal (private) network.
Destination Address Name	Select the source address that you defined in Step 1.
Action	Select <i>ACCEPT</i> .
NAT	Enable.

Source Interface/Zone	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate units.
Source Address Name	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
Destination Interface/Zone	Select the source address that you defined in Step 1.
Destination Address Name	Select the hub’s interface to the internal (private) network.
Action	Select <i>ACCEPT</i> .
NAT	Enable.

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the hub and the spoke. Enter these settings in particular:

Source Interface/Zone	Select the hub’s interface to the internal (private) network.
Source Address Name	Select the source address that you defined in Step 1.
Destination Interface/Zone	Select the hub’s public network interface.
Destination Address Name	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
Action	IPSEC
VPN Tunnel	<p>Select the name of the phase 1 configuration that you created for the spoke in Step 1.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

- 5 In the policy list, arrange the policies in the following order:
 - IPsec policies that control traffic between the hub and the spokes first
 - the default security policy last

Configuring communication between spokes (policy-based VPN)

For a policy-based hub-and-spoke VPN, you define a concentrator to enable communication between the spokes.

To define the VPN concentrator

- 1 At the hub, go to *VPN > IPSEC > Concentrator* and select *Create New*.
- 2 In the *Concentrator Name* field, type a name to identify the concentrator.
- 3 From the *Available Tunnels* list, select a VPN tunnel and then select the right-pointing arrow.



To remove tunnels from the VPN concentrator, select the tunnel in the *Members* list and select the left-pointing arrow.

- 4 Repeat Step 3 until all of the tunnels associated with the spokes are included in the concentrator.
- 5 Select *OK*.

Configuring communication between spokes (route-based VPN)

For a route-based hub-and-spoke VPN, there are several ways you can enable communication between the spokes:

- put all of the IPsec interfaces into a zone and enable intra-zone traffic. This eliminates the need for any security policy for the VPN, but you cannot apply UTM features to scan the traffic for security threats.
- put all of the IPsec interfaces into a zone and create a single zone-to-zone security policy
- create a security policy for each pair of spokes that are allowed to communicate with each other. The number of policies required increases rapidly as the number of spokes increases.

Using a zone as a concentrator

A simple way to provide communication among all of the spokes is to create a zone and allow intra-zone communication. You cannot apply UTM features using this method.

- 1 Go to *System > Network > Interface*.
- 2 Select the down-arrow on the *Create New* button and select *Zone*.
- 3 In the *Zone Name* field, enter a name, such as *Our_VPN_zone*.
- 4 Clear *Block intra-zone traffic*.
- 5 In the *Interface Members* list, select the IPsec interfaces that are part of your VPN.
- 6 Select *OK*.

Using a zone with a policy as a concentrator

If you put all of the hub IPsec interfaces involved in the VPN into a zone, you can enable communication among all of the spokes and apply UTM features with just one security policy.

To create a zone for the VPN

- 1 Go to *System > Network > Interface*.
- 2 Select the down-arrow on the *Create New* button and select *Zone*.
- 3 In the *Zone Name* field, enter a name, such as *Our_VPN_zone*.
- 4 Select *Block intra-zone traffic*.
- 5 In the *Interface Members* list, select the IPsec interfaces that are part of your VPN.
- 6 Select *OK*.

To create a security policy for the zone

- 1 Go to *Policy > Policy > Policy*. Select *Create New* and enter the settings:

Source Interface/Zone	Select the zone you created for your VPN.
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select the zone you created for your VPN.
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Enable.
UTM	If you want to apply UTM features to this traffic, select the appropriate profiles.

- 2 Select *OK*.

Using security policies as a concentrator

To enable communication between two spokes, you need to define an *ACCEPT* security policy for them. To allow either spoke to initiate communication, you must create a policy for each direction. This procedure describes a security policy for communication from Spoke 1 to Spoke 2. Others are similar.

- 1 Define names for the addresses or address ranges of the private networks behind each spoke. For more information, see [“Defining policy addresses” on page 1431](#).
- 2 Go to *Policy > Policy > Policy*. Select *Create New* and enter the settings:

Source Interface/Zone	Select the IPsec interface that connects to Spoke 1.
Source Address Name	Select the address of the private network behind Spoke 1.
Destination Interface/Zone	Select the IPsec interface that connects to Spoke 2.
Destination Address Name	Select the address of the private network behind Spoke 2.
Action	Select <i>ACCEPT</i> .

NAT	Enable.
UTM	If you want to apply UTM features to this traffic, select the appropriate profiles.

- 3 Select OK.

Configure the spokes

Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

Perform these steps at each FortiGate unit that will act as a spoke.

To create the phase 1 and phase_2 configurations

- 1 At the spoke, define the phase 1 parameters that the spoke will use to establish a secure connection with the hub. See [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings:

Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Type the IP address of the interface that connects to the hub.
Enable IPsec Interface Mode	Enable if you are creating a route-based VPN. Clear if you are creating a policy-based VPN.

- 2 Create the phase 2 tunnel definition. See [“Phase 2 parameters” on page 1425](#). Select the set of phase 1 parameters that you defined for the hub. You can select the name of the hub from the *Static IP Address* part of the list.

Configuring security policies for hub-to-spoke communication

- 1 Create an address for this spoke. See [“Defining policy addresses” on page 1431](#). Enter the IP address and netmask of the private network behind the spoke.
- 2 Create an address to represent the hub. See [“Defining policy addresses” on page 1431](#). Enter the IP address and netmask of the private network behind the hub.

3 Define the security policy to enable communication with the hub.

Route-based VPN security policy

Define two security policies to permit communications to and from the hub. Enter these settings:

Source Interface/Zone	Select the virtual IPsec interface you created.
Source Address Name	Select the hub address you defined in Step 1.
Destination Interface/Zone	Select the spoke's interface to the internal (private) network.
Destination Address Name	Select the spoke addresses you defined in Step 2.
Action	Select ACCEPT
NAT	Enable

Source Interface/Zone	Select the spoke's interface to the internal (private) network.
Source Address Name	Select the spoke address you defined in Step 1.
Destination Interface/Zone	Select the virtual IPsec interface you created.
Destination Address Name	Select the hub destination addresses you defined in Step 2.
Action	Select ACCEPT
NAT	Enable

Policy-based VPN security policy

Define an IPsec security policy to permit communications with the hub. See [“Defining VPN security policies” on page 1432](#). Enter these settings in particular:

Source Interface/Zone	Select the spoke's interface to the internal (private) network.
Source Address Name	Select the spoke address you defined in Step 1.
Destination Interface/Zone	Select the spoke's interface to the external (public) network.
Destination Address Name	Select the hub address you defined in Step 2.
Action	Select IPSEC
VPN Tunnel	<p>Select the name of the phase 1 configuration you defined.</p> <p>Select Allow inbound to enable traffic from the remote network to initiate the tunnel.</p> <p>Select Allow outbound to enable traffic from the local network to initiate the tunnel.</p>

Configuring security policies for spoke-to-spoke communication

Each spoke requires security policies to enable communication with the other spokes. Instead of creating separate security policies for each spoke, you can create an address group that contains the addresses of the networks behind the other spokes. The security policy then applies to all of the spokes in the group.

- 1 Define destination addresses to represent the networks behind each of the other spokes. Add these addresses to an address group. For more information, see “Configuring Address Groups” section in the “Firewall Address” chapter of the *FortiGate Administration Guide*.
- 2 Define the security policy to enable communication between this spoke and the spokes in the address group you created.

Policy-based VPN security policy

Define an IPsec security policy to permit communications with the other spokes. See “Defining VPN security policies” on page 1432. Enter these settings in particular:

Route-based VPN security policy

Define two security policies to permit communications to and from the other spokes. Enter these settings in particular:

Source Interface/Zone	Select the virtual IPsec interface you created.
Source Address Name	Select the spoke address group you defined in Step 1.
Destination Interface/Zone	Select the spoke’s interface to the internal (private) network.
Destination Address Name	Select this spoke’s address name.
Action	Select <i>ACCEPT</i>
NAT	Enable

Source Interface/Zone	Select the spoke’s interface to the internal (private) network.
Source Address Name	Select this spoke’s address name.
Destination Interface/Zone	Select the virtual IPsec interface you created.
Destination Address Name	Select the spoke address group you defined in Step 1.
Action	Select <i>ACCEPT</i>
NAT	Enable

Policy-based VPN security policy

Source Interface/Zone	Select this spoke’s internal (private) network interface.
Source Address Name	Select this spoke’s source address.
Destination Interface/Zone	Select the spoke’s interface to the external (public) network.
Destination Address Name	Select the spoke address group you defined in Step 1.


Action	Select <i>IPSEC</i>
VPN Tunnel	<p>Select the name of the phase 1 configuration you defined.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

- 1 Place this policy or policies in the policy list above any other policies having similar source and destination addresses.

Dynamic spokes configuration example

This example demonstrates how to set up a basic route-based hub-and-spoke IPsec VPN that uses preshared keys to authenticate VPN peers.

Figure 134: Example hub-and-spoke configuration



In the example configuration, the protected networks 10.1.0.0/24, 10.1.1.0/24 and 10.1.2.0/24 are all part of the larger subnet 10.1.0.0/16. The steps for setting up the example hub-and-spoke configuration create a VPN among Site 1, Site 2, and the HR Network.

The spokes are dialup. Their addresses are not part of the configuration on the hub, so only one spoke definition is required no matter the number of spokes. For simplicity, only two spokes are shown.

Configure the hub (FortiGate_1)

The phase 1 configuration defines the parameters that FortiGate_1 will use to authenticate spokes and establish secure connections.

For the purposes of this example, one preshared key will be used to authenticate all of the spokes. Each key must contain at least 6 printable characters and best practices dictates that it only be known by network administrators. For optimum protection against currently known attacks, each key must consist of a minimum of 16 randomly chosen alphanumeric characters.

Define the IPsec configuration

To define the phase 1 parameters

- 1 At FortiGate_1, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Define the phase 1 parameters that the hub will use to establish a secure connection to the spokes. Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Type a name (for example, toSpokes).
Remote Gateway	Dialup user
Local Interface	External
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID

The basic phase 2 settings associate IPsec phase 2 parameters with the phase 1 configuration and specify the remote end points of the VPN tunnels.

To define the phase 2 parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Create a phase 2 tunnel definition for the spokes. Select *Create Phase 2*, enter the following information, and select *OK*:

Name	Enter a name for the phase 2 definition (for example, toSpokes_ph2).
Phase 1	Select the Phase 1 configuration that you defined previously (for example, toSpokes).

Define the security policies

security policies control all IP traffic passing between a source address and a destination address. For a route-based VPN, the policies are simpler than for a policy-based VPN. Instead of an IPSEC policy, you use an ACCEPT policy with the virtual IPsec interface as the external interface.

Before you define security policies, you must first define firewall addresses to use in those policies. You need addresses for:

- the HR network behind FortiGate_1
- the aggregate subnet address for the protected networks

To define the IP address of the HR network behind FortiGate_1

- 1 Go to *Firewall Objects > Address > Address*.

- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter an address name (for example, <code>HR_Network</code>).
Subnet/IP Range	Enter the IP address of the HR network behind FortiGate_1 (for example, <code>10.1.0.0/24</code>).

To specify the IP address the aggregate protected subnet

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter an address name (for example, <code>Spoke_net</code>).
Subnet/IP Range	Enter the IP address of the aggregate protected network, <code>10.1.0.0/16</code>

To define the security policy for traffic from the hub to the spokes

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Select the interface to the HR network, <i>port 1</i> .
Address Name	<code>HR_Network</code>
Destination Interface/Zone	Select the virtual IPsec interface that connects to the spokes, <i>toSpokes</i>
Address Name	<code>Spoke_net</code>
Schedule	As required.
Service	As required.
Action	ACCEPT

- 3 Place the policy in the policy list above any other policies having similar source and destination addresses.

Configure communication between spokes

Spokes communicate with each other through the hub. You need to configure the hub to allow this communication. An easy way to do this is to create a zone containing the virtual IPsec interfaces even if there is only one, and create a zone-to-zone security policy.

To create a zone for the VPN

- 1 Go to *System > Network > Interface*.
- 2 Select the down-arrow on the *Create New* button and select *Zone*.
- 3 In the *Zone Name* field, enter a name, such as `Our_VPN_zone`.
- 4 Select *Block intra-zone traffic*.
You could enable intra-zone traffic and then you would not need to create a security policy. But, you would not be able to apply UTM features.
- 5 In *Interface Members*, select the virtual IPsec interface, *toSpokes*.
- 6 Select *OK*.

To create a security policy for the zone

- 1 Go to *Policy > Policy > Policy*. Select *Create New* and enter these settings:

Source Interface/Zone	Select <i>Our_VPN_zone</i> .
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select <i>Our_VPN_zone</i> .
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Enable.
UTM	Select the appropriate UTM profiles.

- 2 Select *OK*.

Configure the spokes

In this example, all spokes have nearly identical configuration, requiring the following:

- phase 1 authentication parameters to initiate a connection with the hub
- phase 2 tunnel creation parameters to establish a VPN tunnel with the hub
- a source address that represents the network behind the spoke. This is the only part of the configuration that is different for each spoke.
- a destination address that represents the aggregate protected network
- a security policy to enable communications between the spoke and the aggregate protected network

Define the IPsec configuration

At each spoke, create the following configuration.

To define the Phase 1 parameters

- 1 At the spoke, go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Type a name, for example, <i>toHub</i> .
Remote Gateway	Static IP Address
IP Address	<i>172.16.10.1</i>
Local Interface	<i>Port2</i>
Mode	<i>Main</i>
Authentication Method	<i>Preshared Key</i>
Pre-shared Key	Enter the preshared key. The value must be identical to the preshared key that you specified previously in the <i>FortiGate_1</i> configuration.
Peer Options	<i>Accept any peer ID</i>
Enable IPsec Interface Mode	Select <i>Advanced</i> to see this option. Enable the option to create a route-based VPN.

To define the Phase 2 parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*, enter the following information, and select *OK*:

Name	Enter a name for the tunnel, for example, <code>toHub_ph2</code> .
Phase 1	Select the name of the phase 1 configuration that you defined previously, for example, <code>toHub</code> .
Advanced	Select to show the following <i>Quick Mode Selector</i> settings.
Source	Enter the address of the protected network at this spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .
Destination	Enter the aggregate protected subnet address, <code>10.1.0.0/16</code> .

Define the security policies

You need to define firewall addresses for the spokes and the aggregate protected network and then create a security policy to enable communication between them.

To define the IP address of the network behind the spoke

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter an address name, for example <code>LocalNet</code> .
Subnet/IP Range	Enter the IP address of the private network behind the spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .

To specify the IP address of the aggregate protected network

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Enter an address nam, for example, <code>Spoke_net</code> .
Subnet/IP Range	Enter the IP address of the aggregate protected network, <code>10.1.0.0/16</code> .

To define the security policy

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and select *OK*:


Source Interface/Zone	Select the virtual IPsec interface, <code>toHub</code> .
Address Name	Select the aggregate protected network address <code>Spoke_net</code>
Destination Interface/Zone	Select the interface to the internal (private) network, <code>port1</code> .

Address Name	Select the address for this spoke's protected network LocalNet
Schedule	As required.
Service	As required.
Action	ACCEPT

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Select the interface to the internal private network, port1.
Address Name	Select the address for this spoke's protected network, LocalNet
Destination Interface/Zone	Select the virtual IPsec interface, toHub.
Address Name	Select the aggregate protected network address, Spoke_net
Schedule	As required.
Service	As required.
Action	ACCEPT

- 4 Place these policies in the policy list above any other policies having similar source and destination addresses.



Dynamic DNS configuration

This section describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a domain name and a dynamic IP address.

The following topics are included in this section:

- [Dynamic DNS over VPN concepts](#)
- [Dynamic DNS topology](#)
- [General configuration steps](#)
- [Configure the dynamically-addressed VPN peer](#)
- [Configure the fixed-address VPN peer](#)
- [Testing](#)

Dynamic DNS over VPN concepts

A typical computer has a static IP address and one or more DNS servers to resolve fully qualified domain names (FQDN) into IP addresses. A domain name assigned to this computer is resolved by any DNS server having an entry for the domain name and its static IP address. The IP address never changes or changes only rarely so the DNS server can reliably say it has the correct address for that domain all the time.

Dynamic DNS (DDNS)

It is different when a computer has a dynamic IP address, such as an IP address assigned dynamically by a DHCP server, and a domain name. Computers that want to contact this computer do not know what its current IP address is. To solve this problem there are dynamic DNS servers. These are public servers that store a DNS entry for your computer that includes its current IP address and associated domain name. These entries are kept up to date by your computer sending its current IP address to the dynamic DNS (DDNS) server to ensure its entry is always up to date. When other computers want to contact your domain, their DNS gets your IP address from your DDNS server. To use DDNS servers, you must subscribe to them and usually pay for their services.

When configuring DDNS on your FortiGate unit, go to *System > Network > DNS* and enable *Use DDNS*. Then select the interface with the dynamic connection, which DDNS server you have an account with, your domain name, and account information. If your DDNS server is not on the list, there is a generic option where you can provide your DDNS server information.

Routing

When an interface has some form of changing IP address (DDNS, PPPoE, or DHCP assigned address), routing needs special attention. The standard static route cannot handle the changing IP address. The solution is to use the dynamic-gateway command in the CLI. Say for example you already have four static routes, and you have a PPPoE connection over the wan2 interface and you want to use that as your default route.

The route is configured on the dynamic address VPN peer trying to access the static address FortiGate unit.

To configure dynamic gateway routing - CLI

```
config router static
  edit 5
    set dst 0.0.0.0 0.0.0.0
    set dynamic-gateway enable
    set device wan2
  next
end
```

handbook chapter

Dynamic DNS over VPN

IPsec VPN expects an IP address for each end of the VPN tunnel. All configuration and communication with that tunnel depends on the IP addresses as reference points. However, when the interface the tunnel is on has DDNS enabled there is no set IP address. The remote end of the VPN tunnel now needs another way to reference your end of the VPN tunnel. This is accomplished using Local ID.

A FortiGate unit that has a domain name and a dynamic IP address can initiate VPN connections anytime—the remote peer can reply to the local FortiGate unit using the source IP address that was sent in the packet header because it is current. Without doing a DNS lookup first, the remote peer runs the risk of the dynamic IP changing before it attempts to connect. To avoid this, the remote peer must perform a DNS lookup for the domain name of to be sure of the dynamic IP address before initiating the connection.

Remote Gateway

When configuring the Phase 1 entry for a VPN tunnel, the Remote Gateway determines the addressing method the remote end of the tunnel uses as one of Static IP Address, Dialup User, or Dynamic DNS. There are different fields for each option.

When you select the Dynamic DNS VPN type there is a related field called Dynamic DNS. The Dynamic DNS field is asking for the FQDN of the remote end of the tunnel. It uses this information to look up the IP address of the remote end of the tunnel through the DDNS server associated with that domain name.

Local ID (peer ID)

The Local ID or peer ID can be used to uniquely identify one end of a VPN tunnel. This enables a more secure connection. Also if you have multiple VPN tunnels negotiating, this ensures the proper remote and local ends connect. When you configure it on your end, it is your Local ID. When the remote end connects to you, they see it as your peer ID.

If you are debugging a VPN connection, the Local ID is part of the VPN negotiations. You can use it to help troubleshoot connection problems.

To configure your Local ID

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create New Phase 1* or edit an existing Phase 1 entry.
- 3 Select *Advanced*.
- 4 In the *P1 Proposal* section, enter your Local ID.
- 5 Select *OK*.

The default configuration is to accept all local IDs (peer IDs). If you have the Local ID set, the remote end of the tunnel must be configured to accept your Local ID.

To accept a specific Peer ID

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create New Phase 1*.
- 3 Select *Aggressive mode*.
- 4 For *Peer Options*, select *Accept this peer ID*. This option becomes visible only when Aggressive mode is selected.
- 5 Enter the string the other end of the tunnel used for its Local ID.
- 6 Configure the rest of the Phase 1 entry as required.
- 7 Select *OK*.

Route-based or policy-based VPN

VPN over dynamic DNS can be configured with either route-based or policy-based VPN settings. Both are valid, but have differences in configuration. Choose the best method based on your requirements. For more information on route-based and policy-based, see [“Types of VPNs” on page 1389](#).


Route-based VPN configuration requires two security policies to be configured (one for each direction of traffic) to permit traffic over the VPN virtual interface, and you must also add a static route entry for that VPN interface or the VPN traffic will not reach its destination. See [“Creating branch_2 route-based security policies” on page 1475](#) and [“Creating branch_1 route-based security policies” on page 1480](#).

Policy-based VPN configuration uses more complex and often more IPsec security policies, but does not require a static route entry. It has the benefit of being able to configure multiple policies for handling multiple protocols in different ways, such as more scanning of less secure protocols or guaranteeing a minimum bandwidth for protocols such as VoIP. See [“Creating branch_2 policy-based security policies” on page 1477](#) and [“Creating branch_1 policy-based security policies” on page 1481](#).

Dynamic DNS topology

In this scenario, two branch offices each have a FortiGate unit and are connected in a gateway-to-gateway VPN configuration. One FortiGate unit has a domain name (example.com) with a dynamic IP address. See [branch_2 in Figure 135](#).

Whenever the [branch_2](#) unit connects to the Internet (and possibly also at predefined intervals set by the ISP), the ISP may assign a different IP address to the FortiGate unit. The unit has its domain name registered with a dynamic DNS service. The [branch_2](#) unit checks in with the DDNS server on a regular basis, and that server provides the DNS information for the domain name, updating the IP address from time to time. Remote peers have to locate the [branch_2](#) FortiGate unit through a DNS lookup each time to ensure the address they get is current and correct.

Figure 135: Example dynamic DNS configuration

When a remote peer (such as the `branch_1` FortiGate unit in Figure 135) initiates a connection to `example.com`, the local DNS server looks up and returns the IP address that matches the domain name `example.com`. The remote peer uses the retrieved IP address to establish a VPN connection with the `branch_2` FortiGate unit.

Assumptions

- You have administrator access to both FortiGate units.
- Both FortiGate units have interfaces named `wan1` and `internal`. (If not, you can use the alias feature to assign these labels as “nicknames” to other interfaces to follow this example.)
- Both FortiGate units have the most recent firmware installed, have been configured for their networks, and are currently passing normal network traffic.
- The `branch_2` FortiGate unit has its `wan1` interface defined as a dynamic DNS interface with the domain name of `example.com`.
- A basic gateway-to-gateway configuration is in place (see “[Gateway-to-gateway configurations](#)” on page 1437) except one of the FortiGate units has a static domain name and a dynamic IP address instead of a static IP address.
- The FortiGate unit with the domain name is subscribed to one of the supported dynamic DNS services. Contact one of the services to set up an account. For more information and instructions about how to configure the FortiGate unit to push its dynamic IP address to a dynamic DNS server, see the [System Administration handbook chapter](#).

General configuration steps

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec phase 1 parameters to establish a secure connection and authenticate the VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

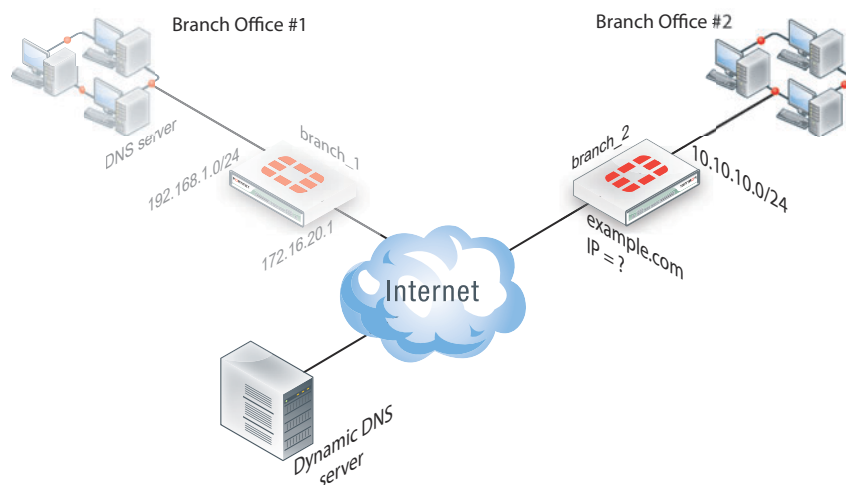
To support these functions, the following general configuration steps must be performed:

- Configure the branch_2 FortiGate unit with the dynamic IP address. This unit uses a Local ID string instead of an IP address to identify itself to the remote peer. See [“Configure the dynamically-addressed VPN peer” on page 1473](#).
 - [Configuring branch_2 VPN tunnel settings](#)
 - [Configuring branch_2 security policies](#)
- Configure the fixed-address VPN peer. To initiate a VPN tunnel with the dynamically-addressed peer, this unit must first retrieve the IP address for the domain from the dynamic DNS service. See [“Configure the fixed-address VPN peer” on page 1478](#).
 - [Configuring branch_1 VPN tunnel settings](#)
 - [Configuring branch_1 security policies](#)

Configure the dynamically-addressed VPN peer

It is assumed that this FortiGate unit (branch_2) has already had its public facing interface, for example the wan1, configured with the proper dynamic DNS configuration.

Figure 136: Configure branch_2, the dynamic address side



Configuring the dynamically-addressed VPN peer includes:

- [Configuring branch_2 VPN tunnel settings](#)
- [Configuring branch_2 security policies](#)

Configuring branch_2 VPN tunnel settings

Define the phase 1 parameters needed to establish a secure connection with the remote peer. See [“Auto Key phase 1 parameters” on page 1407](#). During this procedure you need to choose if you will be using route-based or policy-based VPNs.

To configure branch_2 VPN tunnel settings

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create New Phase 1*.

- 3 Enter the following information and select **OK**.

Name	Enter <code>branch_2</code> , a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies, and the VPN monitor.
Remote Gateway	Select <i>Static IP Address</i> . The remote peer this FortiGate is connecting to has a static IP public address. If the remote interface is PPPoE do not select <i>Retrieve default gateway from server</i> .
IP Address	Enter 172.16.20.1 The IP address of the public interface to the remote peer.
Mode	Select <i>Aggressive</i> .
Advanced	
Enable IPsec Interface Mode	Enable for a route-based VPN and when configuring policies, go to “Creating branch_2 route-based security policies” on page 1475 . Disable for a policy-based VPN and when configuring policies, go to “Creating branch_2 policy-based security policies” on page 1477 . If enabled, default settings are used.
Local ID	Enter <code>example.com</code> A character string used by the <code>branch_2</code> FortiGate unit to identify itself to the remote peer. This value must be identical to the value in the <i>Accept this peer ID</i> field of the phase 1 remote gateway configuration on the <code>branch_1</code> remote peer. See “Configuring branch_1 VPN tunnel settings” on page 1478 .

- 4 Select *Create Phase 2*.

Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. For details on phase 2, see [“Phase 2 parameters” on page 1425](#).

- 5 Enter the following information and select **OK**.

Name	Enter <code>branch_2_phase2</code> . A name to identify this phase 2 configuration.
Phase 1	Select <code>branch_2</code> . The name of the phase 1 configuration that you defined earlier.

Configuring branch_2 security policies

Define security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1432](#).

After defining the two address ranges, select one of [“Creating branch_2 route-based security policies” on page 1475](#) or [“Creating branch_2 policy-based security policies” on page 1477](#) to configure the appropriate VPN policies.

Define address ranges for branch_2 security policies

Define VPN connection names for the address ranges of the private networks. These addresses are used in the security policies that permit communication between the networks. For more information, see [“Defining policy addresses” on page 1431](#).

Define an address name for the IP address and netmask of the private network behind the local FortiGate unit.

To define branch_2 address ranges

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*.

Address Name	Enter <code>branch_2_internal</code> . Enter a meaningful name.
Type	Subnet/IP Range.
Subnet / IP Range	10.10.10.0/24 Include the netmask or specify a specific range.
Interface	internal The interface that will be handling the traffic from the internal network.

Define an address name for the IP address and netmask of the private network behind the remote peer.

- 4 Select *Create New*.
- 5 Enter the following information, and select *OK*.

Address Name	Enter <code>branch_1_internal</code> . A meaningful name for the private network at the remote end of the VPN tunnel.
Type	Subnet/IP Range.
Subnet / IP Range	192.168.1.0/24 Include the netmask. Optionally you can specify a range.
Interface	any The interface that will be handling the remote VPN traffic on this FortiGate unit. If you are unsure, or multiple interfaces may be handling this traffic use <code>any</code> .

Creating branch_2 route-based security policies

Define ACCEPT security policies to permit communication between the branch_2 and branch_1 private networks.

Once the route-based policy is configured a routing entry must be configured to route traffic over the VPN interface.

To create route-based security policies

1 Go to *Policy > Policy > Policy*.

2 Select *Create New*.

Define a policy to permit the branch_2 local FortiGate unit to initiate a VPN session with the branch_1 VPN peer.

3 Enter the following information, and select *OK*.

Source Interface/Zone	Select <code>internal</code> . The interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <code>branch_2_internal</code> . Select the address name for the private network behind this FortiGate unit.
Destination Interface/Zone	Select <code>branch_2</code> . The VPN Tunnel (IPsec Interface).
Destination Address Name	Select <code>branch_1_internal</code> . The address name the private network behind the remote peer.
Action	Select <i>ACCEPT</i> . Accept VPN traffic on this interface pair in this direction.
NAT	Disable.
Comment	route-based: Initiate a branch_2 to branch_1 VPN tunnel

4 Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.

Define a policy to permit the branch_1 remote VPN peer to initiate VPN sessions.

5 Enter the following information, and select *OK*.

Source Interface/Zone	Select <code>branch_2</code> . The VPN Tunnel (IPsec Interface).
Source Address Name	Select <code>branch_1_internal</code> . The address name for the private network behind the remote peer.
Destination Interface/Zone	Select <code>internal</code> . The interface connecting the private network behind this FortiGate unit.
Destination Address Name	Select <code>branch_2_internal</code> . The address name for the private network behind this FortiGate unit.
Action	Select <i>ACCEPT</i> .
NAT	Disable.
Comment	route-based: Initiate a branch_1 to branch_2 internal VPN tunnel.

- 6 Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
- 7 Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

To create routing entry for VPN interface - CLI

```
config router static
edit 5
set dst 0.0.0.0 0.0.0.0
set dynamic-gateway enable
set device wan1
next
end
```



This routing entry must be added in the CLI because the dynamic-gateway option is not available in the web-based manager.

Creating branch_2 policy-based security policies

Define an IPsec policy to permit VPN sessions between the private networks.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.

Define an IPsec policy to permit the VPN sessions between the local branch_2 unit and the remote branch_1 unit.

- 3 Enter the following information, and select *OK*.


Source Interface/Zone	Select <i>internal</i> . The interface connecting the private network behind this FortiGate unit.
Source Address Name	Select <i>branch_2_internal</i> . The address name for the private network behind this local FortiGate unit.
Destination Interface/Zone	Select <i>wan1</i> . The FortiGate unit's public interface.
Destination Address Name	Select <i>branch_1_internal</i> . The address name for the private network behind branch_1, the remote peer.
Action	Select <i>IPSEC</i> .
VPN Tunnel	Select <i>branch_2</i> . The name of the phase 1 tunnel.
Allow Inbound Allow Outbound	Select both <i>Allow inbound</i> and <i>Allow outbound</i> to enable traffic from either direction to initiate the tunnel. If you have problems connecting, try enabling NAT.
Comment	policy-based: allows traffic in either direction to initiate the VPN tunnel.

- 4 Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
- 5 Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

Configure the fixed-address VPN peer

The fixed-address VPN peer, `branch_1`, needs to retrieve the IP address from the dynamic DNS service to initiate communication with the dynamically-addressed peer, `branch_2`. It also depends on the peer ID (local ID) to initiate the VPN tunnel with `branch_2`.

Figure 137: Configure branch_1, the fixed address side



Configuring the fixed-address VPN peer includes:

- [Configuring branch_1 VPN tunnel settings](#)
- [Configuring branch_1 security policies](#)

Configuring branch_1 VPN tunnel settings

Define the phase 1 parameters needed to establish a secure connection with the remote peer. For more information, see [“Auto Key phase 1 parameters” on page 1407](#).

To configure branch_1 phase 1 VPN settings

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create New Phase 1*.
- 3 Enter the following information and select *OK*.

Name	Enter <code>branch_1</code> . A name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select <i>Dynamic DNS</i> . The remote peer this FortiGate is connecting to has a dynamic IP address.

Dynamic DNS	Type the fully qualified domain name of the remote peer (for example, <code>example.com</code>).
Interface	Enter <code>wan1</code> . The public facing interface on the fixed-address FortiGate unit. This interface cannot be a loopback interface.
Mode	Select <i>Aggressive</i> .
Peer Options	Select <i>Accept this peer ID</i> , and enter <code>example.com</code> . This option only appears when the mode is set to Aggressive. The identifier of the FortiGate unit with the dynamic address.
Advanced	
Enable IPsec Interface Mode	Enable for a route-based VPN and when configuring policies, go to “Creating branch_1 route-based security policies” on page 1480 . Disable for a policy-based VPN and when configuring policies, go to “Creating branch_1 policy-based security policies” on page 1481 . If Interface mode is enabled, default settings are used.

- 4 Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Name	Enter <code>branch_1_p2</code> . A name to identify this phase 2 configuration.
Phase 1	Select <code>branch_1</code> . The name of the phase 1 configuration that you defined for the remote peer. You can select the name of the remote gateway from the Dynamic DNS part of the list.

Configuring branch_1 security policies

The `branch_1` FortiGate unit has a fixed IP address and will be connecting to the `branch_2` FortiGate unit that has a dynamic IP address and a domain name of `example.com`.

Remember if you are using route-based security policies that you must add a route for the VPN traffic.

Defining address ranges for branch_1 security policies

As with `branch_2` previously, `branch_1` needs address ranges defined as well. See [“Defining policy addresses” on page 1431](#).

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*.

Address Name	Enter <code>branch_2_internal</code> . A meaningful name for the private network behind the <code>branch_2</code> FortiGate unit.
Type	Subnet/IP Range.

Subnet / IP Range	10.10.10.0/24 Include the netmask or specify a specific range.
Interface	internal This is the interface on this FortiGate unit that will be handling with this traffic.

- 4 Define an address name for the IP address and netmask of the private network behind the remote peer.
- 5 Select *Create New*.
- 6 Enter the following information, and select *OK*.

Address Name	Enter <code>branch_1_internal</code> . A meaningful name for the private network behind the <code>branch_1</code> peer.
Type	Subnet/IP Range.
Subnet / IP Range	192.168.1.0/24 Include the netmask or specify a specific range.
Interface	any The interface on this FortiGate unit that will be handling with this traffic. If you are unsure, or multiple interfaces may be handling this traffic use <code>any</code> .

Creating `branch_1` route-based security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses. See “[Defining VPN security policies](#)” on page 1432.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*.

Source Interface/Zone	Select <code>internal</code> . The interface that connects to the private network behind the <code>branch_1</code> FortiGate unit.
Source Address Name	Select <code>branch_1_internal</code> . The address name that you defined for the private network behind this FortiGate unit.
Destination Interface/Zone	Select <code>branch_1</code> . The VPN Tunnel (IPsec Interface) you configured earlier.
Destination Address Name	Select <code>branch_2_internal</code> . The address name that you defined for the private network behind the <code>branch_2</code> peer.
Action	Select <i>ACCEPT</i> .
NAT	Disable
Comments	Internal -> branch22

To permit the remote client to initiate communication, you need to define a security policy for communication in that direction.

- 4 Select *Create New*.
- 5 Enter the following information, and select *OK*.

Source Interface/Zone	Select <code>branch_1</code> The VPN Tunnel (IPsec Interface) you configured earlier.
Source Address Name	Select <code>branch_2_internal</code> . The address name that you defined for the private network behind the <code>branch_2</code> remote peer.
Destination Interface/Zone	Select <code>internal</code> . The interface that connects to the private network behind this FortiGate unit.
Destination Address Name	Select <code>branch_1_internal</code> . The address name that you defined for the private network behind this FortiGate unit.
Action	Select <i>ACCEPT</i> .
NAT	Disable
Comments	<code>branch_2 -> Internal</code>

Creating `branch_1` policy-based security policies

A policy-based security policy allows you the flexibility to allow inbound or outbound traffic or both through this single policy.

This policy-based IPsec VPN security policy allows both inbound and outbound traffic

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information, and select *OK*.

Source Interface/Zone	Select <code>internal</code> . The interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <code>branch_1_internal</code> . The address name that you defined for the private network behind this FortiGate unit.
Destination Interface/Zone	Select <code>wan1</code> . The FortiGate unit's public interface.
Destination Address Name	Select <code>branch_2_internal</code> . The address name that you defined for the private network behind the remote peer.

Action	Select <i>IPSEC</i> .
VPN Tunnel	<p>Select <i>branch_1</i>. The name of the phase 1 configuration that you created earlier.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

- 4 Place this security policy in the policy list above any other policies having similar source and destination addresses.

Testing

Once both ends are configured, you can test the VPN tunnel.

To test the VPN initiated by branch_2

- 1 On branch_2, go to *VPN > Monitor > IPsec Monitor*.
All IPsec VPN tunnels will be listed on this page, no matter if they are connected or disconnected.
- 2 Select the tunnel listed for branch_2, and select the status column for that entry.
The status will say *Bring Up* and remote port, incoming and outgoing data will all be zero. This indicates an inactive tunnel. When you select *Bring Up*, the FortiGate will try to set up a VPN session over this tunnel. If it is successful, *Bring Up* will change to *Active*, and the arrow icon will change to a green up arrow icon.
- 3 If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting:

To test the VPN initiated by branch_1

- 1 On branch_1, go to *VPN > Monitor > IPsec Monitor*.
- 2 Select the tunnel listed for branch_1, and select the status column.
The difference between branch_2 and branch_1 at this point is that the tunnel entry for branch-1 will not have a remote gateway IP address. It will be resolved when the VPN tunnel is started.
- 3 If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting.

Some troubleshooting ideas include:

- If there was no entry for the tunnel on the monitor page, check the Auto Key (IKE) page to verify the phase 1 and phase 2 entries exist.
- Check the security policy or policies, and ensure there is an outgoing policy as a minimum.
- Check that you entered a local ID in the phase 1 configuration, and that branch_1 has the same local ID.
- Ensure the local DNS server has an up-to-date DNS entry for exmaple.com.

For more information on VPN troubleshooting and testing, see [“VPN troubleshooting tips” on page 1609](#).



FortiClient dialup-client configurations

The FortiClient Endpoint Security application is an IPsec VPN client with antivirus, antispam and firewall capabilities. This section explains how to configure dialup VPN connections between a FortiGate unit and one or more FortiClient Endpoint Security applications.

FortiClient users are usually mobile or remote users who need to connect to a private network behind a FortiGate unit. For example, the users might be employees who connect to the office network while traveling or from their homes.

For greatest ease of use, the FortiClient application can download the VPN settings from the FortiGate unit to configure itself automatically. This section covers both automatic and manual configuration.



The FortiClient configurations in this guide do not apply to the FortiClient Consumer Edition, which does not include the IPsec VPN feature.

The following topics are included in this section:

- [Configuration overview](#)
- [FortiClient-to-FortiGate VPN configuration steps](#)
- [Configure the FortiGate unit](#)
- [Configure the FortiClient Endpoint Security application](#)
- [Adding XAuth authentication](#)
- [FortiClient dialup-client configuration example](#)

Configuration overview


Dialup users typically obtain dynamic IP addresses from an ISP through Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE). Then, the FortiClient Endpoint Security application initiates a connection to a FortiGate dialup server.

By default the FortiClient dialup client has the same IP address as the host PC on which it runs. If the host connects directly to the Internet, this is a public IP address. If the host is behind a NAT device, such as a router, the IP address is a private IP address. The NAT device must be NAT traversal (NAT-T) compatible to pass encrypted packets (see [“NAT traversal” on page 1420](#)). The FortiClient application also can be configured to use a virtual IP address (VIP). For the duration of the connection, the FortiClient application and the FortiGate unit both use the VIP address as the IP address of the FortiClient dialup client.

For a faster and easier method of configuring a FortiGate - to - FortiClient VPN, see [“One button FortiGate - to - FortiClient Phase1 VPN” on page 1485](#).

The FortiClient application sends its encrypted packets to the VPN remote gateway, which is usually the public interface of the FortiGate unit. It also uses this interface to download VPN settings from the FortiGate unit. See [“Automatic configuration of FortiClient dialup clients” on page 1484](#).

Figure 138: Example FortiClient dialup-client configuration



Peer identification

The FortiClient application can establish an IPsec tunnel with a FortiGate unit configured to act as a dialup server. When the FortiGate unit acts as a dialup server, it does not identify the client using the phase 1 remote gateway address. The IPsec tunnel is established if authentication is successful and the IPsec security policy associated with the tunnel permits access. There are several different ways to authenticate dialup clients and restrict access to private networks based on client credentials. For more information, see [“Authenticating remote peers and clients” on page 1412](#).

Automatic configuration of FortiClient dialup clients

The FortiClient application can obtain its VPN settings from the FortiGate VPN server. FortiClient users need to know only the FortiGate VPN server IP address and their user name and password on the FortiGate unit.

The FortiGate unit listens for VPN policy requests from clients on TCP port 8900. When the dialup client connects:

- The client initiates a Secure Sockets Layer (SSL) connection to the FortiGate unit.
- The FortiGate unit requests a user name and password from the FortiClient user. Using these credentials, it authenticates the client and determines which VPN policy applies to the client.
- Provided that authentication is successful, the FortiGate unit downloads a VPN policy to the client over the SSL connection. The information includes IPsec phase 1 and phase 2 settings, and the IP addresses of the private networks that the client is authorized to access.
- The client uses the VPN policy settings to establish an IPsec phase 1 connection and phase 2 tunnel with the FortiGate unit.

One button FortiGate - to - FortiClient Phase1 VPN

On the FortiOS VPN IKE page there is a button to create a Phase1 portion of a VPN tunnel between the FortiGate and FortiClient. Very little information is required for this configuration. No encryption or authentication method is required. This feature is ideal for setting up quick VPN connections with basic settings.



This one button is only compatible with FortiClient 4.3 and higher. Earlier versions of FortiClient need to create IKE Phase-1 object separately, similar to earlier versions of FortiOS.

On the Phase 1 screen (*VPN > IPsec > Phase 1*) is a button called *Create a FortiClient VPN*. This button asks a few basic VPN configuration related questions. Once all the information is added, click *Create Now*. This will create a new dial-up IPsec-interface mode tunnel. Phase 1 and Phase 2 will be added using the default ike settings.

The following Settings will be used when creating a one-button FortiClient VPN Phase1 object:

- Remote Gateway: Dialup User
- Mode: Aggressive
- Enable IPsec Interface Mode
- Default setting for P1 and P2 Proposal
- XAUTH Enable as Server (Auto)
- IKE mode-config will be enabled
- Peer Option set to "Accept any peer ID"
- Rest of the setting use the current defaults (Default value needs to be the same on FCT side)

Once the one button Phase1 is complete, you must create a default Phase2 configuration. This only requires a name for the Phase2 object, and select the one-button Phase1 name.

How the FortiGate unit determines which settings to apply

The FortiGate unit follows these steps to determine the configuration information to send to the FortiClient application:

- 1 Check the virtual domain associated with the connection to determine which VPN policies might apply.
- 2 Select the VPN policy that matches the dialup client's user group and determine which tunnel (phase 1 configuration) is involved.
- 3 Check all IPsec security policies that use the specified tunnel to determine which private networks the dialup clients may access.
- 4 Retrieve the rest of the VPN policy information from the existing IPsec phase 1 and phase 2 parameters in the dialup-client configuration.

Using virtual IP addresses

When the FortiClient host PC is located behind a NAT device, unintended IP address overlap issues may arise between the private networks at the two ends of the tunnel. For example, the client's host might receive a private IP address from a DHCP server on its network that by co-incidence is the same as a private IP address on the network behind the FortiGate unit. A conflict will occur in the host's routing table and the FortiClient Endpoint Security application will be unable to send traffic through the tunnel. Configuring virtual IP (VIP) addresses for FortiClient applications prevents this problem.

Using VIPs ensures that client IP addresses are in a predictable range. You can then define security policies that allow access only to that source address range. If you do not use VIPs, the security policies must allow all source addresses because you cannot predict the IP address for a remote mobile user.

The FortiClient application must not have the same IP address as any host on the private network behind the FortiGate unit or any other connected FortiClient application. You can ensure this by reserving a range of IP addresses on the private network for FortiClient users. Or, you can assign FortiClient VIPs from an uncommonly used subnet such as 10.254.254.0/24 or 192.168.254.0/24.

You can reserve a VIP address for a particular client according to its device MAC address and type of connection. The DHCP server then always assigns the reserved VIP address to the client. For more information about this feature, see the "dhcp reserved-address" section in the "system" chapter of the [FortiGate CLI Reference](#).



On the host computer, you can find out the VIP address that the FortiClient Endpoint Security application is using. For example, On Windows, type `ipconfig /all` at the Windows Command Prompt. On Linux or Mac OS X, type `ifconfig` in a terminal window. The output will also show the IP address that has been assigned to the host Network Interface Card (NIC).

It is best to assign VIPs using DHCP over IPsec. The FortiGate dialup server can act as a DHCP server or relay requests to an external DHCP server. You can also configure VIPs manually on FortiClient applications, but it is more difficult to ensure that all clients use unique addresses.

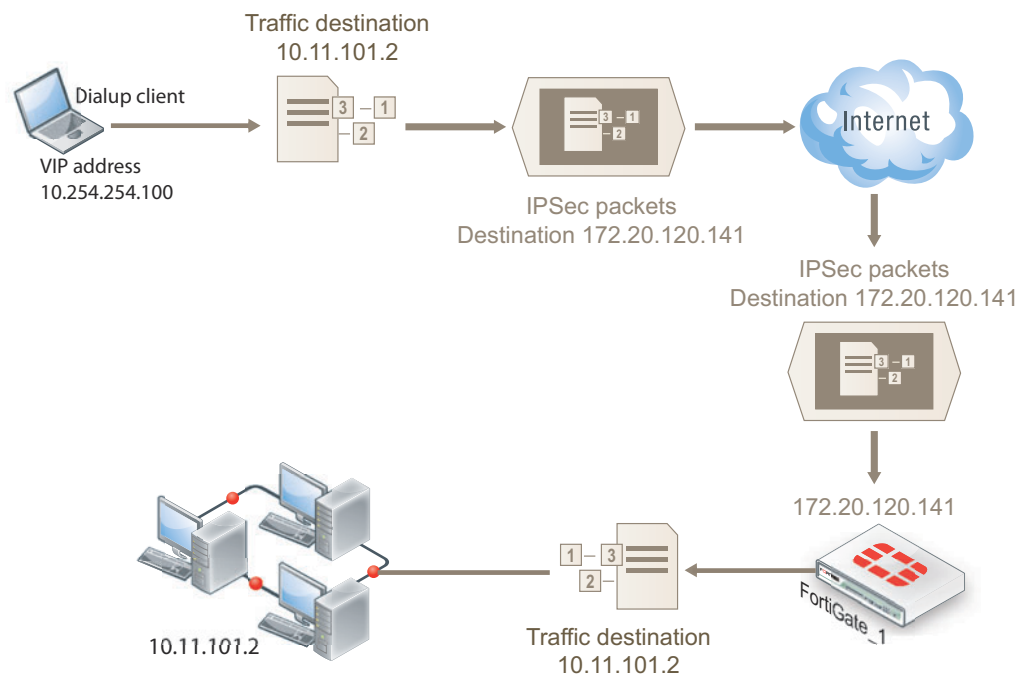


If you assign a VIP on the private network behind the FortiGate unit and enable DHCP-IPsec (a phase 2 advanced option), the FortiGate unit acts as a proxy on the local private network for the FortiClient dialup client. Whenever a host on the network behind the dialup server issues an ARP request for the device MAC address of the FortiClient host, the FortiGate unit answers the ARP request on behalf of the FortiClient host and forwards the associated traffic to the FortiClient host through the tunnel. For more information, see "DHCP-IPsec" on page 1427.



FortiGate units fully support [RFC 3456](#). The FortiGate DHCP over IPsec feature can be enabled to allocate VIP addresses to FortiClient dialup clients using a FortiGate DHCP server.

[Figure 139](#) shows an example of a FortiClient-to-FortiGate VPN where the FortiClient application is assigned a VIP on an uncommonly used subnet. The diagram also shows that while the destination for the information in the encrypted packets is the private network behind the FortiGate unit, the destination of the IPsec packets themselves is the public interface of the FortiGate unit that acts as the end of the VPN tunnel.

Figure 139: IP address assignments in a FortiClient dialup-client configuration

Assigning VIPs by RADIUS user group

If you use XAuth authentication, you can assign users the virtual IP address stored in the Framed-IP-Address field of their record on the RADIUS server. (See [RFC 2865](#) and [RFC 2866](#) for more information about RADIUS fields.) To do this:

- Set the DHCP server *IP Assignment Mode* to *User-group defined method*. This is an Advanced setting. See “[To configure a DHCP server on the FortiGate unit](#)” on [page 1491](#).
- Create a new firewall user group and add the RADIUS server to it.
- In your phase 1 settings, configure the FortiGate unit as an XAuth server and select from *User Group* the new user group that you created. For more information, see “[Using the FortiGate unit as an XAuth server](#)” on [page 1422](#).
- Configure the FortiClient application to use XAuth. See “[Adding XAuth authentication](#)” on [page 1494](#).

FortiClient dialup-client infrastructure requirements

- To support policy-based VPNs, the FortiGate dialup server may operate in either NAT mode or transparent mode. NAT mode is required if you want to create a route-based VPN.
- If the FortiClient dialup clients will be configured to obtain VIP addresses through FortiGate DHCP relay, a DHCP server must be available on the network behind the FortiGate unit and the DHCP server must have a direct route to the FortiGate unit.

- If the FortiGate interface to the private network is not the default gateway, the private network behind the FortiGate unit must be configured to route IP traffic destined for dialup clients back (through an appropriate gateway) to the FortiGate interface to the private network. As an alternative, you can configure the IPsec security policy on the FortiGate unit to perform inbound NAT on IP packets. Inbound NAT translates the source addresses of inbound decrypted packets into the IP address of the FortiGate interface to the local private network.

FortiClient-to-FortiGate VPN configuration steps

Configuring dialup client capability for FortiClient dialup clients involves the following general configuration steps:

- 1 If you will be using VIP addresses to identify dialup clients, determine which VIP addresses to use. As a precaution, consider using VIP addresses that are not commonly used.
- 2 Configure the FortiGate unit to act as a dialup server. See [“Configure the FortiGate unit” on page 1488](#).
- 3 If the dialup clients will be configured to obtain VIP addresses through DHCP over IPsec, configure the FortiGate unit to act as a DHCP server or to relay DHCP requests to an external DHCP server.
- 4 Configure the dialup clients. See [“Configure the FortiClient Endpoint Security application” on page 1493](#).



When a FortiGate unit has been configured to accept connections from FortiClient dialup clients, you can optionally arrange to have an IPsec VPN configuration downloaded to FortiClient dialup clients automatically. For more information, see [“Configuring the FortiGate unit as a VPN policy server” on page 1491](#).

Configure the FortiGate unit

Configuring the FortiGate unit to establish VPN connections with FortiClient Endpoint Security users involves the following steps:

- 1 configure the VPN settings
- 2 if the dialup clients use automatic configuration, configure the FortiGate unit as a VPN policy server
- 3 if the dialup clients obtain VIP addresses by DHCP over IPsec, configure an IPsec DHCP server or relay

The procedures in this section cover basic setup of policy-based and route-based VPNs compatible with FortiClient Endpoint Security. A route-based VPN is simpler to configure.

Configuring FortiGate unit VPN settings

To configure FortiGate unit VPN settings to support FortiClient users, you need to:

- configure the FortiGate Phase 1 VPN settings
- configure the FortiGate Phase 2 VPN settings

- add the security policy
- 1 At the local FortiGate unit, define the phase 1 configuration needed to establish a secure connection with the FortiClient peer. See [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select <i>Dialup User</i> .
Local Interface	Select the interface through which clients connect to the FortiGate unit.
Mode	Select <i>Main (ID Protection)</i> .
Authentication Method	Select <i>Pre-shared Key</i> .
Pre-shared Key	Enter the pre-shared key. This must be the same preshared key provided to the FortiClient users.
Peer option	Select <i>Accept any peer ID</i> .
Enable IPsec Interface Mode	You must select <i>Advanced</i> to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN.

- 2 Define the phase 2 parameters needed to create a VPN tunnel with the FortiClient peer. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Name	Enter a name to identify this phase 2 configuration.
Phase 1	Select the name of the phase 1 configuration that you defined.
Advanced	Select to configure the following optional setting.
DHCP-IPsec	Select if you provide virtual IP addresses to clients using DHCP.

- 3 Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks. For more information, see [“Defining policy addresses” on page 1431](#).

Enter these settings in particular:

- Define an address name for the individual address or the subnet address that the dialup users access through the VPN.
 - If FortiClient users are assigned VIP addresses, define an address name for the subnet to which these VIPs belong.
- 4 Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1432](#).



If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPSec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

Route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses. Enter these settings in particular:

Source Interface/Zone	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Disable.

If you want to allow hosts on the private network to initiate communications with the FortiClient users after the tunnel is established, you need to define a security policy for communication in that direction. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Disable.

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
Destination Interface/Zone	Select the FortiGate unit's public interface.
Destination Address Name	If FortiClient users are assigned VIPs, select the address name that you defined in Step 3 for the VIP subnet. Otherwise, select <i>All</i> .
Action	Select <i>IPSEC</i> .
VPN Tunnel	<p>Select the name of the phase 1 configuration that you created in Step 1.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> if you want to allow hosts on the private network to initiate communications with the FortiClient users after the tunnel is established.</p>

Place VPN policies in the policy list above any other policies having similar source and destination addresses.

Configuring the FortiGate unit as a VPN policy server

When a FortiClient application set to automatic configuration connects to the FortiGate unit, the FortiGate unit requests a user name and password. If the user supplies valid credentials, the FortiGate unit downloads the VPN settings to the FortiClient application.

You must do the following to configure the FortiGate unit to work as a VPN policy server for FortiClient automatic configuration:

- 1 Create user accounts for FortiClient users.
- 2 Create a user group for FortiClient users and the user accounts that you created in step 1.
- 3 Connect to the FortiGate unit CLI and configure VPN policy distribution as follows:

```
config vpn ipsec forticlient
  edit <policy_name>
    set phase2name <tunnel_name>
    set usergroupname <group_name>
    set status enable
  end
```

<tunnel_name> must be the Name you specified in the step 2 of “[Configure the FortiGate unit](#)” on page 1488. <group_name> must be the name of the user group you created for FortiClient users.

Configuring DHCP service on the FortiGate unit

If the FortiClient dialup clients are configured to obtain a VIP address using DHCP, configure the FortiGate dialup server to either:

- relay DHCP requests to a DHCP server behind the FortiGate unit (see “[To configure DHCP relay on the FortiGate unit](#)” below).
- act as a DHCP server (see “[To configure a DHCP server on the FortiGate unit](#)” on page 1491).

To configure DHCP relay on the FortiGate unit

- 1 Go to *System > Network > DHCP Server* and select *Create New*.
- 2 In *Interface Name*, select the interface that connects to the Internet (for example, external or wan1).
- 3 In *Mode*, select *Relay*.
- 4 In *Type* select *IPsec*.
- 5 In the *DHCP Server IP* field, type the IP address of the DHCP server.
- 6 Select *OK*.
- 7 If a router is installed between the FortiGate unit and the DHCP server, define a static route to the DHCP server.

To configure a DHCP server on the FortiGate unit

- 1 Go to *System > DHCP Server* and select *Create New*.
- 2 In *Interface Name*, select the interface that connects to the Internet (for example, external or wan1).
- 3 In *Mode*, select *Server*.

- 4 Select *Enable*.
- 5 Enter the following information and select *OK*:

Type	IPsec
IP Range	<p>Enter the range of VIP addresses that the DHCP server can dynamically assign to dialup clients when they connect. As a precaution, do not assign VIP addresses that match the private network behind the FortiGate unit.</p> <p>If you need to exclude specific IP addresses from the range, you can define an exclusion range (see Advanced... below).</p> <p>Note: If you will use a RADIUS server to assign VIP addresses, these fields are not needed.</p>
Network Mask	Enter the network mask of the IP addresses that you specified in the IP Range fields (for example, 255.255.255.0 for a class C network).
Default Gateway	Enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
DNS Service	<p>Select <i>Use System DNS Setting</i>.</p> <p>If you want to use a different DNS server for VPN clients, select <i>Specify</i> and enter an IP address in <i>DNS Server 0</i>.</p>
Advanced...	Select <i>Advanced</i> to configure any of the following options.
Domain	If you want the FortiGate unit to assign a domain name to dialup clients when they connect, enter the registered domain name.
Lease Time	<p>Specify a lease time:</p> <ul style="list-style-type: none"> Select <i>Unlimited</i> to allow the dialup client to use the assigned IP address for an unlimited amount of time (that is, until the client disconnects). Enter the amount of time (in days, hours, and minutes) that the dialup client may use the assigned IP address, after which the dialup client must request new settings from the DHCP server. The range is from 5 minutes to 100 days.
IP Assignment Mode	<p><i>Server IP Range</i> — assign addresses from <i>IP Range</i> (default)</p> <p><i>User-group defined method</i> — assign addresses from user's record on RADIUS server. See "Assigning VIPs by RADIUS user group" on page 1487.</p>
WINS Server 0 WINS Server 1	Optionally, enter the IP addresses of one or two Windows Internet Service (WINS) servers that dialup clients can access after the tunnel has been established.
Options	Optionally, you can send up to three DHCP options to the dialup client. Select <i>Options</i> and enter the option code in the <i>Code</i> field, and if applicable, type any associated data in the <i>Options</i> field. For more information, see RFC 2132 .
Exclude Ranges	To specify any VIP addresses that must be excluded from the VIP address range, select <i>Exclude Ranges</i> , select the + button and then type the starting and ending IP addresses. You can add multiple ranges to exclude.

Configure the FortiClient Endpoint Security application

The following procedure explains how to configure the FortiClient Endpoint Security application to communicate with a remote FortiGate dialup server using the VIP address that you specify manually.

Configuring FortiClient to work with VPN policy distribution

If the remote FortiGate gateway is configured as a VPN policy server, you can configure the FortiClient software to download the VPN settings from the FortiGate gateway.

For VPNs with automatic configuration, only preshared keys are supported. Certificates are not supported.

To add a VPN with automatic configuration on the FortiClient PC

- 1 Go to *VPN > Connections*.
- 2 Select *Advanced* and then select *Add*.
- 3 In the *New Connection* dialog box, enter a connection name.
- 4 For *Configuration*, select *Automatic*.
- 5 For *Policy Server*, enter the IP address or FQDN of the FortiGate gateway.
- 6 Select *OK*.

Configuring FortiClient manually

This procedure explains how to configure the FortiClient application manually using the default IKE and IPsec settings. For more information, refer to the [FortiClient Endpoint Security User Guide](#).

This procedure includes instructions for configuring a virtual IP for the FortiClient application, either manually or using DHCP over IPsec.

To create a FortiClient VPN configuration

- 1 Go to *VPN > Connections*.
- 2 Select *Advanced* and then select *Add*.
- 3 Enter the following information:

Connection Name	Enter a descriptive name for the connection.
Configuration	Select <i>Manual</i>
Remote Gateway	Enter the IP address or the fully qualified domain name (FQDN) of the remote gateway.
Remote Network	Enter the IP address and netmask of the network behind the FortiGate unit.
Authentication Method	Select <i>Pre-shared Key</i> .
Pre-shared Key	Enter the pre-shared key.

- 4 Follow the remaining steps only if you want to configure a VIP. Otherwise, select *OK*.
- 5 Select *Advanced*.
- 6 Enable *Acquire a virtual IP address* and then select the adjacent *Config* button.

- 7 Enter the following information and select **OK**.

Options	Select one of these options:
DHCP	Obtain virtual IP address from the FortiGate unit using DHCP over IPsec.
Manually Set	Assign the virtual IP address manually using the settings in the <i>Manual VIP</i> section.
Manual VIP	These settings are available only if you select <i>Manually Set</i> in the <i>Options</i> section.
IP	Enter the IP address that the FortiClient dialup client uses. This address must not conflict with any IP address at either end of the VPN tunnel.
Subnet Mask	Enter the subnet for the private network.
DNS Server WINS Server	Optionally, enter the addresses of the DNS and WINS servers that the FortiClient user can access through the VPN.

- 8 Select **OK** twice to close the dialog boxes.
9 Repeat this procedure for each FortiClient dialup client.

Adding XAuth authentication

Extended Authentication (XAuth) increases security by requiring additional user authentication in a separate exchange at the end of the VPN phase 1 negotiation. The FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.

Implementation of XAuth requires configuration at both the FortiGate unit and the FortiClient application. For information about configuring a FortiGate unit as an XAuth server, see "[Using the FortiGate unit as an XAuth server](#)" on page 1422. The following procedure explains how to configure the FortiClient application.



XAuth is not compatible with IKE version 2.

To configure the FortiClient Endpoint Security application

In the FortiClient Endpoint Security application, make the following changes to the VPN configuration to enable XAuth authentication to the FortiGate unit.


- 1 Go to *VPN > Connections*, select the VPN connection you want to modify.
- 2 Select *Advanced > Edit > Advanced*.
- 3 Select the *eXtended Authentication* check box and then select the *Config* button to the right of it.

- 4 In the *Extended Authentication (XAuth)* dialog box, select one of the 2 following choices:
 - To always prompt for username and password, select *Prompt to login*.
This is the default.
 - To never prompt for username and password, clear the *Prompt to login* check box and fill in the *User Name* and *Password* fields.
The FortiClient Endpoint Security application will use these values and automatically responds to the XAuth challenge.
- 5 Select *OK* to close all dialog boxes.

FortiClient dialup-client configuration example

This example demonstrates how to set up a FortiClient dialup-client IPsec VPN that uses preshared keys for authentication purposes. In the example configuration, the DHCP over IPsec feature is enabled in the FortiClient Endpoint Security application so that the FortiClient Endpoint Security application can acquire a VIP address through the FortiGate DHCP server. Both route-based and policy-based solutions are covered.

Figure 140: Example FortiClient dialup-client configuration



In the example configuration:

- VIP addresses that are not commonly used (in this case, 10.254.254.0/24) are assigned to the FortiClient dialup clients using a DHCP server.
- The dialup clients have access to the LAN behind FortiGate_1.
- The other network devices are assigned IP addresses as shown in [Figure 140](#).

Configuring FortiGate_1

When a FortiGate unit receives a connection request from a dialup client, it uses IPsec phase 1 parameters to establish a secure connection and authenticate the client. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec phase 2 parameters and applies the IPsec security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

To support these functions, the following general configuration steps must be performed at the FortiGate unit:

- Define the phase 1 parameters that the FortiGate unit needs to authenticate the dialup clients and establish a secure connection. See [“To define the phase 1 parameters” on page 1496](#).
- Define the phase 2 parameters that the FortiGate unit needs to create a VPN tunnel and enable all dialup clients having VIP addresses on the 10.254.254.0/24 network to connect using the same tunnel definition. See [“To define the phase 2 parameters” on page 1496](#).
- Create security policy to control the permitted services and permitted direction of traffic between the IP source address and the dialup clients. See [“To define the firewall addresses” on page 1496](#).
- Configure the FortiGate unit to service DHCP requests from dialup clients. See [“To configure a DHCP server on the FortiGate unit” on page 1498](#).

To define the phase 1 parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	todialups
Remote Gateway	Dialup User
Local Interface	Port 1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	hardtoguess
Peer Options	Accept any peer ID
Advanced	Select
Enable IPsec Interface Mode	Enable for route-based VPN. Disable for policy-based VPN.

To define the phase 2 parameters

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 2*.
- 2 Select *Advanced*, enter the following information, and select *OK*:

Name	td_2
Phase 1	todialups
Advanced	DHCP-IPsec

To define the firewall addresses

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	internal_net
Subnet/IP Range	10.11.101.0/24
Interface	Port 2

- 3 Select *Create New*, enter the following information, and select *OK*:

Address Name	dialups
Subnet/IP Range	10.254.254.[1-10]
Interface	Route-based VPN: todialups Policy-based VPN: Any

The security policies for route-based and policy-based VPNs are described in separate sections below.

To define security policies - route-based VPN

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	todialups
Source Address Name	dialups
Destination Interface/Zone	Port 2
Destination Address Name	internal_net
Schedule	As required.
Service	As required.
Action	ACCEPT
NAT	Disable

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Port 2
Source Address Name	internal_net
Destination Interface/Zone	todialups
Destination Address Name	dialups
Schedule	As required.
Service	As required.
Action	ACCEPT
NAT	Disable

- 4 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Port 2
Source Address Name	internal_net
Destination Interface/Zone	todialups
Destination Address Name	all
Schedule	As required.
Service	DHCP

Action	ACCEPT
NAT	Disable

- Place these policies in the policy list above any other policies having similar source and destination addresses.



The policy in step 4 is required for DHCP to function properly for policy-based VPNs. You can omit this policy if you change the *Destination Address Name* to `all` in the step before. Route-based policies are not affected by this.

To define the security policy - policy-based VPN

- Go to *Policy > Policy > Policy*.
- Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Port 2
Source Address Name	internal_net
Destination Interface/Zone	Port 1
Destination Address Name	dialups
Schedule	As required.
Service	As required.
Action	IPSEC
VPN Tunnel	todialups.
Allow Inbound	Enable
Allow Outbound	Enable if you want to allow hosts on the private network behind the FortiGate unit to initiate communications with the FortiClient users after the tunnel is established.
Inbound NAT	Disable
Outbound NAT	Disable

- Place the policy in the policy list above any other policies having similar source and destination addresses.

To configure a DHCP server on the FortiGate unit

- Go to *System > DHCP Server* and select *Create New*.
- Enter the following information and select *OK*:

Interface Name	Route-based VPN: select virtual IPsec interface. For example, todialups. Policy-based VPN: select the public interface. For example, Port 1.
Mode	Server
Type	IPSEC.
IP Range	10.254.254.1 - 10.254.254.10

Network Mask	255.255.255.0
Default Gateway	172.20.120.2

Configuring the FortiClient Endpoint Security application

The following procedure explains how to configure the FortiClient Endpoint Security application to connect to FortiGate_1 and broadcast a DHCP request. The dialup client uses the VIP address acquired through FortiGate DHCP relay as its IP source address for the duration of the connection.

To configure FortiClient

- 1 At the remote host, start FortiClient.
- 2 Go to *VPN > Connections* and select *Advanced > Add*.
- 3 Enter the following settings:

Connection Name	Office
VPN Type	Manual IPsec
Remote Gateway	172.20.120.141
Remote Network	10.11.101.0 / 255.255.255.0
Authentication Method	Preshared Key
Preshared Key	hardtoguess

- 4 Select *Advanced*.
- 5 In the *Advanced Settings* dialog box, select *Acquire virtual IP address* and then select *Config*.
- 6 Verify that the *Dynamic Host Configuration Protocol (DHCP) over IPsec* option is selected, and then select *OK*.
- 7 Select *OK* twice to close the dialog boxes.
- 8 Exit FortiClient and repeat this procedure at all other remote hosts.



FortiGate dialup-client configurations

This section explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit having a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.


The following topics are included in this section:

- [Configuration overview](#)
- [FortiGate dialup-client configuration steps](#)
- [Configure the server to accept FortiGate dialup-client connections](#)
- [Configure the FortiGate dialup client](#)

Configuration overview

A dialup client can be a FortiGate unit—the FortiGate dialup client typically obtains a dynamic IP address from an ISP through the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) before initiating a connection to a FortiGate dialup server.

Figure 141: Example FortiGate dialup-client configuration



In a dialup-client configuration, the FortiGate dialup server does not rely on a phase 1 remote gateway address to establish an IPsec VPN connection with dialup clients. As long as authentication is successful and the IPsec security policy associated with the tunnel permits access, the tunnel is established.

Several different ways to authenticate dialup clients and restrict access to private networks based on client credentials are available. To authenticate FortiGate dialup clients and help to distinguish them from FortiClient dialup clients when multiple clients will be connecting to the VPN through the same tunnel, best practices dictate that you assign a unique identifier (local ID or peer ID) to each FortiGate dialup client. For more information, see [“Authenticating remote peers and clients” on page 1412](#).



Whenever you add a unique identifier (local ID) to a FortiGate dialup client for identification purposes, you must select Aggressive mode on the FortiGate dialup server and also specify the identifier as a peer ID on the FortiGate dialup server. For more information, see [“Enabling VPN access with user accounts and pre-shared keys” on page 1415](#).

Users behind the FortiGate dialup server cannot initiate the tunnel because the FortiGate dialup client does not have a static IP address. After the tunnel is initiated by users behind the FortiGate dialup client, traffic from the private network behind the FortiGate dialup server can be sent to the private network behind the FortiGate dialup client.

Encrypted packets from the FortiGate dialup client are addressed to the public interface of the dialup server. Encrypted packets from the dialup server are addressed either to the public IP address of the FortiGate dialup client (if the dialup client connects to the Internet directly), or if the FortiGate dialup client is behind a NAT device, encrypted packets from the dialup server are addressed to the public IP address of the NAT device.



If a router with NAT capabilities is in front of the FortiGate dialup client, the router must be NAT-T compatible for encrypted traffic to pass through the NAT device. For more information, see [“NAT traversal” on page 1420](#).

When the FortiGate dialup server decrypts a packet from the FortiGate dialup client, the source address in the IP header may be one of the following values, depending on the configuration of the network at the far end of the tunnel:

- If the FortiGate dialup client connects to the Internet directly, the source address will be the private IP address of a host or server on the network behind the FortiGate dialup client.
- If the FortiGate dialup client is behind a NAT device, the source address will be the public IP address of the NAT device.


In some cases, computers on the private network behind the FortiGate dialup client may (by co-incidence) have IP addresses that are already used by computers on the network behind the FortiGate dialup server. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent.

In many cases, computers on the private network behind the FortiGate dialup client will most likely obtain IP addresses from a local DHCP server behind the FortiGate dialup client. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and IP-address overlap issues may arise.

To avoid these issues, you can configure FortiGate DHCP relay on the dialup client instead of using a DHCP server on the network behind the dialup client. The FortiGate dialup client can be configured to relay DHCP requests from the local private network to a DHCP server that resides on the network behind the FortiGate dialup server (see [Figure 142 on page 1503](#)). You configure the FortiGate dialup client to pass traffic from the local private network to the remote network by enabling FortiGate DHCP relay on the FortiGate dialup client interface that is connected to the local private network.

Afterward, when a computer on the network behind the dialup client broadcasts a DHCP request, the dialup client relays the message through the tunnel to the remote DHCP server. The remote DHCP server responds with a private IP address for the computer. To avoid ambiguous routing and network overlap issues, the IP addresses assigned to computers behind the dialup client cannot match the network address space used by the private network behind the FortiGate dialup server.

Figure 142: Preventing network overlap in a FortiGate dialup-client configuration



When the DHCP server resides on the private network behind the FortiGate dialup server, the IP destination address specified in the IPsec security policy on the FortiGate dialup client must refer to that network.



You must add a static route to the DHCP server FortiGate unit if it is not directly connected to the private network behind the FortiGate dialup server—its IP address does not match the IP address of the private network. Also, the destination address in the IPsec security policy on the FortiGate dialup client must refer to the DHCP server address. The DHCP server must be configured to assign a range of IP addresses different from the DHCP server's local network, and also different from the private network addresses behind the FortiGate dialup server. See [“Routing” on page 1469](#).

FortiGate dialup-client infrastructure requirements

The requirements are:

- The FortiGate dialup server must have a static public IP address.
- NAT mode is required if you want to create a route-based VPN.
- The FortiGate dialup server may operate in either NAT mode or transparent mode to support a policy-based VPN.

- Computers on the private network behind the FortiGate dialup client can obtain IP addresses either from a DHCP server behind the FortiGate dialup client, or a DHCP server behind the FortiGate dialup server.
- If the DHCP server resides on the network behind the dialup client, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup server.
- If the DHCP server resides on the network behind the FortiGate dialup server, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup client.

FortiGate dialup-client configuration steps

The procedures in this section assume that computers on the private network behind the FortiGate dialup client obtain IP addresses from a local DHCP server. The assigned IP addresses do not match the private network behind the FortiGate dialup server.



In situations where IP-address overlap between the local and remote private networks is likely to occur, FortiGate DHCP relay can be configured on the FortiGate dialup client to relay DHCP requests to a DHCP server behind the FortiGate dialup server. For more information, see [“To configure DHCP relay on the FortiGate unit” on page 1491](#).

Configuring dialup client capability for FortiGate dialup clients involves the following general configuration steps:

- Determine which IP addresses to assign to the private network behind the FortiGate dialup client, and add the IP addresses to the DHCP server behind the FortiGate dialup client. Refer to the software supplier’s documentation to configure the DHCP server.
- Configure the FortiGate dialup server. See [“Configure the server to accept FortiGate dialup-client connections” on page 1504](#).
- Configure the FortiGate dialup client. See [“Configure the FortiGate dialup client” on page 1506](#).

Configure the server to accept FortiGate dialup-client connections

Before you begin, optionally reserve a unique identifier (peer ID) for the FortiGate dialup client. The dialup client will supply this value to the FortiGate dialup server for authentication purposes during the IPsec phase 1 exchange. In addition, the value will enable you to distinguish FortiGate dialup-client connections from FortiClient dialup-client connections. The same value must be specified on the dialup server and on the dialup client.

- 1 At the FortiGate dialup server, define the phase 1 parameters needed to authenticate the FortiGate dialup client and establish a secure connection. See [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel. This name appears in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select <i>Dialup User</i> .
Local Interface	Select the interface through which clients connect to the FortiGate unit.

Mode	If you will be assigning an ID to the FortiGate dialup client, select <i>Aggressive</i> .
Peer Options	If you will be assigning an ID to the FortiGate dialup client, select <i>Accept this peer ID</i> and type the identifier that you reserved for the FortiGate dialup client into the adjacent field.
Enable IPsec Interface Mode	You must select <i>Advanced</i> to see this setting. If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN. After you select <i>OK</i> to create the phase 1 configuration, you cannot change this setting.

- 2 Define the phase 2 parameters needed to create a VPN tunnel with the FortiGate dialup client. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Name	Enter a name to identify this phase 2 configuration.
Phase 1	Select the name of the phase 1 configuration that you defined.

- 3 Define names for the addresses or address ranges of the private networks that the VPN links. See [“Defining policy addresses” on page 1431](#). Enter these settings in particular:
- Define an address name for the server, host, or network behind the FortiGate dialup server.
 - Define an address name for the private network behind the FortiGate dialup client.
- 4 Define the security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1432](#).

Route-based VPN security policy

Define an *ACCEPT* security policy to permit communications between hosts on the private network behind the FortiGate dialup client and the private network behind this FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy.

Enter these settings in particular:

Source Interface/Zone	Select the VPN tunnel (IPsec interface) created in Step 1.
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Disable

Policy-based VPN security policy

Define an IPsec security policy. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
Destination Interface/Zone	Select the FortiGate unit's public interface.
Destination Address Name	Select the address name that you defined in Step 3.
Action	Select <i>IPSEC</i> .
VPN Tunnel	<p>Select the name of the phase 1 configuration that you created in Step 1.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Clear <i>Allow outbound</i> to prevent traffic from the local network from initiating the tunnel after the tunnel has been established.</p>

- 1 Place the policy in the policy list above any other policies having similar source and destination addresses.
- 2 If configuring a route-based policy, configure a default route for VPN traffic on this interface.

Configure the FortiGate dialup client

Configure the FortiGate dialup client as follows:

- 1 At the FortiGate dialup client, define the phase 1 parameters needed to authenticate the dialup server and establish a secure connection. See [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel.
Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Type the IP address of the dialup server's public interface.
Local Interface	Select the interface that connects to the public network.
Mode	The FortiGate dialup client has a dynamic IP address, select <i>Aggressive</i> .
Advanced	Select to view the following options.
Local ID	If you defined a peer ID for the dialup client in the FortiGate dialup server configuration, enter the identifier of the dialup client. The value must be identical to the peer ID that you specified previously in the FortiGate dialup server configuration.

Enable IPsec Interface Mode	<p>If <i>IPsec Interface Mode</i> is enabled, the FortiGate unit creates a virtual IPsec interface for a route-based VPN. Disable this option if you want to create a policy-based VPN.</p> <p>After you select OK to create the phase 1 configuration, you cannot change this setting.</p>
------------------------------------	---

- 2 Define the phase 2 parameters needed to create a VPN tunnel with the dialup server. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Name	Enter a name to identify this phase 2 configuration.
Phase 1	Select the set of phase 1 parameters that you defined in step 1.

- 3 Define names for the addresses or address ranges of the private networks that the VPN links. See [“Defining policy addresses” on page 1431](#). Enter these settings in particular:
- Define an address name for the server, host, or network behind the FortiGate dialup server.
 - Define an address name for the private network behind the FortiGate dialup client.
- 4 Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [“Defining VPN security policies” on page 1432](#).

Route-based VPN security policy

Define an ACCEPT security policy to permit communications between hosts on the private network behind this FortiGate dialup client and the private network behind the FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select <i>All</i> .
Destination Interface/Zone	Select the VPN tunnel (IPsec interface) created in Step 1.
Destination Address Name	Select <i>All</i> .
Action	Select <i>ACCEPT</i> .
NAT	Disable

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses. Enter these settings in particular:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	Select the address name that you defined in Step 3 for the private network behind this FortiGate unit.
Destination Interface/Zone	Select the FortiGate unit's public interface.

Destination Address Name	Select the address name that you defined in Step 3 for the private network behind the dialup server.
Action	Select <i>IPSEC</i> .
VPN Tunnel	<p>Select the name of the phase 1 configuration that you created in Step 1.</p> <p>Clear <i>Allow inbound</i> to prevent traffic from the remote network from initiating the tunnel after the tunnel has been established.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

Place the policy in the policy list above any other policies having similar source and destination addresses.



Supporting IKE Mode config clients

IKE Mode Config is an alternative to DHCP over IPsec. A FortiGate unit can be configured as either an IKE Mode Config server or client. This chapter contains the following sections:

- [Automatic configuration overview](#)
- [IKE Mode Config overview](#)
- [Configuring IKE Mode Config](#)
- [Example: FortiGate unit as IKE Mode Config server](#)
- [Example: FortiGate unit as IKE Mode Config client](#)

Automatic configuration overview

VPN configuration for remote clients is simpler if it is automated. Several protocols support automatic configuration:

- The Fortinet FortiClient Endpoint Security application can completely configure a VPN connection with a suitably configured FortiGate unit given only the FortiGate unit's address. This protocol is exclusive to Fortinet. For more information, see the [“FortiClient dialup-client configurations”](#) chapter.
- DHCP over IPsec can assign an IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms.
- IKE Mode Config can configure host IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms. Several network equipment vendors support IKE Mode Config, which is described in the ISAKMP Configuration Method document [draft-dukes-ike-mode-cfg-02.txt](#).

This chapter describes how to configure a FortiGate unit as either an IKE Mode Config server or client.

IKE Mode Config overview

Dialup VPN clients connect to a FortiGate unit that acts as a VPN server, providing the client the necessary configuration information to establish a VPN tunnel. The configuration information typically includes a virtual IP address, netmask, and DNS server address.

IKE Mode Config is available only for VPNs that are route-based, also known as interface-based. A FortiGate unit can function as either an IKE Configuration Method server or client. IKE Mode Config is configurable only in the CLI.

Configuring IKE Mode Config

IKE Mode Config is configured with the CLI command `vpn ipsec phase1-interface`. The `mode-cfg` variable enables IKE Mode Config. The `type` field determines whether you are creating an IKE Mode Config server or a client. Setting `type` to `dynamic` creates a server configuration, otherwise the configuration is a client.

Configuring an IKE Mode Config client

If the FortiGate unit will connect as a dialup client to a remote gateway that supports IKE Mode Config, the relevant `vpn ipsec phase1-interface` variables are as follows:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is not compatible with IKE v2.
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type {ddns static}</code>	If you set <code>type</code> to <code>dynamic</code> , an IKE Mode Config server is created.
<code>assign-ip {enable disable}</code>	Enable to request an IP address from the server.
<code>interface <interface_name></code>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.
<code>proposal <encryption_combination></code>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the client will accept. For more information, see “Defining IKE negotiation parameters” on page 1417 .
<code>mode-cfg-ip-version {4 6}</code>	Select if the Method client receives an IPv4 or IPv6 IP address. The default is 4. the <code>ip-version</code> setting matches this variable’s value.
<code>ip-version <4 6></code>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to 6 to create a VPN with IPv6 addressing.

Configuring an IKE Mode Config server

If the FortiGate unit will accept connection requests from dialup clients that support IKE Mode Config, the following `vpn ipsec phase1-interface` settings are required before any other configuration is attempted:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is not compatible with IKE v2.
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type dynamic</code>	Any other setting creates an IKE Mode Config client.

Variable	Description
interface <interface_name>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.
proposal <encryption_combination>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the server will accept. For more information, see “Defining IKE negotiation parameters” on page 1417 .
ip-version <4 6>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to <code>6</code> to create a VPN with IPv6 addressing.

After you have enabled the basic configuration, you can configure:

- IP address assignment for clients
- DNS and WINS server assignment

IP address assignment

Usually you will want to assign IP addresses to clients. The simplest method is to assign addresses from a specific range, similar to a DHCP server.

If your clients are authenticated by a RADIUS server, you can obtain the user's IP address assignment from the Framed-IP-Address attribute. The user must be authenticated using XAuth.

To assign IP addresses from an address range

If your VPN uses IPv4 addresses,

```
config vpn ipsec phase1-interface
edit vpn1
set mode-cfg-ipversion 4
set assign-ip enable
set assign-ip-type ip
set assign-ip-from range
set ipv4-start-ip <range_start>
set ipv4-end-ip <range_end>
set ipv4-netmask <netmask>
end
```

If your VPN uses IPv6 addresses,

```
config vpn ipsec phase1-interface
edit vpn1
set mode-cfg-ipversion 6
set assign-ip enable
set assign-ip-type ip
set assign-ip-from range
set ipv6-start-ip <range_start>
set ipv6-end-ip <range_end>
end
```

To assign IP addresses from a RADIUS server

The users must be authenticated by a RADIUS server and assigned to the FortiGate user group <grpname>. Since the IP address will not be static, `type` is set to `dynamic`, and `mode-cfg` is enabled. This is IKE Configuration Method so that compatible clients can configure themselves with settings that the FortiGate unit provides.

```
config vpn ipsec phase1-interface
  edit vpn1
    set type dynamic
    set mode-cfg enable
    set assign-ip enable
    set assign-ip-from usrgrp
    set xauthtype auto
    set authusrgrp <grpname>
  end
```

Example: FortiGate unit as IKE Mode Config server

In this example, the FortiGate unit assigns IKE Mode Config clients addresses in the range of 10.11.101.160 through 10.11.101.180. DNS and WINS server addresses are also provided. The public interface of the FortiGate unit is Port 1.

The `ipv4-split-include` variable specifies a firewall address that represents the networks to which the clients will have access. This destination IP address information is sent to the clients.


Only the CLI fields required for IKE Mode Config are shown here. For detailed information about these variables, see the [FortiGate CLI Reference](#).

```
config vpn ipsec phase1-interface
  edit vpn1
    set ip-version 4
    set type dynamic
    set interface port1
    set proposal 3des-sha1 aes128-sha1
    set mode-cfg enable
    set mode-cfg-ipversion 4
    set assign-ip enable
    set assign-ip-type ip
    set assign-ip-from range
    set ipv4-start-ip 10.11.101.160
    set ipv4-end-ip 10.11.101.180
    set ipv4-netmask 255.255.255.0
    set dns-server1 10.11.101.199
    set dns-server2 66.11.168.195
    set wins-server1 10.11.101.191
    set domain example
    set ipv4-split-include OfficeLAN
  end
```

Example: FortiGate unit as IKE Mode Config client

In this example, the FortiGate unit connects to a VPN gateway with a static IP address that can be reached through Port 1. Only the port, gateway and proposal information needs to be configured. All other configuration information will come from the IKE Mode Config server.

```
config vpn ipsec phase1-interface
edit vpn1
set ip-version 4
set type static
set remote-gw <gw_address>
set interface port 1
set proposal 3des-sha1 aes128-sha1
set mode-cfg enable
set mode-cfg-ipversion 4
set assign-ip enable
end
```

Internet-browsing configuration

This section explains how to support secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.

The following topics are included in this section:


- [Configuration overview](#)
- [Creating an Internet browsing security policy](#)
- [Routing all remote traffic through the VPN tunnel](#)

Configuration overview

A VPN provides secure access to a private network behind the FortiGate unit. You can also enable VPN clients to access the Internet securely. The FortiGate unit inspects and processes all traffic between the VPN clients and hosts on the Internet according to the Internet browsing policy. This is accomplished even though the same FortiGate interface is used for both encrypted VPN client traffic and unencrypted Internet traffic.

In [Figure 143](#), FortiGate_1 enables secure Internet browsing for FortiClient Endpoint Security users such as Dialup_1 and users on the Site_2 network behind FortiGate_2, which could be a VPN peer or a dialup client.

Figure 143: Example Internet-browsing configuration



You can adapt any of the following configurations to provide secure Internet browsing:

- a gateway-to-gateway configuration (see [“Gateway-to-gateway configurations”](#) on page 1437)
- a FortiClient dialup-client configuration (see [“FortiClient dialup-client configurations”](#) on page 1483)

- a FortiGate dialup-client configuration (see [“FortiGate dialup-client configurations” on page 1501](#))

The procedures in this section assume that one of these configurations is in place, and that it is operating properly.

To create an internet-browsing configuration based on an existing gateway-to-gateway configuration, you must edit the gateway-to-gateway configuration as follows:

- On the FortiGate unit that will provide Internet access, create an Internet browsing security policy. See [“Creating an Internet browsing security policy”](#), below.
- Configure the remote peer or client to route all traffic through the VPN tunnel. You can do this on a FortiGate unit or on a FortiClient Endpoint Security application. See [“Routing all remote traffic through the VPN tunnel” on page 1517](#).

Creating an Internet browsing security policy

On the FortiGate unit that acts as a VPN server and will provide secure access to the Internet, you must create an Internet browsing security policy. This policy differs depending on whether your gateway-to-gateway configuration is policy-based or route-based.

To create an Internet browsing policy - policy-based VPN

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information and then select *OK*:

Source Interface	The interface to which the VPN tunnel is bound.
Source Address Name	All
Destination Interface	The interface to which the VPN tunnel is bound.
Destination Address Name	The internal range of address of the remote spoke site.
Schedule	As required.
Service	As required.
Action	IPSEC
VPN Tunnel	Select the tunnel that provides access to the private network behind the FortiGate unit.
UTM	Select the UTM feature profiles that you want to apply to Internet access.
Allow Inbound	Enable
Allow Outbound	Enable
Inbound NAT	Enable

To create an Internet browsing policy - route-based VPN

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information and then select *OK*:

Source Interface	The IPsec VPN interface.
Source Address Name	All
Destination Interface	The interface that connects to the Internet. The virtual IPsec interface is configured on this physical interface.
Destination Address Name	All
Schedule	As required.
Service	As required.
Action	ACCEPT
NAT	Enable
UTM	Select the UTM features that you want to apply to Internet access.

The VPN clients must be configured to route all Internet traffic through the VPN tunnel.

Routing all remote traffic through the VPN tunnel

To make use of the Internet browsing configuration on the VPN server, the VPN peer or client must route all traffic through the VPN tunnel. Usually, only the traffic destined for the private network behind the FortiGate VPN server is sent through the tunnel.

The remote end of the VPN can be a FortiGate unit that acts as a peer in a gateway-to-gateway configuration or a FortiClient Endpoint Security application that protects an individual client such as a notebook PC.

- To configure a remote peer FortiGate unit for Internet browsing via VPN, see [“Configuring a FortiGate remote peer to support Internet browsing”](#).
- To configure a FortiClient Endpoint Security application for Internet browsing via VPN, see [“Configuring a FortiClient application to support Internet browsing” on page 1518](#).

These procedures assume that your VPN connection to the protected private network is working and that you have configured the FortiGate VPN server for Internet browsing as described in [“Creating an Internet browsing security policy” on page 1516](#).

Configuring a FortiGate remote peer to support Internet browsing

The configuration changes to send all traffic through the VPN differ for policy-based and route-based VPNs.

To route all traffic through a policy-based VPN

- 1 At the FortiGate dialup client, go to *Policy > Policy > Policy*.
- 2 Select the IPsec security policy and then select *Edit*.
- 3 From the *Destination Address* list, select *all*.
- 4 Select *OK*.

All packets are routed through the VPN tunnel, not just packets destined for the protected private network.

To route all traffic through a route-based VPN

- 1 At the FortiGate dialup client, go to *Router > Static > Static Route*.
- 2 Select the default route (destination IP 0.0.0.0) and then select *Edit*. If there is no default route, select *Create New*. Enter the following information and select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	Select the IPsec virtual interface.
Distance	Leave at default.

All packets are routed through the VPN tunnel, not just packets destined for the protected private network.

Configuring a FortiClient application to support Internet browsing

By default, the FortiClient application configures the PC so that traffic destined for the remote protected network passes through the VPN tunnel but all other traffic is sent to the default gateway. You need to modify the FortiClient settings so that it configures the PC to route all outbound traffic through the VPN.

To route all traffic through VPN - FortiClient application

- 1 At the remote host, start FortiClient.
- 2 Go to *VPN > Connections*.
- 3 Select the definition that connects FortiClient to the FortiGate dialup server.
- 4 Select *Advanced* and then select *Edit*.
- 5 In the *Edit Connection* dialog box, select *Advanced*.
- 6 In the *Remote Network* group, select *Add*.
- 7 In the *IP* and *Subnet Mask* fields, type 0.0.0.0/0.0.0.0 and select *OK*.

The address is added to the *Remote Network* list. The first destination IP address in the list establishes a VPN tunnel. The second destination address (0.0.0.0/0.0.0.0 in this case) forces all other traffic through the VPN tunnel.

- 8 Select *OK*.



Redundant VPN configurations

This section discusses the options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches.

The following topics are included in this section:

- [Configuration overview](#)
- [General configuration steps](#)
- [Configure the VPN peers - route-based VPN](#)
- [Redundant route-based VPN configuration example](#)
- [Partially-redundant route-based VPN example](#)
- [Creating a backup IPsec interface](#)

Configuration overview

A FortiGate unit with two interfaces connected to the Internet can be configured to support redundant VPNs to the same remote peer. If the primary connection fails, the FortiGate unit can establish a VPN using the other connection.



Redundant tunnels do not support Tunnel Mode or Manual Keys. You must use Interface Mode.

A fully-redundant configuration requires redundant connections to the Internet on both peers. [Figure 144 on page 1520](#) shows an example of this. This is useful to create a reliable connection between two FortiGate units with static IP addresses.

When only one peer has redundant connections, the configuration is partially-redundant. For an example of this, see [“Partially-redundant route-based VPN example” on page 1534](#). This is useful to provide reliable service from a FortiGate unit with static IP addresses that accepts connections from dialup IPsec VPN clients.


In a fully-redundant VPN configuration with two interfaces on each peer, four distinct paths are possible for VPN traffic from end to end. Each interface on a peer can communicate with both interfaces on the other peer. This ensures that a VPN will be available as long as each peer has one working connection to the Internet.

You configure a VPN and an entry in the routing table for each of the four paths. All of these VPNs are ready to carry data. You set different routing distances for each route and only the shortest distance route is used. If this route fails, the route with the next shortest distance is used.

The redundant configurations described in this chapter use route-based VPNs, otherwise known as virtual IPsec interfaces. This means that the FortiGate unit must operate in NAT mode. You must use auto-keying. A VPN that is created using manual keys (see [“Manual-key configurations” on page 1551](#)) cannot be included in a redundant-tunnel configuration.

The configuration described here assumes that your redundant VPNs are essentially equal in cost and capability. When the original VPN returns to service, traffic continues to use the replacement VPN until the replacement VPN fails. If your redundant VPN uses more expensive facilities, you want to use it only as a backup while the main VPN is down. For information on how to do this, see [“Creating a backup IPsec interface” on page 1541](#).

Figure 144: Example redundant-tunnel configuration



A VPN that is created using manual keys (see [“Manual-key configurations” on page 1551](#)) cannot be included in a redundant-tunnel configuration.

General configuration steps

A redundant configuration at each VPN peer includes:

- one phase 1 configuration (virtual IPsec interface) for each path between the two peers. In a fully-meshed redundant configuration, each network interface on one peer can communicate with each network interface on the remote peer. If both peers have two public interfaces, this means that each peer has four paths, for example.
- one phase 2 definition for each phase 1 configuration
- one static route for each IPsec interface, with different distance values to prioritize the routes
- two Accept security policies per IPsec interface, one for each direction of traffic
- dead peer detection enabled in each phase 1 definition

The procedures in this section assume that two separate interfaces to the Internet are available on each VPN peer.

Configure the VPN peers - route-based VPN

VPN peers are configured using Interface Mode for redundant tunnels.

Configure each VPN peer as follows:

- 1 Ensure that the interfaces used in the VPN have static IP addresses.
- 2 Create a phase 1 configuration for each of the paths between the peers. Enable IPsec Interface mode so that this creates a virtual IPsec interface. Enable dead peer detection so that one of the other paths is activated if this path fails.

Enter these settings in particular:

Path 1

Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Type the IP address of the primary interface of the remote peer.
Local Interface	Select the primary public interface of this peer.
Enable IPsec Interface Mode	Enable
Dead Peer Detection	Enable

Other settings as required by VPN.

Path 2

Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Type the IP address of the secondary interface of the remote peer.
Local Interface	Select the primary public interface of this peer.
Enable IPsec Interface Mode	Enable
Dead Peer Detection	Enable

Other settings as required by VPN.

Path 3

Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Type the IP address of the primary interface of the remote peer.
Local Interface	Select the secondary public interface of this peer.
Enable IPsec Interface Mode	Enable
Dead Peer Detection	Enable

Other settings as required by VPN.

Path 4

Remote Gateway	Select Static IP Address.
IP Address	Type the IP address of the secondary interface of the remote peer.
Local Interface	Select the secondary public interface of this peer.
Enable IPsec Interface Mode	Enable
Dead Peer Detection	Enable

Other settings as required by VPN.

For more information, see [“Auto Key phase 1 parameters” on page 1407](#).

- 3 Create a phase 2 definition for each path. See [“Phase 2 parameters” on page 1425](#). Enter these settings in particular:

Phase 1	Select the phase 1 configuration (virtual IPsec interface) that you defined for this path. You can select the name from the Static IP Address part of the list.
----------------	---

- 4 Create a route for each path to the other peer. If there are two ports on each peer, there are four possible paths between the peer devices.

Destination IP/Mask	The IP address and netmask of the private network behind the remote peer.
Device	One of the virtual IPsec interfaces on the local peer.
Distance	For each path, enter a different value to prioritize the paths.

- 5 Define the security policy for the local primary interface. See [“Defining VPN security policies” on page 1432](#). You need to create two policies for each path to enable communication in both directions. Enter these settings in particular:

Source Interface/Zone	Select the local interface to the internal (private) network
Source Address Name	All
Destination Interface/Zone	Select one of the virtual IPsec interfaces you created in Step 2.
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

Source Interface/Zone	Select one of the virtual IPsec interfaces you created in Step 2.
Source Address Name	All
Destination Interface/Zone	Select the local interface to the internal (private) network.
Destination Address Name	All

Schedule	Always
Service	Any
Action	ACCEPT


- 6 Place the policy in the policy list above any other policies having similar source and destination addresses.
- 7 Repeat this procedure at the remote FortiGate unit.

Redundant route-based VPN configuration example

This example demonstrates a fully redundant site-to-site VPN configuration using route-based VPNs. At each site, the FortiGate unit has two interfaces connected to the Internet through different ISPs. This means that there are four possible paths for communication between the two units. In this example, these paths, listed in descending priority, are:

- FortiGate_1 WAN 1 to FortiGate_2 WAN 1
- FortiGate_1 WAN 1 to FortiGate_2 WAN 2
- FortiGate_1 WAN 2 to FortiGate_2 WAN 1
- FortiGate_1 WAN 2 to FortiGate_2 WAN 2

Figure 145: Example redundant route-based VPN configuration



For each path, VPN configuration, security policies and routing are defined. By specifying a different routing distance for each path, the paths are prioritized. A VPN tunnel is established on each path, but only the highest priority one is used. If the highest priority path goes down, the traffic is automatically routed over the next highest priority path. You could use dynamic routing, but to keep this example simple, static routing is used.

Configuring FortiGate_1

You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the four possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the four possible paths
- configure routes for the four IPsec interfaces, assigning the appropriate priorities

- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

To configure the network interfaces

- 1 Go to *System > Network > Interface*.
- 2 Select the the Internal interface and select *Edit*.
- 3 Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	10.21.101.0/255.255.255.0

- 4 Select the WAN1 interface and select *Edit*, enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	192.168.10.2/255.255.255.0

- 5 Select the WAN2 interface and select *Edit*, enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	172.16.20.2/255.255.255.0

To configure the IPsec interfaces (phase 1 configurations)

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_A
Remote Gateway	Static IP Address
IP Address	192.168.20.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 3 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_B
Remote Gateway	Static IP Address

IP Address	172.16.30.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 4 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_C
Remote Gateway	Static IP Address
IP Address	192.168.20.2
Local Interface	WAN2
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 5 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_D
Remote Gateway	Static IP Address
IP Address	172.16.30.2
Local Interface	WAN2
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

To define the phase 2 configurations for the four VPNs

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_A.
Phase 1	Site_1_A

- 3 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_B.
Phase 1	Site_1_B

- 4 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_C.
Phase 1	Site_1_C

- 5 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_D.
Phase 1	Site_1_D

To configure routes

- 1 Go to *Router > Static > Static Route*.
- 2 Select *Create New*, enter the following default gateway information and then select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	WAN1
Gateway	192.168.10.1
Distance	10

- 3 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.31.101.0/255.255.255.0
Device	Site_1_A
Distance	1

- 4 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.31.101.0/255.255.255.0
Device	Site_1_B
Distance	2

- 5 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.31.101.0/255.255.255.0
Device	Site_1_C
Distance	3

- 6 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.31.101.0/255.255.255.0
Device	Site_1_D
Distance	4

To configure security policies

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and then select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_A
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_1_A
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 4 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_B
Destination Address Name	All
Schedule	Always

Service	Any
Action	ACCEPT

- 5 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_1_B
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 6 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_C
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 7 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_1_C
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 8 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_D
Destination Address Name	All
Schedule	Always

Service	Any
Action	ACCEPT

- 9 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_1_D
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

Configuring FortiGate_2

The configuration for FortiGate_2 is very similar that of FortiGate_1. You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the four possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the four possible paths
- configure routes for the four IPsec interfaces, assigning the appropriate priorities
- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

To configure the network interfaces

- 1 Go to *System > Network > Interface*.
- 2 Select the Internal interface and then select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	10.31.101.0/255.255.255.0

- 3 Select the WAN1 interface and then select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	192.168.20.2/255.255.255.0

- 4 Select the WAN2 interface and then select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	172.16.30.2/255.255.255.0

To configure the IPsec interfaces (phase 1 configurations)

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.

- 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_A
Remote Gateway	Static IP Address
IP Address	192.168.10.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 3 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_B
Remote Gateway	Static IP Address
IP Address	172.16.20.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 4 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_C
Remote Gateway	Static IP Address
IP Address	192.168.10.2
Local Interface	WAN2
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	

Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 5 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_D
Remote Gateway	Static IP Address
IP Address	172.16.20.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

To define the phase 2 configurations for the four VPNs

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_A.
Phase 1	Site_2_A

- 3 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_B.
Phase 1	Site_2_B

- 4 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_C.
Phase 1	Site_2_C

- 5 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_D.
Phase 1	Site_2_D

To configure routes

- 1 Go to *Router > Static > Static Route*.

- 2 Select *Create New*, enter the following default gateway information and then select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	WAN1
Gateway	192.168.10.1
Distance	10

- 3 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_A
Distance	1

- 4 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_B
Distance	2

- 5 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_C
Distance	3

- 6 Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_D
Distance	4

To configure security policies

- 1 Go to *Policy > Policy > Policy*.
 2 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_A

Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_2_A
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 4 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_B
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 5 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_2_B
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 6 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_C

Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 7 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_2_C
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 8 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_D
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 9 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Site_2_D
Source Address Name	All
Destination Interface/Zone	Internal
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

Partially-redundant route-based VPN example

This example demonstrates how to set up a partially redundant IPsec VPN between a local FortiGate unit and a remote VPN peer that receives a dynamic IP address from an ISP before it connects to the FortiGate unit. For more information about FortiGate dialup-client configurations, see [“FortiGate dialup-client configurations” on page 1501](#).

When a FortiGate unit has more than one interface to the Internet (see FortiGate_1 in Figure 146), you can configure redundant routes—if the primary connection fails, the FortiGate unit can establish a VPN using the redundant connection.

In this case, FortiGate_2 has only one connection to the Internet. If the link to the ISP were to go down, the connection to FortiGate_1 would be lost, and the tunnel would be taken down. The tunnel is said to be partially redundant because FortiGate_2 does not support a redundant connection.

In the configuration example:


- Both FortiGate units operate in NAT mode.
- Two separate interfaces to the Internet (192.168.10.2 and 172.16.20.2) are available on FortiGate_1. Each interface has a static public IP address.
- FortiGate_2 has a single connection to the Internet and obtains a dynamic public IP address (for example, 172.16.30.1) when it connects to the Internet.
- FortiGate_2 forwards IP packets from the SOHO network (10.31.101.0/24) to the corporate network (10.21.101.0/24) behind FortiGate_1 through a partially redundant IPsec VPN. Encrypted packets from FortiGate_2 are addressed to the public interface of FortiGate_1. Encrypted packets from FortiGate_1 are addressed to the public IP address of FortiGate_2.

There are two possible paths for communication between the two units. In this example, these paths, listed in descending priority, are:

- FortiGate_1 WAN 1 to FortiGate_2 WAN 1
- FortiGate_1 WAN 2 to FortiGate_2 WAN 1

For each path, VPN configuration, security policies and routing are defined. By specifying a different routing distance for each path, the paths are prioritized. A VPN tunnel is established on each path, but only the highest priority one is used. If the highest priority path goes down, the traffic is automatically routed over the next highest priority path. You could use dynamic routing, but to keep this example simple, static routing is used.

Figure 146: Example partially redundant route-based configuration



Configuring FortiGate_1

You must

- configure the interfaces involved in the VPN
- define the phase 1 configuration for each of the two possible paths, creating a virtual IPsec interface for each one
- define the phase 2 configuration for each of the two possible paths
- configure incoming and outgoing security policies between the internal interface and each of the virtual IPsec interfaces

To configure the network interfaces

- 1 Go to *System > Network > Interface*.
- 2 Select the Internal interface and select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	10.21.101.2/255.255.255.0

- 3 Select the WAN1 interface and select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	192.168.10.2/255.255.255.0

- 4 Select the WAN2 interface and select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	172.16.20.2/255.255.255.0

To configure the IPsec interfaces (phase 1 configurations)

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_A
Remote Gateway	Dialup User
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	

Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 3 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_1_B
Remote Gateway	Dialup User
Local Interface	WAN2
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

To define the phase 2 configurations for the two VPNs

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
- 2 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_A.
Phase 1	Site_1_A

- 3 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_B.
Phase 1	Site_1_B

To configure routes

- 1 Go to *Router > Static > Static Route*.
- 2 Select *Create New*, enter the following default gateway information and then select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	WAN1
Gateway	192.168.10.1
Distance	10

To configure security policies

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_A
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_1_B
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

Configuring FortiGate_2

The configuration for FortiGate_2 is similar to that of FortiGate_1. You must

- configure the interface involved in the VPN
- define the phase 1 configuration for the primary and redundant paths, creating a virtual IPsec interface for each one
- define the phase 2 configurations for the primary and redundant paths, defining the internal network as the source address so that FortiGate_1 can automatically configure routing
- configure the routes for the two IPsec interfaces, assigning the appropriate priorities
- configure security policies between the internal interface and each of the virtual IPsec interfaces

To configure the network interfaces

- 1 Go to *System > Network > Interface*.
- 2 Select the Internal interface and select *Edit*. Enter the following information and then select *OK*:

Addressing mode	Manual
IP/Netmask	10.31.101.2/255.255.255.0

- 3 Select the WAN1 interface and select *Edit*. Enter the following information and then select *OK*:

Addressing mode	DHCP
------------------------	------

To configure the two IPsec interfaces (phase 1 configurations)

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.
 2 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_A
Remote Gateway	Static IP Address
IP Address	192.168.10.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

- 3 Select *Create Phase 1*, enter the following information, and select *OK*:

Name	Site_2_B
Remote Gateway	Static IP Address
IP Address	172.16.20.2
Local Interface	WAN1
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Accept any peer ID
Advanced	
Enable IPsec Interface Mode	Select
Dead Peer Detection	Select

To define the phase 2 configurations for the two VPNs

- 1 Go to *VPN > IPsec > Auto Key (IKE)*.

- 2 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_A.
Phase 1	Site_2_A
Advanced	
Source Address	10.31.101.0/24

- 3 Select *Create Phase 2*, enter the following information and select *OK*:

Name	Route_B.
Phase 1	Site_2_B
Advanced	
Source Address	10.31.101.0/24

To configure routes

- Go to *Router > Static > Static Route*.
- Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_A
Distance	1

- Select *Create New*, enter the following information and then select *OK*:

Destination IP/Mask	10.21.101.0/255.255.255.0
Device	Site_2_B
Distance	2

To configure security policies

- Go to *Policy > Policy > Policy*.
- Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_A
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	Internal
Source Address Name	All
Destination Interface/Zone	Site_2_B
Destination Address Name	All
Schedule	Always
Service	Any
Action	ACCEPT

Creating a backup IPsec interface

You can configure a route-based VPN that acts as a backup facility to another VPN. It is used only while your main VPN is out of service. This is desirable when the redundant VPN uses a more expensive facility.

You can configure a backup IPsec interface only in the CLI. The backup feature works only on interfaces with static addresses that have dead peer detection enabled. The `monitor-phase1` option creates a backup VPN for the specified phase 1 configuration.

In the following example, `backup_vpn` is a backup for `main_vpn`.

```
config vpn ipsec phase1-interface
  edit main_vpn
    set dpd on
    set interface port1
    set nattraversal enable
    set psksecret "hard-to-guess"
    set remote-gw 192.168.10.8
    set type static
  end
  edit backup_vpn
    set dpd on
    set interface port2
    set monitor-phase1 main_vpn
    set nattraversal enable
    set psksecret "hard-to-guess"
    set remote-gw 192.168.10.8
    set type static
  end
```




Transparent mode VPNs

This section describes transparent VPN configurations, in which two FortiGate units create a VPN tunnel between two separate private networks transparently.


The following topics are included in this section:

- [Configuration overview](#)
- [Configure the VPN peers](#)


Configuration overview

In transparent mode, all interfaces of the FortiGate unit except the management interface (which by default is assigned IP address 10.10.10.1/255.255.255.0) are invisible at the network layer. Typically, when a FortiGate unit runs in transparent mode, different network segments are connected to the FortiGate interfaces. [Figure 147](#) shows the management station on the same subnet. The management station can connect to the FortiGate unit directly through the web-based manager.

Figure 147: Management station on internal network



An edge router typically provides a public connection to the Internet and one interface of the FortiGate unit is connected to the router. If the FortiGate unit is managed from an external address (see [Figure 148 on page 1544](#)), the router must translate (NAT) a routable address to direct management traffic to the FortiGate management interface.

Figure 148: Management station on external network

In a transparent VPN configuration, two FortiGate units create a VPN tunnel between two separate private networks transparently. All traffic between the two networks is encrypted and protected by FortiGate security policies.

Both FortiGate units may be running in transparent mode, or one could be running in transparent mode and the other running in NAT mode. If the remote peer is running in NAT mode, it must have a static public IP address.




VPNs between two FortiGate units running in transparent mode do not support inbound/outbound NAT (supported through CLI commands) within the tunnel. In addition, a FortiGate unit running in transparent mode cannot be used in a hub-and-spoke configuration.


Encrypted packets from the remote VPN peer are addressed to the management interface of the local FortiGate unit. If the local FortiGate unit can reach the VPN peer locally, a static route to the VPN peer must be added to the routing table on the local FortiGate unit. If the VPN peer connects through the Internet, encrypted packets from the local FortiGate unit must be routed to the edge router instead. For information about how to add a static route to the FortiGate routing table, see the [Advanced Routing](#).

In the example configuration shown in [Figure 148](#), Network Address Translation (NAT) is enabled on the router. When an encrypted packet from the remote VPN peer arrives at the router through the Internet, the router performs inbound NAT and forwards the packet to the FortiGate unit. Refer to the software supplier's documentation to configure the router.

If you want to configure a VPN between two FortiGate units running in transparent mode, each unit must have an independent connection to a router that acts as a gateway to the Internet, and both units must be on separate networks that have a different address space. When the two networks linked by the VPN tunnel have different address spaces (see [Figure 149 on page 1545](#)), at least one router must separate the two FortiGate units, unless the packets can be redirected using ICMP (see [Figure 150 on page 1545](#)).


Figure 149: Link between two FortiGate units in transparent mode

In [Figure 150](#), interface C behind the router is the default gateway for both FortiGate units. Packets that cannot be delivered on Network_1 are routed to interface C by default. Similarly, packets that cannot be delivered on Network_2 are routed to interface C. In this case, the router must be configured to redirect packets destined for Network_1 to interface A and redirect packets destined for Network_2 to interface B.

Figure 150: ICMP redirecting packets to two FortiGate units in transparent mode

If there are additional routers behind the FortiGate unit (see [Figure 151 on page 1546](#)) and the destination IP address of an inbound packet is on a network behind one of those routers, the FortiGate routing table must include routes to those networks. For example, in [Figure 151](#), the FortiGate unit must be configured with static routes to interfaces A and B in order to forward packets to Network_1 and Network_2 respectively.

Figure 151: Destinations on remote networks behind internal routers



Transparent VPN infrastructure requirements

- The local FortiGate unit must be operating in transparent mode.
- The management IP address of the local FortiGate unit specifies the local VPN gateway. The management IP address is considered a static IP address for the local VPN peer.
- If the local FortiGate unit is managed through the Internet, or if the VPN peer connects through the Internet, the edge router must be configured to perform inbound NAT and forward management traffic and/or encrypted packets to the FortiGate unit.
- If the remote peer is operating in NAT mode, it must have a static public IP address.

A FortiGate unit operating in transparent mode requires the following basic configuration to operate as a node on the IP network:

- The unit must have sufficient routing information to reach the management station.
- For any traffic to reach external destinations, a default static route to an edge router that forwards packets to the Internet must be present in the FortiGate routing table.
- When all of the destinations are located on the external network, the FortiGate unit may route packets using a single default static route. If the network topology is more complex, one or more static routes in addition to the default static route may be required in the FortiGate routing table.

Only policy-based VPN configurations are possible in transparent mode.

Before you begin

An IPsec VPN definition links a gateway with a tunnel and an IPsec policy. If your network topology includes more than one virtual domain, you must choose components that were created in the same virtual domain. Therefore, before you define a transparent VPN configuration, choose an appropriate virtual domain in which to create the required interfaces, security policies, and VPN components. For more information, see the [Virtual Domains](#) chapter of *The Handbook*.

Configure the VPN peers

- 1 The local VPN peer need to operate in transparent mode.

To determine if your FortiGate unit is in transparent mode, go to *System > Dashboard > Status to the System Information* widget. Select *[change]*. Select transparent for the Operation Mode— two new fields will appear to enter the *Management IP/Netmask*, and the *Default Gateway*.



In transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. You only have to configure a management IP address so that you can make configuration changes.

The remote VPN peer may operate in NAT mode or transparent mode.

- 2 At the local FortiGate unit, define the phase 1 parameters needed to establish a secure connection with the remote peer. See [“Auto Key phase 1 parameters” on page 1407](#). Select *Advanced* and enter these settings in particular:

Remote Gateway	Select Static IP Address.
IP Address	Type the IP address of the public interface to the remote peer. If the remote peer is a FortiGate unit running in transparent mode, type the IP address of the remote management interface.
Advanced	Select Nat-traversal, and type a value into the Keepalive Frequency field. These settings protect the headers of encrypted packets from being altered by external NAT devices and ensure that NAT address mappings do not change while the VPN tunnel is open. For more information, see “NAT traversal” on page 1420 and “NAT keepalive frequency” on page 1421 .

- 3 Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. See [“Phase 2 parameters” on page 1425](#).

Phase 1	Select the set of phase 1 parameters that you defined for the remote peer. The name of the remote peer can be selected from the Static IP Address list.
----------------	---

- 4 Define the source and destination addresses of the IP packets that are to be transported through the VPN tunnel. See [“Defining policy addresses” on page 1431](#). Enter these settings in particular:
 - For the originating address (source address), enter the IP address and netmask of the private network behind the local peer network. for the management interface, for example, 10.10.10.0/24. This address needs to be a range to allow traffic from your network through the tunnel. Optionally select `any` for this address.
 - For the remote address (destination address), enter the IP address and netmask of the private network behind the remote peer (for example, 192.168.10.0/24). If the remote peer is a FortiGate unit running in transparent mode, enter the IP address of the remote management interface instead.
- 5 Define an IPsec security policy to permit communications between the source and destination addresses. See [“Defining VPN security policies” on page 1432](#). Enter these settings in particular:

Source Interface/Zone	Select the local interface to the internal (private) network.
Source Address Name	Select the source address that you defined in Step 4.
Destination Interface/Zone	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
Destination Address Name	Select the destination address that you defined in Step 4.
Action	IPSEC
VPN Tunnel	<p>Select the name of the phase 2 tunnel configuration that you created in Step 3.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

- 6 Place the policy in the policy list above any other policies having similar source and destination addresses.
- 7 Define another IPsec security policy to permit communications between the source and destination addresses in the opposite direction. This security policy and the previous one form a bi-directional policy pair. See [“Defining VPN security policies” on page 1432](#). Enter these settings in particular:

Source Interface/Zone	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
Source Address Name	Select the destination address that you defined in Step 4.
Destination Interface/Zone	Select the local interface to the internal (private) network.
Destination Address Name	Select the source address that you defined in Step 4.

Action	IPSEC
VPN Tunnel	<p>Select the name of the phase 2 tunnel configuration that you created in Step 3.</p> <p>Select <i>Allow inbound</i> to enable traffic from the remote network to initiate the tunnel.</p> <p>Select <i>Allow outbound</i> to enable traffic from the local network to initiate the tunnel.</p>

- 8 Repeat this procedure at the remote FortiGate unit to create bidirectional security policies. Use the local interface and address information local to the remote FortiGate unit.

For more information on transparent mode, see the [System Administration handbook chapter](#).



Manual-key configurations

This section explains how to manually define cryptographic keys to establish an IPsec VPN, either policy-based or route-based.

For more information on web-based manual key configuration, see [“Manual Key” on page 1403](#).

The following topics are included in this section:

- [Configuration overview](#)
- [Specify the manual keys for creating a tunnel](#)

Configuration overview

You manually define cryptographic keys where prior knowledge of the encryption and/or authentication key is required (that is, one of the VPN peers requires a specific IPsec encryption and/or authentication key). In this case, you do not specify IPsec phase 1 and phase 2 parameters; you define manual keys by going to *VPN > IPsec > Manual Key*.

If one VPN peer uses specific authentication and encryption keys to establish a tunnel, both VPN peers must be configured to use the same encryption and authentication algorithms and keys.



It may not be safe or practical to define manual keys because network administrators must be trusted to keep the keys confidential, and propagating changes to remote VPN peers in a secure manner may be difficult.

It is essential that both VPN peers be configured with matching encryption and authentication algorithms, matching authentication and encryption keys, and complementary Security Parameter Index (SPI) settings.

You can define either the encryption or the authentication as NULL (disabled), but not both.

Each SPI identifies a Security Association (SA). The value is placed in ESP datagrams to link the datagrams to the SA. When an ESP datagram is received, the recipient refers to the SPI to determine which SA applies to the datagram. An SPI must be specified manually for each SA. Because an SA applies to communication in one direction only, you must specify two SPIs per configuration (a local SPI and a remote SPI) to cover bidirectional communications between two VPN peers.



If you are not familiar with the security policies, SAs, selectors, and SA databases for your particular installation, do not attempt the following procedure without qualified assistance.

Specify the manual keys for creating a tunnel


Specify the manual keys for creating a tunnel as follows:

- 1 Go to *VPN > IPsec > Manual Key* and select *Create New*.
- 2 Include appropriate entries as follows:

Name	Type a name for the VPN tunnel.
Local SPI	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles outbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the <i>Remote SPI</i> value in the manual key configuration at the remote peer.
Remote SPI	Type a hexadecimal number (up to 8 characters, 0-9, a-f) that represents the SA that handles inbound traffic on the local FortiGate unit. The valid range is from 0x100 to 0xffffffff. This value must match the <i>Local SPI</i> value in the manual key configuration at the remote peer.
Remote Gateway	Type the IP address of the public interface to the remote peer. The address identifies the recipient of ESP datagrams.
Local Interface	Select the name of the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound. The FortiGate unit obtains the IP address of the interface from <i>System > Network > Interface</i> settings. This is available in NAT mode only.
Encryption Algorithm	Select one of the following symmetric-key encryption algorithms: <ul style="list-style-type: none"> • DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key. • 3DES — Triple-DES, in which plain text is encrypted three times by three keys. • AES128 — A 128-bit block algorithm that uses a 128-bit key. • AES192 — A 128-bit block algorithm that uses a 192-bit key. • AES256 — A 128-bit block algorithm that uses a 256-bit key.
Encryption Key	If you selected: <ul style="list-style-type: none"> • DES, type a 16-character hexadecimal number (0-9, a-f). • 3DES, type a 48-character hexadecimal number (0-9, a-f) separated into three segments of 16 characters. • AES128, type a 32-character hexadecimal number (0-9, a-f) separated into two segments of 16 characters. • AES192, type a 48-character hexadecimal number (0-9, a-f) separated into three segments of 16 characters. • AES256, type a 64-character hexadecimal number (0-9, a-f) separated into four segments of 16 characters.
Authentication Algorithm	Select one of the following message digests: <ul style="list-style-type: none"> • MD5 — Message Digest 5 algorithm, which produces a 128-bit message digest. • SHA1 — Secure Hash Algorithm 1, which produces a 160-bit message digest.

Authentication Key	<p>If you selected:</p> <ul style="list-style-type: none">• MD5, type a 32-character hexadecimal number (0-9, a-f) separated into two segments of 16 characters.• SHA1, type 40-character hexadecimal number (0-9, a-f) separated into one segment of 16 characters and a second segment of 24 characters.
IPsec Interface Mode	<p>Select to create a route-based VPN. A virtual IPsec interface is created on the Local Interface that you selected. This option is available only in NAT mode.</p>

3 Select OK.



IPv6 IPsec VPNs

This chapter describes how to configure your FortiGate unit's IPv6 IPsec VPN functionality. To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6* in the section, *Display Options on GUI*.

The following topics are included in this section:

- [Overview of IPv6 IPsec support](#)
- [Configuring IPv6 IPsec VPNs](#)
- [Site-to-site IPv6 over IPv6 VPN example](#)
- [Site-to-site IPv4 over IPv6 VPN example](#)
- [Site-to-site IPv6 over IPv4 VPN example](#)

Overview of IPv6 IPsec support

FortiOS supports route-based IPv6 IPsec, but not policy-based. This section describes how IPv6 IPsec support differs from IPv4 IPsec support. FortiOS 4.0 MR3 is Pv6 Ready Logo Program Phase 2 certified.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

IPv4 over IPv6	The VPN gateways have IPv6 addresses. The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.
IPv6 over IPv4	The VPN gateways have IPv4 addresses. The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported. This is because FortiOS 3.0 does not support IPv6 DNS.
- You cannot use RSA certificates in which the common name (cn) is a domain name that resolves to an IPv6 address. This is because FortiOS 3.0 does not support IPv6 DNS.
- DHCP over IPsec is not supported, because FortiOS 3.0 does not support IPv6 DHCP.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies and routing.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6* in the section, *Display Options on GUI*.

Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the Advanced settings. Enable the IPv6 Version check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses. For example:

```
config vpn ipsec phase1-interface
edit tunnel6
set ip-version 6
set remote-gw6 0:123:4567::1234
set interface port3
set proposal 3des-md5
end
```

Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default “0.0.0.0/0” address for Source address and Destination address to the IPv6 value “::/0”. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, “::/0” for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like this:

```
config vpn ipsec phase2-interface
edit tunnel6_p2
set phase1name tunnel6
set proposal 3des-md5
set src-addr-type subnet6
set dst-addr-type subnet6
end
```

Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network’s port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

To create a static route in the web-based manager, go to *Router > Static > Static Route*. Select the drop-down arrow on the Create New button and select IPv6 Route. Enter the information and select OK. In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

```
config router static6
  edit 1
    set device port2
    set dst 0::/0
  next
  edit 2
    set device toB
    set dst fec0:0000:0000:0004::/64
  next
end
```


If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

Site-to-site IPv6 over IPv6 VPN example

In this example, computers on IPv6-addressed private networks communicate securely over public IPv6 infrastructure.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6* in the section, *Display Options on GUI*.

Figure 152: Example IPv6-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the public network and port 3 connects to the local network.

```
config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f2/64
    end
  next
```

```

edit port3
  config ipv6
    set ip6-address fec0::0000:209:0fff:fe83:25f3/64
  end
next
end

```

Configure FortiGate A IPsec settings

The phase 1 configuration creates a virtual IPsec interface on port 2 and sets the remote gateway to the public IP address FortiGate B. This configuration is the same as for an IPv4 route-based VPN, except that `ip-version` is set to 6 and the `remote-gw6` keyword is used to specify an IPv6 remote gateway address.

```

config vpn ipsec phase1-interface
edit toB
  set ip-version 6
  set interface port2
  set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
  set dpd enable
  set psksecret maryhadalittlelamb
  set proposal 3des-md5 3des-sha1
end

```

By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are 0.0.0.0/0 for IPv4, `::/0` for IPv6.

```

config vpn ipsec phase2-interface
edit toB2
  set phase1name toB
  set proposal 3des-md5 3des-sha1
  set pfs enable
  set replay enable
  set src-addr-type subnet6
  set dst-addr-type subnet6
end

```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```

config firewall policy6
edit 1
  set srcintf port3
  set dstintf toB
  set srcaddr all6
  set dstaddr all6
  set action accept
  set service ANY
  set schedule always
next
edit 2
  set srcintf toB
  set dstintf port3

```

```
set srcaddr all6
set dstaddr all6
set action accept
set service ANY
set schedule always
end
```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB. A default route sends all IPv6 traffic out on port2.

```
config router static6
edit 1
set device port2
set dst 0::/0
next
edit 2
set device toB
set dst fec0:0000:0000:0004::/64
end
```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. Security policies enable traffic to pass between the private network and the IPsec interface. Routing ensures traffic for the private network behind FortiGate A goes through the VPN and that all IPv6 packets are routed to the public network.

```
config system interface
edit port2
config ipv6
set ip6-address fec0::0003:209:0fff:fe83:25c7/64
end
next
edit port3
config ipv6
set ip6-address fec0::0004:209:0fff:fe83:2569/64
end
end
config vpn ipsec phase1-interface
edit toA
set ip-version 6
set interface port2
set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
set dpd enable
set psksecret maryhadalittlelamb
set proposal 3des-md5 3des-sha1
end
config vpn ipsec phase2-interface
edit toA2
set phase1name toA
set proposal 3des-md5 3des-sha1
set pfs enable
set replay enable
```

```


        set src-addr-type subnet6
        set dst-addr-type subnet6
    end
config firewall policy6
    edit 1
        set srcintf port3
        set dstintf toA
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toA
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end
config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toA
        set dst fec0:0000:0000:0000::/64
    end

```

Site-to-site IPv4 over IPv6 VPN example

In this example, two private networks with IPv4 addressing communicate securely over IPv6 infrastructure.

Figure 153: Example IPv4-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the IPv6 public network and port 3 connects to the IPv4 LAN.

```
config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f2/64
    end
  next
  edit port3
    set 192.168.2.1/24
  end
```

Configure FortiGate A IPsec settings

The phase 1 configuration is the same as in the IPv6 over IPv6 example.

```
config vpn ipsec phase1-interface
  edit toB
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
```

The phase 2 configuration is the same as you would use for an IPv4 VPN. By default, phase 2 selectors are set to accept all subnet addresses for source and destination.

```
config vpn ipsec phase2-interface
  edit toB2
    set phase1name toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
  end
```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. These are IPv4 security policies.

```
config firewall policy
  edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toB
    set dstintf port3
    set srcaddr all
    set dstaddr all
```

```

    set action accept
    set service ANY
    set schedule always
end

```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv4 static route. A default route sends all IPv6 traffic, including the IPv6 IPsec packets, out on port2.

```

config router static6
  edit 1
    set device port2
    set dst 0::/0
  next
  edit 2
    set device toB
    set dst 192.168.3.0/24
  end
end

```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. The IPsec phase 2 configuration has IPv4 selectors.

IPv4 security policies enable traffic to pass between the private network and the IPsec interface. An IPv4 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv6 static route ensures that all IPv6 packets are routed to the public network.

```

config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0003:fe83:25c7/64
    end
  next
  edit port3
    set 192.168.3.1/24
  end
config vpn ipsec phase1-interface
  edit toA
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
config vpn ipsec phase2-interface
  edit toA2
    set phase1name toA
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
  end
end

```



```


config firewall policy
  edit 1
    set srcintf port3
    set dstintf toA
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toA
    set dstintf port3
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
config router static6
  edit 1
    set device port2
    set dst 0::/0
  next
  edit 2
    set device toA
    set dst 192.168.2.0/24
  end

```

Site-to-site IPv6 over IPv4 VPN example

In this example, IPv6-addressed private networks communicate securely over IPv4 public infrastructure.

Figure 154: Example IPv6-over-IPv4 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the IPv4 public network and port 3 connects to the IPv6 LAN.

```

config system interface
  edit port2
    set 10.0.0.1/24
  next
  edit port3
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f3/64
    end
  end

```

Configure FortiGate A IPsec settings

The phase 1 configuration uses IPv4 addressing.

```

config vpn ipsec phase1-interface
  edit toB
    set interface port2
    set remote-gw 10.0.1.1
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end

```

The phase 2 configuration uses IPv6 selectors. By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```

config vpn ipsec phase2-interface
  edit toB2
    set phase1name toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
    set src-addr-type subnet6
    set dst-addr-type subnet6
  end

```

Configure FortiGate A security policies

IPv6 security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```

config firewall policy6
  edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toB
    set dstintf port3
    set srcaddr all6
    set dstaddr all6
  end

```

```
set action accept
set service ANY
set schedule always
end
```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv6 static route. A default route sends all IPv4 traffic, including the IPv4 IPsec packets, out on port2.

```
config router static6
edit 1
set device toB
set dst fec0:0000:0000:0004::/64
end
config router static
edit 1
set device port2
set dst 0.0.0.0/0
set gateway 10.0.0.254
end
```


Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the IPv4 public IP address of FortiGate A. The IPsec phase 2 configuration has IPv6 selectors.

IPv6 security policies enable traffic to pass between the private network and the IPsec interface. An IPv6 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv4 static route ensures that all IPv4 packets are routed to the public network.

```
config system interface
edit port2
set 10.0.1.1/24
next
edit port3
config ipv6
set ip6-address fec0::0004:209:0fff:fe83:2569/64
end
config vpn ipsec phase1-interface
edit toA
set interface port2
set remote-gw 10.0.0.1
set dpd enable
set psksecret maryhadalittlelamb
set proposal 3des-md5 3des-sha1
end
config vpn ipsec phase2-interface
edit toA2
set phase1name toA
set proposal 3des-md5 3des-sha1
set pfs enable
set replay enable
set src-addr-type subnet6
```

```
        set dst-addr-type subnet6
    end
config firewall policy6
    edit 1
        set srcintf port3
        set dstintf toA
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toA
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end
config router static6
    edit 1
        set device toA
        set dst fec0:0000:0000:0000::/64
    end
config router static
    edit 1
        set device port2
        set gateway 10.0.1.254
    end
```



L2TP and IPsec (Microsoft VPN)

This section describes how to set up a VPN that is compatible with the Microsoft Windows native VPN, which is Layer 2 Tunneling Protocol (L2TP) with IPsec encryption.


The following topics are included in this section:

- [Overview](#)
- [Assumptions](#)
- [Configuring the FortiGate unit](#)
- [Configuring the Windows PC](#)
- [Troubleshooting](#)

Overview

The topology of a VPN for Microsoft Windows dialup clients is very similar to the topology for FortiClient Endpoint Security clients.

Figure 155: Example FortiGate VPN configuration with Microsoft clients



For users, the difference is that instead of installing and using the FortiClient application, they configure a network connection using the software built into the Microsoft Windows operating system. Starting in FortiOS 4.0 MR2, you can configure a FortiGate unit to work with unmodified Microsoft VPN client software.

Layer 2 Tunneling Protocol (L2TP)

L2TP is a tunneling protocol published in 1999 that is used with VPNs, as the name suggests. Microsoft Windows operating system has a built-in L2TP client starting since Windows 2000. Mac OS X 10.3 system and higher also have a built-in client.

L2TP provides no encryption and used UDP port 1701. IPsec is used to secure L2TP packets. The initiator of the L2TP tunnel is called the L2TP Access Concentrator (LAC).



L2TP and IPsec is supported for native Windows XP, Windows Vista and Mac OSX native VPN clients. However, in Mac OSX (OSX 10.6.3, including patch releases) the L2TP feature does not work properly on the Mac OS side.

Assumptions

The following assumptions have been made for this example:

- L2TP protocol traffic is allowed through network firewalls (TCP and UDP port 1701)
- User has Microsoft Windows 2000 or higher — a Windows version that supports L2TP

Configuring the FortiGate unit

To configure the FortiGate unit, you need to:

- configure L2TP users and firewall user group;
- configure the L2TP VPN, including the IP address range it assigns to clients;
- configure an IPsec VPN with encryption and authentication settings that match the Microsoft VPN client;
- configure security policies.

Configuring L2TP users and firewall user group

Remote users must be authenticated before they can request services and/or access network resources through the VPN. The authentication process can use a password defined on the FortiGate unit or an established external authentication mechanism such as RADIUS or LDAP.

Creating user accounts

You need to create user accounts and then add these users to a firewall user group to be used for L2TP authentication. The Microsoft VPN client can automatically send the user's Windows network logon credentials. You might want to use these for their L2TP user name and password.

To create a user account - web-based manager

- 1 Go to *User > User > User* and select *Create New*.
- 2 Enter the *User Name*.
- 3 Do one of the following:
 - Select *Password* and enter the user's assigned password.
 - Select *LDAP*, *RADIUS*, or *TACACS+* and select the authentication server from the list. The authentication server must be already configured on the FortiGate unit.
- 4 Select *OK*.

To create a user account - CLI

To create a user account called `user1` with the password `123_user`, you would enter:

```
config user local
  edit user1
    set type password
    set passwd "123_user"
    set status enable
  end
```

Creating a user group

When clients connect using the L2TP-over-IPsec VPN, the FortiGate unit checks their credentials against the user group you specify for L2TP authentication. You need to create a firewall user group to use for this purpose.

To create a user group - web-based manager

- 1 Go to *User > User Group > User Group*, select *Create New*, and enter the following information:

Name	Type or edit the user group name (for example, <code>L2TP_group</code>).
Type	Select <i>Firewall</i> .
Available Users/Groups	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that can be added to the user group. To add a member to this list, select the name and then select the right arrow button.
Members	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that belong to the user group. To remove a member, select the name and then select the left arrow button.

- 2 Select *OK*.

To create a user group - CLI

To create the user group `L2TP_group` and add members `User_1`, `User_2`, and `User_3`, you would enter:

```
config user group
  edit L2TP_group
    set group-type firewall
    set member User_1 User_2 User_3
  end
```

Configuring L2TP

You can only configure L2TP settings in the CLI. As well as enabling L2TP, you set the range of IP address values that are assigned to L2TP clients and specify the user group that can access the VPN. For example, to allow access to users in the `L2TP_group` and assign them addresses in the range 192.168.0.50 to 192.168.0.59, you would enter

```
config vpn l2tp
  set sip 192.168.0.50
  set eip 192.168.0.59
  set status enable
  set usrgrp "L2TP_group"
end
```

One of the security policies for the L2TP over IPsec VPN uses the client address range, so you need also need to create a firewall address for that range. For example,

```
config firewall address
  edit L2TPclients
    set type iprange
    set end-ip 192.168.6.88
    set start-ip 192.168.6.85
  end
```

Alternatively, you could define this range in the web-based manager.

Configuring IPsec

The Microsoft VPN client uses IPsec for encryption. The configuration needed on the FortiGate unit is the same as for any other IPsec VPN with the following exceptions.

- Transport mode is used instead of tunnel mode.
- The encryption and authentication proposals must be compatible with the Microsoft client.

L2TP over IPsec is supported on the FortiGate unit using policy-based, not route-based configurations.

Configuring phase 1 - web-based manager

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
- 2 Enter the following information and then select *OK*.

Name	Enter a name for this VPN, dialup_p1 for example.
Remote Gateway	<i>Dialup User</i>
Local Interface	Select the network interface that connects to the Internet. For example, port1.
Mode	<i>Main (ID protection)</i>
Authentication Method	<i>Preshared Key</i>
Pre-shared Key	Enter the preshared key. This key must also be entered in the Microsoft VPN client.
Advanced	Select <i>Advanced</i> to enter the following information.
Enable IPsec Interface Mode	This must not be selected.
P1 Proposal	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
DH Group	2
NAT Traversal	Enable
Dead Peer Detection	Enable

Configuring phase 1 - CLI

To create a phase 1 configuration called dialup_p1 on a FortiGate unit that has port1 connected to the Internet, you would enter:

```
config vpn ipsec phase1
  edit dialup_p1
    set type dynamic
    set interface port1
    set mode main
    set psksecret *****
    set proposal aes256-md5 3des-sha1 aes192-sha1
    set dhgrp 2
    set nattraversal enable
    set dpd enable
  end
```

Configuring phase 2 - web-based manager

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 2*.
- 2 Enter the following information and then select *OK*.

Name	Enter a name for this phase 2 configuration.
Phase 1	Select the name of the phase 1 configuration.
Advanced	Select <i>Advanced</i> to enter the following information.
P2 Proposal	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
Enable replay detection	Enable
Enable perfect forward secrecy (PFS)	Disable
Keylife	3600 seconds

- 3 Make this a transport-mode VPN. You must use the CLI to do this. If your phase 2 name is dialup_p2, you would enter:

```
config vpn ipsec phase2
  edit dialup_p2
    set encapsulation transport-mode
  end
```

Configuring phase 2 - CLI

To configure a phase 2 to work with your phase_1 configuration, you would enter:

```
config vpn ipsec phase2
  edit dialup_p2
    set phase1name dialup_p1
    set proposal aes256-md5 3des-sha1 aes192-sha1
    set replay enable
    set pfs disable
    set keylifeseconds 3600
    set encapsulation transport-mode
  end
```

Configuring security policies

The security policies required for L2TP over IPsec VPN are:

- an IPSEC policy, as you would create for any policy-based IPsec VPN
- a regular ACCEPT policy to allow traffic from the L2TP clients to access the protected network

Configuring the IPSEC security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and then select OK:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address	<i>all</i>
Destination Interface/Zone	Select the FortiGate unit's public interface.
Destination Address	<i>all</i>
Action	<i>IPSEC</i>
VPN Tunnel	Select the name of the phase 1 configuration that you created. For example, dialup_p1. See “Configuring IPsec” on page 1570 .
Allow inbound	Enable
Allow outbound	Enable
UTM	Optional settings for UTM features.

Configuring the IPSEC security policy - CLI

If your VPN tunnel (phase 1) is called dialup_p1, your protected network is on port2, and your public interface is port1, you would enter:

```
config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action ipsec
    set schedule always
    set service ANY
    set inbound enable
    set outbound enable
    set vpngroup dialup_p1
  end
```

Configuring the ACCEPT security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and then select *OK*:

Source Interface/Zone	Select the FortiGate unit's public interface.
Source Address	Select the firewall address that you defined for the L2TP clients.
Destination Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address	<i>all</i>
Action	<i>ACCEPT</i>
UTM	Optionally, select UTM feature profiles.

Configuring the ACCEPT security policy - CLI

If your public interface is port1, your protected network is on port2, and L2TPclients is the address range that L2TP clients use, you would enter:

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr L2TPclients
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  end
```

Configuring the Windows PC

Configuration of the Windows PC for a VPN connection to the FortiGate unit consists of the following:

- In Network Connections, configure a Virtual Private Network connection to the FortiGate unit.
- Ensure that the IPSEC service is running.
- Ensure that IPsec has not been disabled for the VPN client. It may have been disabled to make the Microsoft VPN compatible with an earlier version of FortiOS.

The instructions in this section are based on Windows XP SP3. Other versions of Windows may vary slightly.

To configure the network connection

- 1 Open *Network Connections*.
This is available through the Control Panel.
- 2 Double-click *New Connection Wizard* and *Select Next*.
- 3 Select *Connect to the network at my workplace*.
- 4 Select *Next*.

- 5 Select *Virtual Private Network connection* and then select *Next*.
- 6 In the *Company Name* field, enter a name for the connection and then select *Next*.
- 7 Select *Do not dial the initial connection* and then select *Next*.
- 8 Enter the public IP address or FQDN of the FortiGate unit and then select *Next*.
- 9 Optionally, select *Add a shortcut to this connection to my desktop*.
- 10 Select *Finish*.
The *Connect* dialog opens on the desktop.
- 11 Select *Properties* and then select the *Security* tab.
- 12 Select IPsec Settings.
- 13 Select *Use pre-shared key for authentication*, enter the preshared key that you configured for your VPN, and select *OK*.
- 14 Select *OK*.

To check that the IPSEC service is running

- 1 Open *Administrative Tools*.
This is available through the Control Panel.
- 2 Double-click *Services*.
- 3 Look for IPSEC Services. Confirm that the *Startup Type* is *Automatic* and *Status* is set to *Started*. If needed, double-click *IPSEC Services* to change these settings.

To check that IPsec has not been disabled

- 1 Select *Start > Run*.
- 2 Enter *regedit* and select *OK*.
- 3 Find the Registry key
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters
- 4 If there is a *ProhibitIPSec* value, it must be set to 0.

Troubleshooting

This section describes some checks and tools you can use to resolve issues with L2TP-over-IPsec VPNs.

This section includes:

- [Quick checks](#)
- [Mac OS X and L2TP](#)
- [Setting up logging](#)
- [Understanding the log messages](#)
- [Using the FortiGate unit debug commands](#)

Quick checks

The table below is a list of common L2TP over IPsec VPN problems and the possible solutions.

Problem	What to check
IPsec tunnel does not come up.	<p>Check the logs to determine whether the failure is in Phase 1 or Phase 2.</p> <p>Check the settings, including encapsulation setting, which must be transport-mode.</p> <p>Check the user password.</p> <p>Confirm that the user is a member of the user group assigned to L2TP.</p> <p>On the Windows PC, check that the IPsec service is running and has not been disabled. See “Configuring the Windows PC” on page 1573.</p>
Tunnel connects, but there is no communication.	<p>Did you create an ACCEPT security policy from the public network to the protected network for the L2TP clients? See “Configuring security policies” on page 1572.</p>

Mac OS X and L2TP

FortiOS allows L2TP connections with empty AVP host names and therefore Mac OS X L2TP connections can connect to the FortiGate.

Prior to FortiOS 4.0 MR3, FortiOS refused L2TP connections with empty AVP host names in compliance with RFC 2661 and RFC 3931.

Setting up logging

L2TP logging must be enabled to record L2TP events. Alert email can be configured to report L2TP errors.

To configure FortiGate logging for L2TP over IPsec

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select *Event Log*.
- 3 Select the *L2TP/PPTP/PPPoE service event* and *IPsec negotiation event* check boxes.
- 4 Select *Apply*.

To configure FortiGate logging level

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select the *Local Logging & Archiving* check box.
- 3 Set *Minimum log level* to *Information*.
- 4 Select *Apply*.

To view FortiGate logs

- 1 Go to *Log&Report > Log & Archive Access > Event Log*.
- 2 Select the storage location if required.
- 3 After each attempt to start the L2TP over IPsec VPN, select *Refresh* to view any logged events.

Understanding the log messages

Successful startup of an L2TP over IPsec VPN follows a well-defined sequence. If you compare your logs to the sequence shown in [Table 91](#), you will be able to determine at what stage your configuration is failing.

Table 91: Sequence of log messages for L2TP over IPsec VPN connection startup

	ID	Sub Type	Action	Message
1	37127	ipsec	negotiate	progress IPsec phase 1
2	37127	ipsec	negotiate	progress IPsec phase 1
3	37127	ipsec	negotiate	progress IPsec phase 1
4	37127	ipsec	negotiate	progress IPsec phase 1
5	37129	ipsec	negotiate	progress IPsec phase 2
6	37133	ipsec	install_sa	install IPsec SA
7	37139	ipsec	phase2-up	IPsec phase 2 status change
8	37138	ipsec	tunnel-up	IPsec connection status change
9	37129	ipsec	negotiate	progress IPsec phase 2
10	37122	ipsec	negotiate	negotiate IPsec phase 2
11	31008	ppp	connect	Client 172.20.120.151 control connection started (id 743), assigned ip 192.168.6.85
12	29013	ppp		
13	29002	ppp	auth_success	User 'user1' using l2tp with authentication protocol MSCHAP_V2, succeeded
14	31101	ppp	tunnel-up	L2TP tunnel established



This table lists messages in top down chronological order. In the web-based manager log viewer, you need to read the messages from the bottom up. The newest message appears at the top of that list.

In [Table 91](#), messages 1 through 4 show the IKE phase 1 negotiation stages that result in the creation of the Security Association (SA) shown in message 6. Phase 2 negotiation in messages 5, 9, 10 produce the tunnel-up condition reported in message 8.

With IPsec communication established, the L2TP connection is established (message 11), the pppd daemon starts (message 12), the user is authenticated (message 13), and the L2TP tunnel is now ready to use.

Using the FortiGate unit debug commands

To view debug output for IKE and L2TP

- 1 Start an SSH or Telnet session to your FortiGate unit.
- 2 Enter the following CLI commands


```
diagnose debug application ike -1
diagnose debug application l2tp -1
diagnose debug enable
```
- 3 Attempt to use the VPN and note the debug output in the SSH or Telnet session.
- 4 Enter the following command to reset debug settings to default:


```
diagnose debug reset
```

To use the packet sniffer

- 1 Start an SSH or Telnet session to your FortiGate unit.
- 2 Enter the following CLI command


```
diagnose sniffer packet any icmp 4
```
- 3 Attempt to use the VPN and note the debug output.
- 4 Enter `Ctrl-C` to end sniffer operation.

Typical L2TP over IPsec session startup log entries - raw format

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 1" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success
init=remote mode=main dir=outbound stage=1 role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 1" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success
init=remote mode=main dir=outbound stage=2 role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 1" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success
init=remote mode=main dir=inbound stage=3 role=responder result=DONE
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 1" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success
init=remote mode=main dir=outbound stage=3 role=responder result=DONE
```

2010-01-11 16:39:58 log_id=0101037129 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 2" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success
init=remote mode=quick dir=outbound stage=1 role=responder result=OK

2010-01-11 16:39:58 log_id=0101037133 type=event subtype=ipsec pri=notice
vd="root" msg="install IPsec SA" action="install_sa" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" role=responder
in_spi=61100fe2 out_spi=bd70fca1

2010-01-11 16:39:58 log_id=0101037139 type=event subtype=ipsec pri=notice
vd="root" msg="IPsec phase 2 status change" action="phase2-up"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0"
phase2_name=dialup_p2

2010-01-11 16:39:58 log_id=0101037138 type=event subtype=ipsec pri=notice
vd="root" msg="IPsec connection status change" action="tunnel-up"
rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500
out_intf="port1" cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A"
group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0"
tunnel_ip=172.20.120.151 tunnel_id=1552003005 tunnel_type=ipsec duration=0 sent=0
rcvd=0 next_stat=0 tunnel=dialup_p1_0

2010-01-11 16:39:58 log_id=0101037129 type=event subtype=ipsec pri=notice
vd="root" msg="progress IPsec phase 2" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success
init=remote mode=quick dir=inbound stage=2 role=responder result=DONE

2010-01-11 16:39:58 log_id=0101037122 type=event subtype=ipsec pri=notice
vd="root" msg="negotiate IPsec phase 2" action="negotiate" rem_ip=172.20.120.151
loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1"
cookies="5f6da1c0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A"
xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success
role=responder esp_transform=ESP_3DES esp_auth=HMAC_SHA1

2010-01-11 16:39:58 log_id=0103031008 type=event subtype=ppp vd=root
pri=information action=connect status=success msg="Client 172.20.120.151 control
connection started (id 805), assigned ip 192.168.6.85"

2010-01-11 16:39:58 log_id=0103029013 type=event subtype=ppp vd=root pri=notice
pppd is started

2010-01-11 16:39:58 log_id=0103029002 type=event subtype=ppp vd=root pri=notice
user="user1" local=172.20.120.141 remote=172.20.120.151 assigned=192.168.6.85
action=auth_success msg="User 'user1' using l2tp with authentication protocol
MSCHAP_V2, succeeded"

2010-01-11 16:39:58 log_id=0103031101 type=event subtype=ppp vd=root
pri=information action=tunnel-up tunnel_id=1645784497 tunnel_type=l2tp
remote_ip=172.20.120.151 tunnel_ip=192.168.6.85 user="user1" group="L2TPusers"
msg="L2TP tunnel established"



GRE over IPsec (Cisco VPN) configurations

This section describes how to configure a FortiGate VPN that is compatible with Cisco-style VPNs that use GRE in an IPsec tunnel.

The following topics are included in this section:


- [Overview](#)
- [Configuring the FortiGate unit](#)
- [Configuring the Cisco router](#)
- [Troubleshooting](#)

Overview

Cisco products that include VPN support often use Generic Routing Encapsulation (GRE) protocol tunnel over IPsec encryption. This chapter describes how to configure a FortiGate unit to work with this type of Cisco VPN.

Cisco VPNs can use either transport mode or tunnel mode IPsec. Before FortiOS 4.0 MR2, the FortiGate unit was compatible only with tunnel mode IPsec.

Figure 156: Example FortiGate to Cisco GRE-over-IPsec VPN



In this example, users on LAN-1 are provided access to LAN-2.

Configuring the FortiGate unit

There are several steps to the GRE-over-IPsec configuration:

- Enable overlapping subnets. This is needed because the IPsec and GRE tunnels will use the same addresses.
- Configure a route-based IPsec VPN on the external interface.
- Configure a GRE tunnel on the virtual IPsec interface. Set its local gateway and remote gateway addresses to match the local and remote gateways of the IPsec tunnel.
- Configure security policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- Configure security policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.
- Configure a static route to direct traffic destined for the network behind the Cisco router into the GRE-over-IPsec tunnel.

Enabling overlapping subnets

By default, each FortiGate unit network interface must be on a separate network. The configuration described in this chapter assigns an IPsec tunnel end point and the external interface to the same network. Enable subnet overlap as follows:

```
config system settings
    set allow-subnet-overlap enable
end
```

Configuring the IPsec VPN

A route-based VPN is required. It must use encryption and authentication algorithms compatible with the Cisco equipment to which it connects. In this chapter, preshared key authentication is shown.

To configure the IPsec VPN - web-based manager

- 1 Define the phase 1 configuration needed to establish a secure connection with the remote Cisco device. Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel, tocsico for example. This is the name of the virtual IPsec interface. It appears in phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Enter the IP address of the Cisco device public interface. For example, 192.168.5.113
Local Interface	Select the FortiGate unit's public interface. For example, 172.20.120.141
Mode	Select <i>Main (ID Protection)</i> .
Authentication Method	<i>Preshared Key</i>
Pre-shared Key	Enter the preshared key. It must match the preshared key on the Cisco device.
Advanced	Select the Advanced button to see the following settings.

Enable IPsec Interface Mode	Enable.
P1 Proposal	3DES-MD5 At least one proposal must match the settings on the Cisco unit.

For more information about these settings, see [“Auto Key phase 1 parameters” on page 1407](#).

- 2 Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. For compatibility with the Cisco router, Quick Mode Selectors must be entered, which includes specifying protocol 47, the GRE protocol. Enter these settings in particular:

Name	Enter a name to identify this phase 2 configuration.
Phase 1	Select the name of the phase 1 configuration that you defined in Step 1.
Advanced	Select <i>Advanced</i> to view the following fields.
P2 Proposal	3DES-MD5 At least one proposal must match the settings on the Cisco unit.
Quick Mode Selector	
Source Address	Enter the GRE local tunnel end IP address. For example 172.20.120.141
Source Port	0
Destination Address	Enter the GRE remote tunnel end IP address. For example 192.168.5.113
Destination Port	0
Protocol	47

For more information about these settings, see [“Phase 2 parameters” on page 1425](#).

- 3 If the Cisco device is configured to use transport mode IPsec, you need to use transport mode on the FortiGate VPN. You can configure this only in the CLI. In your phase 2 configuration, set `encapsulation` to `transport-mode` (default is `tunnel-mode`) as follows:

```
config vpn phase2-interface
  edit to_cisco_p2
    set encapsulation transport-mode
  end
```

To configure the IPsec VPN - CLI

```
config vpn ipsec phase1-interface
  edit tocisco
    set interface port1
    set proposal 3des-sha1 aes128-sha1
    set remote-gw 192.168.5.113
    set psksecret xxxxxxxxxxxxxxxxx
  end
config vpn ipsec phase2-interface
  edit tocisco_p2
    set phase1name "tocisco"
    set proposal 3des-md5
    set encapsulation tunnel-mode      // if tunnel mode
```

```
set encapsulation transport-mode // if transport mode
set protocol 47
set src-addr-type ip
set dst-start-ip 192.168.5.113
set src-start-ip 172.20.120.141
end
```

Adding IPsec tunnel end addresses

The Cisco configuration requires an address for its end of the IPsec tunnel. The addresses are set to match the GRE gateway addresses. Use the CLI to set the addresses, like this:

```
config system interface
edit tocisco
set ip 172.20.120.141 255.255.255.255
set remote-ip 192.168.5.113
end
```

Configuring the GRE tunnel

The GRE tunnel runs between the virtual IPsec public interface on the FortiGate unit and the Cisco router. You must use the CLI to configure a GRE tunnel. In the example, you would enter:

```
config system gre-tunnel
edit gre1
set interface tocisco
set local-gw 172.20.120.141
set remote-gw 192.168.5.113
end
```

`interface` is the virtual IPsec interface, `local-gw` is the FortiGate unit public IP address, and `remote-gw` is the remote Cisco device public IP address

Adding GRE tunnel end addresses

You will also need to add tunnel end addresses. The Cisco router configuration requires an address for its end of the GRE tunnel. Using the CLI, enter tunnel end addresses that are not used elsewhere on the FortiGate unit, like this:

```
config system interface
edit gre1
set ip 10.0.1.1 255.255.255.255
set remote-ip 10.0.1.2
end
```

Configuring security policies

Two sets of security policies are required:

- policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.

To configure security policies - web-based manager

- 1 Define an ACCEPT security policy to permit communications between the protected network and the GRE tunnel:

Source Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Source Address Name	All
Destination Interface/Zone	Select the GRE tunnel virtual interface you configured.
Destination Address Name	All
Action	ACCEPT
NAT	Disable

- 2 To permit the remote client to initiate communication, you need to define a security policy for communication in that direction:

Source Interface/Zone	Select the GRE tunnel virtual interface you configured.
Source Address Name	All
Destination Interface/Zone	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address Name	All
Action	ACCEPT.
NAT	Disable

- 3 Define a pair of ACCEPT security policies to permit traffic to flow between the GRE virtual interface and the IPsec virtual interface:

Source Interface/Zone	Select the GRE virtual interface. See “Configuring the GRE tunnel” on page 1582 .
Source Address Name	All
Destination Interface/Zone	Select the virtual IPsec interface you created. See “Configuring the IPsec VPN” on page 1580 .
Destination Address Name	All
Action	ACCEPT.
NAT	Disable

Source Interface/Zone	Select the virtual IPsec interface you created. See “Configuring the IPsec VPN” on page 1580 .
Source Address Name	All

Destination Interface/Zone	Select the GRE virtual interface. See “Configuring the GRE tunnel” on page 1582 .
Destination Address Name	All
Action	Select <i>ACCEPT</i> .
NAT	Disable.

To configure security policies - CLI

```

config firewall policy
  edit 1                                // LAN to GRE tunnel
    set srcintf port2
    set dstintf gre1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
  edit 2                                // GRE tunnel to LAN
    set srcintf gre1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
  edit 3                                // GRE tunnel to IPsec interface
    set srcintf "gre1"
    set dstintf "tocisco"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
  next
  edit 4                                // IPsec interface to GRE tunnel
    set srcintf "tocisco"
    set dstintf "gre1"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
end

```

Configuring routing

Traffic destined for the network behind the Cisco router must be routed to the GRE tunnel. To do this, create a static route as follows:

Destination IP/Mask	Enter the IP address and netmask for the network behind the Cisco router. For example 10.21.101.0 255.255.255.0
Device	Select the GRE virtual interface.
Distance	Leave setting at default value.

In the CLI, using the example values, you would enter

```
config router static
edit 0
set device gre1
set dst 10.21.101.0 255.255.255.0
end
```

Configuring the Cisco router

Using Cisco IOS, you would configure the Cisco router as follows, using the addresses from the example:

```
config ter
crypto ipsec transform-set myset esp-3des esp-md5-hmac
no mode
exit
no ip access-list extended tunnel
ip access-list extended tunnel
permit gre host 192.168.5.113 host 172.20.120.141
exit
interface Tunnel1
ip address 10.0.1.2 255.255.255.0
tunnel source 192.168.5.113
tunnel destination 172.20.120.141
!
ip route 10.11.101.0 255.255.255.0 Tunnel1
end
clea crypto sa
clea crypto isakmp
```

For transport mode, change `no mode` to `mode transport`.

This is only the portion of the Cisco router configuration that applies to the GRE-over-IPsec tunnel. For more information, refer to the Cisco documentation.

Troubleshooting

This section describes some checks and tools you can use to resolve issues with the GRE-over-IPsec VPN.

Quick checks

Here is a list of common problems and what to verify.

Problem	What to check
No communication with remote network.	Use the <code>execute ping</code> command to ping the Cisco device public interface. Use the FortiGate VPN Monitor page to see whether the IPsec tunnel is up or can be brought up.
IPsec tunnel does not come up.	Check the logs to determine whether the failure is in Phase 1 or Phase 2. Check that the encryption and authentication settings match those on the Cisco device. Check the encapsulation setting: tunnel-mode or transport-mode. Both devices must use the same mode.
Tunnel connects, but there is no communication.	Check the security policies. See “Configuring security policies” on page 1582 . Check routing. See “Configuring routing” on page 1585 .

Setting up logging

To configure FortiGate logging for IPsec

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select the *Event Logging*.
- 3 Select *IPsec negotiation event*.
- 4 Select *Apply*.

To configure FortiGate logging level

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select *Local Logging & Archiving*.
- 3 Set *Minimum log level* to *Information*.
- 4 Optionally, enable and configure disk logging.
- 5 Select *Apply*.

To view FortiGate logs

- 1 Go to *Log & Report > Log & Archive Access > Event Log*.
- 2 Select the log storage type.
- 3 Select *Refresh* to view any logged events.

Understanding the log messages

Successful startup of an IPsec VPN follows a well-defined sequence. If you compare your logs to the sequence shown in [Table 92](#), you will be able to determine at what stage your configuration is failing.

Table 92: Typical sequence of log messages for IPsec VPN connection startup

	ID	Sub Type	Action	Message
1	37127	ipsec	negotiate	progress IPsec phase 1
2	37127	ipsec	negotiate	progress IPsec phase 1
3	37127	ipsec	negotiate	progress IPsec phase 1
4	37127	ipsec	negotiate	progress IPsec phase 1
5	37129	ipsec	negotiate	progress IPsec phase 2
6	37133	ipsec	install_sa	install IPsec SA
7	37139	ipsec	phase2-up	IPsec phase 2 status change
8	37138	ipsec	tunnel-up	IPsec connection status change



This table lists messages in top down chronological order. In the web-based manager log viewer, you need to read the messages from the bottom up. The newest message appears at the top of that list.

Using diagnostic commands

There are some diagnostic commands that can provide useful information.



When using diagnostic commands, it is best practice that you connect to the CLI using a terminal program, such as puTTY, that allows you to save output to a file. This will allow you to review the data later on at your own speed without worry about missed data as the diag output scrolls by.

To use the packet sniffer

- 1 Enter the following CLI command:

```
diag sniff packet any icmp 4
```
- 2 Ping an address on the network behind the FortiGate unit from the network behind the Cisco router.

The output will show packets coming in from the GRE interface going out of the interface that connects to the protected network (LAN) and vice versa. For example:

```
114.124303 gre1 in 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124367 port2 out 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124466 port2 in 10.11.101.10 -> 10.0.1.2: icmp: echo reply
114.124476 gre1 out 10.11.101.10 -> 10.0.1.2: icmp: echo reply
```


- 3 Enter CTRL-C to stop the sniffer.

To view debug output for IKE

- 1 Enter the following CLI commands

```
diagnose debug application ike -1
diagnose debug enable
```
- 2 Attempt to use the VPN or set up the VPN tunnel and note the debug output.

- 3 Enter CTRL-C to stop the debug output.
- 4 Enter the following command to reset debug settings to default:
`diagnose debug reset`



Protecting OSPF with IPsec

For enhanced security, OSPF dynamic routing can be carried over IPsec VPN links.


The following topics are included in this section:

- [Overview](#)
- [OSPF over IPsec configuration](#)
- [Creating a redundant configuration](#)

Overview

This chapter shows an example of OSPF routing conducted over an IPsec tunnel between two FortiGate units. The network shown in [Figure 157](#) is a single OSPF area. FortiGate_1 is an Area border router that advertises a static route to 10.22.10.0/24 in OSPF. FortiGate_2 advertises its local LAN as an OSPF internal route.

Figure 157: OSPF over an IPsec VPN tunnel



The section [“OSPF over IPsec configuration”](#) describes the configuration with only one IPsec VPN tunnel, tunnel_wan1. Then, the section [“Creating a redundant configuration”](#) on [page 1595](#) describes how you can add a second tunnel to provide a redundant backup path. This is shown in [Figure 157](#) as VPN tunnel “tunnel_wan2”.

Only the parts of the configuration concerned with creating the IPsec tunnel and integrating it into the OSPF network are described. It is assumed that security policies are already in place to allow traffic to flow between the interfaces on each FortiGate unit.

OSPF over IPsec configuration

There are several steps to the OSPF-over-IPsec configuration:

- Configure a route-based IPsec VPN on an external interface. It will connect to a corresponding interface on the other FortiGate unit. Define the two tunnel-end addresses.
- Configure a static route to the other FortiGate unit.
- Configure the tunnel network as part of the OSPF network and define the virtual IPsec interface as an OSPF interface.

This section describes the configuration with only one VPN, tunnel_wan1. The other VPN is added in the section [“Creating a redundant configuration” on page 1595](#).

Configuring the IPsec VPN

A route-based VPN is required. In this chapter, preshared key authentication is shown. Certificate authentication is also possible. Both FortiGate units need this configuration.

To configure Phase 1

- 1 Define the phase 1 configuration needed to establish a secure connection with the other FortiGate unit. For more information, see [“Auto Key phase 1 parameters” on page 1407](#). Enter these settings in particular

Name	Enter a name to identify the VPN tunnel, tunnel_wan1 for example. This becomes the name of the virtual IPsec interface.
Remote Gateway	Select <i>Static IP Address</i> .
IP Address	Enter the IP address of the other FortiGate unit's public (Port 2) interface.
Local Interface	Select this FortiGate unit's public (Port 2) interface.
Mode	Select <i>Main (ID Protection)</i> .
Authentication Method	<i>Preshared Key</i>
Pre-shared Key	Enter the preshared key. It must match the preshared key on the other FortiGate unit.
Advanced	Select <i>Advanced</i> .
Enable IPsec Interface Mode	Enable

To assign the tunnel end IP addresses

- 1 Go to *System > Network > Interface*, select the virtual IPsec interface that you just created on Port 2 and select *Edit*.
- 2 In the *IP* and *Remote IP* fields, enter the following tunnel end addresses:

	FortiGate_1	FortiGate_2
IP	10.1.1.1	10.1.1.2
Remote_IP	10.1.1.2	10.1.1.1

These addresses are from a network that is not used for anything else.

To configure Phase 2

- 1 Define the phase 2 parameters needed to create a VPN tunnel with the remote peer. For more information, see “Phase 2 parameters” on page 1425. Enter these settings in particular:

Name	Enter a name to identify this phase 2 configuration, twan1_p2, for example.
Phase 1	Select the name of the phase 1 configuration that you defined in Step 1, tunnel_wan1 for example.

Configuring static routing

You need to define the route for traffic leaving the external interface. Go to *Router > Static > Static Route*, select *Create New*, and enter the following information.

Destination IP/Mask	Leave as 0.0.0.0 0.0.0.0.
Device	Select the external interface.
Gateway	Enter the IP address of the next hop router.
Distance	Leave setting at default value.

Configuring OSPF

This section does not attempt to explain OSPF router configuration. It focusses on the integration of the IPsec tunnel into the OSPF network. This is accomplished by assigning the tunnel as an OSPF interface, creating an OSPF route to the other FortiGate unit.

This configuration uses loopback interfaces to ease OSPF troubleshooting. The OSPF router ID is set to the loopback interface address. The loopback interface ensures the router is always up. Even though technically the router ID doesn't have to match a valid IP address on the FortiGate unit, having an IP that matches the router ID makes troubleshooting a lot easier.

The two FortiGate units have slightly different configurations. FortiGate_1 is an AS border router that advertises its static default route. FortiGate_2 advertises its local LAN as an OSPF internal route.

Setting the router ID for each FortiGate unit to the lowest possible value is useful if you want the FortiGate units to be the designated router (DR) for their respective ASes. This is the router that broadcasts the updates for the AS.

Leaving the IP address on the OSPF interface at 0.0.0.0 indicates that all potential routes will be advertised, and it will not be limited to any specific subnet. For example if this IP address was 10.1.0.0, then only routes that match that subnet will be advertised through this interface in OSPF.

FortiGate_1 OSPF configuration

When configuring FortiGate_1 for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

To create the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.1, you would enter:

```
config system interface
  edit lback1
    set vdom root
    set ip 10.0.0.1 255.255.255.255
    set type loopback
  end
```

The loopback addresses and corresponding router IDs on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

To configure OSPF area, networks, and interfaces - web-based manager

- 1 On FortiGate_1, go to *Router > Dynamic > OSPF* and enter the following information to define the router, area, and interface information.

Router ID	Enter 10.0.0.1. Select <i>Apply</i> before entering the remaining information.
Advanced Options	
Redistribute	Select the <i>Connected</i> and <i>Static</i> check boxes. Use their default metric values.
Areas	Select <i>Create New</i> , enter the <i>Area</i> and <i>Type</i> and then select <i>OK</i> .
Area	0.0.0.0
Type	Regular
Interfaces	
Name	Enter a name for the OSPF interface, ospf_wan1 for example.
Interface	Select the virtual IPsec interface, tunnel_wan1.
IP	0.0.0.0

- 2 For *Networks*, select *Create New*.
- 3 Enter the following information.

IP/Netmask	10.1.1.0/255.255.255.0.
Area	0.0.0.0

- 4 For *Networks*, select *Create New*.
- 5 Enter the following information:

IP/Netmask	10.0.0.1/255.255.255.255
Area	0.0.0.0

- 6 Select *Apply*.

To configure OSPF area and interfaces - CLI

Your loopback interface is 10.0.0.1, your tunnel ends are on the 10.1.1.0/24 network, and your virtual IPsec interface is named `tunnel_wan1`. Enter the following CLI commands:

```
config router ospf
  set router-id 10.0.0.1
  config area
    edit 0.0.0.0
  end
  config network
    edit 4
      set prefix 10.1.1.0 255.255.255.0
    next
    edit 2
      set prefix 10.0.0.1 255.255.255.255
    end
  config ospf-interface
    edit ospf_wan1
      set cost 10
      set interface tunnel_wan1
      set network-type point-to-point
    end
  config redistribute connected
    set status enable
  end
  config redistribute static
    set status enable
  end
end
```

FortiGate_2 OSPF configuration

When configuring FortiGate_2 for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

Configuring FortiGate_2 differs from FortiGate_1 in that three interfaces are defined instead of two. The third interface is the local LAN that will be advertised into OSPF.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

To create the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.2, you would enter:

```
config system interface
  edit lback1
    set vdom root
    set ip 10.0.0.2 255.255.255.255
    set type loopback
  end
```

The loopback addresses on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

To configure OSPF area and interfaces - web-based manager

- 1 On FortiGate_2, go to *Router > Dynamic > OSPF*.
- 2 For *Router ID*, enter 10.0.0.2.

Router ID	10.0.0.2
Areas	Select <i>Create New</i> , enter the <i>Area</i> and <i>Type</i> and then select <i>OK</i> .
Area	0.0.0.0
Type	Regular
Interfaces	
Name	Enter a name for the OSPF interface, ospf_wan1 for example.
Interface	Select the virtual IPsec interface, tunnel_wan1
IP	0.0.0.0

- 3 For *Networks*, select *Create New*.
- 4 Enter the following information for the loopback interface:

IP/Netmask	10.0.0.2/255.255.255.255
Area	0.0.0.0

- 5 For *Networks*, select *Create New*.
- 6 Enter the following information for the tunnel interface:

IP/Netmask	10.1.1.0/255.255.255.0
Area	0.0.0.0

- 7 For *Networks*, select *Create New*.
- 8 Enter the following information for the local LAN interface:

IP/Netmask	10.31.101.0/255.255.255.0
Area	0.0.0.0

- 9 Select *Apply*.

To configure OSPF area and interfaces - CLI

If for example, your loopback interface is 10.0.0.2, your tunnel ends are on the 10.1.1.0/24 network, your local LAN is 10.31.101.0/24, and your virtual IPsec interface is named tunnel_wan1, you would enter:

```
config router ospf
  set router-id 10.0.0.2
  config area
    edit 0.0.0.0
  end
  config network
    edit 1
      set prefix 10.1.1.0 255.255.255.0
    next
    edit 2
      set prefix 10.31.101.0 255.255.255.0
```



```
next
edit 2
set prefix 10.0.0.2 255.255.255.255
end
config ospf-interface
edit ospf_wan1
set interface tunnel_wan1
set network-type point-to-point
end
end
```

Creating a redundant configuration

You can improve the reliability of the OSPF over IPsec configuration described in the previous section by adding a second IPsec tunnel to use if the default one goes down. Redundancy in this case is not controlled by the IPsec VPN configuration but by the OSPF routing protocol.

To do this you:

- Create a second route-based IPsec tunnel on a different interface and define tunnel end addresses for it.
- Add the tunnel network as part of the OSPF network and define the virtual IPsec interface as an additional OSPF interface.
- Set the OSPF cost for the added OSPF interface to be significantly higher than the cost of the default route.

Adding the second IPsec tunnel

The configuration is the same as in [“Configuring the IPsec VPN” on page 1590](#), but the interface and addresses will be different. Ideally, the network interface you use is connected to a different Internet service provider for added redundancy.

When adding the second tunnel to the OSPF network, choose another unused subnet for the tunnel ends, 10.1.2.1 and 10.1.2.2 for example.

Adding the OSPF interface

OSPF uses the metric called cost when determining the best route, with lower costs being preferred. Up to now in this example, only the default cost of 10 has been used. Cost can be set only in the CLI.

The new IPsec tunnel will have its OSPF cost set higher than that of the default tunnel to ensure that it is only used if the first tunnel goes down. The new tunnel could be set to a cost of 200 compared to the default cost is 10. Such a large difference in cost will ensure this new tunnel will only be used as a last resort.

If the new tunnel is called tunnel_wan2, you would enter the following on both FortiGate units:

```
config router ospf
config ospf-interface
edit ospf_wan2
set cost 200
set interface tunnel_wan2
set network-type point-to-point
end
end
```




Hardware offloading and acceleration

FortiGate units incorporate proprietary FortiASIC NP2 network processors that can provide accelerated processing for IPsec VPN traffic. This section describes how to configure offloading and acceleration.

The following topics are included in this section:

- [Overview](#)
- [IPsec offloading configuration examples](#)

Overview

Fortinet's NP2 network processors contain features to improve IPsec tunnel performance. For example, network processors can encrypt and decrypt packets, offloading cryptographic work from the FortiGate unit's main processing resources.

On FortiGate units with the appropriate hardware, you can configure offloading of both IPsec sessions and HMAC checking.

IPsec session offloading requirements

Sessions must be fast path ready. Fast path ready session requirements are:

- Layer 2 type/length must be 0x0800 (IEEE 802.1q VLAN specification is supported); link aggregation between any network interfaces sharing the same network processor(s) may be used (IEEE 802.3ad specification is supported)
- Layer 3 protocol must be IPv4
- Layer 4 protocol must be UDP, TCP or ICMP
- Layer 3 / Layer 4 header or content modification must not require a session helper (for example, SNAT, DNAT, and TTL reduction are supported, but application layer content modification is not supported)
- FortiGate unit security policy must not require antivirus or IPS inspection, although hardware accelerated anomaly checks are acceptable.
- The session must not use an aggregated link or require QoS, including rate limits and bandwidth guarantees (NP1 processor only).
- Ingress and egress network interfaces are both attached to the same network processor(s)
- In Phase I configuration, Local Gateway IP must be specified as an IP address of a network interface attached to a network processor

- In Phase II configuration:
 - encryption algorithm must be DES, 3DES, AES-128, AES-192, AES-256, or null (for NP1 processor, only 3DES is supported)
 - authentication must be MD5, SHA1, or null (for NP1 processor, only MD5 is supported)
 - if replay detection is enabled, encryption and decryption options must be enabled in the CLI (see “[IPsec encryption offloading](#)”, below)

If the IPsec session meets the above requirements, the FortiGate unit sends the IPsec security association (SA) and configured processing actions to the network processors.

Packet offloading requirements

In addition to the session requirements, the packets themselves must meet fast-path requirements:

- Incoming packets must not be fragmented.
- Outgoing packets must be 385 bytes or larger after any fragmentation. This means the configured MTU (Maximum Transmission Unit) for the network processors’ interfaces must have an MTU of 385 bytes or larger.

If packet offloading requirements are not met, an individual packet will use the FortiGate unit main processing resources, regardless of whether other packets in the session are offloaded to the specialized network processors.

IPsec encryption offloading

Network processing unit (NPU) settings configure offloading behavior for IPsec VPNs. Configured behavior applies to all network processors contained by the FortiGate unit itself or any installed AMC modules.

If replay detection is not enabled (IPsec Phase 2 settings), encryption is always offloaded.

To enable offloading of encryption even when replay detection is enabled

```
config system npu
    set enc-offload-antireplay enable
end
```

To enable offloading of decryption even when replay detection is enabled

```
config system npu
    set dec-offload-antireplay enable
end
```

HMAC check offloading

The Hash-based Message Authentication Code (HMAC) check can also be offloaded to hardware. SHA-256, SHA-384, or SHA-512 cannot be off-loaded to hardware, and must be processed using only software resources.


To enable HMAC check offloading

```
configure system global
    set ipsec-hmac-offload (enable|disable)
end
```

IPsec offloading configuration examples

The following examples configure two FortiASIC NP2 accelerated VPNs, one route-based, the other policy based. In both cases, the network topology is as shown in Figure 158.

Figure 158: Hardware accelerated IPsec VPN topology



Accelerated route-based VPN configuration

This example uses the accelerated ports on FortiGate-ASM-FB4 modules in each FortiGate unit. These accelerated ports on the modules are paired interfaces that have their own network processor (NPU) to offload work from the FortiGate unit CPU. Beyond this fact, the example is normal VPN example.

Configuring the FortiGate units require the same basic steps:

- Configure VPN Phase 1
- Configure VPN Phase 2
- Create security policies to allow traffic to flow
- Create a static route to allow traffic to flow

When both FortiGates are have the VPN tunnel configured, test to ensure it is working properly.

To configure FortiGate_1

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
- 2 Configure Phase 1 settings (name *FGT_1_IPsec*), plus
 - Select *Advanced*.
 - Select *Enable IPsec Interface Mode*.
 - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1's FortiGate-ASM-FB4 module on port 2.
- 3 Select *OK*.
- 4 Select *Create Phase 2* and configure Phase 2 settings, including
 - Select *Enable replay detection*.
 - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.

- 5 Go to *Policy > Policy > Policy*.
- 6 Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.
- 7 Go to *Router > Static > Static Route*.
- 8 Configure a static route to route traffic destined for FortiGate_2's protected network to the virtual IPsec interface, FGT_1_IPsec.

To add the static route from the CLI:

```
config router static
edit 2
set device "FGT_1_IPsec"
set dst 2.2.2.0 255.255.255.0
end
```

To configure FortiGate_2

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
- 2 Configure Phase 1 settings (name FGT_2_IPsec), plus
 - Select *Advanced*.
 - Select *Enable IPsec Interface Mode*.
 - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2's FortiGate-ASM-FB4 module on port 2.
- 3 Select *OK*.
- 4 Select *Create Phase 2* and configure Phase 2 settings, including
 - Select *Enable replay detection*.
 - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.
- 5 Go to *Policy > Policy > Policy*.
- 6 Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.
- 7 Go to *Router > Static > Static Route*.
- 8 Configure a static route to route traffic destined for FortiGate_1's protected network to the virtual IPsec interface, FGT_2_IPsec.

To add the static route from the CLI:

```
config router static
edit 2
set device "FGT_2_IPsec"
set dst 1.1.1.0 255.255.255.0
end
```

To test the VPN

- 1 Activate the IPsec tunnel by sending traffic between the two protected networks.
- 2 To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.

Accelerated policy-based VPN configuration

To configure FortiGate_1

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
- 2 Configure Phase 1 settings (name FGT_1_IPsec), plus
 - Select *Advanced*.
 - Ensure that the *Enable IPsec Interface Mode* check box is not selected.
 - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1's FortiGate-ASM-FB4 module on port 2.
- 3 Select *OK*.
- 4 Select *Create Phase 2* and configure Phase 2 settings, including
 - Select *Enable replay detection*.
 - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.
- 5 Go to *Policy > Policy > Policy*.
- 6 Configure an IPSEC policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.
- 7 Go to *Router > Static > Static Route*.
- 8 Configure a static route to route traffic destined for FortiGate_2's protected network to FortiGate_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

To add the static route from the CLI:

```
config router static
  edit 0
    set device "AMC-SW1/2"
    set dst 2.2.2.0 255.255.255.0
    set gateway 3.3.3.1
  end
```

To configure FortiGate_2

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and select *Create Phase 1*.
- 2 Configure Phase 1 settings (name FGT_2_IPsec), plus
 - Select *Advanced*.
 - Select *Enable IPsec Interface Mode*.
 - In *Local Gateway IP*, select *Specify* and enter the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2's FortiGate-ASM-FB4 module on port 2.
- 3 Select *OK*.
- 4 Select *Create Phase 2* and configure Phase 2 settings, including
 - Select *Enable replay detection*.
 - set `enc-offload-antireplay` to enable using the `config system npu` CLI command.
- 5 Go to *Policy > Policy > Policy*.
- 6 Configure an IPSEC policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.
- 7 Go to *Router > Static > Static Route*.

- 8 Configure a static route to route traffic destined for FortiGate_1's protected network to FortiGate_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

To add the static route from the CLI:

```
config router static
edit 0
set device "AMC-SW1/2"
set dst 1.1.1.0 255.255.255.0
set gateway 3.3.3.2
end
```

To test the VPN

- 1 Activate the IPsec tunnel by sending traffic between the two protected networks.
- 2 To verify tunnel activation, go to *VPN > Monitor > IPsec Monitor*.



Monitoring and troubleshooting

This section provides some general maintenance and monitoring procedures for VPNs.

The following topics are included in this section:

- [Monitoring VPN connections](#)
- [Testing VPN connections](#)
- [Testing VPN connections](#)
- [Logging VPN events](#)
- [VPN troubleshooting tips](#)

Monitoring VPN connections

You can use the monitor to view activity on IPsec VPN tunnels and to start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels. See [“IPsec Monitor” on page 1405](#).

Monitoring connections to remote peers

The list of tunnels provides information about VPN connections to remote peers that have static IP addresses or domain names. You can use this list to view status and IP addressing information for each tunnel configuration. You can also start and stop individual tunnels from the list.

To view the list of static-IP and dynamic-DNS tunnels go to *VPN > Monitor > IPsec Monitor*.

Monitoring dialup IPsec connections

The list of dialup tunnels provides information about the status of tunnels that have been established for dialup clients. The list displays the IP addresses of dialup clients and the names of all active tunnels. The number of tunnels shown in the list can change as dialup clients connect and disconnect.

To view the list of dialup tunnels go to *VPN > Monitor > IPsec Monitor*.



If you take down an active tunnel while a dialup client such as FortiClient is still connected, FortiClient will continue to show the tunnel connected and idle. The dialup client must disconnect before another tunnel can be initiated.

The list of dialup tunnels displays the following statistics:

- The Name column displays the name of the tunnel.
- The meaning of the value in the Remote gateway column changes, depending on the configuration of the network at the far end:

- When a FortiClient dialup client establishes a tunnel, the Remote gateway column displays either the public IP address and UDP port of the remote host device (on which the FortiClient Endpoint Security application is installed), or if a NAT device exists in front of the remote host, the Remote gateway column displays the public IP address and UDP port of the remote host.
- When a FortiGate dialup client establishes a tunnel, the Remote gateway column displays the public IP address and UDP port of the FortiGate dialup client.
- The Username column displays the peer ID, certificate name, or XAuth user name of the dialup client (if a peer ID, certificate name, or XAuth user name was assigned to the dialup client for authentication purposes).
- The Timeout column displays the time before the next key exchange. The time is calculated by subtracting the time elapsed since the last key exchange from the keylife.
- The Proxy ID Source column displays the IP addresses of the hosts, servers, or private networks behind the FortiGate unit. A network range may be displayed if the source address in the security encryption policy was expressed as a range of IP addresses.
- The meaning of the value in the Proxy ID Destination column changes, depending on the configuration of the network at the far end:
 - When a FortiClient dialup client establishes a tunnel:
 - If VIP addresses are not used and the remote host connects to the Internet directly, the Proxy ID Destination field displays the public IP address of the Network Interface Card (NIC) in the remote host.
 - If VIP addresses are not used and the remote host is behind a NAT device, the Proxy ID Destination field displays the private IP address of the NIC in the remote host.
 - If VIP addresses were configured (manually or through FortiGate DHCP relay), the Proxy ID Destination field displays either the VIP address belonging to a FortiClient dialup client, or a subnet address from which VIP addresses were assigned.
 - When a FortiGate dialup client establishes a tunnel, the Proxy ID Destination field displays the IP address of the remote private network.

Testing VPN connections

A VPN connection has multiple stages that can be confirmed to ensure the connection is working properly. It is easiest to see if the final stage is successful first since if it is successful the other stages will be working properly. Otherwise, you will need to work back through the stages to see where the problem is located.

- [Testing VPN connection](#)
- [Troubleshooting VPN connections](#)

Testing VPN connection

When a VPN connection is properly established, traffic will flow from one end to the other as if both ends were physically in the same place. If you can determine the connection is working properly then any problems are likely problems with your applications.

If the connection is not working properly, you can move on to [“Troubleshooting VPN connections” on page 1605](#) to determine the exact problem.

LAN interface connection

To confirm whether a VPN connection over LAN interfaces has been configured correctly, issue a ping or traceroute command on the network behind the FortiGate unit to test the connection to a computer on the remote network. If the connection is properly configured, a VPN tunnel will be established automatically when the first data packet destined for the remote network is intercepted by the FortiGate unit.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel. You can confirm this by going to *VPN > Monitor > IPsec Monitor* where you will be able to see your connection. A green arrow means the tunnel is up and currently processing traffic. A red arrow means the tunnel is not processing traffic, and this VPN connection has a problem.

If the connection has problems, see [“Troubleshooting VPN connections” on page 1605](#).

Dialup connection

A dialup VPN connection has additional steps. To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel, or dialup client. As with the LAN connection, confirm the VPN tunnel is established by checking *VPN > Monitor > IPsec Monitor*.

If the connection has problems, see [Troubleshooting](#).

Troubleshooting VPN connections

If you have determined that your VPN connection is not working properly through [“Testing VPN connection” on page 1604](#), the next step is to verify that you have a phase2 connection.



If traffic is not passing through the FortiGate unit as you expect, ensure the traffic does not contain IPcomp packets (IP protocol 108, RFC 3173). FortiGate units do not allow IPcomp packets, they compress packet payload, preventing it from being scanned.

Testing phase 1 and 2 connections is a bit more difficult than testing the working VPN. This is because they require diagnose CLI commands. These commands are typically used by Fortinet customer support to discover more information about your FortiGate unit and its current configuration.

Before you start troubleshooting you need to:

- configure FortiGate units on both ends for interface VPN
- record the information in your VPN phase 1 and phase 2 configurations - for our example here the remote IP address is 10.101.101.101 and the names of the phases are Phase1 and Phase2
- install a telnet or SSH client such as putty that allows logging of output
- ensure that the admin interface supports your chosen connection protocol so you can connect to your FortiGate unit admin interface.
- For this example, default values were used unless stated otherwise.

To get diagnose information for the VPN connection - CLI

- 1 Log into the CLI as admin with the output being logged to a file.
- 2 Stop any diag debug sessions that are currently running with the CLI command
`diag debug disable`
- 3 Clear any existing log-filters by running
`diag debug log-filter clear`
- 4 Set the log-filter to the IP address of the remote computer (10.11.101.10). This filters out all VPN connections except ones to the IP address we are concerned with. The command is
`diag debug log-filter dst-addr4 10.11.101.10.`
- 5 Set up the commands to output the VPN handshaking. The commands are:
`diag debug app ike 255`
`diag debug enable`
- 6 Have the remote FortiGate initiate the VPN connection in the web-based manager by going to *VPN > Monitor* and selecting *Bring up*.
This makes the remote FortiGate the initiator and the local FortiGate becomes the responder. Establishing the connection in this manner means the local FortiGate will have its configuration information as well as the information the remote computer sends. Having both sets of information locally makes it easier to troubleshoot your VPN connection.
- 7 Watch the screen for output, and after roughly 15 seconds enter the following CLI command to stop the output.
`diag debug disable`
- 8 If needed, save the log file of this output to a file on your local computer.
Saving the output to a file can make it easier to search for a particular phrase, or save the output for future comparison with different output.

To troubleshoot a phase1 VPN connection

Using the output from [“To get diagnose information for the VPN connection - CLI” on page 1606](#), search for the word `proposal` in the output. It may occur once indicating a successful connection, or it will occur two or more times for an unsuccessful connection — there will be one proposal listed for each end of the tunnel and each possible combination in their settings. For example if 10.11.101.10 selected both DH Group 1 and 5, that would be at least 2 proposals set.

A successful negotiation proposal will look similar to

- XXX insert output sample here XXX

Note the phrases “initiator: main mode is sending 1st message...” and “” which show you the handshake between the ends of the tunnel is in progress. Initiator shows the remote unit is sending the first message.



In the proposal you will see the Diffie-Hellman Group (DH Group) listed as OAKLEY and vn=<a number>. This is normal.

Logging VPN events

You can configure the FortiGate unit to log VPN events. For IPsec VPNs, phase 1 and phase 2 authentication and encryption events are logged. For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

To log VPN events

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Enable the storage of log messages to one or more of the following locations:
 - a FortiLog unit
 - the FortiGate system memory
 - a remote computer running a syslog server



If available on your FortiGate unit, you can enable the storage of log messages to a system hard disk. In addition, as an alternative to the options listed above, you may choose to forward log messages to a remote computer running a WebTrends firewall reporting server. For more information about enabling either of these options through CLI commands, see the “log” chapter of the [FortiGate CLI Reference](#).

- 3 If the options are concealed, select the blue arrow beside each option to reveal and configure associated settings.
- 4 If logs will be written to system memory, from the Log Level list, select Information. For more information, see the [Logging and Reporting](#) chapter of *The Handbook*.
- 5 Select *Apply*.

To filter VPN events

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Verify that the *IPsec negotiation event* option is selected.
- 3 Select *Apply*.

To view event logs

- 1 Go to *Log&Report > Log & Archive Access > Event Log*.
- 2 If the option is available from the *Type* list, select the log file from disk or memory.

Entries similar to the following indicate that a tunnel has been established. The following log messages concern a VPN tunnel called `vpn_test` on `port2` interface in the root VDOM. Pay attention to the `status` and `msg` values.

```
2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/0000000000000000 action=negotiate
init=local mode=main stage=1 dir=outbound status=success
msg="Initiator: sent 172.16.62.11 main mode message #1 (OK)"

2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
init=local mode=main stage=2 dir=outbound status=success
msg="Initiator: sent 172.16.62.11 main mode message #2 (OK)"
```

```
2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
init=local mode=main stage=3 dir=outbound status=success
msg="Initiator: sent 172.16.62.11 main mode message #3 (OK)"

2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
init=local mode=main stage=3 dir=inbound status=success
msg="Initiator: parsed 172.16.62.11 main mode message #3 (DONE)"

2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
init=local mode=quick stage=1 dir=outbound status=success
msg="Initiator: sent 172.16.62.11 quick mode message #1 (OK)"

2005-03-31 15:38:29 log_id=0101023006 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=install_sa
in_spi=66867f2b out_spi=e22de275 msg="Initiator: tunnel
172.16.62.10/172.16.62.11 install ipsec sa"

2005-03-31 15:38:29 log_id=0101023004 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
init=local mode=quick stage=2 dir=outbound status=success
msg="Initiator: sent 172.16.62.11 quick mode message #2 (DONE)"

2005-03-31 15:38:29 log_id=0101023002 type=event subtype=ipsec
pri=notice vd=root loc_ip=172.16.62.10 loc_port=500
rem_ip=172.16.62.11 rem_port=500 out_if=port2 vpn_tunnel=vpn_test
cookies=151c3a5c6dd93c54/5ed26a81fb7a2d0c action=negotiate
status=success msg="Initiator: tunnel 172.16.62.11,
transform=ESP_3DES, HMAC_SHA1"
```

Entries similar to the following indicate that phase 1 negotiations broke down because the preshared keys belonging to the VPN peers were not identical. A tunnel was not established. Pay attention to the `status` and `msg` values.

```
2005-03-31 16:06:39 log_id=0101023003 type=event subtype=ipsec
pri=error vd=root loc_ip=192.168.70.2 loc_port=500
rem_ip=192.168.80.2 rem_port=500 out_if=port2 vpn_tunnel=vpn_test2
cookies=3896343ae575f210/0a7ba199149e31e9 action=negotiate
status=negotiate_error msg="Negotiate SA Error: probable pre-
shared secret mismatch"
```

For more information about how to interpret error log messages, see the [FortiGate Log Message Reference](#).

VPN troubleshooting tips

More in-depth VPN troubleshooting can be found in the [Troubleshooting handbook chapter](#).

The VPN proposal is not connecting

One side may attempt to initiate the VPN tunnel unsuccessful. There are a number of potential reasons for this problem.

Attempting hardware offloading beyond SHA1

If you are trying to off-load VPN processing to a network processing unit (NPU), remember that only SHA1 authentication is supported. For high levels of authentication such as SHA256, SHA384, and SHA512 hardware offloading is not an option — all VPN processing must be done in software.

Check Phase 1 proposal settings

Ensure that both sides have at least one Phase 1 proposal in common. Otherwise they will not connect. If there are many proposals in the list, this will slow down the negotiating of Phase 1. If its too slow, the connection may timeout before completing. If this happens, try removing some of the unused proposals.

Ensure you are not using a loopback for the local VPN interface

The Local Interface field in Phase 1 VPN configuration cannot not be configured to point to a loopback interface, or the IPSec tunnel will not be established.

Check your routing

If routing is not properly configured with an entry for the remote end of the VPN tunnel, traffic will not flow properly. You may need static routes on both ends of the tunnel. If routing is the problem, the proposal will likely setup properly but no traffic will flow.

Try enabling XAuth

If one end of an attempted VPN tunnel is using XAuth and the other end is not, the connection attempt will fail. The log messages for the attempted connection will not mention XAuth is the reason, but when connections are failing it is a good idea to ensure both ends have the same XAuth settings. If you do not know the other end's settings enable or disable XAuth on your end to see if that is the problem.

General troubleshooting tips

Most connection failures are due to a configuration mismatch between the FortiGate unit and the remote peer. In general, begin troubleshooting an IPSec VPN connection failure as follows:

- 1 Ping the remote network or client to verify whether the connection is up. See [“Testing VPN connections” on page 1604](#).
- 2 Traceroute the remote network or client. If DNS is working, you can use domain names. Otherwise use IP addresses.
- 3 Check the routing behind the dialup client. Routing problems may be affecting DHCP. If this appears to be the case, configure a DHCP relay service to enable DHCP requests to be relayed to a DHCP server on or behind the FortiGate server.
- 4 Verify the configuration of the FortiGate unit and the remote peer. Check the following IPSec parameters:

- The mode setting for ID protection (main or aggressive) on both VPN peers must be identical.
- The authentication method (preshared keys or certificates) used by the client must be supported on the FortiGate unit and configured properly.
- If preshared keys are being used for authentication purposes, both VPN peers must have identical preshared keys.
- The remote client must have at least one set of phase 1 encryption, authentication, and Diffie-Hellman settings that match corresponding settings on the FortiGate unit.
- Both VPN peers must have the same NAT traversal setting (enabled or disabled).
- The remote client must have at least one set of phase 2 encryption and authentication algorithm settings that match the corresponding settings on the FortiGate unit.
- If you are using manual keys to establish a tunnel, the *Remote SPI* setting on the FortiGate unit must be identical to the *Local SPI* setting on the remote peer, and vice versa.

5 To correct the problem, see the following table.

Table 93: VPN trouble-shooting tips

Configuration problem	Correction
Mode settings do not match.	Select complementary mode settings. See “Choosing main mode or aggressive mode” on page 1408.
Peer ID or certificate name of the remote peer or dialup client is not recognized by FortiGate VPN server.	Check Phase 1 configuration. Depending on the Remote Gateway and Authentication Method settings, you have a choice of options to authenticate FortiGate dialup clients or VPN peers by ID or certificate name (see “Authenticating remote peers and clients” on page 1412). If you are configuring authentication parameters for FortiClient dialup clients, refer to the Authenticating FortiClient Dialup Clients Technical Note .
Preshared keys do not match.	Reenter the preshared key. See “Authenticating remote peers and clients” on page 1412.
Phase 1 or phase 2 key exchange proposals are mismatched.	Make sure that both VPN peers have at least one set of proposals in common for each phase. See “Defining IKE negotiation parameters” on page 1417 and “Configure the phase 2 parameters” on page 1428.
NAT traversal settings are mismatched.	Select or clear both options as required. See “NAT traversal” on page 1420 and “NAT keepalive frequency” on page 1421.
SPI settings for manual key tunnels are mismatched.	Enter complementary SPI settings. See “Manual-key configurations” on page 1551.

A word about NAT devices

When a device with NAT capabilities is located between two VPN peers or a VPN peer and a dialup client, that device must be NAT traversal (NAT-T) compatible for encrypted traffic to pass through the NAT device. For more information, see [“NAT traversal”](#) on page 1420.



Chapter 9 SSL VPN

- [Introduction to SSL VPN](#) provides useful general information about VPN and SSL, how the FortiGate unit implements them, and gives guidance on how to choose between SSL and IPsec.
- [Basic Configuration](#) explains how to configure the FortiGate unit and the web portal. Along with these configuration details, this chapter also explains how to grant unique access permissions, configure the SSL virtual interface (`ssl.root`), and describes the SSL VPN OS Patch Check feature that allows a client with a specific OS patch to access SSL VPN services.
- [The SSL VPN client](#) provides an overview of the FortiClient software required for tunnel mode, where to obtain the software, install it and the configuration information required for remote users to connect to the internal network.
- [Setup examples](#) explores several configuration scenarios with step-by-step instructions. While the information provided is enough to set up the described SSL VPN configurations, these scenarios are not the only possible SSL VPN setups.



Introduction to SSL VPN

Over the past several years, as organizations have grown and become more complex, secure remote access to network resources has become critical for day-to-day operations. In addition, businesses are expected to provide clients with efficient, convenient services including knowledge bases and customer portals. Employees travelling across the country or around the world require timely and comprehensive access to network resources. As a result of the growing need for providing remote/mobile clients with easy, cost-effective and secure access to a multitude of resources, the concept of a Virtual Private Network was developed.

SSL VPNs establish connectivity using SSL, which functions at Levels 4 - 5 (Transport and Session). Information is encapsulated at Levels 6 - 7 (Presentation and Application), and SSL VPNs communicate at the highest levels in the OSI model. SSL is not strictly a Virtual Private Network (VPN) technology allows clients to connect to remote networks in a secure way. A VPN is a secure logical network created from physically separate networks. VPNs use encryption and other security methods to ensure that only authorized users can access the network. VPNs also ensure that the data transmitted between computers cannot be intercepted by unauthorized users. When data is encoded and transmitted over the Internet, the data is said to be sent through a "VPN tunnel". A VPN tunnel is a non-application oriented tunnel that allows the users and networks to exchange a wide range of traffic regardless of application or protocol.

The advantages of a VPN over an actual physical private network are two-fold. Rather than utilizing expensive leased lines or other infrastructure, you use the relatively inexpensive, high-bandwidth Internet. Perhaps more important though is the universal availability of the Internet - in most areas, access to the Internet is readily obtainable without any special arrangements or long wait times.

SSL (Secure Sockets Layer) as HTTPS is supported by most web browsers for exchanging sensitive information securely between a web server and a client. SSL establishes an encrypted link, ensuring that all data passed between the web server and the browser remains private and secure. SSL protection is initiated automatically when a user (client) connects to a web server that is SSL-enabled. Once the successful connection is established, the browser encrypts all the information before it leaves the computer. When the information reaches its destination, it is decrypted using a secret (private) key. Any data sent back is first encrypted, and is decrypted when it reaches the client.

SSL VPN modes of operation

When a remote client connects to the FortiGate unit, the FortiGate unit authenticates the user based on user name, password, and authentication domain. A successful login determines the access rights of remote users according to user group. The user group settings specify whether the connection will operate in web-only mode or tunnel mode.

Web-only mode

Web-only mode provides remote users with a fast and efficient way to access server applications from any thin client computer equipped with a web browser. Web-only mode offers true clientless network access using any web browser that has built-in SSL encryption and the Sun Java runtime environment.

Support for SSL VPN web-only mode is built into the FortiOS operating system. The feature comprises of an SSL daemon running on the FortiGate unit, and a web portal, which provides users with access to network services and resources including HTTP/HTTPS, telnet, FTP, SMB/CIFS, VNC, RDP and SSH.

In web-only mode, the FortiGate unit acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group. After successful authentication, the FortiGate unit redirects the web browser to the web portal home page and the user can access the server applications behind the FortiGate unit.

When the FortiGate unit provides services in web-only mode, a secure connection between the remote client and the FortiGate unit is established through the SSL VPN security in the FortiGate unit and the SSL security in the web browser. After the connection has been established, the FortiGate unit provides access to selected services and network resources through a web portal.

FortiGate SSL VPN web portals have a 1- or 2-column page layout and portal functionality is provided through small applets called widgets. Widget windows can be moved or minimized. The controls within each widget depend on its function. There are predefined web portals and the administrator can create additional portals.

Configuring the FortiGate unit involves selecting the appropriate web portal configuration in the user group settings. These configuration settings determine which server applications can be accessed. SSL encryption is used to ensure traffic confidentiality.

For information about client operating system and browser requirements, see the Release Notes for your FortiGate firmware.

Tunnel mode

Tunnel mode offers remote users the freedom to connect to the internal network using the traditional means of web-based access from laptop computers, as well as from airport kiosks, hotel business centers, and Internet cafés. If the applications on the client computers used by your user community vary greatly, you can deploy a dedicated SSL VPN client to any remote client through its web browser. The SSL VPN client encrypts all traffic from the remote client computer and sends it to the FortiGate unit through an SSL VPN tunnel over the HTTPS link between the web browser and the FortiGate unit. Another option is split tunneling, which ensures that only the traffic for the private network is sent to the SSL VPN gateway. Internet traffic is sent through the usual unencrypted route. This conserves bandwidth and alleviates bottlenecks.

In tunnel mode, remote clients connect to the FortiGate unit and the web portal login page using Microsoft Internet Explorer, Mozilla Foundation/Firefox, Mac OS, or Linux. The FortiGate unit acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group. After successful authentication, the FortiGate unit redirects the web browser to the web portal home page dictated by the user group settings. If the user does not have the SSL VPN client installed, they will be prompted to download the SSL VPN client (an ActiveX or Java plugin) and install it using controls provided through the web portal. SSL VPN tunnel mode can also be initiated from a standalone application on Windows, Mac OS, and Linux.

When the user initiates a VPN connection with the FortiGate unit through the SSL VPN client, the FortiGate unit establishes a tunnel with the client and assigns the client a virtual IP address from a range of reserved addresses. The client uses the assigned IP address as its source address for the duration of the connection. After the tunnel has been established, the user can access the network behind the FortiGate unit.

Configuring the FortiGate unit to establish a tunnel with remote clients involves enabling the feature through SSL VPN configuration settings and selecting the appropriate web portal configuration for tunnel-mode access in the user group settings. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely



The user account used to install the SSL VPN client on the remote computer must have administrator privileges.



If you are using Windows Vista, you must disable UAC (User Account Control) before installing the SSL VPN tunnel client. This UAC setting must be disabled before the SSL VPN tunnel client is installed. IE7 in Windows Vista runs in Protected Mode by default. To install SSL VPN client ActiveX, you need to launch IE7 by using 'Run as administrator' (right-click the IE7 icon and select 'Run as administrator').

For information about client operating system requirements, see the Release Notes for your FortiGate firmware.

For information on configuring tunnel mode, see [“Tunnel mode settings” on page 1625](#).

Port forwarding mode

While tunnel mode provides a Layer 3 tunnel that users can run any application over it, the user needs to install the tunnel client, and have the required administrative rights to do so. In some situations, this may not be desirable, yet the simple web mode does not provide enough flexibility for application support. For example, using an email client that needs to communicate with a POP3 server. The port forward mode, or proxy mode, provides this middle ground between web mode and tunnel mode.

SSL VPN port forwarding listens on local ports on the user's computer. When it receives data from a client application, the port forward module encrypts and sends the data to the FortiGate unit, which then forwards the traffic to the application server.

The port forward module is implemented with a Java applet, which is downloaded and runs on the user's computer. The applet provides the up-to-date status information such as addressing and bytes sent and received.

On the user end, the user logs into the FortiGate SSL VPN portal, and selects a port forward bookmark configured for a specific application. The bookmark defines the server address and port as well as which port to listen to on the user's computer.



The user must configure the application on the PC to point to the local proxy instead of the application server. For information on this configuration change, see the application documentation.

This mode only supports client/server applications that are using a static TCP port. It will not support client/server applications using dynamic ports or traffic over UDP.

For information on configuring a port forward tunnel, see [“Port forward tunnel” on page 1628](#).

Application support

With Citrix application servers, the server downloads an ICA configuration file to the user's PC. The client application uses this information to connect to the Citrix server. The FortiGate unit will read this file and append a SOCKS entry to set the SOCKS proxy to localhost. The Citrix client will then be able to connect to the SSL VPN port forward module to provide the connection. When configuring the port forwarding module, an selection is available for Citrix servers.

For Windows Remote Desktop Connections, when selecting the RDP option, the tunnel will launch the RDP client and connect to the local loopback address after the port forward module has been initiated.

SSL VPN and IPv6

FortiOS supports SSL VPN using IPv6 addressing using IPv6 configurations for security policies and addressing including:

- Policy matching for IPv6 addresses
- Support for DNS resolving in SSL VPN
- Support IPv6 for ping
- FTP applications
- SMB
- Support IPV6 for all the java applets (Telnet, VNC, RDP and so on)

Traveling and security

Because SSL VPN provides a means for “on-the-go” users to dial in to the network while away from the office, you need to ensure that wherever and however they choose to dial in is secure, and not potentially compromising the corporate network.

When setting up the portal, you can include two options to ensure corporate data is safe; a host check for antivirus software, and a cache cleaner.

Host check

You can enable a host integrity checker to scan the remote client. The integrity checker probes the remote client computer to verify that it is safe before access is granted. Security attributes recorded on the client computer (for example, in the Windows registry, in specific files, or held in memory due to running processes) are examined and uploaded to the FortiGate unit.

For more information, see [“Host Check” on page 1640](#).

Cache cleaning

You can enable a cache cleaner to remove any sensitive data that would otherwise remain on the remote computer after the session ends. For example, all cache entries, browser history, cookies, encrypted information related to user authentication, and any temporary data generated during the session are removed from the remote computer. If the client's browser cannot install and run the cache cleaner, the user is not allowed to access the SSL-VPN portal.

For more information, see [“Cache cleaning” on page 1616](#).



Basic Configuration

Configuring SSL VPN involves a number of configurations within FortiOS that you need to complete to make it all come together. This chapter describes the components required, and how and where to configure them to set up the FortiGate unit as an SSL VPN server. The configurations and steps are high level, to show you the procedures needed, and where in FortiOS they are located. For real-world examples, see the chapter, [“Setup examples” on page 1655](#).

There are three or four key steps to configuring an SSL VPN tunnel. The first three in the points below are mandatory, while the other is optional. This chapter will outline these four key steps, as well as additional configuration you can do for tighter security and monitoring.

The key steps are:

- Create user accounts and user groups for the remote clients.
([“User accounts and groups” on page 1617](#))
- Create a web portal to define user access to network resources.
([“Configuring SSL VPN web portals” on page 1620](#))
- Configure the security policies.
([“Configuring security policies” on page 1631](#))
- For tunnel-mode operation, add routing to ensure that client tunnel-mode packets reach the SSL VPN interface.
([“Routing in tunnel mode” on page 1638](#))
- Setup logging of SSL VPN activities.
([“SSL VPN logs” on page 1647](#))

User accounts and groups

The first step for an SSL VPN tunnel is to add the users and user groups that will access the tunnel. You may already have users defined for other authentication-based security policies. These users and groups are identified when creating the security policy when defining the authentication rules.

The user group is associated with the web portal that the user sees after logging in. If you have multiple portals, you will need multiple user groups. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

To create a user account

- in the web-based manager, go to *User > User*, and select *Create New*.
- in the CLI, use the commands in `config user local`.

All users accessing the SSL tunnel must be in a firewall user group. As part of configuring the user group, you select the SSL VPN web portal that the members of the group access after authentication.

To create user groups

- in the web-based manager, go to *User > User Group > User Group* and select *Create New*.
- in the CLI, use the commands in `config user group`.

Authentication

Remote users must be authenticated before they can request services and/or access network resources through the web portal. The authentication process can use a password defined on the FortiGate unit or optionally use established external authentication mechanisms such as RADIUS or LDAP.

To authenticate users, you can use a plain text password on the FortiGate unit (Local domain), forward authentication requests to an external RADIUS, LDAP or TACACS+ server, or utilize PKI certificates.

For information about how to create RADIUS, LDAP, TACACS+ or PKI user accounts and certificates, see the *User Authentication* chapter of *The Handbook*.



FortiOS supports LDAP password renewal notification and updates through SSL VPN. Configuration is enabled using the CLI commands:

```
config user ldap
  edit <username>
    set password-expiry-warning enable
    set password-renewal enable
  end
```

For more information, see the *User Authentication* chapter of *The Handbook*.

IP addresses for users

After the FortiGate unit authenticates a request for a tunnel-mode connection, the FortiGate unit assigns the SSL VPN client an IP address for the session. The address is assigned from an address range (IP Pool) which is a firewall address that defines an IP address range.



Take care to prevent overlapping IP addresses. Do not assign to clients any IP addresses that are already in use on the private network. As a precaution, consider assigning IP addresses from a network that is not commonly used (for example, 10.254.254.0/24).

To set tunnel-mode client IP address range - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.
- 2 Enter an *Address Name*, for example, `SSL_VPN_tunnel_range`.
- 3 In the *Subnet/IP Range* field, enter the starting and ending IP addresses that you want to assign to SSL VPN clients, for example `10.254.254.[80-100]`.
- 4 In *Interface*, select *Any*.
- 5 Select *OK*.

To set tunnel-mode client IP address range - CLI

If your SSL VPN tunnel range is for example 10.254.254.80 - 10.254.254.100, you could enter

```
config firewall address
  edit SSL_tunnel_users
    set type iprange
    set end-ip 10.254.254.100
    set start-ip 10.254.254.80
  end
end
```

You can select the tunnel-mode IP Pools in two places:

- The **VPN > SSL > Config** page *IP Pools* setting applies to all web portals that do not specify their own IP Pools.
- The web portal Tunnel Mode widget *IP Pools* setting, if used, applies only to the web portal and overrides the setting in **VPN > SSL > Config**. See [“Tunnel mode settings” on page 1625](#).

Authentication of remote users

When remote users connect to the SSL VPN tunnel, they must perform authentication before being able to use the internal network resources. This can be as simple as assigning users with their own passwords, connecting to an LDAP server or using more secure options. FortiOS provides a number of options for authentication as well as security option for those connected users.

The web portal can include bookmarks to connect to internal network resources. A web (HTTP/HTTPS) bookmark can include login credentials so that the FortiGate unit automatically logs the user into the web site. This means that the user logs into the SSL VPN and then does not have to enter any more credentials to visit preconfigured web sites.

Both the administrator and the end user can configure bookmarks, including SSO bookmarks. To add bookmarks as a web portal user, see [“The Bookmarks widget” on page 1628](#).

Setting the client authentication timeout

The client authentication timeout controls how long an authenticated user will remain connected. When this time expires, the system forces the remote client to authenticate again. As with the idle timeout, a shorter period of time is more secure. The default value is 28800 seconds (8 hours). You can only modify this timeout value in the CLI.

For example, to change the authentication timeout to 18 000 seconds, enter the following commands:

```
config vpn ssl settings
  set auth-timeout 18000
end
```

Strong authentication with security certificates

The FortiGate unit supports strong (two-factor) authentication through X.509 security certificates (version 1 or 3). The FortiGate unit can require clients to authenticate using a certificate. Similarly, the client can require the FortiGate unit to authenticate using a certificate.

For information about obtaining and installing certificates, see the [User Authentication](#) chapter of *The Handbook*.

You can select the *Require Client Certificate* option in *SSL VPN settings* so that clients must authenticate using certificates. The client browser must have a local certificate installed, and the FortiGate unit must have the corresponding CA certificate installed.

When the remote client initiates a connection, the FortiGate unit prompts the client browser for its client-side certificate as part of the authentication process.

To require client authentication by security certificates - web-based manager

- 1 Go to *VPN > SSL > Config*.
- 2 Select *Require Client Certificate*.
- 3 Select *Apply*.

To require client authentication by security certificates - CLI

```
config vpn ssl settings
    set reqclientcert enable
end
```

If your SSL VPN clients require strong authentication, the FortiGate unit must offer a CA certificate that the client browser has installed.

In the FortiGate unit SSL VPN settings, you can select which certificate the FortiGate offers to authenticate itself. By default, the FortiGate unit offers its factory installed (self-signed) certificate from Fortinet to remote clients when they connect.

To enable FortiGate unit authentication by certificate - web-based manager

- 1 Go to *VPN > SSL > Config*.
- 2 From the *Server Certificate* list, select the certificate that the FortiGate unit uses to identify itself to SSL VPN clients.
- 3 Select *Apply*.

To enable FortiGate unit authentication by certificate - CLI

For example, to use the `example_cert` certificate

```
config vpn ssl settings
    set servercert example_cert
end
```

Configuring SSL VPN web portals

The SSL VPN portal enables remote users to access internal network resources through a secure channel using a web browser. FortiGate administrators can configure log in privileges for system users and which network resources are available to the users.

This step in the configuration of the SSL VPN tunnel sets up the infrastructure; the addressing, encryption, certificates needed to make the initial connection to the FortiGate unit. This step also is where you set up what the remote user sees when the connection is successful. The portal view defines what resources are available to the remote users and what functionality they have on the network.

SSL connection configuration

To configure the basic SSL VPN settings for encryption and log in options, go to **VPN > SSL > Config**.

IP Pools	Select <i>Edit</i> to select the range or subnet firewall addresses that represent IP address ranges reserved for tunnel-mode SSL VPN clients. The IP Pool that you select will be the one created in the previous steps.
Server Certificate	Select the signed server certificate to use for authentication. If you leave the default setting (Self-Signed), the FortiGate unit offers its factory installed certificate from Fortinet, to remote clients when they connect.
Require Client Certificate	Select to use group certificates for authenticating remote clients. When the remote client initiates a connection, the FortiGate unit prompts the client for its client-side certificate as part of the authentication process. For information on using PKI to provide client certificate authentication, see the User Authentication Guide .
Encryption Key Algorithm	Select the algorithm for creating a secure SSL connection between the remote client web browser and the FortiGate unit. This will depend on what the web browser of the client can support. The FortiGate unit supports a range of cryptographic cipher suites to match the capabilities of various web browsers. The web browser and the FortiGate unit negotiate a cipher suite before any information is transmitted over the SSL link.
Idle Timeout	Type the period of time (in seconds) that the connection can remain idle before the user must log in again. The range is from 10 to 28800 seconds. Setting the value to 0 will disable the idle connection timeout. This setting applies to the SSL VPN session. The interface does not time out when web application sessions or tunnels are up.
Login Port	Enter the port number for HTTPS access.
Advanced (DNS and WINS Servers)	Enter up to two DNS servers and/r two WINS servers to be provided for the use of clients.

Portal configuration

The portal configuration determines what the remote user sees when they log in to the portal. Both the system administrator and the user have the ability to customize the SSL VPN portal.

There are three pre-defined default web portal configurations available:

- *full-access*
- *tunnel-access*
- *web-access*

To view the portals settings page, go to *VPN > SSL > Portal*.

Portal Settings page	
Edit Settings window	Provides general, virtual desktop and security control settings for the SSL VPN Service portal page. This window appears when you select <i>Settings</i> . This window also appears when you select <i>Create New</i> .
Settings	Select to edit the settings for the SSL VPN web portal.
Add Widget	Select to add a new widget to the page.
Widgets	The widgets that will appear on the SSL VPN Service page. By default when starting a new portal, all four of the widgets available for the portal appear on the screen. You can add widgets from the Add Widgets drop-down list.

Portal settings

The portal settings determine what SSL VPN users see when they log in to the FortiGate unit. Both the FortiGate administrator and the SSL VPN user have the ability to customize the web portal settings.

Portal settings are configured by going to *VPN > SSL > Portal* and select *Settings*.

General settings

The general settings tab enables you to set up the portal container - what the remote user sees when they log in. It also is the location where you define what applications the remote user can use when connecting. The applications selected affect how you configure the widgets later on. For example, if you do not select the HTTP/HTTPS option, you cannot add bookmarks in the bookmark widget.

Virtual Desktop settings

The virtual desktop options, available for Windows XP and Windows Vista client PCs, are configured to completely isolate the SSL VPN session from the client computer's desktop environment. All data is encrypted, including cached user credentials, browser history, cookies, temporary files, and user files created during the session. When the SSL VPN session ends normally, the files are deleted. If the session ends unexpectedly, any files that may remain will be encrypted.

Virtual desktop requires the Fortinet host check plug in. If the plug in is not present, it is automatically downloaded to the client computer.

Security Control settings

Security control options provide cache cleaning and host checking to the clients of your web portal. Cache cleaning clears information from the client browser cache just before the SSL VPN session ends. The cache cleaner is effective only if the session terminates normally. The cache is not cleaned if the session ends unexpectedly.

Host checking enforces the client's use of antivirus or firewall software. Each client is checked for security software that is recognized by the Windows Security Center. As an alternative, you can create a custom host check that looks for specific security software selected from the Host Check list located at *VPN > SSL > Host Check*. See ["Host Check" on page 1640](#).

Portal widgets

Portal widgets are sections of information that the user will view when they log in to the portal. By default, all widgets are shown. You can modify or remove widgets as required.

Session Information

The *Session Information* widget displays the login name of the user, the amount of time the user has been logged in and the inbound and outbound traffic statistics.

Bookmarks

Bookmarks are used as links to internal network resources. When a bookmark is selected from a bookmark list, a pop-up window appears with the web page. Telnet, VNC, and RDP require a browser plug-in. FTP and Samba replace the bookmarks page with an HTML file-browser.

A web bookmark can include login credentials to automatically log the SSL VPN user into the web site. When the administrator configures bookmarks, the web site credentials must be the same as the user's SSL VPN credentials. Users configuring their own bookmarks can specify alternative credentials for the web site.

Connection Tool

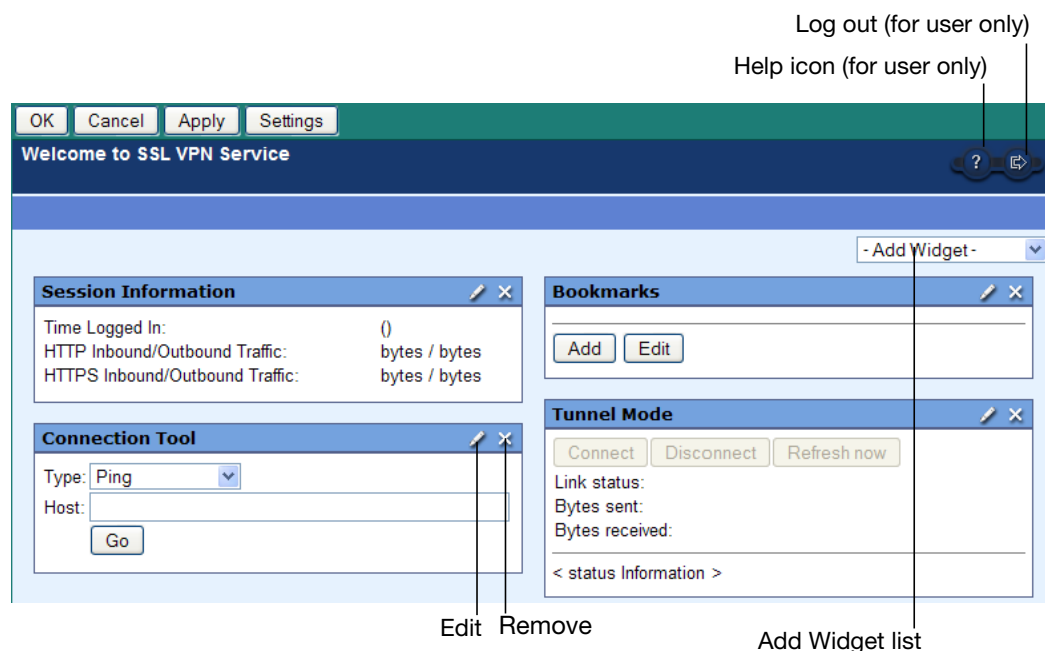
Use the *Connection Tool* widget to connect to a internal network resource without adding a bookmark to the bookmark list. You select the type of resource and specify the URL or IP address of the host computer.

Tunnel Mode

If your web portal provides tunnel mode access, you need to configure the *Tunnel Mode* widget. These settings determine how tunnel mode clients are assigned IP addresses.

Configuring the web portal page layout

You can determine which widgets are displayed on the web portal page and adjust the layout.

Figure 159: Configuring the SSL VPN web portal page**To configure the web portal page - web-based manager**

On the web portal page itself, you can make several adjustments to the appearance of the portal:

- Arrange widgets on the page by dragging them by their title bar.
- Add a widget by choosing a widget from the *Add Widget* list.
- Remove a widget by selecting the *Remove* icon in the widget title bar.
- Configure a widget by selecting the *Edit* icon in the widget title bar. For configuration information about each widget type, see the following sections:
 - “[Tunnel mode settings](#)” on page 1625
 - “[The Session Information widget](#)” on page 1628
 - “[The Connection Tool widget](#)” on page 1630
- To modify the color scheme and other basic settings, select the *Settings* button. You can also configure several advanced features. For more information, see
 - “[The Connection Tool widget](#)” on page 1630
 - “[Configuring cache cleaning](#)” on page 1643
 - “[Configuring virtual desktop](#)” on page 1644
 - “[Configuring client OS Check](#)” on page 1646 (CLI only)

When you have finished configuring the web portal page, select *Apply* to save the modifications.

To configure the web portal page - CLI

You can also define a portal layout using CLI commands, although due its complexity, is not recommended. Unlike configuring with the web-based manager, a new portal created in the CLI has by default no heading and no widgets. Also, the widgets do not have default names. You must specify all of this information.

For example, to create the portal layout shown in [Figure 159 on page 1624](#), you would enter:

```
config vpn ssl web portal
  set heading "Welcome to SSL VPN Service"
  set page-layout double-column
  set theme blue
  edit myportal
    config widget
      edit 0
        set type info
        set name "Session Information"
        set column one
      next
      edit 0
        set type bookmark
        set name "Bookmarks"
        set column one
      next
      edit 0
        set type tunnel
        set name "Tunnel Mode"
        set column two
      next
      edit 0
        set type tool
        set name "Connection Tool"
        set column two
    end
```



When you use `edit 0`, as in this example, the CLI automatically assigns an unused index value when you exit the edit shell by typing `end`.

Tunnel mode settings

If your web portal provides tunnel mode access, the *Tunnel Mode* widget is included automatically when creating a new portal, with the *Split Tunneling* option enabled. These settings determine how tunnel mode clients are assigned IP addresses.

If this web portal will assign a different range of IP addresses to clients than the IP Pools you specified on the *VPN > SSL > Config* page, you need to define a firewall address for the IP address range that you want to use. You will then need to specify this address in the Tunnel Mode widget *IP Pools* setting.



If the tunnel mode and session information widgets are the only widgets configured, the user will automatically be logged into the SSL-VPN tunnel.

Optionally, you can enable a split tunneling configuration so that the VPN carries only the traffic for the networks behind the FortiGate unit. The user's other traffic follows its normal route.

To configure tunnel mode settings - web-based manager

- 1 Go to *VPN > SSL > Portal* and select *Create New*.
- 2 Select the *Edit* icon in the *Tunnel Mode* widget title bar.

Figure 160: Tunnel Mode widget - edit mode

The screenshot shows the 'Tunnel Mode' configuration window. The top section, labeled 'Tunnel mode settings', includes the Name field (set to 'Tunnel Mode'), IP Mode selection (Radio buttons for 'Range' and 'User Group', with 'Range' selected), IP Pools (displaying 'SSLVPN_TUNNEL_ADDR1' with an 'Edit' link), and a 'Split Tunneling' checkbox (unchecked). The bottom section, labeled 'Tunnel controls (for users only)', includes 'Connect', 'Disconnect', and 'Refresh' buttons, followed by 'Link status:', 'Bytes sent:', 'Bytes received:', and a '< status information >' link.

- 3 Enter the following information:

Name	Enter a name for the widget.
IP Mode	Select the mode by which the IP address is assigned to the user.
Range	Use the IP address ranges specified by <i>IP Pools</i> .
User Group	The user is assigned the IP address specified in the Framed-IP-Address field of the user's record on the RADIUS server. This option is valid only for users authenticated by a RADIUS server.
IP Pools	If you want to specify an IP address range for clients of this portal only, select <i>Edit</i> . From the <i>Available</i> list, select the appropriate firewall address. You must configure the desired IP address range as a firewall address before you can select it here.
Split Tunneling	<p>Split tunneling is enabled by default. When enabled, only traffic that requires the SSL VPN is sent through the tunnel. Other traffic follows the user's regular routing.</p> <p>When split tunneling is disabled, all of the user's traffic with other networks passes through the tunnel. This does not affect traffic between the user's computer and hosts on the local network.</p> <p>For enhanced security, some administrators prefer to force all traffic through the SSL VPN tunnel, including traffic between the user and the user's local network. To do this, use the CLI tunnel mode settings to enable <code>exclusive-routing</code>.</p>

The remaining items in the widget are controls that are available to the user during an SSL VPN session.

- 4 Select *OK* in the *Tunnel Mode* widget.
- 5 Select *Apply*.

To configure tunnel mode settings - CLI

To enable tunnel mode operation for portal2 portal users and assign them addresses from the `SSLVPN_TUNNEL_ADDR2` range, you would enter:

```
config vpn ssl web portal
edit portal2
config widget
edit 0
set type tunnel
set ip-mode range
set ip-pools SSLVPN_TUNNEL_ADDR2
end
end
```

The preceding example applies to a web portal that does not already have a tunnel mode widget. To modify the settings on an existing tunnel mode widget, you need to determine the widget's number. Enter:

```
config vpn ssl web portal
edit portal1
config widget
show
```

In the output, you will see, for example,

```
edit 3
set name "Tunnel Mode"
set type tunnel
...
```

You can now enter `edit 3` and modify the tunnel mode widget's settings.

To force all traffic through the tunnel - CLI

If you disable split tunneling, all of the user's traffic to other networks passes through the SSL VPN tunnel. But, this does not apply to traffic between the user and the user's local network. For enhanced security, some administrators prefer to force all of the user's traffic, including traffic with the local network, to pass through the SSL VPN tunnel. To do this, enable `exclusive-routing` in the tunnel widget settings. For example:

```
config vpn ssl web portal
edit portal2
config widget
edit 0
set type tunnel
set ip-mode range
set ip-pools SSLVPN_TUNNEL_ADDR2
set split-tunneling disable
set exclusive-routing enable
end
end
```

Port forward tunnel

Port forwarding provides a method of connecting to application servers without configuring a tunnel mode connection, and requiring the installation of tunnel mode client. Set up the portal as described at [“Configuring SSL VPN web portals” on page 1620](#). To configure the application, create a bookmark with the *Type* of *PortForward*.



Ensure that *Port Forward* is enabled in the *Applications* list of the *General* settings, by selecting the *Settings* button in the portal configuration window.

The Session Information widget

The *Session Information* widget displays the login name of the user, the amount of time the user has been logged in, and the inbound and outbound traffic statistics of HTTP and HTTPS. You can change the widget name.

To edit the session information, in the *Session Information* widget select *Edit* and enter the session name.

To configure Session Information settings - CLI

To change the name of the web-access Session Information widget to “My Session”, you would enter:

```
config vpn ssl web portal
edit web-access
config widget
edit 4
set name "My Session"
end
```

The Bookmarks widget

Bookmarks are used as links to specific resources on the network. When a bookmark is selected from a bookmark list, a pop-up window appears with the requested web page. Telnet, VNC, and RDP all pop up a window that requires a browser plug-in. FTP and Samba replace the bookmarks page with an HTML file-browser.

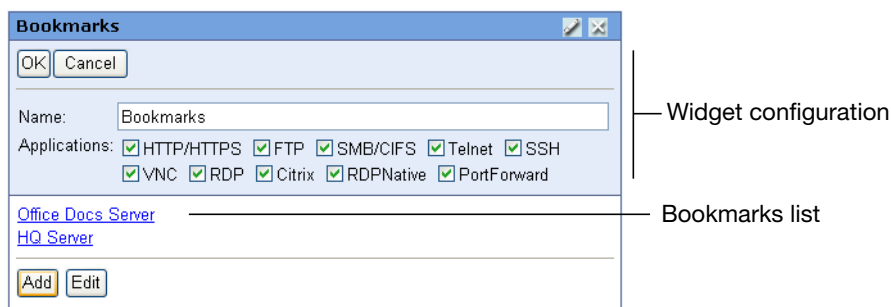


Ensure that *HTTP/HTTPS* is enabled in the *Applications* list of the *General* settings, by selecting the *Settings* button in the portal configuration window.

To configure the Bookmarks widget

- 1 Go to *VPN > SSL > Portal*, and select *Create New*.

- 2 Select the *Edit* icon in the *Bookmarks* widget title bar.



- 3 Select the *Applications* check boxes for the types of bookmarks that you want to support.
- 4 To add a bookmark, select *Add*.
- 5 Enter or edit the following information:

Name	Enter a name for the bookmark.
Type	Select the type of application to which the bookmark links. For example, select HTTP/HTTPS for a web site. Only the application types that you configured in the settings are in the list. You can select Edit in the widget title bar to enable additional application types.
Location	Enter the destination of the bookmark.
Description	Enter a descriptive tooltip for the bookmark.
SSO	A Single Sign-On (SSO) bookmark automatically enters the login credentials for the bookmark destination. Select one of: Disabled — This is not an SSO bookmark. Automatic — Use the user's SSL VPN credentials for login. Static — Use the login credentials defined below.
Single Sign-On settings available when SSO is Static	
Field Name	Enter a required login page field name, "User Name" for example.
Value	Enter the value to enter in the field identified by <i>Field Name</i> . If you are an administrator configuring a bookmark for users: Enter %username% to represent the user's SSL VPN user name. Enter %passwd% to represent the user's SSL VPN password.
Add	Enter another <i>Field Name</i> / <i>Value</i> pair, for the password for example. A new set of <i>Field Name</i> / <i>Value</i> fields is added.

- 6 Select *OK*.
- 7 Select *Apply* at the top of the web portal page to save the changes that you made.

To configure the Bookmarks widget and add/edit bookmarks - CLI

To allow only FTP and web connections on the web-access portal and to configure a bookmark to example.com, you would enter:

```
config vpn ssl web portal
  edit web-access
    config widget
      edit 1
        set type bookmark
        set allow-apps ftp web
        config bookmarks
          edit "example"
            set apptype web
            set description "example bookmark"
            set url "http://example.com"
          end
        end
      end
    end
  end
```

To delete bookmarks - CLI

To delete the bookmark added above, you would enter:

```
config vpn ssl web portal
  edit web-access
    config widget
      edit 1
        config bookmarks
          delete example
        end
      end
    end
  end
```

The Connection Tool widget

The *Connection Tool* enables a user to connect to resources for which there are no bookmarks.



Ensure that what you want remote users to connect to is enabled in the *Applications* list of the *General* settings, by selecting the *Settings* button in the portal configuration window.

To configure the Connection Tool widget

- 1 Go to *VPN > SSL > Portal*, and select *Create New*.
- 2 Select the *Edit* icon in the *Connection Tool* widget title bar.

- 3 Enter a new *Name* for the widget.
- 4 Select the types of *Applications* that the Connection Tool is enabled to access.
- 5 Select *OK*.

To configure the Connection Tool widget - CLI

To change, for example, the full-access portal Connection Tool widget to allow all application types except Telnet, you would enter:

```
config vpn ssl web portal
edit full-access
config widget
edit 3
set allow-apps ftp rdp smb ssh vnc web
end
end
end
```

Configuring security policies

You will need at least one SSL VPN security policy. This is an identity-based policy that authenticates users and enables them to access the SSL VPN web portal. The SSL VPN user groups named in the policy determine who can authenticate and which web portal they will use. From the web portal, users can access protected resources or download the SSL VPN tunnel client application.

This section contains the procedures needed to configure security policies for web-only mode operation and tunnel-mode operation. These procedures assume that you have already completed the procedures outlined in [“User accounts and groups” on page 1617](#).

If you will provide tunnel mode access, you will need a second security policy — an ACCEPT tunnel mode policy to permit traffic to flow between the SSL VPN tunnel and the protected networks.

Firewall addresses

Before you can create security policies, you need to define the firewall addresses you will use in those policies. For both web-only and tunnel mode operation, you need to create firewall addresses for all of the destination networks and servers to which the SSL VPN client will be able to connect.

For tunnel mode, you will already have defined firewall addresses for the IP address ranges that the FortiGate unit will assign to SSL VPN clients. See [“Windows OS check” on page 1643](#).

The source address for your SSL VPN security policies will be the predefined “all” address. Both the address and the netmask are 0.0.0.0. The “all” address is used because VPN clients will be connecting from various addresses, not just one or two known networks. For improved security, if clients will be connecting from one or two known locations you should configure firewall addresses for those locations, instead of using the “all” address.

To create a firewall address, in the web-based manager, go to *Firewall Objects > Address > Address*, and select *Create New*. Using the CLI, use the commands in `config firewall address`.

Create an SSL VPN security policy

At minimum, you need one SSL VPN security policy to authenticate users and provide access to the protected networks. You will need additional security policies only if you have multiple web portals that provide access to different resources.

The user group is associated with the web portal that the user sees after logging in. If you have multiple portals, you will need multiple user groups. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

The SSL VPN security policy specifies:

- the source address that corresponds to the IP address of the remote user.
- the destination address that corresponds to the IP address or addresses that remote clients need to access.

The destination address may correspond to an entire private network, a range of private IP addresses, or the private IP address of a server or host.

- the level of SSL encryption to use and the authentication method.
- which SSL VPN user groups can use the security policy.
- the times (schedule) and types of services that users can access.
- the UTM features and logging that are applied to the connection.



Do not use ALL as the destination address. If you do, you will see the “Destination address of Split Tunneling policy is invalid” error when you enable Split Tunneling.

To create an SSL-VPN security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information:

Source Interface/Zone	Select the name of the FortiGate network interface to that connects to the Internet.
Source Address	Select <i>all</i> .
Destination Interface/Zone	Select the FortiGate network interface that connects to the protected network.

Destination Address	Select the firewall address you created that represents the networks and servers to which the SSL VPN clients will connect. If you want to associate multiple firewall addresses or address groups with the <i>Destination Interface/Zone</i> , from <i>Destination Address</i> , select the plus symbol. In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i> .
Action	Select <i>SSL-VPN</i> .
SSL Client Certificate Restrictive	Allow access only to holders of a (shared) group certificate. The holders of the group certificate must be members of an SSL VPN user group, and the name of that user group must be present in the <i>Allowed</i> field. See “Strong authentication with security certificates” on page 1619 .
Cipher Strength	Select the bit level of SSL encryption. The web browser on the remote client must be capable of matching the level that you select.
Configure SSL-VPN Users	A security policy for an SSL VPN is automatically an identity-based policy.
Add	Add a user group to the policy. The Edit Authentication Rule window opens on top of the security policy. Enter the following information and then select <i>OK</i> . You can select <i>Add</i> again to add more groups.
User Group	Select user groups in the left list and use the right arrow button to move them to the right list.
Service	Select service in the left list and use the right arrow button to move them to the right list. Select the <i>ANY</i> service to allow the user group access to all services.

3 Select *OK*.

Your identity-based policies are listed in the security policy table. The FortiGate unit searches the table from the top down to find a policy to match the client's user group. Using the move icon in each row, you can change the order of the policies in the table to ensure the best policy will be matched first. You can also use the icons to edit or delete policies.

To create an SSL VPN security policy - CLI

To create the security policy by entering the following CLI commands.

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr OfficeLAN
    set action ssl-vpn
    set nat enable
  config identity-based-policy
    edit 0
      set groups SSL-VPN
      set schedule always
```

```

        set service ANY
    end
end

```

Create a tunnel mode security policy

If your SSL VPN will provide tunnel mode operation, you need to create a security policy to enable traffic to pass between the SSL VPN virtual interface and the protected networks. This is in addition to the SSL VPN security policy that you created in the preceding section.

The SSL VPN virtual interface is the FortiGate unit end of the SSL tunnel that connects to the remote client. It is named `ssl.<vdom_name>`. In the root VDOM, for example, it is named `ssl.root`. If VDOMs are not enabled on your FortiGate unit, the SSL VPN virtual interface is also named `ssl.root`.

To configure the tunnel mode security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*.

Source Interface/Zone	Select the virtual SSL VPN interface, such as <i>ssl.root</i> .
Source Address	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients, such as <i>SSL_VPN_tunnel_users</i> .
Destination Interface/Zone	Select the interface that connects to the protected network.
Destination Address	Select the firewall address that represents the networks and servers the SSL VPN clients will connect to. To select multiple firewall addresses or address groups, select the plus sign next to the drop-down list.
Action	Select <i>Accept</i> .
NAT	Select <i>Enable NAT</i>

To configure the tunnel mode security policy - CLI

```

config firewall policy
    edit <id>
        set srcintf ssl.root
        set dstintf <dst_interface_name>
        set srcaddr <tunnel_ip_address>
        set dstaddr <protected_network_address_name>
        set schedule always
        set service ANY
        set nat enable
    end

```

This policy enables the SSL VPN client to initiate communication with hosts on the protected network. If you want to enable hosts on the protected network to initiate communication with the SSL VPN client, you should create another Accept policy like the preceding one but with the source and destination settings reversed.

You must also add a static route for tunnel mode operation.

Routing for tunnel mode

If your SSL VPN operates in tunnel mode, you must add a static route so that replies from the protected network can reach the remote SSL VPN client.

To add the tunnel mode route - web-based manager

- 1 Go to *Router > Static > Static Route* and select *Create New*.
- 2 Enter the *Destination IP/Mask* of the tunnel IP address that you assigned to the users of the web portal.
- 3 Select the SSL VPN virtual interface for the *Device*.
- 4 Select *OK*.

To add the tunnel mode route - CLI

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
edit <id>
set device ssl.root
set dst 10.11.254.0/24
set gateway <gateway_IP>
end
```

Split tunnel Internet browsing policy

With split tunneling disabled, all of the SSL VPN client's requests are sent through the SSL VPN tunnel. But the tunnel mode security policy provides access only to the protected networks behind the FortiGate unit. Clients will receive no response if they attempt to access Internet resources. You can enable clients to connect to the Internet through the FortiGate unit.

To add an Internet browsing policy

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*.

Source Interface/Zone	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
Source Address	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients.
Destination Interface/Zone	Select the FortiGate network interface that connects to the Internet.
Destination Address	Select <i>all</i> .
Action	Select <i>Accept</i> .
NAT	Enable.
Leave other settings at their default values.	

To configure the Internet browsing security policy - CLI

To enable browsing the Internet through port1, you would enter:

```
config firewall policy
edit 0
set srcintf ssl.root
```

```

set dstintf port1
set srcaddr SSL_tunnel_users
set dstaddr all
set schedule always
set service ANY
set nat enable
end

```

Enabling a connection to an IPsec VPN

You might want to provide your SSL VPN clients access to another network, such as a branch office, that is connected by an IPsec VPN. To do this, you need only to add the appropriate security policy. For information about route-based and policy-based IPsec VPNs, see the [IPsec VPN Guide](#).

Route-based Connection

To configure interconnection with a route-based IPsec VPN - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*.

Source Interface/Zone	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
Source Address	Select the firewall address that represents the IP address range assigned to SSL VPN clients.
Destination Interface/Zone	Select the virtual IPsec interface for your IPsec VPN.
Destination Address	Select the address of the IPsec VPN remote protected subnet.
Action	Select <i>ACCEPT</i> .
NAT	Enable.
Leave other settings at their default values.	

To configure interconnection with a route-based IPsec VPN - CLI

If, for example, you want to enable SSL VPN users to connect to the private network (address name *OfficeAnet*) through the *toOfficeA* IPsec VPN, you would enter:

```

config firewall policy
edit 0
set srcintf ssl.root
set dstintf toOfficeA
set srcaddr SSL_tunnel_users
set dstaddr OfficeAnet
set action accept
set nat enable
set schedule always
set service ANY
end

```

Policy-based connection

To configure interconnection with a policy-based IPsec VPN - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*.

Source Interface/Zone	Select the virtual SSL VPN interface, <i>ssl.root</i> , for example.
Source Address	Select the firewall address that represents the IP address range assigned to SSL VPN clients.
Destination Interface/Zone	Select the FortiGate network interface that connects to the Internet.
Destination Address	Select the address of the IPsec VPN remote protected subnet.
Action	Select <i>IPSEC</i> .
VPN tunnel	Select the Phase 1 configuration name of your IPsec VPN.
Allow inbound	Enable
Allow outbound	Enable
NAT inbound	Enable
Leave other settings at their default values.	

To configure interconnection with a policy-based IPsec VPN - CLI

If, for example, you want to enable SSL VPN users to connect to the private network (address name *OfficeAnet*) through the *OfficeA* IPsec VPN, you would enter:

```
config firewall policy
  edit 0
    set srcintf ssl.root
    set dstintf port1
    set srcaddr SSL_tunnel_users
    set dstaddr OfficeAnet
    set action ipsec
    set schedule always
    set service ANY
    set inbound enable
    set outbound enable
    set natinbound enable
    set vpntunnel toOfficeA
  end
```

In this example, *port1* is connected to the Internet.

Additional configuration options

Beyond the basics of setting up the SSL VPN, you can configure a number of other options that can help to ensure your internal network is secure and limit the possibility of attacks and viruses entering the network from an outside source.

Routing in tunnel mode

If are creating a SSL VPN connection in tunnel mode, you need to add a static route so that replies from the protected network can reach the remote SSL VPN client.

To add the tunnel mode route - web-based manager

- 1 Go to *Router > Static > Static Route* and select *Create New*.
- 2 Enter the *Destination IP/Mask* of the tunnel IP address that you assigned to the users of the web portal.
- 3 Select the SSL VPN virtual interface for the *Device*.
- 4 Select *OK*.

To add the tunnel mode route - CLI

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
edit <id>
set device ssl.root
set dst 10.11.254.0/24
set gateway <gateway_IP>
end
```

Changing the port number for web portal connections

You can specify a different TCP port number for users to access the web portal login page through the HTTPS link. By default, the port number is 10443 and users can access the web portal login page using the following default URL:

`https://<FortiGate_IP_address>:10443/remote/login`

where <FortiGate_IP_address> is the IP address of the FortiGate interface that accepts connections from remote users.

To change the SSL VPN port - web-based manager

- 1 If *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
- 2 Go to *VPN > SSL > Config*.
- 3 Type an unused port number in *SSLVPN Login Port*, and select *Apply*.

To change the SSL VPN port - CLI

This is a global setting. For example, to set the SSL VPN port to 10443, enter:

```
config global
config system global
set sslvpn-sport 10443
end
```

SSL offloading

Configuring SSL offloading that allows or denies client renegotiation, is configured in the CLI. This feature helps to resolve the issues that affect all SSL and TLS servers that support renegotiation, identified by the Common Vulnerabilities and Exposures system in CVE-2009-3555. The IETF is currently working on a TLS protocol change that will permanently resolve the issue. The SSL offloading renegotiation feature is considered a workaround until the IETF permanently resolves the issue.

The CLI command is `ssl-client-renegotiation` and is found in `config firewall vip` command.

Customizing the web portal login page

The default web portal login page shows only the *Name* and *Password* fields and the Login button, centred in the web browser window. You can customize the page with your company name or other information.

The login page is a form of replacement message, in HTML format. You can modify the content to display a customized message. Note that there are specific fields that must remain in the code to ensure the page appears correctly in the user's browser.



Before you begin, copy the default web portal login page text to a separate text file for safe-keeping. Afterward, if needed you can restore the text to the original version.

To configure the SSL VPN login page - web-based manager

- 1 If you want to edit the global login page and *Current VDOM* appears at the bottom left of the screen, select *Global* from the list of VDOMs.
- 2 Go to *System > Config > Replacement Messages*.
- 3 Expand the *SSL VPN* row and select the *Edit* icon for the *SSL VPN login message*.
- 4 Edit the HTML text. Note the following content that must remain on the page:
 - The login page must contain a form with `ACTION="%%SSL_ACT%%"` and `METHOD="%%SSL_METHOD%%"`
 - The form must contain the `%%SSL_LOGIN%%` tag to provide the login form.
 - The form must contain the `%%SSL_HIDDEN%%` tag.
- 5 Select *OK*.

To configure the SSL VPN login page - CLI

Do one of the following:

- If VDOMs are enabled and you want to modify the global login page, enter:

```
config global
config system replacemsg sslvpn sslvpn-login
```

- If you want to modify the login page for a VDOM, enter:

```
config vdom
edit <vdom_name>
  config system replacemsg-group
  edit default
    config sslvpn
    edit sslvpn-login
```

To change the login page content, enter the modified page content as a string. In this example, the page title is changed to "Secure Portal login" and headings are added above the login dialog which say "example.com Secure Portal":

```

set buffer "<html><head><title>Secure Portal login</title>
<meta http-equiv='Pragma' content='no-cache'><meta http-
equiv='cache-control' content='no-cache'> <meta http-
equiv='cache-control' content='must-revalidate'><link
href='/sslvpn/css/login.css' rel='stylesheet'
type='text/css'><script type='text/javascript'>if (top &&
top.location != window.location) top.location =
top.location;if (window.opener && window.opener.top) {
window.opener.top.location = window.opener.top.location;
self.close(); }</script></head><body class='main'>
<center><table width='100%' height='100%' align='center'
class='container' valign='middle' cellpadding='0'
cellspacing='0'><tr valign=top><td align=center>
<h1>example.com</h1><h3>Secure Portal</h3></td></tr><tr><td><form action='%%SSL_ACT%%'
method='%%SSL_METHOD%%' name='f'><table class='list'
cellpadding=10 cellspacing=0 align=center width=400
height=180>%%SSL_LOGIN%%</table>%%SSL_HIDDEN%%</td></tr></tr></form></center></body><script>document.forms[0].use
rname.focus();</script></html>"
end

```

Your console application determines how the text wraps. It is easier to edit the code in a separate text editor and then paste the finished code into the `set buffer` command. Be sure to enclose the entire string in quotation (") marks.

Host Check

When you enable AV, FW, or AV-FW host checking in the web portal Security Control settings, each client is checked for security software that is recognized by the Windows Security Center. As an alternative, you can create a custom host check that looks for security software selected from the Host Check list. For more information, see [“Portal settings” on page 1622](#).

The Host Check list includes default entries for many security software products.

To configure host checking, go to *VPN > SSL > Host Check*. To add to the list, select *Create New*.



Host integrity checking is only possible with client computers running Microsoft Windows platforms.

Host Check Software page	
Name	Enter a name for the host check list.
Type	Select the type of host checking, either AV or FW.
GUID	Enter the Globally Unique Identifier (GUID) for the host check application, if known. Windows uses GUIDs to identify applications in the Windows Registry. The GUID can be found in the Windows registry in the HKEY_CLASSES_ROOT section.
Version	Enter the software's version.
Create New	Creates a new check item to add to the list below.

Edit	Modifies the settings within the software.
Delete	Removes the check software item from the list on the Host Check Software page. To remove multiple check software items from the list, select the check box for each row to remove, and select <i>Delete</i> .
#	The order in which each item is listed.
Target	The type of target that you chose.
Type	The type of check that you chose.
Action	The type of action that you chose.
Host Check Software window	
Type	Select how the FortiGate unit checks for the correct version of the application.
Action	Select one of the following: <i>Require</i> – If the item is found, the client meets the check item condition. <i>Deny</i> – If the item is found, the client is considered to not meet the check item condition. Use this option if it is necessary to prevent use of a particular security product.
File/Path	Enter the file name and path. This field appears when you select the <i>Type of File</i> .
Process	Enter the application's executable file name. This field appears when you select the <i>Type of Process</i> .
Registry	Enter the registry number of the application. This field appears when you select the <i>Type of Registry</i> .
Version	Enter the application's version.
MD5 Signatures (one per line)	Enter the MD5 signature for application executable file. You can enter more than one but each one needs to be on a separate line. You can use a third-party utility to calculate MD5 signatures or hashes for the file. Entering multiple MD5 signatures helps to match multiple versions of the application.

To configure host checking - CLI

To configure the full-access portal to check for AV and firewall software on client Windows computers, you would enter the following:

```
config vpn ssl web portal
edit full-access
set host-check av-fw
end
```

To configure the full-access portal to perform a custom host check for FortiClient Host Security AV and firewall software, you would enter the following:

```
config vpn ssl web portal
edit full-access
set host-check custom
set host-check-policy FortiClient-AV FortiClient-FW
end
```

Creating a custom host check list

If you configure a custom host check for your web portal (see [“Host Check” on page 1640](#)), you choose security applications from the list on the *VPN > SSL > Host Check* page. The *Host Check* list includes default entries for many security software products. You can add, remove, or modify entries in this list.



Host integrity checking is only possible with client computers running Microsoft Windows platforms.

To add an entry to the Host Check list - web-based manager

- 1 Go to *VPN > SSL > Host Check*.
- 2 Select *Create New* and enter the following information:

Name	Enter a name of the application. The name does not need to match the actual application name.
Type	Select the type of security application. Can be AV for antivirus or FW for firewall.
GUID	Enter the Globally Unique Identifier (GUID) for the host check application, if known. Windows uses GUIDs to identify applications in the Windows Registry. The GUID can be found in the Windows registry in the HKEY_CLASSES_ROOT section.
Version	The version of the security application. To get the exact versioning, in Windows right-click on the .EXE file of the application and select <i>Properties</i> . Select the <i>Version</i> tab.
Create New	If you do not know the GUID, add alternative checks for the application. The security software is considered found only if all checks succeed.
Type	Select how to check for the application: <ul style="list-style-type: none"> • File — Look for a file. This could be the application's executable file or any other file that would confirm the presence of the application. In <i>File/Path</i>, enter the full path to the file. Where applicable, you can use environment variables enclosed in percent (%) marks. For example, <code>%ProgramFiles%\Fortinet\FortiClient\FortiClient.exe</code> • Process — Look for the application as a running process. In <i>Process</i>, enter the application's executable file name. • Registry — Search for a Windows Registry entry. In <i>Registry</i>, enter a registry item, for example <code>HKLM\SOFTWARE\Fortinet\FortiClient\Misc</code>

Action	Select one of <ul style="list-style-type: none"> • Require — If the item is found, the client meets the check item condition. • Deny — If the item is found, the client is considered to not meet the check item condition. Use this option if it is necessary to prevent use of a particular security product.
MD5 Signatures	If <i>Type</i> is <i>File</i> or <i>Process</i> , enter one or more known MD5 signatures for the application executable file. You can use a third-party utility to calculate MD5 signatures or hashes for any file. You can enter multiple signatures to match multiple versions of the application.

3 Select OK.

Windows OS check

The Windows patch check enables you to define the minimum Windows version and patch level allowed when connecting to the SSL VPN portal. When the user attempts to connect to the web portal, FortiOS performs a query on the version of Windows the user has installed. If it does not match the minimum requirement, the connection is denied. The Windows patch check is configured in the CLI.

The following example shows how you would add an OS check to the g1portal web portal. This OS check accepts all Windows XP users and Windows 2000 users running patch level 3.

To specify the acceptable patch level, you set the `latest-patch-level` and the `tolerance`. The lowest acceptable patch level is `latest-patch-level` minus `tolerance`. In this case, `latest-patch-level` is 3 and `tolerance` is 1, so 2 is the lowest acceptable patch level.

```
config vpn ssl web portal
edit g1portal
set os-check enable
config os-check-list windows-2000
set action check-up-to-date
set latest-patch-level 3
set tolerance 1
end
config os-check-list windows-xp
set action allow
end
end
```

Configuring cache cleaning

When the SSL VPN session ends, the client browser cache may retain some information. To enhance security, cache cleaning clears this information just before the SSL VPN session ends.



The cache cleaner is effective only if the session terminates normally. The cache is not cleaned if the session ends due to a malfunction, such as a power failure.

To enable cache cleaning - web-based manager

- 1 Go to *VPN > SSL > Portal*, select the web portal and then select *Edit*.
- 2 Select the *Settings* and then select the *Security Control* tab.
- 3 Select *Clean Cache*.
- 4 Select *OK* then *Apply*.

To enable cache cleaning - CLI

To enable cache cleaning on the full-access portal, you would enter:

```
config vpn ssl web portal
edit full-access
set cache-cleaner enable
end
```

Cache cleaning requires a browser plug-in. If the user does not have the plug-in, it is automatically downloaded to the client computer.

Configuring virtual desktop

Available for Windows XP, Windows Vista, and Windows 7 client PCs, the virtual desktop feature completely isolates the SSL VPN session from the client computer's desktop environment. All data is encrypted, including cached user credentials, browser history, cookies, temporary files, and user files created during the session. When the SSL VPN session ends normally, the files are deleted. If the session ends due to a malfunction, files might remain, but they are encrypted, so the information is protected.

When the user starts an SSL VPN session which has virtual desktop enabled, the virtual desktop replaces the user's normal desktop. When the virtual desktop exits, the user's normal desktop is restored.

Virtual desktop requires the Fortinet cache cleaner plug in. If the plug in is not present, it is automatically downloaded to the client computer.

To enable virtual desktop - web-based manager

- 1 Go to *VPN > SSL > Portal*, select the web portal and then select *Edit*.
- 2 Select the *Settings* and then select the *Virtual Desktop* tab.
- 3 Select *Enable Virtual Desktop*.
- 4 Enable the other options as needed.
- 5 Optionally, select an Application Control List.
See ["Configuring virtual desktop application control"](#).
- 6 Select *OK*, then select *Apply*.

To enable virtual desktop - CLI

To enable virtual desktop on the full-access portal and apply the application control list List1, for example, you would enter:

```
config vpn ssl web portal
edit full-access
set virtual-desktop enable
set virtual-desktop-app-list List1
end
```

Configuring virtual desktop application control

You can control which applications users can run on their virtual desktop. To do this, you create an Application Control List of either allowed or blocked applications. When you configure the web portal, you select the list to use.

Configuration is located in *VPN > SSL > Virtual Desktop Application Control* and select *Create New*.

Virtual Desktop Application Control List page	
Name	Enter a name for the virtual desktop application list.
Allow the applications on the list and block all others	Select to allow the applications on this list.
Block the applications on the list and allow all others	Select to block the applications on the list.
Create New	Creates a new application signature.
Edit	Modifies the settings within the application signature.
Delete	Removes an application signature from the list on the Virtual Desktop Application Control List page. To remove multiple signatures from the list, on the Virtual Desktop Application Control List page, select the check box for the applications and select <i>Delete</i> .
Applications	The name of the application.
Application Signatures (window)	
Name	Enter the name of the application.
MD5 Signatures (one per line)	Enter the MD5 signature for application executable file. You can enter more than one but each one requires to be on a separate line. Entering multiple MD5 signatures helps to match multiple versions of the application.

There are two types of application control list:

- allow the listed applications and block all others
- or
- block the listed applications and allow all others.

You can create multiple application control lists, but each in web portal you can select only one list to use.

To create an Application Control List - web-based manager

- 1 Go to *VPN > SSL > Virtual Desktop Application Control* and select *Create New*.
- 2 Enter a *Name* for the list.

- 3 Select one of the following:
 - *Allow the applications on this list and block all others*
 - *Block the applications on this list and allow all others*
- 4 Select *Add*.
- 5 Enter a *Name* for the application.
This can be any name and does not have to match the official name of the application.
- 6 Enter one or more known *MD5 Signatures* for the application executable file.
You can use a third-party utility to calculate MD5 signatures or hashes for any file. You can enter multiple signatures to match multiple versions of the application.
- 7 Select *OK*.
- 8 Repeat steps 4 through 7 for each additional application.
- 9 Select *OK*.

To create an Application Control List - CLI

If you want to add BannedApp to List1, a list of blocked applications, you would enter:

```
config vpn ssl web virtual-desktop-app-list
edit "List1"
set action block
config apps
edit "BannedApp"
set md5s "06321103A343B04DF9283B80D1E00F6B"
end
end
```

Configuring client OS Check

The SSLVPN client OS Check feature can determine if clients are running the Windows 2000, Windows XP, Windows Vista or Windows 7 operating system. You can configure the OS Check to do any of the following:

- allow the client access
- allow the client access only if the operating system has been updated to a specified patch (service pack) version
- deny the client access

The OS Check has no effect on clients running other operating systems.

To configure OS Check - CLI

OS Check is configurable only in the CLI.

```
config vpn ssl web portal
edit <portal_name>
set os-check enable
config os-check-list {windows-2000 | windows-xp
| windows-vista | windows-7}
set action {allow | check-up-to-date | deny}
set latest-patch-level {disable | 0 - 255}
set tolerance {tolerance_num}
end
end
```

Adding WINS and DNS services for clients

You can specify the WINS or DNS servers that are made available to SSL-VPN clients.

DNS servers provide the IP addresses that browsers need to access web sites. For Internet sites, you can specify the DNS server that your FortiGate unit uses. If SSL VPN users will access intranet sites using URLs, you need to provide them access to the intranet's DNS server. You specify a primary and a secondary DNS server.

A WINS server provides IP addresses for named servers in a Windows domain. If SSL VPN users will access a Windows network, you need to provide them access to the domain WINS server. You specify a primary and a secondary WINS server.

To specify WINS and DNS services for clients - web-based manager

- 1 Go to *VPN > SSL > Config*.
- 2 Select the *Expand Arrow* to display the *Advanced* section.
- 3 Enter the IP addresses of DNS servers in the *DNS Server* fields as needed.
- 4 Enter the IP addresses of WINS servers in the *WINS Server* fields as needed.
- 5 Select *Apply*.

To specify WINS and DNS services for clients - CLI

```
config vpn ssl settings
  set dns-server1 <address_ipv4>
  set dns-server2 <address_ipv4>
  set wins-server1 <address_ipv4>
  set wins-server2 <address_ipv4>
end
```

Setting the idle timeout setting

The idle timeout setting controls how long the connection can remain idle before the system forces the remote user to log in again. For security, keep the default value of 300 seconds (5 minutes) or less.

To set the idle timeout - web-based manager

- 1 Go to *VPN > SSL > Config*.
- 2 In the *Idle Timeout* field, enter the timeout value.
The valid range is from 10 to 28800 seconds.
- 3 Select *Apply*.

To set the idle timeout - CLI

```
config vpn ssl settings
  set idle-timeout <seconds_int>
end
```

SSL VPN logs

Logging is available for SSP VPN traffic so you can monitor users connected to the FortiGate unit and their activity.

For more information on configuring logs on the FortiGate unit, see the [Logging and Reporting](#) chapter of *The Handbook*.

To enable logging of SSL VPN events - web-based manager

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select *Enable*, and then select one or more of the following options:
 - SSL VPN user authentication event
 - SSL VPN administration event
 - SSL VPN session event
- 3 Select *Apply*.

To enable logging of SSL VPN events - CLI

```
config log {fortianalyzer | memory | syslog} filter
    set event enable
    set sslvpn-log-adm enable
    set sslvpn-log-auth enable
    set sslvpn-log-session enable
end
```

Logging of SSL VPN traffic is configured when you create the security policy, by selecting the *Log Allowed Traffic* option in the web-based manager or using the CLI command `set logtraffic enable` under `config firewall policy`.

Viewing log data

To view the SSL VPN log data, in the web-based manager, go to *Log&Report > Log & Archive Access* and select either the *Event Log* or *Traffic Log*.

In event log entries, look for the sub-types “sslvpn-session” and “sslvpn-user”.

In the traffic logs, look for the sub-type “allowed”. For web-mode traffic, the source is the host IP address. For tunnel-mode traffic, the source is the address assigned to the host from the SSL VPN address pool.

For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

Monitoring active SSL VPN sessions

You can go to *User > Monitor* to view a list of active SSL VPN sessions. The list displays the user name of the remote user, the IP address of the remote client, and the time the connection was made. You can also see which services are being provided, and delete an active web session from the FortiGate unit.

To monitor SSL VPNs - web-based manager

To view the list of active SSL VPN sessions, go to *VPN > SSL > Monitor*.

When a tunnel-mode user is connected, the *Description* field displays the IP address that the FortiGate unit assigned to the remote host.

If required, you can end a session/connection by selecting its check box and then selecting the *Delete* icon.

To monitor SSL VPNs - CLI

To list all of the SSL VPN sessions and their index numbers:

```
get vpn ssl monitor
```

To delete tunnel-mode or web-mode sessions:

```
execute vpn sslvpn del-tunnel <index_int>
```

```
execute vpn sslvpn del-web <index_int>
```

Troubleshooting

Here is a list of common SSL VPN problems and the likely solutions.

No response from SSL VPN URL	Check SSL VPN port assignment (default 10443). Verify the SSL VPN security policy.
Error: "The web page cannot be found."	Check URL: <code>https://<FortiGate_IP>:<SSLVPN_port>/remote/login</code>
Tunnel connects, but there is no communication.	Check that there is a static route to direct packets destined for the tunnel users to the SSL VPN interface. See "Routing for tunnel mode" on page 1635 .
Tunnel-mode connection shuts down after a few seconds	This issue occurs when there are multiple interfaces connected to the Internet, for example, a dual WAN configuration. Upgrade to the latest firmware then use the following CLI command: <pre>config vpn ssl settings set route-source-interface enable end</pre>
Error: "Destination address of Split Tunneling policy is invalid."	The SSL VPN security policy uses the ALL address as its destination. Specify the address of the protected network instead.
When trying to connect using FortiClient the error message "Unable to logon to the server. Your user name or password may not be configured properly for this connection. (-12)" appears. When trying to login to the web portal, login and password are entered and login page will be sent back.	Cookies must be enabled for SSL VPN to function in Web portal or with FortiClient. Access to the web portal or tunnel will fail if Internet Explorer has the privacy Internet Options set to High. If set to High, Internet Explorer will: Block cookies that do not have a compact privacy policy. Block cookies that use personally identifiable information without your explicit consent.



The SSL VPN client

The remote client connects to the SSL VPN tunnel in various ways, depending on the VPN configuration.

- Web mode requires nothing more than a web browser. Microsoft Internet Explorer, Mozilla Firefox, and Apple Safari browsers are supported. For detailed information about supported browsers see the Release Notes for your FortiOS firmware.
- Tunnel mode establishes a connection to the remote protected network that any application can use. This requires FortiClient SSL VPN application that sends and receives data through the SSL VPN tunnel.

If the client computer runs Microsoft Windows, they can download the tunnel mode client from the web portal Tunnel Mode widget. After installing the client, they can start and stop tunnel operation from the Tunnel Mode widget, or open the tunnel mode client as a standalone application. The tunnel mode client is available on the Start menu at *All Programs > FortiClient > FortiClient SSL VPN*.

If the client computer runs Linux or Mac OS X, the user needs to download the tunnel mode client application from the Fortinet Support web site. See the Release Notes for your FortiOS firmware for the specific operating system versions that are supported. On Linux and Mac OS X platforms, tunnel mode operation cannot be initiated from the web portal Tunnel Mode widget. The remote user must use the standalone tunnel client application.

- The virtual desktop application creates a virtual desktop on a user's PC and monitors the data read/write activity of the web browser running inside the virtual desktop. When the application starts, it presents a 'virtual desktop' to the user. The user starts the web browser from within the virtual desktop and connects to the SSL VPN web portal. The browser file/directory operation is redirected to a new location, and the data is encrypted before it is written to the local disk. When the virtual desktop application exits normally, all the data written to the disk is removed. If the session terminates abnormally (power loss, system failure), the data left behind is encrypted and unusable to the user. The next time you start the virtual desktop, the encrypted data is removed.

FortiClient

Remote users can use FortiClient software to initiate an SSL VPN tunnel to connect to the internal network. FortiClient uses local port TCP 1024 to initiate an SSL encrypted connection to the FortiGate unit, on port TCP 10443. When connection using FortiClient, the FortiGate unit authenticates the FortiClient SSL VPN request based on the user group options. The FortiGate unit establishes a tunnel with the client and assigns a virtual IP address to the client PC. Once the tunnel has been established, the user can access the network behind the FortiGate unit.

There are three FortiClient application options available, depending on your requirements:

- FortiClient/FortiClient Premium
- FortiClient Lite (free) - http://download.cnet.com/FortiClient-Lite/3000-2239_4-75532356.html
- Standalone client (free)
- iPhone and iPad app available for free from the iTunes App Store.

For details on configuring FortiClient for SSL VPN, see the FortiClient documentation.

Downloading the SSL VPN tunnel mode client

SSL VPN standalone tunnel client applications are available for Windows, Linux, and Mac OS X systems (see the Release Notes for your FortiOS firmware for the specific versions that are supported). There are separate download files for each operating system.



Windows users can also download the tunnel mode client from an SSL VPN web portal that contains the Tunnel Mode widget.

The most recent version of the SSL VPN standalone client applications can be found at: <http://support.fortinet.com/>

To download the SSL VPN tunnel client

- 1 Log in to Fortinet Support at <http://support.fortinet.com/>.
- 2 In the Download area, select Firmware Images.
- 3 Select FortiGate.
- 4 Select v4.00 and then select the latest firmware release.
- 5 Select *SSL VPN Clients*.
- 6 Select the appropriate client.

Windows: `SslvpnClient.exe` or `SslvpnClient.msi`

Linux: `forticlientsslvpn_linux_<version>.tar.gz`

Mac OS X: `forticlientsslvpn_macosx_<version>.dmg`



The location of the SSL VPN tunnel client on the Support web site is subject to change. If you have difficulty finding the appropriate file, contact Customer Support.

Tunnel mode client configuration

The FortiClient SSL VPN tunnel client requires basic configuration by the remote user to connect to the SSL VPN tunnel. When distributing the FortiClient software, provide the following information for the remote user to enter once the client software has been started. Once entered, they can select *Connect* to begin a SSL VPN session.

Connection Name	If you have pre-configured the connection settings, select the connection from the list and then select <i>Connect</i> . Otherwise, enter the settings in the fields below.
Server Address	Enter the IP address or FQDN of the FortiGate unit that hosts the SSL VPN.
Username	Enter your user name.
Password	Enter the password associated with your user account.
Client Certificate	Use this field if the SSL VPN requires a certificate for authentication. Select the required certificate from the drop-down list. The certificate must be installed in the Internet Explorer certificate store.
Settings...	Select to open the <i>Settings</i> dialog and select the <i>Keep connection alive until manually stopped</i> check box to prevent tunnel connections from closing due to inactivity.

Uninstalling the tunnel mode client

If you want to remove the tunnel mode client application, follow the instructions for your operating system.

To uninstall from Windows


- 1 In the *Control Panel*, select *Programs and Features* (*Add or Remove Programs* in Windows XP).
- 2 Select *FortiClient SSL VPN* and then *Remove*.

To uninstall from Linux

Remove/delete the folder containing all the SSL VPN client application files.

To uninstall from Mac OS X

In the Applications folder, select `forticlientsslvpn.app` and drag it into the Trash.



Setup examples


The examples in this chapter demonstrate the basic configurations needed for common connections to the SSL VPN tunnel and portals, applying the steps outlined in the chapter “Basic Configuration” on page 1617.

The example included are:

- Secure internet browsing
- Split Tunnel
- Multiple user groups with different access permissions example

Secure internet browsing

This example sets up an SSL VPN tunnel to provide remote users the ability to access the Internet while travelling, and ensure that they are not subjected to malware and other dangers, by using the corporate firewall to filter all of their Internet traffic. Essentially, the remote user will connect to the corporate FortiGate unit to surf the Internet.



Using SSL VPN and FortiClient SSL VPN software, you create a means to use the corporate FortiGate to browse the web safely.

Creating an SSL VPN IP pool and SSL VPN web portal

- 1 Go to *VPN > SSL > Config* and for *IP Pools* select *Edit* and add *SSLVPN_TUNNEL_ADDR1* to the *Selected* table.
- 2 Create the SSL VPN portal to by going to *VPN > SSL > Portal* and selecting *tunnel-access*.

- 3 Select the *Edit* pencil icon for the Tunnel Mode widget and enter the following:

Name	Browsing
IP Mode	User Group
IP Pools	SSLVPN_TUNNEL_ADDR1

- 4 Select *OK*.

Creating the SSL VPN user and user group

Create the SSL VPN user and add the user to a user group configured for SSL VPN use.

- 1 Go to *User > User > User* and select *Create New* to add the user:

User Name	twhite
Password	password

- 2 Select *OK*.
- 3 Go to *User > User Group > User Group* and select *Create New* to add *twhite* to the SSL VPN user group:

Name	Tunnel
Type	Firewall
Allow SSL-VPN Access	tunnel-access



Make sure you select the *Allow SSL VPN Access* option. If not selected, the *Tunnel* user group will not appear in the group list when configuring the authentication security policy.

- 4 Move *twhite* to the *Members* list.
- 5 Select *OK*.

Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

- 1 Go to *Router > Static > Static* and select *Create New* to add the static route:

Destination IP/Mask	10.212.134.0/255.255.255.0
Device	ssl.root



The *Destination IP/Mask* matches the network address of the remote SSL VPN user.

- 2 Select *OK*.

Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit. Create a normal security policy from ssl.root to wan1 to allow SSL VPN traffic to connect to the Internet.

- 1 Go to **Policy > Policy > Policy** and select **Create New** to add the SSL VPN security policy:

Source Interface/Zone	wan1
Source Address	all
Destination Interface/Zone	ssl.root
Destination Address	all
Action	SSL-VPN

- 2 Select *Configure SSL-VPN Users* and select *Add* to add an authentication rule for the remote user:

Selected User Groups	Tunnel
Selected Services	ANY
Schedule	always



If the *Tunnel* user group does not appear in the *User Group* list, ensure you select the *SSL VPN Access* option when creating the user group. If that option is not selected, the *Tunnel* user group will not appear in the user group list when configuring the authentication security policy.

- 3 Select *OK*.
- 4 Select *Create New* to add a security policy that allows remote SSL VPN users to connect to the Internet:

Source Interface/Zone	ssl.root
Source Address	all
Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 5 Select *OK*.

Results


Using FortiClient SSLVPN application, log into the VPN using the address `https://172.20.120.136:10443/` and log in as `twhite`. Once connected, you can browse the Internet.

From the FortiGate web-based manager go to *VPN > Monitor > SSL-VPN Monitor* to view the list of users connected using SSL VPN. The *Subsession* entry indicates the split tunnel which redirects to the Internet.

Split Tunnel

For this example, the remote users are configured to be able to securely access head office internal network servers, and browse the Internet through the head office firewall. This will enable the remote user to use the FortiGate security to connect to the internal network and the web.

This solution describes how to configure FortiGate SSL VPN split tunnelling using the FortiClient SSL VPN software, available from the Fortinet Support site.



Using split tunneling, all communication from remote SSL VPN users to the head office internal network and to the Internet uses an SSL VPN tunnel between the user's PC and the head office FortiGate unit. Connections to the Internet are routed back out the head office FortiGate unit to the Internet. Replies come back into the head office FortiGate unit before being routed back through the SSL VPN tunnel to the remote user.

Creating a firewall address for the head office server

- 1 Go to *Firewall Objects > Address > Address* and select *Create New* and add the head office server address:

Address Name	Head office server
Type	Subnet / IP Range
Subnet / IP Range	192.168.1.12
Interface	Internal

- 2 Select *OK*.

Creating an SSL VPN IP pool and SSL VPN web portal

- 1 Go to *VPN > SSL > Config* and for *IP Pools* select *Edit* and add *SSLVPN_TUNNEL_ADDR1* to the *Selected* table.
- 2 Create the SSL VPN portal to by going to *VPN > SSL > Portal* and selecting *tunnel-access*.

- 3 Select the *Edit* pencil icon for the *Tunnel Mode* widget and enter the following:

Name	Connect to head office server
IP Mode	User Group
IP Pools	SSLVPN_TUNNEL_ADDR1
Split Tunneling	Enable

- 4 Select *OK*.

Creating the SSL VPN user and user group

Create the SSL VPN user and add the user to a user group configured for SSL VPN use.

- 1 Go to *User > User > User*, select *Create New* and add the user:

User Name	twhite
Password	password

- 2 Select *OK*.
- 3 Go to *User > User Group > User Group* and select *Create New* to add *twhite* to the SSL VPN user group:

Name	Tunnel
Type	Firewall
Allow SSL-VPN Access	tunnel-access



Make sure you select the *Allow SSL-VPN Access* option. If not selected, the *Tunnel* user group will not appear in the group list when configuring the authentication security policy.

- 4 Move *twhite* to the *Members* list.
- 5 Select *OK*.

Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

- 1 Go to *Router > Static > Static* and select *Create New* to add the static route:

Destination IP/Mask	10.212.134.0/255.255.255.0
Device	ssl.root



The *Destination IP/Mask* matches the network address of the remote SSL VPN user.

- 2 Select *OK*.

Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit. Create a normal security policy from ssl.root to wan1 to allow SSL VPN traffic to connect to the Internet.

- 1 Go to *Policy > Policy > Policy* and select *Create New* to add the SSL VPN security policy:

Source Interface/Zone	wan1
Source Address	all
Destination Interface/Zone	internal
Destination Address	Head office server
Action	SSL-VPN

- 2 Select *Configure SSL-VPN Users* and select *Add* to add an authentication rule for the remote user:

Selected User Groups	Tunnel
Selected Services	ANY
Schedule	always



If the *Tunnel* user group does not appear in the *User Group* list, ensure you select the *SSL VPN Access* option when creating the user group. If that option is not selected, the *Tunnel* user group will not appear in the user group list when configuring the authentication security policy.

- 3 Select *OK*.
- 4 Select *Create New* to add a security policy that allows remote SSL VPN users to connect to the Internet:

Source Interface/Zone	ssl.root
Source Address	all
Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 5 Select *OK*.

Results


Using the FortiClient SSL VPN application on the remote PC, connect to the VPN using the address <https://172.20.120.136:10443/> and log in with the *twhite* user account. Once connected, you can connect to the head office server or browse to web sites on the Internet.

From the web-based manager go to *VPN > Monitor > SSL-VPN Monitor* to view the list of users connected using SSL VPN. The *Subsession* entry indicates the split tunnel which redirects SSL VPN sessions to the Internet.

Multiple user groups with different access permissions example

You might need to provide access to several user groups with different access permissions. Consider the following example topology in which users on the Internet have controlled access to servers and workstations on private networks behind a FortiGate unit.

Figure 161: SSL VPN configuration for different access permissions by user group



In this example configuration, there are two users:

- user1 can access the servers on Subnet_1
- user2 can access the workstation PCs on Subnet_2

You could easily add more users to either user group to provide them access to the user group's assigned web portal.

General configuration steps

- 1 Create firewall addresses for
 - the destination networks
 - two non-overlapping tunnel IP address ranges that the FortiGate unit will assign to tunnel clients in the two user groups
- 2 Create two web portals.
- 3 Create two user accounts, user1 and user2.
- 4 Create two user groups. For each group, add a user as a member and select a web portal. In this example, user1 will belong to group1, which will be assigned to portal1.
- 5 Create security policies:
 - two SSL VPN security policies, one to each destination
 - two tunnel-mode policies to allow each group of users to reach its permitted destination network
- 6 Create the static route to direct packets for the users to the tunnel.

Creating the firewall addresses

Security policies do not accept direct entry of IP addresses and address ranges. You must define firewall addresses in advance.

Creating the destination addresses

SSL VPN users in this example can access either Subnet_1 or Subnet_2.

To define destination addresses - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Subnet_1
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0/24
Interface	port2

- 3 Select *Create New*, enter the following information, and select *OK*.

Address Name	Subnet_2
Type	Subnet / IP Range
Subnet / IP Range	10.11.201.0/24
Interface	port3

To define destination addresses - CLI

```
config firewall address
  edit Subnet_1
    set type ipmask
    set subnet 10.11.101.0/24
    set associated-interface port2
  next
  edit Subnet_2
    set type ipmask
    set subnet 10.11.201.0/24
    set associated-interface port3
  end
```

Creating the tunnel client range addresses

To accommodate the two groups of users, split an otherwise unused subnet into two ranges. The tunnel client addresses must not conflict with each other or with other addresses in your network.

To define tunnel client addresses - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Address Name	Tunnel_group1
Type	Subnet / IP Range

Subnet / IP Range	10.11.254.[1-50]
Interface	Any

- 3 Select *Create New*, enter the following information, and select *OK*.

Address Name	Tunnel_group2
Type	Subnet / IP Range
Subnet / IP Range	10.11.254.[51-100]
Interface	Any

To define tunnel client addresses - CLI

```
config firewall address
  edit Tunnel_group1
    set type iprange
    set end-ip 10.11.254.50
    set start-ip 10.11.254.1
  next
  edit Tunnel_group2
    set type iprange
    set end-ip 10.11.254.100
    set start-ip 10.11.254.51
end
```

Creating the web portals

To accommodate two different sets of access permissions, you need to create two web portals, portal1 and portal2, for example. Later, you will create two SSL VPN user groups, one to assign to portal1 and the other to assign to portal2.

To create the portal1 web portal

- 1 Go to *VPN > SSL > Portal* and select *Create New*.
- 2 Enter `portal1` in the *Name* field and select *OK*.
- 3 In *Applications*, select all of the application types that the users can access.
- 4 Select the *Edit* icon on the *Tunnel Mode* widget.
- 5 In *IP Pools*, select *Edit*.
- 6 In the *Available* list, select *Tunnel_group1* and then select the down arrow button. Select *OK*.
- 7 Select *OK* in the *Tunnel Mode* widget.
- 8 Select *OK*.

To create the portal2 web portal

- 1 Go to *VPN > SSL > Portal* and select *Create New*.
- 2 Enter `portal2` in the *Name* field and select *OK*.
- 3 In *Applications*, select all of the application types that the users can access.
- 4 Select the *Edit* icon on the *Tunnel Mode* widget.
- 5 In *IP Pools*, select *Edit*.
- 6 In the *Available* list, select *Tunnel_group2* and then select the down arrow button. Select *OK*.

7 Select *OK* in the *Tunnel Mode* widget.

8 Select *OK*.

To create the web portals - CLI

```
config vpn ssl web portal
edit portal1
set allow-access ftp ping rdp smb ssh telnet vnc web
config widget
edit 0
set type tunnel
set tunnel-status enable
set ip-pools "Tunnel_group1"
end
next
edit portal2
set allow-access ftp ping rdp smb ssh telnet vnc web
config widget
edit 0
set type tunnel
set tunnel-status enable
set ip-pools "Tunnel_group2"
end
end
end
```

Later, you can configure these portals with bookmarks and enable connection tool capabilities for the convenience of your users.

Creating the user accounts and user groups

After enabling SSL VPN and creating the web portals that you need, you need to create the user accounts and then the user groups that require SSL VPN access.

Go to *User > User* and create *user1* and *user2* with password authentication. After you create the users, create the SSL VPN user groups.

To create the user groups - web-based manager

- 1 Go to *User > User Group > User Group*.
- 2 Select *Create New* and enter the following information:

Name	group1
Type	SSL VPN
Portal	portal1

- 3 From the *Available* list, select *user1* and move it to the *Members* list by selecting the right arrow button.
- 4 Select *OK*.
- 5 Repeat steps 2 through 4 to create *group2*, assigned to *portal2*, with *user2* as its only member.

To create the user groups - CLI

```
config user group
edit group1
```

```

set group-type sslvpn
set member user1
set sslvpn-portal portal1
next
edit group2
set group-type sslvpn
set member user2
set sslvpn-portal portal2
end

```

Creating the security policies

You need to define security policies to permit your SSL VPN clients, web-mode or tunnel-mode, to connect to the protected networks behind the FortiGate unit. Before you create the security policies, you must define the source and destination addresses to include in the policy. See [“Creating the firewall addresses” on page 1662](#).

Two types of security policy are required:

- An SSL VPN policy enables clients to authenticate and permits a web-mode connection to the destination network. In this example, there are two destination networks, so there will be two SSL VPN policies. The authentication, ensures that only authorized users access the destination network.
- A tunnel-mode policy is a regular ACCEPT security policy that enables traffic to flow between the SSL VPN tunnel interface and the protected network. Tunnel-mode policies are required if you want to provide tunnel-mode connections for your clients. In this example, there are two destination networks, so there will be two tunnel-mode policies.

To create the SSL VPN security policies - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New and* enter the following information:

Source Interface/Zone	port1
Source Address	All
Destination Interface/Zone	port2
Destination Address	Subnet_1
Action	SSL-VPN

- 3 Select *Add* and enter the following information:

User Group	group1
Service	Any

- 4 Select *OK*, and then select *OK* again.
- 5 Select *Create New and* enter the following information:

Source Interface/Zone	port1
Source Address	All
Destination Interface/Zone	port3

Destination Address	Subnet_2
Action	SSL-VPN

- 6 Select *Add* and enter the following information:

User Group	group2
Service	Any

- 7 Select *OK*, and then select *OK* again.

To create the SSL VPN security policies - CLI

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr Subnet_1
    set action ssl-vpn
    set nat enable
    config identity-based-policy
      edit 1
        set groups group1
        set schedule always
        set service ANY
      end
    next
  edit 0
    set srcintf port1
    set dstintf port3
    set srcaddr all
    set dstaddr Subnet_2
    set action ssl-vpn
    set nat enable
    config identity-based-policy
      edit 1
        set groups group2
        set schedule always
        set service ANY
      end
    end
  end
```

To create the tunnel-mode security policies - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	sslvpn tunnel interface (ssl.root)
Source Address	Tunnel_group1
Destination Interface/Zone	port2
Destination Address	Subnet_1
Action	ACCEPT
NAT	Enable

- 3 Select *Create New*, enter the following information, and select *OK*:

Source Interface/Zone	sslvpn tunnel interface (ssl.root)
Source Address	Tunnel_group2
Destination Interface/Zone	port3
Destination Address	Subnet_2
Action	ACCEPT
NAT	Enable

To create the tunnel-mode security policies - CLI

```
config firewall policy
  edit 0
    set srcintf ssl.root
    set dstintf port2
    set srcaddr Tunnel_group1
    set dstaddr Subnet_1
    set action accept
    set schedule always
    set service ANY
    set nat enable
  next
  edit 0
    set srcintf ssl.root
    set dstintf port3
    set srcaddr Tunnel_group2
    set dstaddr Subnet_2
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
end
```

Create the static route to tunnel mode clients

Reply packets destined for tunnel mode clients must pass through the SSL VPN tunnel. You need to define a static route to accomplish this.

To add a route to SSL VPN tunnel mode clients - web-based manager

- 1 Go to *Router > Static > Static Route* and select *Create New*.
- 2 Enter the following information and select *OK*.

Destination IP/Mask	10.11.254.0/24 This IP address range covers both ranges that you assigned to SSL VPN tunnel-mode users. See “Creating the tunnel client range addresses” on page 1662 .
Device	Select the SSL VPN virtual interface, <i>ssl.root</i> for example.
Leave other settings at their default values.	

To add a route to SSL VPN tunnel mode clients - CLI

```
config router static
edit 0
    set device ssl.root
    set dst 10.11.254.0/24
end
```

Enabling SSL VPN operation

By default, SSL VPN is not enabled. SSL VPN is configured in the CLI only.

To enable SSL VPN and set tunnel address range - CLI

```
config vpn ssl settings
    set sslvpn-enable enable
    set tunnel-ip-pools SSL_tunnel_users
end
```



In this example, the *IP Pools* field on the *VPN > SSL > Config* page is not used because each web portal specifies its own tunnel IP address range.



Chapter 10 Advanced Routing

This chapter describes advanced static routing concepts and how to implement dynamic routing on FortiGate units.

This FortiOS Handbook chapter contains the following sections:

[Advanced Static routing](#) explains routing concepts, equal cost multipath (ECMP) and load balancing, policy routing, and routing in transparent mode.

[Dynamic Routing Overview](#) provides some basic routing concepts needed to explain dynamic routing, compares static and dynamic routing, and walks you through deciding which dynamic routing protocol is best for you.

[Routing Information Protocol \(RIP\)](#), [Border Gateway Protocol \(BGP\)](#), [Open Shortest Path First \(OSPF\)](#), and [Intermediate System To Intermediate System Protocol \(IS-IS\)](#) provide background on the protocol, explains the terms used, how the protocol works, looks at some troubleshooting, and examples on configuring the protocols in different situations.

[Router Reference](#) describes the web-based manager Router pages, providing explanations for each field.



Advanced Static routing

Advanced static routing includes features and concepts that are used in more complex networks. Dynamic routing is not addressed in this section.

This section includes:

- [Routing concepts](#)
- [ECMP route failover and load balancing](#)
- [Static routing tips](#)
- [Policy Routing](#)
- [Transparent mode static routing](#)

Routing concepts

Many routing concepts apply to static routing. However without first understanding these basic concepts, it is difficult to understand the more complex dynamic routing.

This section includes:

- [Routing in VDOMs](#)
- [Default route](#)
- [Routing table](#)
- [Building the routing table](#)
- [Static routing security](#)
- [Multipath routing and determining the best route](#)

Routing in VDOMs

Routing on FortiGate units is configured per-VDOM. This means if VDOMs are enabled, you must enter a VDOM to do any routing configuration. This allows each VDOM to operate independently, with its own default routes and routing configuration.

In this guide, the procedures assume your FortiGate unit has VDOMs disabled. This is stated in the assumptions for the examples. If you have VDOMs enabled you will need to perform the following steps in addition to the procedure's steps.

To route in VDOMs - web-based manager

Select the VDOM that you want to view or configure at the bottom of the main menu.

To route in VDOMs - CLI

Before following any CLI routing procedures with VDOMs enabled, enter the following commands. For this example, it is assumed you will be working in the root VDOM. Change root to the name of your selected VDOM as needed.

```
config vdom
edit root
```

Following these commands, you can enter any routing CLI commands as normal.

Default route

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

All routers, including FortiGate units, are shipped with default routes in place. This allows customers to set up and become operational more quickly. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address.

Routing table

When two computers are directly connected, there is no need for routing because each computer knows exactly where to find the other computer. They communicate directly.

Networking computers allows many computers to communicate with each other. This requires each computer to have an IP address to identify its location to the other computers. This is much like a mailing address - you will not receive your postal mail at home if you do not have an address for people to send mail to. The routing table on a computer is much like an address book used to mail letters to people in that the routing table maintains a list of how to reach computers. Routing tables may also include information about the quality of service (QoS) of the route, and the interface associated with the route if the device has multiple interfaces.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information, but also the source information to ensure that it exists. If there is no source to be found, then that packet is dropped because the router assumes it to be an error or an attack on the network.

Looking at routing as delivering letters is more simple than reality. In reality, routers loose power or have bad cabling, network equipment is moved without warning, and other such events happen that prevent static routes from reaching their destinations. When any changes such as these happen along a static route, traffic can no longer reach the destination — the route goes down. Dynamic routing can address these changes to ensure traffic still reaches its destination. The process of realizing there is a problem, backtracking and finding a route that is operational is called convergence. If there is fast convergence in a network, users won't even know that re-routing is taking place.

The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes—the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information, but also the source information to ensure that it exists. If there is no source to be found, then that packet is dropped because the router assumes it to be an error or an attack on the network.

The routing table is used to store routes that are learned. The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes — the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Some actions you can perform on the routing table include:

- [Viewing the routing table in the web-based manager](#)
- [Viewing the routing table in the CLI](#)
- [Searching the routing table](#)

Viewing the routing table in the web-based manager

VDOM

By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of “any/all” packets.

To display the routes in the routing table, go to *Router > Monitor > Routing Monitor*.

[Figure 162](#) shows the Routing Monitor list belonging to a FortiGate unit that has interfaces named “port1”, “port4”, and “lan”. The names of the interfaces on your FortiGate unit may be different.

[Figure 163](#) shows the Routing Monitor list when IPv6 has been selected. Note that the information available for IPv6 is limited.

Figure 162: Routing Monitor list - IPv4

IP Version: IPv4 Type: All Network:		Gateway:		Apply Filter			
Type	Subtype	Network	Distance	Metric	Gateway	Interface	Up Time (d h:m:s)
Connected		10.10.10.0/24	0	0	0.0.0.0	port4	
Connected		172.20.120.0/24	0	0	0.0.0.0	port1	

Figure 163: Routing Monitor list - IPv6

IP Version: IPv6		
Interface	Network	Gateway
hvdlink1	fe80::/10	
hvdlink1	::/8	

IP version:	Select IPv4 or IPv6. This is available only when IPv6 is enabled in the web-based manager. The fields displayed in the table depend on which IP version is selected.
Type:	<p>Select one of the following route types to search the routing table and display routes of the selected type only:</p> <p><i>All</i> — all routes recorded in the routing table.</p> <p><i>Connected</i> — all routes associated with direct connections to FortiGate unit interfaces.</p> <p><i>Static</i> — the static routes that have been added to the routing table manually.</p> <p><i>RIP</i> — all routes learned through RIP. For more information see “Routing Information Protocol (RIP)” on page 1715.</p> <p><i>BGP</i> — all routes learned through BGP. For more information see “Border Gateway Protocol (BGP)” on page 1751.</p> <p><i>OSPF</i> — all routes learned through OSPF. For more information see “Open Shortest Path First (OSPF)” on page 1787.</p> <p><i>IS-IS</i> — all routes learned through IS-IS. For more information see “Intermediate System To Intermediate System Protocol (IS-IS)” on page 1827.</p> <p><i>HA</i> — RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you are viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster.</p> <p>Not displayed when IP version IPv6 is selected.</p> <p>For details about HA routing synchronization, see the FortiGate HA User Guide.</p>
Network:	<p>Enter an IP address and netmask (for example, 172.16.14.0/24) to search the routing table and display routes that match the specified network.</p> <p>Not displayed when IP version IPv6 is selected.</p>
Gateway:	<p>Enter an IP address and netmask (for example, 192.168.12.1/32) to search the routing table and display routes that match the specified gateway.</p> <p>Not displayed when IP version IPv6 is selected.</p>
Apply Filter	<p>Select to search the entries in the routing table based on the specified search criteria and display any matching routes.</p> <p>Not displayed when IP version IPv6 is selected.</p>

Type	<p>The type values assigned to FortiGate unit routes (Static, Connected, RIP, OSPF, or BGP).</p> <p>Not displayed when IP version IPv6 is selected.</p>
Subtype	<p>If applicable, the subtype classification assigned to OSPF routes.</p> <p>An empty string implies an intra-area route. The destination is in an area to which the FortiGate unit is connected.</p> <p><i>OSPF inter area</i> — the destination is in the OSPF AS, but the FortiGate unit is not connected to that area.</p> <p><i>External 1</i> — the destination is outside the OSPF AS. This is known as OSPF E1 type. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.</p> <p><i>External 2</i> — the destination is outside the OSPF AS. This is known as OSPF E2 type. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost.</p> <p><i>OSPF NSSA 1</i> — same as External 1, but the route was received through a not-so-stubby area (NSSA).</p> <p><i>OSPF NSSA 2</i> — same as External 2, but the route was received through a not-so-stubby area.</p> <p>For more information on OSPF subtypes, see “OSPF Background and concepts” on page 1787.</p> <p>Not displayed when IP version 6 is selected.</p>
Network	<p>The IP addresses and network masks of destination networks that the FortiGate unit can reach.</p>
Distance	<p>The administrative distance associated with the route. A value of 0 means the route is preferable compared to other routes to the same destination.</p> <p>Modifying this distance for dynamic routes is route distribution.</p>
Metric	<p>The metric associated with the route type. The metric of a route influences how the FortiGate unit dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to.</p> <p><i>Hop count</i> — routes learned through RIP.</p> <p><i>Relative cost</i> — routes learned through OSPF.</p> <p><i>Multi-Exit Discriminator (MED)</i> — routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network. For more information on BGP attributes, see “BGP attributes” on page 1757.</p>
Gateway	<p>The IP addresses of gateways to the destination networks.</p>
Interface	<p>The interface through which packets are forwarded to the gateway of the destination network.</p>
Up Time	<p>The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable.</p> <p>Not displayed when IP version IPv6 is selected.</p>

Viewing the routing table in the CLI

In the CLI, you can easily view the static routing table just as in the web-based manager or you can view the full routing table.

When viewing the list of static routes using the CLI command `get route static`, it is the configured static routes that are displayed. When viewing the routing table using the CLI command `get router info routing-table all`, it is the entire routing table information that is displayed including configured and learned routes of all types. The two are different information in different formats.



If VDOMs are enabled on your FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

To view the routing table

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
    inter area
* - candidate default

S*      0.0.0.0/0 [10/0] via 192.168.183.254, port2
S       1.0.0.0/8 [10/0] via 192.168.183.254, port2
S       2.0.0.0/8 [10/0] via 192.168.183.254, port2
C       10.142.0.0/23 is directly connected, port3
B       10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C       192.168.182.0/23 is directly connected, port2
```

Examining an entry:

```
B       10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route including netmask.
[20/0]	20 indicates and administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
10.142.0.74	The gateway, or next hop.
port3	The interface used by this route.
2d18h02m	How old this route is, in this case almost three days old.

To view the kernel routing table

```
# get router info kernel

tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0-
>10.11.201.0/24 pref=10.11.201.4 gwy=0.0.0.0 dev=5(external1)
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0-
>172.20.120.0/24 pref=172.20.120.146 gwy=0.0.0.0 dev=6(internal)
```

The parts of the routing table entry are:

tab	table number. This will be either 254 (unicast) or 255 (multicast).
vf	virtual domain of the firewall. This is the vdom index number. If vdoms are not enabled, this number will be 0.
type	type of routing connection. Valid values include: 0 - unspecified 1 - unicast 2 - local 3 - broadcast 4 - anycast 5 - multicast 6 - blackhole 7 - unreachable 8 - prohibited
proto	type of installation. This indicates where the route came from. Valid values include: 0 - unspecified 2 - kernel 11 - ZebOS routing module 14 - FortiOS 15 - HA 16 - authentication based 17 - HA1
prio	priority of the route. Lower priorities are preferred.
->10.11.201.0/24 (->x.x.x.x/mask)	the IP address and subnet mask of the destination
pref	preferred next hop along this route
gwy	gateway - the IPv4 address of the gateway this route will use
dev	outgoing interface index. This number is associated with the interface for this route, and if VDOMs are enabled the VDOM will be included here as well. If an interface alias is set for this interface it will also be displayed here.

Searching the routing table

You can apply a filter to search the routing table and display certain routes only. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed — an implicit AND condition is applied to all of the search parameters you specify.

For example, if the FortiGate unit is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select *Connected* from the *Type* list, type 172.16.14.0/24 in the *Network* field, and then select *Apply Filter* to display the associated routing table entry or entries. Any entry that contains the word “Connected” in its *Type* field and the specified value in the *Gateway* field will be displayed.

In this example, you will apply a filter to search for an entry for static route to 10.10.10.10/24

To search the FortiGate unit routing table in the web-based manager

- 1 Go to *Router > Monitor > Routing Monitor*.
- 2 From the *Type* list, select the type of route to display. In our example, select *Static*.
- 3 If you want to display routes to a specific network, type the IP address and netmask of the network in the *Networks* field. In our example, enter 10.10.10.10/24.
- 4 If you want to display routes to a specific gateway, type the IP address of the gateway in the *Gateway* field.
- 5 Select *Apply Filter*.



All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.

To search the FortiGate unit routing table in the CLI

```
FGT # get router info routing-table details 10.10.10.10
Routing entry for 10.10.10.10/24
  Known via "static", distance 10, metric 0, best
```

If there are multiple routes that match your filter, they will all be listed, with the best match at the top of the list as indicated by the word *best*.

Building the routing table

In the factory default configuration, the FortiGate unit routing table contains a single static default route. You can add routing information to the routing table by defining additional static routes.

It is possible that the routing table is faced with several different routes to the same destination — the IP addresses of the next-hop router specified in those routes or the FortiGate interfaces associated with those routes may vary. In this situation, the “best” route is selected from the table.

The FortiGate unit selects the “best” route for a packet by evaluating the information in the routing table. The “best” route to a destination is typically associated with the shortest distance between the FortiGate unit and the closest gateway, also known as a next-hop router. In some cases, the next best route may be selected if the best route is unavailable.

The FortiGate unit installs the best available routes in the unit's forwarding table, which is a subset of the unit's routing table. Packets are forwarded according to the information in the forwarding table.

Static routing security

Securing the information on your company network is a top priority for network administrators. Security is also required as the routing protocols used are internationally known standards that typically provide little or no inherent security by themselves.

The two reasons for securing your network are the sensitive and proprietary information on your network, and also your external bandwidth. Hackers not only can steal your information, but they can also steal your bandwidth. Routing is a good low level way to secure your network, even before UTM features are applied.

Routing provides security to your network in a number of ways including obscuring internal network addresses with NAT and blackhole routing, using RPF to validate traffic sources, and maintaining an access control list (ACL) to limit access to the network.

This section includes:

- [Network Address Translation \(NAT\)](#)
- [Access Control List \(ACL\)](#)
- [Blackhole Route](#)
- [Reverse path lookup](#)

Network Address Translation (NAT)

Network address translation (NAT) is a method of changing the address traffic appears to originate from. This practice is used to hide the IP address on company's internal networks, and helps prevent malicious attacks that use those specific addresses.

This is accomplished by the router connected to that local network changing all the IP addresses to its externally connected IP address before sending the traffic out to the other networks, such as the Internet. Incoming traffic uses the established sessions to determine which traffic goes to which internal IP address. This also has the benefit of requiring only the router to be very secure against external attacks, instead of the whole internal network as would be the case without NAT. Securing one computer is much cheaper and easier to maintain.

Configuring NAT on your FortiGate unit includes the following steps.

- 1 Configure your internal network. For example use the `10.11.101.0` subnet.
- 2 Connect your internal subnet to an interface on your FortiGate unit. For example use `port1`.
- 3 Connect your external connection, for example an ISP gateway of `172.20.120.2`, to another interface on your FortiGate unit, for example `port2`.
- 4 Configure security policies to allow traffic between `port1` and `port2` on your FortiGate unit, ensuring that the NAT feature is enabled.

The above steps show that traffic from your internal network will originate on the `10.11.101.0` subnet and pass on to the `172.20.120.0` network. The FortiGate unit moves the traffic to the proper subnet. In doing that, the traffic appears to originate from the FortiGate unit interface on that subnet — it does not appear to originate from where it actually came from.

NAT “hides” the internal network from the external network. This provides security through obscurity. If a hacker tries to directly access your network, they will find the Fortigate unit, but will not know about your internal network. The hacker would have to get past the security-hardened FortiGate unit to gain access to your internal network. NAT will not prevent hacking attempts that piggy back on valid connections between the internal network and the outside world. However other UTM security measures can deal with these attempts.

Another security aspect of NAT is that many programs and services have problems with NAT. Consider if someone on the Internet tries to initiate a chat with someone on the internal network. The outsider only can access the FortiGate unit’s external interface unless the security policy allows the traffic through to the internal network. If allowed in, the proper internal user would respond to the chat. However if its not allowed, the request to chat will be refused or time-out. This is accomplished in the security policy by allowing or denying different protocols.

Access Control List (ACL)

An access control list (ACL) is a table of addresses that have permission to send and receive data over a router’s interface or interfaces. The router maintains an ACL, and when traffic comes in on a particular interface it is buffered, while the router looks up in the ACL if that traffic is allowed over that port or not. If it is allowed on that incoming interface, then the next step is to check the ACL for the destination interface. If the traffic passes that check as well the buffered traffic is delivered to its accentuation. If either of those steps fail the ACL check, the traffic is dropped and an error message may be sent to the sender. The ACL ensures that traffic follows expected paths, and any unexpected traffic is not delivered. This stops many network attacks. However, to be effective the ACL must be kept up to date —when employees or computers are removed from the internal network their IP addresses must also be removed from the ACL. For more information on the ACL, see the router chapter of the [FortiGate CLI Reference](#).

Blackhole Route

A blackhole route is a route that drops all traffic sent to it. It is very much like /dev/null in Linux programming.

Blackhole routes are used to dispose of packets instead of responding to suspicious inquiries. This provides added security since the originator will not discover any information from the target network.

Blackhole routes can also limit traffic on a subnet. If some subnet addresses are not in use, traffic to those addresses (traffic which may be valid or malicious) can be directed to a blackhole for added security and to reduce traffic on the subnet.

The loopback interface, a virtual interface that does not forward traffic, was added to enable easier configuration of blackhole routing. Similar to a normal interface, this loopback interface has fewer parameters to configure, and all traffic sent to it stops there. Since it cannot have hardware connection or link status problems, it is always available, making it useful for other dynamic routing roles. Once configured, you can use a loopback interface in security policies, routing, and other places that refer to interfaces. You configure this feature only from the CLI. For more information, see the system chapter of the [FortiGate CLI Reference](#).

Reverse path lookup

Whenever a packet arrives at one of the FortiGate unit's interfaces, the unit determines whether the packet was received on a legitimate interface by doing a reverse lookup using the source IP address in the packet header. This is also called anti-spoofing. If the FortiGate unit cannot communicate with the computer at the source IP address through the interface on which the packet was received, the FortiGate unit drops the packet as it is likely a hacking attempt.

If the destination address can be matched to a local address (and the local configuration permits delivery), the FortiGate unit delivers the packet to the local network. If the packet is destined for another network, the FortiGate unit forwards the packet to a next-hop router according to a policy route and the information stored in the FortiGate forwarding table.

Multipath routing and determining the best route

Multipath routing occurs when more than one entry to the same destination is present in the routing table. When multipath routing happens, the FortiGate unit may have several possible destinations for an incoming packet, forcing the FortiGate unit to decide which next-hop is the best one.

It should be noted that some IP addresses will be rejected by routing protocols. These are called Martian addresses. They are typically IP addresses that are invalid and not routable because they have been assigned an address by a misconfigured system, or are spoofed addresses.

Two methods to manually resolve multiple routes to the same destination are to lower the administrative distance of one route or to set the priority of both routes. For the FortiGate unit to select a primary (preferred) route, manually lower the administrative distance associated with one of the possible routes. Setting the priority on the routes is a FortiGate unit feature and may not be supported by non-Fortinet routers.

Administrative distance is based on the expected reliability of a given route. It is determined through a combination of the number of hops from the source and the protocol used. A hop is when traffic moves from one router to the next. More hops from the source means more possible points of failure. The administrative distance can be from 1 to 255, with lower numbers being preferred. A distance of 255 is seen as infinite and will not be installed in the routing table.

Here is an example to illustrate how administration distance works — if there are two possible routes traffic can take between two destinations with administration distances of 5 (always up) and 31 (sometimes not available), the traffic will use the route with an administrative distance of 5. If for some reasons the preferred route (admin distance of 5) is not available, the other route will be used as a backup.

Different routing protocols have different default administrative distances. These different administrative distances are based on a number of factors of each protocol such as reliability, speed, and so on. The default administrative distances for any of these routing protocols are configurable.

Table 94: Default administrative distances for routing protocols and connections

Routing protocol	Default administrative distance
Direct physical connection	1
Static	10
EBGP	20
OSPF	110

Table 94: Default administrative distances for routing protocols and connections

RIP	120
IBGP	200

Another method to determine the best route is to manually change the priority of both routes in question. If the next-hop administrative distances of two routes on the FortiGate unit are equal, it may not be clear which route the packet will take. Manually configuring the priority for each of those routes will make it clear which next-hop will be used in the case of a tie. The priority for a route can only be set from the CLI. Lower priorities are preferred. Priority is a Fortinet value that may or may not be present in other brands of routers.

All entries in the routing table are associated with an administrative distance. If the routing table contains several entries that point to the same destination (the entries may have different gateways or interface associations), the FortiGate unit compares the administrative distances of those entries first, selects the entries having the lowest distances, and installs them as routes in the FortiGate unit forwarding table. As a result, the FortiGate unit forwarding table contains only those routes having the lowest distances to every possible destination. While only static routing uses administrative distance as its routing metric, other routing protocols such as RIP can use metrics that are similar to administrative distance.

Route priority

After the FortiGate unit selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can configure the priority field through the CLI or the web-based manager. Priority values can range from 0 to 4 294 967 295. The route with the lowest value in the priority field is considered the best route. It is also the primary route.

To change the priority of a route - web-based manager

- 1 Go to *Router > Static > Static Route*.
- 2 Select the route entry, and select *Edit*.
- 3 Select *Advanced*.
- 4 Enter the *Priority* value.
- 5 Select *OK*.

To change the priority of a route - CLI

The following command changes the priority to 5 for a route to the address 10.10.10.1 on the port1 interface.

```
config router static
edit 1
set device port1
set gateway 10.10.10.10
set dst 10.10.10.1
set priority 5
end
```

If there are other routes set to priority 10, the route set to priority 5 will be preferred. If there are routes set to priorities less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal cost multipath (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, then different sessions will resolve this problem by using different routes to the same address.

Troubleshooting static routing

When there are problems with your network that you believe to be static routing related, there are a few basic tools available to locate the problem.

These tools include:

- [Ping](#)
- [Traceroute](#)
- [Examine routing table contents](#)

Ping

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is no packet loss detected, your basic network connectivity is OK.

If there is some packet loss detected, you should investigate:

- possible ECMP, split horizon, network loops
- cabling to ensure no loose connections

If there is total packet loss, you should investigate:

- hardware - ensure cabling is correct, and all equipment between the two locations is accounted for
- addresses and routes - ensure all IP addresses and routing information along the route is configured as expected
- firewalls - ensure all firewalls are set to allow PING to pass through

To ping from a Windows PC

- 1 Go to a DOS prompt. Typically you go to *Start > Run*, enter `cmd` and select OK.
- 2 Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate unit with four packets.

To ping from a Linux PC

- 1 Go to a command line prompt.
- 2 Enter `"/bin/etc/ping 10.11.101.101"`.

Traceroute

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, traceroute can be used to locate exactly where the problem is.

To use traceroute on an MS Windows PC

- 1 Go to a DOS prompt. Typically you go to *Start > Run*, enter “cmd” and select OK.
- 2 Enter “tracert fortinet.com” to trace the route from the PC to the Fortinet website.

To perform a traceroute on a Linux PC

- 1 Go to a command line prompt.
- 2 Enter “/bin/etc/traceroute fortinet.com”.

The Linux traceroute output is very similar to the MS Windows traceroute output.

Examine routing table contents

The first place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route isn't used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. Note that if your FortiGate unit is in Transparent mode, you are unable to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the web-based manager, use the Routing Monitor — go to *Router > Monitor > Routing Monitor*. In the CLI, use the command `get router info routing-table all`.

To examine the firewall session list in the web-based manager

- 1 Go to *System > Dashboard > Status > Top Sessions*.
- 2 Select *Detach*, and then *Details*.
- 3 Expand the session window to full screen to display the information.
- 4 Change filters, view associated security policy, column ordering, and so on to analyze the sessions in the table.
- 5 Select the delete icon to terminate the session.

ECMP route failover and load balancing

Equal Cost Multi-Path (ECMP) load balancing and failover are methods that extend basic static routing. They allow you to use your network bandwidth more effectively and with less down time than if you used basic static routing alone.

The concepts in this section include:

- [Route priority](#)
- [Equal-Cost Multi-Path \(ECMP\)](#)
- [Configuring interface status detection for gateway load balancing](#)
- [Configuring spillover or usage-based ECMP](#)
- [Configuring weighted static route load balancing](#)

Route priority

After the FortiGate unit selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can only configure the priority field through the CLI. Priority values can range from 0 to 255. The route with the lowest value in the priority field is considered the best route, and it is also the primary route.

For example, use the following command to change the priority of a route to 5 for a route to the address 10.10.10.1 on the port1 interface.

```
config router static
edit 1
    set device port1
    set gateway 10.10.10.10
    set dst 10.10.10.1
    set priority 5
end
```

If there are other routes at priority 10, this route will be preferred. If there are routes at priority less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal cost multipath (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, then different sessions will resolve this problem by using different routes to the same address.

Equal-Cost Multi-Path (ECMP)

FortiOS uses equal-cost multi-path (ECMP) to distribute traffic to the same destination such as the Internet or another network. Using ECMP you can add multiple routes to the destination and give each of those routes the same distance and priority.



If multiple routes to the same destination have the same priority but different distances, the route with the lowest distance is used. If multiple routes to the same destination have the same distance but different priorities, the route with the lowest priority is used. Distance takes precedence over priority. If multiple routes to the same destination have different distances and different priorities, the route with the lowest distance is always used even if it has the highest priority.

Using ECMP, if more than one ECMP route is available you can configure how the FortiGate unit selects the route to be used for a communication session. If only one ECMP route is available (for example, because an interface cannot process traffic because interface status detection does not receive a reply from the configured server) then all traffic uses this route.

Previous versions of FortiOS provided source IP-based load balancing for ECMP routes, but now FortiOS includes three configuration options for ECMP route failover and load balancing:

Source based (also called source IP based)	The FortiGate unit load balances sessions among ECMP routes based on the source IP address of the sessions to be load balanced. This is the default load balancing method. No configuration changes are required to support source IP load balancing.
Weighted Load Balance (also called weight-based)	The FortiGate unit load balances sessions among ECMP routes based on weights added to ECMP routes. More traffic is directed to routes with higher weights. After selecting weight-based you must add weights to static routes.
Spillover (also called usage-based)	<p>The FortiGate unit distributes sessions among ECMP routes based on how busy the FortiGate interfaces added to the routes are.</p> <p>After selecting spill-over you add route Spillover Thresholds to interfaces added to ECMP routes. The FortiGate unit sends all ECMP-routed sessions to the lowest numbered interface until the bandwidth being processed by this interface reaches its spillover threshold. The FortiGate unit then spills additional sessions over to the next lowest numbered interface.</p> <p>The Spillover Thresholds range is 0-2097000 KBps.</p>

You can configure only one of these ECMP route failover and load balancing methods in a single VDOM. If your FortiGate unit is configured for multiple VDOM operation, each VDOM can have its own ECMP route failover and load balancing configuration.

To configure the ECMP load balancing method from the web-based manager

- 1 Go to *Router > Static > Settings*.
- 2 Set *ECMP Load Balance Method* to *Source IP based*, *Weighted Load Balance*, or *Spillover*.

To configure the ECMP load balancing method from the CLI

For example, to set the load balancing method to usage-based, enter the following:

```
config system settings
    set v4-ecmp-mode usage-based
end
```

ECMP routing of simultaneous sessions to the same destination IP address

When the FortiGate unit selects an ECMP route for a session, a route cache is created that matches the route with the destination IP address of the session. All new sessions to the same destination IP address use the same route until the route is flushed from the cache. Routes are flushed from the cache after a period of time when no new sessions to the destination IP address are received.

The route cache improves FortiGate unit routing performance by reducing how often the FortiGate unit looks up routes in the routing table.

If the FortiGate unit receives a large number of sessions with the same destination IP address, because all of these sessions will be processed by the same route, it may appear that sessions are not distributed according to the ECMP route failover and load balancing configuration.

Configuring interface status detection for gateway load balancing

Interface status detection is used for ECMP route failover and load balancing. Interface status detection consists of the unit confirming that packets sent from an interface result in a response from a server. You can use up to three different protocols to confirm that an interface can connect to the server. Usually the server is the next-hop router that leads to an external network or the Internet. Interface status detection sends a packet using the configured protocols. If a response is received from the server, the unit assumes the interface can connect to the network. If a response is not received, the unit assumes that the interface cannot connect to the network.

Since it is possible that a response may not be received, even if the server and the network are operating normally, the dead gateway detection configuration controls the time interval between testing the connection to the server and the number of times the test can fail before the unit assumes that the interface cannot connect to the server.



As long as the unit receives responses for at least one of the protocols that you select, the unit assumes the server is operating and can forward packets. Responding to more than one protocol does not enhance the status of the server or interface.

To configure gateway failover detection for an interface

- 1 Go to *Router > Static > Settings*.
- 2 Under *Dead Gateway Detection*, select *Create New*.

3 Enter the following information:

Interface		Select the interface to test.
Ping Server		Enter the IP address of the server to test.
Detect Protocol		Select one of the following protocols.
	Ping	Use standard ICMP ping to confirm that the server is responding. Ping confirms that the server can respond to an ICMP ping request.
	TCP Echo	<p>Use TCP echo to confirm that the server is responding. Select this option if the server is configured to provide TCP echo services. In some cases a server may be configured to reply to TCP echo requests but not to reply to ICMP pings.</p> <p>TCP echo uses TCP packets on port number 7 to send a text string to the server and expect an echo reply back from the server. The echo reply just echoes back the same text to confirm that the server can respond to TCP requests.</p> <p>FortiGate units do not recognize RST (reset) packets from TCP Echo servers as normal TCP echo replies. If the unit receives an RST response to a TCP echo request, the unit assumes the server is unreachable.</p>
	UDP Echo	<p>Use UDP echo to detect the server. Select this option if the server is configured to provide UDP echo services. In some cases a server may be configured to reply to UDP echo requests but not to reply ICMP pings.</p> <p>UDP echo uses UDP packets on port number 7 to send a text string to the server and expects an echo reply from the server. The echo reply just echoes back the same text to confirm that the server can respond to UDP requests.</p>
Ping Interval		Enter the interval between pings, in seconds.
Failover Threshold		Enter the number of times the test can fail before the unit assumes that the interface cannot connect to the server.
HA Priority		



For more information about TCP echo and UDP echo, see RFC 862.

4 Select OK.**To configure gateway failover detection for an interface - CLI**

```
config router gwdetect
  edit port1
    set protocol ping
    set server 10.10.10.1
    set interval 5
    set failtime 5
  end
```

Configuring spillover or usage-based ECMP

Spill-over or usage-based ECMP routes new sessions to interfaces that have not reached a configured bandwidth limit (called the *Spillover Threshold* or a route-spillover threshold). To configure spill-over or usage-based ECMP routing, you enable spill-over ECMP, add ECMP routes, and add a *Spillover Threshold* to the interfaces used by the ECMP routes. Set the *Spillover Thresholds* to limit the amount of bandwidth processed by each interface. The range is 0 to 2 097 000 Kbps. The threshold counts only outgoing traffic.

With spill-over ECMP routing configured, the FortiGate unit routes new sessions to an interface used by an ECMP route until that interface reaches its *Spillover Threshold*. Then, when the threshold of that interface is reached, new sessions are routed to one of the other interfaces used by the ECMP routes.

To add Spillover Thresholds to interfaces - web-based manager

Use the following steps to enable usage based ECMP routing, add Spillover Thresholds to FortiGate interfaces port3 and port4, and then to configure EMCP routes with device set to port3 and port4.

- 1 Go to *Router > Static > Settings*.
- 2 Set *ECMP Load Balance Method* to *Spillover*.
- 3 Go to *Router > Static > Static Route*.
- 4 Add ECMP routes for port3 and port4.

Destination IP/Mask	192.168.20.0/24
Device	port3
Gateway	172.20.130.3
Advanced	
Distance	10

Destination IP/Mask	192.168.20.0/24
Device	port4
Gateway	172.20.140.4
Advanced	
Distance	10

- 5 Go to *System > Network > Interface*.
- 6 Edit port3 and port4 and add the following spillover-thresholds:

Interface	port3
Spillover Threshold	100

Interface	port4
Spillover Threshold	200

To add Spillover Thresholds to interfaces - CLI

```
config system settings
    set v4-ecmp-mode usage-based
```

```
end
config router static
  edit 1
    set device port3
    set dst 192.168.20.0 255.255.255.0
    set gateway 172.20.130.3
  next
  edit 2
    set device port4
    set dst 192.168.20.0 255.255.255.0
    set gateway 172.20.140.4
  end
config system interface
  edit port3
    set spillover-threshold 100
  next
  edit port4
    set spillover-threshold 200
  end
```

Detailed description of how spill-over ECMP selects routes

When you add ECMP routes they are added to the routing table in the order displayed by the routing monitor or by the `get router info routing-table static` command. This order is independent of the configured bandwidth limit.

The FortiGate unit selects an ECMP route for a new session by finding the first route in the routing table that sends the session out a FortiGate unit interface that is not processing more traffic than its configured route spill-over limit.



A new session to a destination IP address that already has an entry in the routing cache is routed using the route already added to the cache for that destination address. See [“ECMP routing of simultaneous sessions to the same destination IP address” on page 1687](#).

For example, consider a FortiGate unit with interfaces port3 and port4 both connected to the Internet through different ISPs. ECMP routing is set to usage-based and route spillover for to 100 KBps for port3 and 200 KBps for port4. Two ECMP default routes are added, one for port3 and one for port4.

If the route to port3 is higher in the routing table than the route to port4, the FortiGate unit sends all default route sessions out port3 until port3 is processing 10Mbps of data. When port3 reaches its configured bandwidth limit, the FortiGate unit sends all default route sessions out port4. When the bandwidth usage of port3 falls below 10Mbps, the FortiGate again sends all default route sessions out port3.

New sessions with destination IP addresses that are already in the routing cache; however, use the cached routes. This means that even if port3 is exceeding its bandwidth limit, new sessions can continue to be sent out port3 if their destination addresses are already in the routing cache. As a result, new sessions are sent out port4 only if port3 exceeds its bandwidth limit and if the routing cache does not contain a route for the destination IP address of the new session.

Also, the switch over to port4 does not occur as soon as port3 exceeds its bandwidth limit. Bandwidth usage has to exceed the limit for a period of time before the switch over takes place. If port3 bandwidth usage drops below the bandwidth limit during this time period, sessions are not switched over to port4. This delay reduces route flapping.

FortiGate usage-based ECMP routing is not actually load balancing, since routes are not distributed evenly among FortiGate interfaces. Depending on traffic volumes, most traffic would usually be processed by the first interface with only spillover traffic being processed by other interfaces.

If you are configuring usage-based ECMP in most cases you should add spillover thresholds to all of the interfaces with ECMP routes. The default spillover threshold is 0 which means no bandwidth limiting. If any interface has a spillover threshold of 0, no sessions will be routed to interfaces lower in the list unless the interface goes down or is disconnected. An interface can go down if *Detect interface status for Gateway Load Balancing* does not receive a response from the configured server.

Determining if an interface has exceeded its Spillover Threshold

You can use the `diagnose netlink dstmac list` CLI command to determine if an interface is exceeding its Spillover Threshold. If the command displays `over_bps=1` the interface is exceeding its threshold. If `over_bps=0` the interface has not exceeded its threshold.

Configuring weighted static route load balancing

Configure weighted load balancing to control how the FortiGate unit distributes sessions among ECMP routes by adding weights for each route. Add higher weights to routes that you want to load balance more sessions to.

With the ECMP load balancing method set to weighted, the FortiGate unit distributes sessions with different destination IPs by generating a random value to determine the route to select. The probability of selecting one route over another is based on the weight value of each route. Routes with higher weights are more likely to be selected.

Large numbers of sessions are evenly distributed among ECMP routes according to the route weight values. If all weights are the same, sessions are distributed evenly. The distribution of a small number of sessions; however, may not be even. For example, it's possible that if there are two ECMP routes with the same weight; two sessions to different IP addresses could use the same route. On the other hand, 10,000 sessions with different destination IPs should be load balanced evenly between two routes with equal rates. The distribution could be 5000:5000 or 50001:4999. Also, 10,000 sessions with different destination IP addresses should be load balanced as 3333:6667 if the weights for the two routes are 100 and 200.

Weights only affect how routes are selected for sessions to new destination IP addresses. New sessions to IP addresses already in the routing cache are routed using the route for the session already in the cache. So in practice sessions will not always be distributed according to the routing weight distribution.

To add weights to static routes from the web-based manager

- 1 Go to *Router > Static > Settings*.
- 2 Set *ECMP Load Balance Method* to *Weighted Load Balance*.
- 3 Go to *Router > Static > Static Route*.

- 4 If needed, add new static routes, for example:

Destination IP/Mask	192.168.20.0/24
Device	port1
Gateway	172.20.110.1
Distance	10

Destination IP/Mask	192.168.20.0/24
Device	port2
Gateway	172.20.120.2
Distance	10

- 5 Go to *System > Network > Interface*.
- 6 Edit each route's interface and set the *Weight* value.
For example, set the weight of port1 to 100 and the weight of port2 to 200.

Static routing tips

When your network goes beyond basic static routing, here are some tips to help you plan and manage your static routing.

Always configure a default route

The first thing configured on a router on your network should be the default route. And where possible the default routes should point to either one or very few gateways. This makes it easier to locate and correct problems in the network. By comparison, if one router uses a second router as its gateway which uses a fourth for its gateway and so on, one failure in that chain will appear as an outage for all the devices downstream. By using one or very few addresses as gateways, if there is an outage on the network it will either be very localized or network-wide — either is easy to troubleshoot.

Have an updated network plan

A network plan lists different subnets, user groups, and different servers. Essentially it puts all your resources on the network, and shows how the parts of your network are connected. Keeping your plan updated will also help you troubleshoot problems more quickly when they arise.

A network plan helps your static routing by eliminating potential bottlenecks, and helping troubleshoot any routing problems that come up. Also you can use it to plan for the future and act on any changes to your needs or resources more quickly.

Plan for expansion

No network remains the same size. At some time, all networks grow. If you take future growth into account, there will be less disruption to your existing network when that growth happens. For example allocating a block of addresses for servers can easily prevent having to re-assign IP addresses to multiple servers due to a new server.

With static routing, if you group parts of your network properly you can easily use network masks to address each part of your network separately. This will reduce the amount of administration required both to maintain the routing, and to troubleshoot any problems.

Configure as much security as possible

Securing your network through static routing methods is a good low level method to defend both your important information and your network bandwidth.

- Implement NAT to obscure your IP address is an excellent first step.
- Implement black hole routing to hide which IP addresses are in use or not on your local network.
- Configure and use access control list (ACL) to help ensure you know only valid users are using the network.

All three features limit access to the people who should be using your network, and obscure your network information from the outside world and potential hackers.

Policy Routing

Policy routing enables you to redirect traffic away from a static route. This can be useful if you want to route certain types of network traffic differently. You can use incoming traffic's protocol, source address or interface, destination address, or port number to determine where to send the traffic. For example, generally network traffic would go to the router of a subnet, but you might want to direct SMTP or POP3 traffic directly to the mail server on that subnet.

If you have configured the FortiGate unit with routing policies and a packet arrives at the FortiGate unit, the FortiGate unit starts at the top of the Policy Route list and attempts to match the packet with a policy. If a match is found and the policy contains enough information to route the packet (a minimum of the IP address of the next-hop router and the FortiGate interface for forwarding packets to it), the FortiGate unit routes the packet using the information in the policy. If no policy route matches the packet, the FortiGate unit routes the packet using the routing table.



Most policy settings are optional, so a matching policy alone might not provide enough information for forwarding the packet. The FortiGate unit may refer to the routing table in an attempt to match the information in the packet header with a route in the routing table. For example, if the outgoing interface is the only item in the policy, the FortiGate unit looks up the IP address of the next-hop router in the routing table. This situation could happen when the interfaces are dynamic (such as DHCP or PPPoE) and you do not want or are unable to specify the IP address of the next-hop router.

Policy route options define which attributes of a incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the FortiGate unit routes the packet through the specified interface to the specified gateway.

To view policy routes go to *Router > Static > Policy Route*.

Create New	Add a policy route. See “Adding a policy route” on page 1694 .
Delete icon	Delete the selected policy route.
Edit icon	Edit the selected policy route.
Move To icon	Move the selected policy route. Enter the new position and select OK. For more information, see “Moving a policy route” on page 1696 .
#	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.

Incoming	The interfaces on which packets subjected to route policies are received.
Outgoing	The interfaces through which policy routed packets are routed.
Source	The IP source addresses and network masks that cause policy routing to occur.
Destination	The IP destination addresses and network masks that cause policy routing to occur.

Adding a policy route

To add a policy route, go to *Router > Static > Policy Route* and select *Create New*.

Protocol	<p>Enter the protocol number to match. The Internet Protocol Number is found in the IP packet header. RFC 5237 describes protocol numbers and you can find a list of the assigned protocol numbers here. The range is from 0 to 255. A value of 0 disables the feature.</p> <p>Commonly used <i>Protocol</i> settings include 6 for TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.</p>
Incoming Interface	Select the name of the interface through which incoming packets subjected to the policy are received.
Source Address / Mask	To perform policy routing based on IP source address, type the source address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Address / Mask	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Ports	<p>To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature.</p> <p>The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.</p>
Type of Service	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see “Type of Service” on page 1695 .
Outgoing Interface	Select the name of the interface through which packets affected by the policy will be routed.
Gateway Address	Type the IP address of the next-hop router that the FortiGate unit can access through the specified interface.

Example policy route

Configure the following policy route to send all FTP traffic received at `port1` out the `port10` interface and to a next hop router at IP address `172.20.120.23`. To route FTP traffic set protocol to 6 (for TCP) and set both of the destination ports to 21, the FTP port.

Protocol	6
Incoming interface	port1
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0
Destination Ports	From 21 to 21
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)
Outgoing interface	port10
Gateway Address	172.20.120.23

Type of Service

Type of service (TOS) is an 8-bit field in the IP header that enables you to determine how the IP datagram should be delivered, with such qualities as delay, priority, reliability, and minimum cost.

Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority TOS is 0, the highest is 7 - when bits 3, 4, and 5 are all set to 1. The router tries to match the TOS of the datagram to the TOS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero TOS route.

Using increased quality may increase the cost of delivery because better performance may consume limited network resources. For more information, see RFC 791 and RFC 1349.

Table 95: The role of each bit in the IP header TOS 8-bit field

bits 0, 1, 2	Precedence	Some networks treat high precedence traffic as more important traffic. Precedence should only be used within a network, and can be used differently in each network. Typically you do not care about these bits.
bit 3	Delay	When set to 1, this bit indicates low delay is a priority. This is useful for such services as VoIP where delays degrade the quality of the sound.
bit 4	Throughput	When set to 1, this bit indicates high throughput is a priority. This is useful for services that require lots of bandwidth such as video conferencing.
bit 5	Reliability	When set to 1, this bit indicates high reliability is a priority. This is useful when a service must always be available such as with DNS servers.
bit 6	Cost	When set to 1, this bit indicates low cost is a priority. Generally there is a higher delivery cost associated with enabling bits 3, 4, or 5, and bit 6 indicates to use the lowest cost route.
bit 7	Reserved for future use	Not used at this time.

For example, if you want to assign low delay, and high reliability, say for a VoIP application where delays are unacceptable, you would use a bit pattern of xxx1x1xx where an 'x' indicates that bit can be any value. Since all bits are not set, this is a good use for the bit mask; if the mask is set to 0x14, it will match any TOS packets that are set to low delay and high reliability.

Moving a policy route

A routing policy is added to the bottom of the routing table when it is created. If you prefer to use one policy over another, you may want to move it to a different location in the routing policy table.

The option to use one of two routes happens when both routes are a match, for example 172.20.0.0/255.255.0.0 and 172.20.120.0/255.255.255.0. If both of these routes are in the policy table, both can match a route to 172.20.120.112 but you consider the second one as a better match. In that case the best match route should be positioned before the other route in the policy table.

To change the position of a policy route in the table, go to *Router > Static > Policy Route* and select *Move To* for the policy route you want to move.

Before/After	Select Before to place the selected Policy Route before the indicated route. Select After to place it following the indicated route.
Policy route ID	Enter the Policy route ID of the route in the Policy route table to move the selected route before or after.

Transparent mode static routing

FortiOS operating modes allow you to change the configuration of your FortiGate unit depending on the role it needs to fill in your network.

NAT/Route operating mode is the standard mode where all interfaces are accessed individually, and traffic can be routed between ports to travel from one network to another.

In transparent operating mode, all physical interfaces act like one interface. The FortiGate unit essentially becomes a bridge — traffic coming in over any interface is broadcast back out over all the interfaces on the FortiGate unit.

In transparent mode, there is no entry for routing at the main level of the menu on the web-based manager display as there is in NAT/Route mode. Routing is instead accessed through the network menu option.

To view the routing table in transparent mode, go to *System > Network > Routing Table*.

When viewing or creating a static route entry in transparent mode there are only three fields available.

Destination IP /Mask	The destination of the traffic being routed. The first entry is attempted first for a match, then the next, and so on until a match is found or the last entry is reached. If no match is found, the traffic will not be routed. Use 0.0.0.0 to match all traffic destinations. This is the default route.
-----------------------------	--

Gateway	Specifies the next hop for the traffic. Generally the gateway is the address of a router on the edge of your network.
Priority	<p>The priority is used if there is more than one match for a route. This allows multiple routes to be used, with one preferred. If the preferred route is unavailable the other routes can be used instead.</p> <p>Valid range of priority can be from 0 to 4 294 967 295.</p> <p>If more than one route matches and they have the same priority it becomes an ECMP situation and traffic is shared among those routes. See “Route priority” on page 1685.</p>

When configuring routing on a FortiGate unit in transparent mode, remember that all interfaces must be connected to the same subnet. That means all traffic will be coming from and leaving on the same subnet. This is important because it limits your static routing options to only the gateways attached to this subnet. For example, if you only have one router connecting your network to the Internet then all static routing on the FortiGate unit will use that gateway. For this reason static routing on FortiGate units in transparent mode may be a bit different, but it is not as complex as routing in NAT/Route mode.



Dynamic Routing Overview

This section provides an overview of dynamic routing, and how it compares to static routing. For details on various dynamic routing protocols, see the following chapters for detailed information.

The following topics are included in this section:

- [What is dynamic routing?](#)
- [Comparison of dynamic routing protocols](#)
- [Choosing a routing protocol](#)
- [Dynamic routing terminology](#)
- [IPv6 in dynamic routing](#)

What is dynamic routing?

Dynamic routing uses a dynamic routing protocol to automatically select the best route to put into the routing table. So instead of manually entering static routes in the routing table, dynamic routing automatically receives routing updates, and dynamically decides which routes are best to go into the routing table. Its this intelligent and hands-off approach that makes dynamic routing so useful.

Dynamic routing protocols vary in many ways and this is reflected in the various administrative distances assigned to routes learned from dynamic routing. These variations take into account differences in reliability, speed of convergence, and other similar factors. For more information on these administrative distances, see [“Multipath routing and determining the best route” on page 1681](#).

This section includes:

- [Comparing static and dynamic routing](#)
- [Dynamic routing protocols](#)
- [Minimum configuration for dynamic routing](#)

Comparing static and dynamic routing

A common term used to describe dynamic routing is convergence. Convergence is the ability to work around a network problems and outages — for the routing to come together despite obstacles. For example if the main router between two end points goes down, convergence is the ability to find a way around that failed router and reach the destination. Static routing has zero convergence beyond trying the next route in its limited local routing table — if a network administrator doesn't fix a routing problem manually, it will never be fixed resulting in a downed network. Dynamic routing solves this

problem by involving routers along the route to the destination in decision making about the route, and using the routing tables of these routes for potential routes around the outage. In general dynamic routing has better scalability, robustness, and convergence. However, the cost of these added benefits include more complexity and some overhead — bandwidth that is used by the routing protocol for its own administration.

Table 96: Comparing static and dynamic routing

Feature	Static Routing	Dynamic Routing
Hardware support	Supported by all routing hardware	May require special, more expensive routers
Router Memory Required	Minimal	Can require considerable memory for larger tables
Complexity	Simple	Complex
Overhead	None	Varying amounts of bandwidth used for routing protocol updates
Scalability	Limited to small networks	Very scalable, better for larger networks
Robustness	None - if a route fails it has to be fixed manually	Robust - traffic routed around failures automatically
Convergence	None	Varies from good to excellent

Dynamic routing protocols

A dynamic routing protocol is an agreed on method of routing that the sender, receiver, and all routers along the path (route) support. Typically the routing protocol involves a process running on all computers and routers along that route to enable each router to handle routes in the same way as the others. The routing protocol determines how the routing tables are populated along that route, how the data is formatted for transmission, and what information about a route is included with that route. For example RIP, and BGP use distance vector algorithms, where OSPF uses a shortest path first algorithm. Each routing protocol has different strengths and weaknesses — one protocol may have fast convergence, while another may be very reliable, and a third is very popular for certain businesses like Internet Service Providers (ISPs).

Dynamic routing protocols are different from each other in a number of ways, such as:

- [Classful versus classless routing protocols](#)
- [Interior versus exterior routing protocols](#)
- [Distance vector versus link-state protocols](#)

Classful versus classless routing protocols

Classful or classless routing refers to how the routing protocol handles the IP addresses. In classful addresses there is the specific address, and the host address of the server that address is connected to. Classless addresses use a combination of IP address and netmask.

Classless Inter-Domain Routing (CIDR) was introduced in 1993 (originally with RFC 1519 and most recently with RFC 4632) to keep routing tables from getting too large. With Classful routing, each IP address requires its own entry in the routing table. With Classless routing, a series of addresses can be combined into one entry potentially saving vast amounts of space in routing tables.

Current routing protocols that support classless routing out of necessity include RIPv2, BGP, IS-IS, and OSPF. Older protocols such as RIPv1 do not support CIDR addresses.

Interior versus exterior routing protocols

The names interior and exterior are very descriptive. Interior routing protocols are designed for use within a contained network of limited size, where exterior routing protocols are designed to link multiple networks together. For example, only border routers of a network run the exterior routing protocol, where all the routers on the network run the interior protocol. This overlap is required for the exterior routers to communicate with the interior routers — border routers almost always run multiple routing protocols.

Nearly all routing protocols are interior routing protocols. Only BGP is commonly used as an exterior routing protocol.

You may see interior gateway protocol (IGP) used to refer to interior routing protocols, and exterior gateway protocol (EGP) used to refer to exterior routing protocols.

Distance vector versus link-state protocols

Every routing protocol determines the best route between two addresses using a different method. However, there are two main algorithms for determining the best route — Distance vector and Link-state.

Distance vector protocols

In distance vector protocols, routers are told about remote networks through neighboring routers. The distance part refers to the number of hops to the destination, and in more advanced routing protocols these hops can be weighted by factors such as available bandwidth and delay. The vector part determines which router is the next step along the path for this route. This information is passed along from neighboring routers with routing update packets that keep the routing tables up to date. Using this method, an outage along a route is reported back along to the start of that route, ideally before the outage is encountered.

On distance vector protocols, RFC 1058 which defines RIP v1 states the following:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system.

Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

There are four main weaknesses inherent in the distance vector method. Firstly, the routing information is not discovered by the router itself, but is instead reported information that must be relied on to be accurate and up-to-date. The second weakness is that it can take a while for the information to make its way to all the routers who need the information — in other words it can have slow convergence. The third weakness is the amount of overhead involved in passing these updates all the time. The number of updates between routers in a larger network can significantly reduce the available bandwidth. The fourth weakness is that distance vector protocols can end up with routing-loops. Routing loops are when packets are routed for ever around a network, and often occur with slow convergence. The bandwidth required by these infinite loops will slow your network to a halt. There are methods of preventing these loops however, so this weakness is not as serious as it may first appear.

Link-state protocols

Link-state protocols are also known as shortest path first protocols. Where distance vector uses information passed along that may or may not be current and accurate, in link-state protocols each router passes along only information about networks and devices directly connected to it. This results in a more accurate picture of the network topology around your router, allowing it to make better routing decisions. This information is passed between routers using link-state advertisements (LSAs). To reduce the overhead, LSAs are only sent out when information changes, compared to distance vector sending updates at regular intervals even if no information has changed. The more accurate network picture in link-state protocols greatly speed up convergence and avoid problems such as routing-loops.

Minimum configuration for dynamic routing

Dynamic routing protocols do not pay attention to routing updates from other sources, unless you specifically configure them to do so using CLI redistribute commands within each routing protocol.

The minimum configuration for any dynamic routing to function is dynamic routing configured on one interface the FortiGate unit and one other router configured as well. Some protocols require more

Table 97: Minimum configuration based on dynamic protocol

	BGP	RIP	OSPF
Interface	yes	yes	yes
Network	yes	yes	yes
AS	local and neighbor	no	yes
Neighbors	at least one	at least one	at least one
Version	no	yes	no
Router ID	no	no	yes

Comparison of dynamic routing protocols

Each dynamic routing protocol was designed to meet a specific routing need. Each protocol does some things well, and other things not so well. For this reason, choosing the right dynamic routing protocol for your situation is not an easy task.

Features of dynamic routing protocols

Each protocol is better suited for some situations over others.

Choosing the best dynamic routing protocol depends on the size of your network, speed of convergence required, the level of network maintenance resources available, what protocols the networks you connect to are using, and so on. For more information on these dynamic routing protocols, see [“Routing Information Protocol \(RIP\)” on page 1715](#), [“Border Gateway Protocol \(BGP\)” on page 1751](#), or [“Open Shortest Path First \(OSPF\)” on page 1787](#).

Table 98: Comparing RIP, BGP, and OSPF dynamic routing protocols

Protocol	RIP	BGP	OSPF
Routing algorithm	Distance Vector, basic	Distance Vector, advanced	Link-state
Common uses	Small non-complex networks	Network backbone, ties multinational offices together	Common in large, complex enterprise networks
Strengths	Fast and simple to implement Near universal support Good when no redundant paths	Graceful restart BFD support Only needed on border routers Summarize routes	Fast convergence Robust Little management overhead No hop count limitation Scalable
Weakness	Frequent updates can flood network Slow convergence Maximum 15 hops may limit network configuration	Required full mesh in large networks can cause floods Route flap Load-balance multi-homed networks Not available on low end routers	Complex No support for unequal cost multipath routing Route summary can require network changes
Authentication	Optional authentication using text string or MD5 password. (RIP v1 has no authentication)		
IPv6 Support	Only in RIPng	Only in BGP4+	Only in OSPF6

Routing protocols

Routing Information Protocol (RIP) uses classful routing, as well as incorporating various methods to stop incorrect route information from propagating, such as poisoned horizon. However, on larger networks its frequent updates can flood the network and its slow convergence can be a problem.

Border Gateway Protocol (BGP) has been the core Internet backbone routing protocol since the mid 1990s, and is the most used interior gateway protocol (IGP). However, some configurations require full mesh connections which flood the network, and there can be route flap and load balancing issues for multihomed networks.

Open Shortest Path First (OSPF) is commonly used in large enterprise networks. It is the protocol of choice mainly due to its fast convergence. However, it can be complicated to setup properly.

Multicast addressing is used to broadcast from one source to many destinations efficiently. Protocol Independent Multicast (PIM) is the protocol commonly used in enterprises, multimedia content delivery, and stock exchanges. For more information, see [“Multicast forwarding” on page 471](#) in the System Administration chapter.

Routing algorithm

Each protocol uses a slightly different algorithm for choosing the best route between two addresses on the network. The algorithm is the “intelligent” part of a dynamic protocol because the algorithm is responsible for deciding which route is best and should be added to the local routing table. RIP and BGP use distance vector algorithms, where OSPF uses link-state or a shortest path first algorithm.

Vector algorithms are essentially based on the number of hops between the originator and the destination in a route, possibly weighting hops based on how reliable, fast, and error-free they are.

The link-state algorithm used by OSPF is called the Dijkstra algorithm. Link-state treats each interface as a link, and records information about the state of the interface. The Dijkstra algorithm creates trees to find the shortest paths to the routes it needs based on the total cost of the parts of the routes in the tree.

For more information on the routing algorithm used, see [“Distance vector versus link-state protocols” on page 1701](#).

Authentication

If an attacker gains access to your network, they can masquerade as a router on your network to either gain information about your network or disrupt network traffic. If you have a high quality firewall configured, it will help your network security and stop many of this type of threat. However, the main method for protecting your routing information is to use authentication in your routing protocol. Using authentication on your FortiGate unit and other routers prevents access by attackers — all routers must authenticate with passwords, such as MD5 hash passwords, to ensure they are legitimate routers.

When configuring authentication on your network, ensure you configure it the same on all devices on the network. Failure to do so will create errors and outages as those forgotten devices fail to connect to the rest of the network.

For example, to configure an MD5 key of 123 on an OSPF interface called `ospf_test`, enter the following CLI command:

```
config router ospf
  config ospf-interface
    edit ospf_test
      set authentication md5
      set md5-key 123
    end
  end
```

Convergence

Convergence is the ability of a networking protocol to re-route around network outages. Static routing cannot do this. Dynamic routing protocols can all converge, but take various amounts of time to do this. Slow convergence can cause problems such as network loops which degrade network performance.

You may also hear robustness and redundancy used to describe networking protocols. In many ways they are the same thing as convergence. Robustness is the ability to keep working even though there are problems, including configuration problems as well as network outages. Redundancy involves having duplicate parts that can continue to function in the event of some malfunction, error, or outage. It is relatively easy to configure dynamic routing protocols to have backup routers and configurations that will continue to function no matter the network problem short of a total network failure.

IPv6 Support

IPv4 addressing is in common use everywhere around the world. IPv6 has much larger addresses and it is used by many large companies and government departments. IPv6 is not as common as IPv4 yet, but more companies are adopting it.

If your network uses IPv6, your dynamic routing protocol must support it. None of the dynamic routing protocols originally supported IPv6, but they all have additions, expansions, or new versions that do support IPv6. For more information, see [“RIP and IPv6” on page 1716](#), [“BGP and IPv6” on page 1752](#), or [“OSPF and IPv6” on page 1788](#).

When to adopt dynamic routing

Static routing is more than enough to meet your networking needs when you have a small network. However, as your network grows, the question you need to answer is at what point do you adopt dynamic routing in your networking plan and start using it in your network? The main factors in this decision are typically:

- [Budget](#)
- [Current network size and topology](#)
- [Expected network growth](#)
- [Available resources for ongoing maintenance](#)

Budget

When making any business decision, the budget must always be considered. Static routing does not involve special hardware, fancy software, or expensive training courses.

Dynamic routing can include all of these extra expenses. Any new routing hardware such as routers and switches need to support your chosen protocols. Network management software to help configure and maintain your more complex network, and routing protocol drivers may be necessary as well. If the network administrators are not well versed in dynamic routing, either a training course or some hands on learning time must be budgeted so they can administer the new network with confidence. Together, these factors will use up your budget quickly.

Additionally people account for network starting costs in the budgets, but usually leave out the ongoing cost of network maintenance. Any budget must provide for the hours that will be spent on updating the network routing equipment, and fixing any problems. Without that money in the budget, you may end up back at static routing before you know it.

Current network size and topology

As stated earlier static routing works well on small networks. At those networks get larger, routing takes longer, routing tables get very large, and general performance isn't what it could be.

Topology is a concern as well. If all your computers are in one building, it's much easier to stay with static routing longer. However, connecting a number of locations will be easier with the move to dynamic routing.

If you have a network of 20 computers, you can still likely use static routing. If those computers are in two or three locations, static routing will still be a good choice for connecting them. Also, if you just connect to your ISP and don't worry about any special routing to do that, you are likely safe with just static routing.

If you have a network of 100 computers in one location, you can use static routing but it will be getting slower, more complex, and there won't be much room for expansion. If those 100 computers are spread across three or more locations, dynamic routing is the way to go.

If you have 1000 computers, you definitely need to use dynamic routing no matter how many locations you have.

Hopefully this section has given you an idea of what results you will likely experience from different sized networks using different routing protocols. Your choice of which dynamic routing protocol to use is partly determined by the network size, and topology.

Expected network growth

You may not be sure if your current network is ready for dynamic routing. However, if you are expecting rapid growth in the near future, it is a good idea to start planning for that growth now so you are ready for the coming expansion.

Static routing is very labor intensive. Each network device's routing table needs to be configured and maintained manually. If there is a large number of new computers being added to the network, they each need to have the static routing table configured and maintained. If devices are being moved around the network frequently, they must also be updated each time.

Instead, consider putting dynamic routing in place before those new computers are installed on the network. The installation issues can be worked out with a smaller and less complex network, and when those new computers or routers are added to the network there will be nowhere near the level of manual configuration required. Depending on the level of growth, this labor savings can be significant. For example, in an emergency you can drop a new router into a network or AS wait for it to receive the routing updates from its neighbors, and then remove one of the neighbors. While the routes will not be the most effective possible, this method is much less work than static routing in the same situation with less chance of mistakes.

Also as your network grows and you add more routers, those new routers can help share the load in most dynamic routing configurations. For example if you have 4 OSPF routers and 20,000 external routes those few routers will be overwhelmed. But in a network with 15 OSPF routers they will better be able to handle that number of routes. Be aware though that adding more routers to your network will increase the amount of updates sent between the routers, which will take up some of your bandwidth.

Available resources for ongoing maintenance

As touched on in the budget section, there must be resources dedicated to ongoing network maintenance, upgrades, and troubleshooting. These resources include administrator hours to configure and maintain the network, training for the administrator if needed, extra hardware and software as needed, and possible extra staff to help the administrator in emergencies. Without these resources, you will quickly find the network reverting to static routing out of necessity. This is because:

- Routing software updates will require time.

- Routing hardware updates will require time.
- Office reorganizations or significant personnel movement will require time from a networking point of view.
- Networking problems that occur, such as failed hardware, require time to locate and fix the problem.

If the resources to accomplish these tasks are not budgeted, they will either not happen or not happen at the required level to continue operation. This will result in both the network administration staff and the network users being very frustrated.

A lack of maintenance budget will also result in increasingly heavy reliance on static routing as the network administrators are forced to use quick fixes for problems that come up. This invariably involves going to static routing, and dropping the more complex and time consuming dynamic routing.

Choosing a routing protocol

One of the hardest decisions in routing can be choosing which routing protocol to use on your network. It can be easy to decide when static routing will not meet your needs, but how can you tell which dynamic routing protocol is best for your network and situation?

Here is a brief look at the routing protocols including their strongest and weakest points. The steps to choosing your routing protocol are:

- 1 [Answer questions about your network](#)
- 2 [Dynamic routing terminology](#)
- 3 [Evaluate your chosen protocol](#)
- 4 [Implement your dynamic routing protocol](#)

Answer questions about your network

Before you can decide what is best for your situation, you need to examine what the details of your situation are such as what you have for budget, equipment, and users.

The following questions will help you form a clear idea of your routing needs:

How many computers or devices are on your network?

It matters if you only have a few computers, or if you have many and if they are all at one location or not as well. All routing protocols can be run on any sized network, however it can be inefficient to run some on very small networks. However, routers and network hardware that support dynamic routing can be more expensive than more generic routers for static routing.

What applications typically run over the network?

Finding out what application your users are running will help you determine their needs and the needs of the network regarding bandwidth, quality of service, and other such issues.

What level of service do the users expect from the network?

Different network users have different expectations of the network. It's not critical for someone surfing the Internet to have 100% uptime, but it is required for a stock exchange network or a hospital.

Is there network expansion in your near future?

You may have a small network now, but if it will be growing quickly, you should plan for the expected size so you don't have to change technologies again down the road.

What routing protocols do your networks connect to?

This is most often how routing protocol decisions are made. You need to be able to communicate easily with your service provider and neighbors, so often people simply use what everyone else is using.

Is security a major concern?

Some routing protocols have levels of authentication and other security features built in. Others do not. If security is important to you, be aware of this.

What is your budget — both initial and maintenance?

More robust and feature laden routing protocols generally mean more resources are required to keep them working well. Also more secure configurations require still more resources. This includes both set up costs, as well as ongoing maintenance costs. Ignore these costs at the risk of having to drop the adoption of the new routing protocol mid-change.

Evaluate your chosen protocol

Once you have examined the features of the routing protocols listed above and chosen the one that best meets your needs, you can set up an evaluation or test install of that protocol.

The test install is generally set up in a sandbox configuration so it will not affect critical network traffic. The aim of the test install is to prove that it will work on a larger scale on your network. So be sure that the test install mirrors your larger network well enough for you to discover any problems. If its too simplistic, these problems may not appear.

If your chosen protocol does not meet your goals choose a different protocol and repeat the evaluation process until either a protocol meets your needs, or you change your criteria.

Implement your dynamic routing protocol

You have examined your needs, selected the best matching dynamic routing protocol, tested it, and now you are ready to implement it with confidence.

This guide will help you configure your FortiGate unit to support your chosen dynamic routing protocol. Refer to the various sections in this guide as needed during your implementation to help ensure a smooth transition. Examples for each protocol have been included to show proper configurations for different types of networks.

Dynamic routing terminology

Dynamic routing is a complex subject. There are many routers on different networks and all can be configured differently. It become even more complicated when you add to this each routing protocol having slightly different names for similar features, and many configurable features for each protocol.

To better understand dynamic routing, here are some explanations of common dynamic routing terms.

- [Aggregated routes and addresses](#)
- [Autonomous system \(AS\)](#)
- [Area border router \(ABR\)](#)
- [Neighbor routers](#)
- [Route maps](#)

- [Access lists](#)
- [Bi-directional forwarding detection \(BFD\)](#)

For more details on a term as it applies to a dynamic routing protocol, see one of “[Border Gateway Protocol \(BGP\)](#)” on page 1751, “[Routing Information Protocol \(RIP\)](#)” on page 1715, or “[Open Shortest Path First \(OSPF\)](#)” on page 1787.

Aggregated routes and addresses

Just as an aggregate interface combines multiple interfaces into one virtual interface, an aggregate route combines multiple routes into one. This reduces the amount of space those routes require in the routing tables of the routers along that route. The trade-off is a small amount of processing to aggregate and de-aggregate the routes at either end.

The benefit of this method is that you can combine many addresses into one, potentially reducing the routing table size immensely. The weakness of this method is if there are holes in the address range you are aggregating you need to decide if its better to break it into multiple ranges, or accept the possibility of failed routes to the missing addresses.

To manually aggregate the range of IP addresses from 192.168.1.100 to 192.168.1.103

- 1 Convert the addresses to binary

```
192.168.1.100 = 11000000 10101000 00000001 01100100
192.168.1.101 = 11000000 10101000 00000001 01100101
192.168.1.102 = 11000000 10101000 00000001 01100110
192.168.1.103 = 11000000 10101000 00000001 01100111
```

- 2 Determine the maximum number of matching bits common to the addresses.

There are 30-bits in common, with only the last 2-bits being different.

- 3 Record the common part of the address.

```
11000000 10101000 00000001 0110010X = 192.168.1.100
```

- 4 For the netmask, assume all the bits in the netmask are 1 except those that are different which are 0.

```
11111111 11111111 11111111 11111100 = 255.255.255.252
```

- 5 Combine the common address bits and the netmask.

```
192.168.1.100/255.255.255.252
```

Alternately the IP mask may be written as a single number:

```
192.168.1.100/2
```

- 6 As required, set variables and attributes to declare the routes have been aggregated, and what router did the aggregating.

Autonomous system (AS)

An Autonomous System (AS) is one or more connected networks that use the same routing protocol, and appear to be a single unit to any externally connected networks. For example an ISP may have a number of customer networks connected to it, but to any networks connected externally to the ISP it appears as one system or AS. An AS may also be referred to as a routing domain.

It should be noted that while OSPF routing takes place within one AS, the only part of OSPF that deals with the AS is the AS border router (ASBR).

There are multiple types of AS defined by how they are connected to other ASes. A multihomed AS is connected to at least two other ASes and has the benefit of redundancy — if one of those ASes goes down, your AS can still reach the Internet through its other connection. A stub AS only has one connection, and can be useful in specific configurations where limited access is desirable.

Each AS has a number assigned to it, known as an ASN. In an internal network, you can assign any ASN you like (a private AS number), but for networks connected to the Internet (public AS) you need to have an officially registered ASN from Internet Assigned Numbers Authority (IANA). ASNs are typically 16-bit numbers — ASNs from 1 - 64,511 are designated for public use.



As of January 2010, AS numbers are 4 bytes long instead of the former 2 bytes. RFC 4893 introduced 32-bit ASNs, which FortiGate units support.

Do you need your own AS?

The main factors in deciding if you need your own AS or if you should be part of someone else's are:

- exchanging external routing information
- many prefixes should exist in one AS as long as they use the same routing policy
- when you use a different routing protocol than your border gateway peers (for example your ISP uses BGP, and you use OSPF)
- connected to multiple other AS (multi-homed)

You should not create an AS for each prefix on your network. Neither should you be forced into an AS just so someone else can make AS-based policy decisions on your traffic.

There can be only one AS for any prefix on the Internet. This is to prevent routing issues.

What AS number to use?

In addition to overseeing IP address allocation and Domain Name Systems (DNS), the Internet Assigned Numbers Authority (IANA) assigns public AS numbers. The public AS numbers are from 1 to 64,511. The ASNs 0, 54272–64511, and 65535 are reserved by the IANA. These ASNs should not be used.

ASNs are assigned in blocks by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIRs) who then assign ASNs to companies within that RIRs geographic area. Usually these companies are ISPs, and to receive an ASN you must complete the application process of the local RIR and be approved before being assigned an ASN. The RIRs names and regions are:

AFRINIC	Serves the African continent
APNIC	Asia-Pacific including China, India, and Japan
ARIN	American registry including Canada and United States
LACNIC	Latin America, including Mexico, Caribbean, Central and South America
RIPE NCC	Europe, the Middle East, former USSR, and parts of Central Asia

AS numbers from 64512 to 65534 are reserved for private use. Private AS numbers can be used for any internal networks with no outside connections to the Internet such as test networks, classroom labs, or other internal-only networks that do not access the outside world. You can also configure border routers to filter out any private ASNs before routing traffic to the outside world. If you must use private ASNs with public networks, this is the only way to configure them. However, it is risky because many other private networks could be using the same ASNs and conflicts will happen. It would be very much like your local 192.168.0.0 network being made public — the resulting problems would be widespread.

In 1996, when RFC 1930 was written only 5,100 ASes had been allocated and a little under 600 ASes were actively routed in the global Internet. Since that time many more public ASNs have been assigned, leaving only a small number. For this reason 32-bit ASNs (four-octet ASNs) were defined to provide more public ASNs. RFC 4893 defines 32-bit ASNs, and FortiGate units support these larger ASNs as of FortiOS version 4.2

Area border router (ABR)

Routers within an AS advertise updates internally and only to each other. However, routers on the edge of the AS must communicate both with routers inside their AS and with routers external to their AS, often running a different routing protocol. These routers are called Area Border Routers (ABRs) or edge routers. Often ABRs run multiple routing protocols to be able to redistribute traffic between different ASes that are running different protocols, such as the edge between an ISP's IS-IS routing network and an large company's OSPF network.

OSPF defines ABRs differently from other routers. In OSPF, an ABR is an OSPF router that connects another AS to the backbone AS, and is a member of all the areas it connects to. An OSPF ABR maintains a LSA database for each area that it is connected to. The concept of the edge router is present, but its the edge of the backbone instead of the edge of the OSPF supported ASes.

Neighbor routers

Routing involves routers communicating with each other. To do this, routers need to know information about each other. These routers are called neighbor routers, and are configured in each routing protocol. Each neighbor has custom settings since some routers may have functionality others routers lack. Neighbour routers are sometimes called peers.

Generally neighbor routers must be configured, and discovered by the rest of the network before they can be integrated to the routing calculations. This is a combination of the network administrator configuring the new router with its neighbor router addresses, and the routing network discovering the new router, such as the hello packets in OSPF. That discovery initiates communication between the new router and the rest of the network.

Route maps

Route maps are a way for the FortiGate unit to evaluate optimum routes for forwarding packets or suppressing the routing of packets to particular destinations. Compared to access lists, route maps support enhanced packet-matching criteria. In addition, route maps can be configured to permit or deny the addition of routes to the FortiGate unit routing table and make changes to routing information dynamically as defined through route-map rules.

Route maps can be used for limiting both received route updates, and sent route updates. This can include the redistribution of routes learned from other types of routing. For example if you don't want to advertise local static routes to external networks, you could use a route map to accomplish this.

The FortiGate unit compares the rules in a route map to the attributes of a route. The rules are examined in ascending order until one or more of the rules in the route map are found to match one or more of the route attributes.

As an administrator, route maps allow you to group a set of addresses together and assign them a meaningful name. Then during your configuration, you can use these route-maps to speed up configuration. The meaningful names ensure fewer mistakes during configuration as well.

The default rule in the route map (which the FortiGate unit applies last) denies all routes. For a route map to take effect, it must be called by a FortiGate unit routing process.

The syntax for route maps are:

```
config router route-map
edit <route_map_name>
    set comments
    config rule
    edit <route_map_rule_id>
        set action
        set match-*
        set set-*
```

The `match-*` commands allow you to match various parts of a route. The `set-*` commands allow you to set routing information once a route is matched.

For an example of how route maps can be used to create receiving or sending "groups" in routing, see ["Redistributing and blocking routes in BGP" on page 1780](#).

Access lists

Use this command to add, edit, or delete access lists. Access lists are filters used by FortiGate unit routing processes. For an access list to take effect, it must be called by a FortiGate unit routing process (for example, a process that supports RIP or OSPF). Use `access-list6` for IPv6 routing.

Access lists can be used to filter which updates are passed between routers, or which routes are redistributed to different networks and routing protocols. You can create lists of rules that will match all routes for a specific router or group of routers.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.



If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

The syntax for access lists is:

```
config router access-list, access-list6
edit <access_list_name>
    set comments
```

```
config rule
edit <access_list_id>
set action
set exact-match
set prefix
set prefix6
set wildcard
```

For an example of how access lists can be used to create receiving or sending “groups” in routing, see [“Redistributing and blocking routes in BGP” on page 1780](#).

Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD send packets to each other at a negotiated rate. If packets from a BFD-protected router fail to arrive, then that router is declared down. BFD communicates this information to the routing protocol and the routing information is updated.

BFD neighbors establish if BFD is enabled in OSPF or BGP routers that establish as neighbors.

The CLI commands associated with BFD include:

```
config router bgp
config neighbor
set bfd

config router ospf
set bfd
```

Per-VDOM configuration:

```
config system settings
set bfd
set bfd-desired-min-tx
set bfd-required-min-rx
set bfd-detect-mult
set bfd-dont-enforce-src-port
```

Per-interface (override) configuration:

```
config system interface
edit <interface_name>
set bfd enable
set bfd-desired-min-tx
set bfd-detect-mult
set bfd-required-min-rx
```

For more information about BFD in BGP, see [“Bi-directional forwarding detection \(BFD\)” on page 1766](#).

IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming popular and new versions of the dynamic routing protocols have been introduced.

As of FortiOS v4.1, dynamic routing supports IPv6 on your FortiGate unit. The new versions of these protocols and the corresponding RFCs are:

- **BGP4+** — RFC 2545, and RFC 2858 Multiprotocol Extensions for IPv6 Inter-Domain Routing, and Multiprotocol Extensions for BGP-4 (MP-BGP) respectively. See [“BGP and IPv6” on page 1752](#)
- **RIP next generation (RIPng)** — RFC 2080 - Routing Information Protocol next generation (RIPng). See [“RIP and IPv6” on page 1716](#).
- **OSPFv3** — RFC 2740 Open Shortest Path First version 3 (OSPFv3) for IPv6 support. See [“OSPF and IPv6” on page 1788](#).

As with most advanced routing features on your Fortigate unit, IPv6 settings for dynamic routing protocols are CLI-only. To configure IPv6 for RIP, BGP, or OSPF protocols you must use the CLI commands.



Routing Information Protocol (RIP)

This section describes the Routing Information Protocol (RIP).

The following topics are included in this section:

- [RIP background and concepts](#)
- [Troubleshooting RIP](#)
- [RIP routing examples](#)

RIP background and concepts

This section contains:

- [Background](#)
- [Parts and terminology of RIP](#)
- [How RIP works](#)

Background

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. Its widespread use started when an early version of RIP was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by RIP, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

RIP benefits include being well suited to smaller networks, is in widespread use, near universal support on routing hardware, quick to configure, and works well if there are no redundant paths. However, RIP updates are sent out node-by-node so it can be slow to find a path around network outages. RIP also lacks good authentication, can not choose routes based on different quality of service methods, and can create network loops if you are not careful.

The FortiGate implementation of RIP supports RIP version 1 (see RFC 1058), RIP version 2 (see RFC 2453), and the IPv6 version RIPng (see RFC 2080).

RIP v1

In 1988 RIP version 1, defined in RFC 1058, was released. The RFC even states that RIP v1 is based on Linux routed due to it being a “defacto standard”.

It uses classful addressing and uses broadcasting to send out updates to router neighbors. There is no subnet information included in the routing updates in classful routing, and it does not support CIDR addressing — subnets must all be the same size. Also, route summarization is not possible.

RIP v1 has no router authentication method, so it is vulnerable to attacks through packet sniffing, and spoofing.

RIP v2

In 1993, RIP version 2 was developed to deal with the limitations of RIP v1. It was not standardized until 1998. This new version supports classless routing, and subnets of various sizes.

Router authentication was added in RIP v2 — it supports MD5. MD5 hashes are an older encryption method, but this is much improved over no security at all.

In RIP v2 the hop count limit remained at 15 to be backwards compatible with RIP v1.

RIP v2 uses multicasting to send the entire routing table to router neighbors, thereby reducing the traffic for devices that are not participating in RIP routing.

Routing tags were added as well, which allow internal routes or redistributed routes to be identified as such.

RIPng

RIPng, defined in RFC 2080, is an extension of RIP2 designed to support IPv6. However, RIPng varies from RIPv2 in that it is not fully backwards compatible with RIPv1.

- RIPng does not support RIPv1 update authentication, it relies on IPsec
- RIPng does not allow attaching tags to routes as in RIPv2
- RIPng requires specific encoding of the next hop for a set of route entries, unlike RIPv2 that encodes the next-hop into each route entry .

Parts and terminology of RIP

Before you can understand how RIP functions, you need to understand some of the main concepts and parts of RIP.

This section includes:

- [RIP and IPv6](#)
- [Default information originate option](#)
- [Garbage, timeout, and update timers](#)
- [Authentication and key-chain](#)
- [Access Lists](#)

RIP and IPv6

RIP Next Generation (RIPng) is a new version of RIP was released that includes support for IPv6.

The FortiGate unit command `config router ripng` is almost the same as `config router rip`, except that IPv6 addresses are used. Also if you are going to use prefix or access lists with RIPng, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to troubleshoot RIPng, it is the same as with RIP but specify the different protocol, and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table, or other related information.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ip6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in Transparent mode.

For example, you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01:: where it will need access to an IPv4 network again. Use the following command:

```
config system ipv6-tunnel
  edit test_tunnel
    set destination 2002:A0A:A01::
    set interface port1
    set source 2002:C0A8:3201::
  end
end
```

The CLI commands associated with RIPng include:

```
config router ripng
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

Default information originate option

This is the second advanced option for RIP in the web-based manager, right after metric. Enabling default-information-originate will generate and advertise a default route into the FortiGate unit's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. RIP does not create the default route unless you use the always option.

Select *Disable* if you experience any issues or if you wish to advertise your own static routes into RIP updates.

The CLI commands associated with default information originate include:

```
config router rip
  set default-information-originate
end
```

Garbage, timeout, and update timers

RIP uses various timers to regulate its performance including a garbage timer, timeout timer, and update timer. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations — if you change these settings, ensure that the new settings are compatible with local routers and access servers.



The Timeout period should be at least three times longer than the Update period. If the Update timer is smaller than Timeout or Garbage timers, you will experience an error.

The CLI commands associated with garbage, timeout, and update timers include:

```
config router rip
  set garbage-timer
  set timeout-timer
  set update-timer
end
```

Garbage timer

The garbage timer is the amount of time (in seconds) that the FortiGate unit will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This will result in a smaller routing table which is useful if you have a very large network, or if your network changes frequently.

Update timer

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, otherwise you will experience an error.

If you are experiencing significant RIP traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience time outs that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the FortiGate unit will keep a reachable route in the routing table while no updates for that route are received. If the FortiGate unit receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the depute period, otherwise you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods — it may be considerable time before the time the FortiGate unit is done waiting for all the timers to expire on unresponsive routes.

Authentication and key-chain

RIP version 2 uses authentication keys to ensure that the routing information exchanged between routers is reliable. RIP version 1 has no authentication. For authentication to work both the sending and receiving routers must be set to use authentication, and must be configured with the same keys.

The sending and receiving routers need to have their system dates and times synchronized to ensure both ends are using the same keys at the proper times. However, you can overlap the key lifetimes to ensure that a key is always available even if there is some difference in the system times.

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. The FortiGate unit migrates from one key to the next according to the scheduled send and receive lifetimes.

Key-chain is a CLI router command. You use this command to manage RIP version 2 authentication keys. You can add, edit or delete keys identified by the specified key number.

This example shows how to configure a key-chain with two keys that are valid sequentially in time. This example creates a key-chain called "rip_key" that has a password of "fortinet". The accepted and send lifetimes are both set to the same values — a start time of 9:00am February 23, 2010 and an end time of 9:00am March 17, 2010. A second key is configured with a password of "my_fortigate" that is valid from March 17, 2010 9:01am to April 1 2010 9:00am. This "rip_key" keychain is then used on the port1 interface in RIP.

```
config router key-chain
  edit "rip_key"
    config key
      edit 1
        set accept-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
        set key-string "fortinet"
        set send-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
      next
      edit 2
        set accept-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
        set key-string "my_fortigate"
        set send-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
      next
    end
  end
config router rip
  config interface
    edit port1
      set auth-keychain "rip_key"
    end
  end
end
```

The CLI commands associated with authentication keys include:

```
config router key-chain

config router rip
  config interface
    edit <interface>
      set auth-keychain
      set auth-mode
      set auth-string
    end
  end
end
```

Access Lists

Access lists are filters used by FortiGate unit RIP and OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also its easier to troubleshoot since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the RIPng or OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of 10.10.10.10 and 11.11.11.11, enter the command:

```
config router access-list
  edit test_list
    config rule
      edit 1
        set prefix 10.10.10.10 255.255.255.255
        set action allow
        set exact-match enable
      next
      edit 2
        set prefix 11.11.11.11 255.255.255.255
        set action allow
        set exact-match enable
      end
    end
  end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the command access-list6 as follows:

```
config router access-list6
  edit test_list_ip6
    config rule
      edit 1
        set prefix6 2002:A0A:A0A:0:0:0:0:0/48
        set action deny
      next
      edit 2
        set prefix6 2002:B0B:B0B:0:0:0:0:0/48
        set action deny
      end
    end
  end
```

To use an access_list, you must call it from a routing protocol such as RIP. The following example uses the access_list from the earlier example called `test_list` to match routes coming in on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially increase . Enter the following command:

```
config router rip
  config offset-list
    edit 5
      set access-list test_list
```

```
set direction in
set interface port1
set offset 3
set status enable
end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose

How RIP works

As one of the original modern dynamic routing protocols, RIP is straight forward. It's routing algorithm is not complex, there are some options to allow fine tuning, and its straight forward to configure RIP on FortiGate units.

From RFC 1058:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

This section includes:

- [RIP versus static routing](#)
- [RIP metric — hop count](#)
- [The Bellman–Ford routing algorithm](#)
- [Passive versus active RIP interfaces](#)
- [RIP packet structure](#)

RIP versus static routing

RIP was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, RIP is a big step forward from simple static routing.

While RIP may be slow in response to network outages, static routing has zero response. The same is true for convergence — static routing has zero convergence. Both RIP and static routing have the limited hop count, so its not a strength or a weakness. Count to infinity can be a problem, but typically can be fixed as it happens or is the result of a network outage that would cause even worse problems on static routing network.

Overall, RIP is a large step forward when compared to static routing.

RIP metric — hop count

RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiGate unit, while a hop count of 16 represents a network that cannot be reached. Each network that a packet travels through to reach its destination usually counts as one hop. When the FortiGate unit compares two routes to the same destination, it adds the route having the lowest hop count to the routing table. As you can see in [“RIP packet structure” on page 1725](#), the hop count is part of a RIP v2 packet making it very important.

Similarly, when RIP is enabled on an interface, the FortiGate unit sends RIP responses to neighboring routers on a regular basis. The updates provide information about the routes in the FortiGate unit's routing table, subject to the rules that you specify for advertising those routes. You can specify how often the FortiGate unit sends updates, the period of time a route can be kept in the routing table without being updated, and for routes that are not updated regularly you can specify the period of time that the unit advertises a route as unreachable before it is removed from the routing table.

If hops are weighted higher than one, it becomes very easy to reach the upper limit. This higher weighting will effectively limit the size of your network depending on the numbers used. Merely changing from the default of 1.0 to 1.5 will lower the effective hop count from 15 to 10. This is acceptable for smaller networks, but can be a problem as your network expands over time.

In RIP, you can use the `offset` command to artificially increase the hop count of a route. Doing this will make this route less preferred, and in turn it will get less traffic. Offsetting routes is useful when you have network connections of different bandwidths, different levels of reliability, or different costs. In each of these situations you still want the redundancy of multiple route access, but you don't want the bulk of your traffic using these less preferred routes. For an example of RIP offset, see ["Access Lists" on page 1719](#).

The Bellman-Ford routing algorithm

The routing algorithm used by RIP was first used in 1967 as the initial routing algorithm of the ARPANET. The Bellman-Ford algorithm is distributed because it involves a number of nodes (routers) within an Autonomous system, and consists of the following steps:

- 1 Each node calculates the distances between itself and all other nodes within the AS and stores this information as a table.
- 2 Each node sends its table to all neighboring nodes.
- 3 When a node receives distance tables from its neighbors, it calculates the shortest routes to all other nodes and updates its own table to reflect any changes.


To examine how this algorithm functions let's look at a network with 4 routers — routers 1 through 4. The distance from router1 to router2 is 2 hops, 1 to 3 is 3 hops, and 2 to 3 is 4 hops. Router4 is only connected to routers 2 and 3, each distance being 2 hops.

- 1 Router1 finds all the distance to the other three routers — router 2 is 2, router 3 is 3. Router1 doesn't have a route to router 4.
- 2 Routers 2 through 4 do the same calculations from their point of views.
- 3 Once router 1 gets an update from router 2 or 3, it will get their route to router 4. At that point it now has a route to router 4 and installs that in its local table.
- 4 If router1 gets an update from router3 first, it has a hop count of 5 to reach router4. But when router2 sends its update, router1 will go with router2's shorter 4 hops to reach router4. Future updates don't change this unless they are shorter than 4 hops, or the routing table route goes down.

Figure 164: RIP algorithm example in 4 steps


Step 1. Router1 finds the distance to other routers. It has no route to router4.

Router1 table:
Distance to router2 = 2 hops
Distance to router3 = 3 hops




Step 2. All routers do the same as router1, and send out updates with the table of routes.

Note that router1 and router4 do not update each other, but rely on router2 and router3 to pass along deputed.




Step 3. Each router looks at the updates it receives, and adds any new or shorter routes to its table.

Router1 updated table:
Distance to router2 = 2 hops
Distance to router3 = 3 hops
Distance to router4 = 4 hops
Distance to router4 = 5 hops



Step 4. The shortest route to router4 is installed, and the other routes to router4 are removed from the table.

Router1 updated table:
Distance to router2 = 2 hops
Distance to router3 = 3 hops
Distance to router4 = 4 hops



The good part about the Bellman-Ford algorithm in RIP is that the router only uses the information it needs from the update. If there are no newer, better routes than the ones the router already has in its routing table, there is no need to change its routing table. And no change means no additional update, so less traffic. But even when there is update traffic, the RIP packets are very small so it takes many updates to affect overall network bandwidth. For more information about RIP packets, see [“RIP packet structure” on page 1725](#).

The main disadvantage of the Bellman-Ford algorithm in RIP is that it doesn't take weightings into consideration. While it is possible to assign different weights to routes in RIP, doing so severely limits the effective network size by reducing the hop count limit. Also other dynamic routing protocols can take route qualities, such as reliability or delay, into consideration to provide not only the physically shortest but also the fastest or more reliable routes as you choose.

Another disadvantage of the Bellman-Ford algorithm is due to the slow updates passed from one RIP router to the next. This results in a slow response to changes in the network topology, which in turn results in more attempts to use routes that are down which wastes time and network resources.

Passive versus active RIP interfaces

Normally the FortiGate unit's routing table is kept up to date by periodically asking the neighbors for routes, and sending your routing updates out. This has the downside of generating a lot of extra traffic for large networks. The solution to this problem is passive interfaces.

An standard interface that supports RIP is active by default — it both sends and receives updates by actively communicating with its neighbors. A passive RIP interface does not send out updates — it just listens to the updates of other routers. This is useful in reducing network traffic, and if there are redundant routers in the network that would be sending out essentially the same updates all the time.

The following example shows how to create a passive RIP v2 interface on port1, using MD5 authentication and a key-chain called `passiveRIPv2` that has already been configured. Note that in the CLI, you enable passive by disabling `send-version2-broadcast`.

To create a passive RIP interface - web-based manager

- 1 Go to *Router > Dynamic > RIP*.
- 2 Under *Interfaces*, select *Create New*.
- 3 Select port1 as the *Interface*.
- 4 Select 2 as both the *Send Version* and *Receive Version*.
- 5 Select MD5 for *Authentication*.
- 6 Select the `passiveRIPv2` *Key-chain*.
- 7 Select *Passive Interface*.
- 8 Select OK to accept this configuration, and return to the main RIP display page.

To create a passive RIP v2 interface on port1 using MD5 authentication- CLI

```
config router rip
  config interface
    edit port1
      set send-version2-broadcast disable
      set auth-keychain "passiveRIPv2"
```

```

set auth-mode md5
set receive-version 2
set send-version 2
end
end

```

RIP packet structure

It is hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how will help you better understand the RIP protocol, and better configure your network for it.

This section provides information on the contents of RIP 1 and RIP 2 packets.

RIP version 1

RIP version 1, or RIP IP packets are 24 bytes in length. The empty areas were left for future expansion.

Table 99: RIP IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

The following descriptions summarize the RIP version 1 packet fields.

Command — Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.

Version — Specifies the RIP version used. This field can signal different potentially incompatible versions.

Zero field — This field defaults to zero, and is not used by RFC 1058 RIP.

Address-family identifier (AFI) — Specifies the address family used. RIP is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.

IP Address — Specifies the IP address for the entry.

Metric — This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable the metric is 16.

RIP version 2

RIP version 2 has more features than RIP 1 and this is reflected in its packets. RIP 2 packets are similar in format to RIP 1, but carry more information. All but one of the empty zero fields in RIP 1 packets now contain information.

Table 100: RIP 2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag	4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric
----------------	----------------	---------------	------------	------------------	-------------------	---------------	-----------------	---------------

The following descriptions summarize the fields RIP 2 adds to the RIP IP header. The other fields have been described above for RIP 1.

Unused — Has a value set to zero, and is intended for future use

Route tag — Provides a method for distinguishing between internal routes learned by RIP and external routes learned from other protocols.

Subnet mask — Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.

Next hop — Indicates the IP address of the next hop to which packets for the entry should be forwarded.

Troubleshooting RIP

This section is about troubleshooting RIP. For general troubleshooting information, see the Troubleshooting chapter.

This section includes:

- [Routing Loops](#)
- [Split horizon and Poison reverse updates](#)
- [Debugging IPv6 on RIPng](#)

Routing Loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

Routing loops' effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the RIP routers connected to it. In a network loop, it's possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- [Check your logs](#)

- [Use SNMP network monitoring](#)
- [Use dead gateway detection and e-mail alerts](#)
- [Look at the packet flow](#)

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it. Ping, traceroute, and other basic troubleshooting tools are covered in ["" on page 1714](#).

Check your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report > Log & Archive Access*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

- 1 Go to *System > Config > SNMP*.
- 2 Enable *SMTP Agent* and select *Apply*.

Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.

- 3 Under *SNMP v1/v2* or *SNMP v3* as appropriate, select *Create New*.

SNMP v3

User Name	Enter the SNMP user ID.
Security Level	Select authentication or privacy as desired. Select the authentication or privacy algorithms to use and enter the required passwords.
Notification Host	Enter the IP addresses of up to 16 hosts to notify.
Enable Query	Select. The <i>Port</i> should be 161. Ensure that your security policies allow ports 161 and 162 (SNMP queries and traps) to pass.

SNMP v1/v2

Hosts	Enter the IP addresses of up to 8 hosts to notify. You can also specify the network <i>Interface</i> , or leave it as <i>ANY</i> .
Queries	Enable v1 and/or v2 as needed. The <i>Port</i> should be 161. Ensure that your security policies allow port 161 to pass.
Traps	Enable v1 and/or v2 as needed. The <i>Port</i> should be 162. Ensure that your security policies allow port 162 to pass.

- 4 Select the events for which you want notification. For routing loops this should include *CPU usage is high*, *Memory is low*, and possibly *Log disk space is low*. If there are problems the log will be filling up quickly, and the FortiGate unit's resources will be overused.
- 5 Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Use dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with dead gateway detection and e-mail alerts

- 1 To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
- 2 Enter the *Ping Server* IP address and select the *Interface* that connects to it.
- 3 Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.

To configure notification of failed gateways

- 1 Go to *Log&Report > Log Config > Alert E-mail*.
- 2 Enter your email details.
- 3 Select the *Configuration changes* event.
- 4 Select *Apply*.

You might also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email about the outage.

Look at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. This is same idea as police pulling over a car and asking the driver where they have been, and what the conditions were like.

The method used in the troubleshooting sections “[Debugging IPv6 on RIPng](#)” on [page 1729](#) and on debugging the packet flow apply here as well. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable.

Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

- 1 Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
- 2 Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will not work. However, if B tries to use it this begins an endless loop.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse “poisons” routes away from the destination that use the current router in their route to the destination. This “poisoned” route is marked as unreachable for routers that cannot use it. In RIP this means that route is marked with a distance of 16.

Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
FGT# diagnose debug enable
FGT# diagnose ipv6 router rip level info
FGT# diagnose ipv6 router rip all enable
```

These three commands will:

- turn on debugging in general

- set the debug level to information, a verbose reporting level
- turn on all rip router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

For more information, see [“Testing the IPv6 RIPng information” on page 1750](#)

RIP routing examples

The following examples for RIP:

- [Simple RIP example](#)
- [RIPng — RIP and IPv6](#)

Simple RIP example

This is an example of a typical medium sized network configuration using RIP routing.

Your company has 3 small local networks, one for each department. These networks are connected by RIP, and then connected to the Internet. Each subnet has more than one route, for redundancy. There are two central routers that are both connected to the internet, and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running RIP, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running those will need to be redistributed through the RIP network.

To keep the example simple, there will be no authentication of router traffic.

With RIP properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. RIP is good for a smaller network due to its lack of complex configurations.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units system information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

Basic network layout

Your company has 3 departments each with their own network — Sales, R&D, and Accounting. Each network has routers that are not running RIP as well as FortiGate units running RIP.

The R&D network has two RIP routers, and each is connected to both other departments as well as being connected to the Internet through the ISP router. The links to the Internet are indicated in black.


The three internal networks do not run RIP. They use static routing because they are small networks. This means the FortiGate units have to redistribute any static routes they learn so that the internal networks can communicate with each other.

Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows. Note that the Interfaces that connect Router2 and Router3 also connect to the R&D network.

Table 101: Rip example network topology

Network	Router	Interface & Alias	IP address
Sales	Router1	port1 (internal)	10.11.101.101
		port2 (router2)	10.11.201.101
		port3 (router3)	10.11.202.101
R&D	Router2	port1 (internal)	10.12.101.102
		port2 (router1)	10.11.201.102
		port3 (router4)	10.14.201.102
		port4 (ISP)	172.20.120.102
	Router3	port1 (internal)	10.12.101.103
		port2 (router1)	10.11.201.103
		port3 (router4)	10.14.202.103
		port4 (ISP)	172.20.120.103
Accounting	Router4	port1 (internal)	10.14.101.104
		port2 (router2)	10.14.201.104
		port3 (router3)	10.14.202.104

Figure 165: Network topology for the simple RIP example

Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 4.0 MR3 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 through port4 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- Only FortiGate units are running RIP on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.
- Router2 and Router3 each have their own connection to the Internet, indicated in black in [Figure 165](#).

General configuration steps

This example is very straight forward. The only steps involved are:

- [Configuring the FortiGate units system information](#)
- [Configuring FortiGate unit RIP router information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Configuring the FortiGate units system information

Each FortiGate unit needs their hostname, and interfaces configured.

For IP numbering, Router2 and Router3 use the other routers numbering where needed.

Router2 and Router3 have dead gateway detection enabled on the ISP interfaces using Ping. Remember to contact the ISP and confirm their server has ping enabled.

Configure the hostname, interfaces, and default route

To configure Router1 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router1".
- 3 Go to *Router > Static > Static Route*.
- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 5 Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 6 Go to *System > Network > Interface*.
- 7 Edit port1 (internal) interface.
- 8 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal sales network
Administrative Status	Up

- 9 Edit port2 (router2) interface.
- 10 Set the following information, and select *OK*.

Alias	router2
IP/Netmask	10.11.201.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network & internet through Router2
Administrative Status	Up

- 11 Edit port3 (router3) interface.

12 Set the following information, and select OK.

Alias	router3
IP/Netmask	10.11.202.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

To configure Router1 system information - CLI

```

config system global
    set hostname Router1
end

config router static
    edit 1
        set device "port2"
        set distance 45
        set gateway 10.11.201.102
    next
    edit 2
        set device "port3"
        set distance 45
        set gateway 10.11.202.103
    end
end

config system interface
    edit port1
        set alias internal
        set ip 10.11.101.101/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal sales network"
    next
    edit port2
        set alias ISP
        set allowaccess https ssh ping
        set ip 10.11.201.101/255.255.255.0
        set description "Link to R&D network & internet through
            Router2"
    next
    edit port3
        set alias router3
        set ip 10.11.202.101/255.255.255.0
        set allowaccess https ssh ping
        set description "Link to R&D network & internet through
            Router2"
    end
end
end

```

To configure Router2 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router2".
- 3 Go to *Router > Static > Static Route*.
- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

- 5 Go to *System > Network > Interface*.
- 6 Edit port1 (internal) interface.
- 7 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.12.101.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router3
Administrative Status	Up

- 8 Edit port2 (router1) interface.
- 9 Set the following information, and select *OK*.

Alias	router1
IP/Netmask	10.12.201.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and the Sales network
Administrative Status	Up

- 10 Edit port3 (router4) interface.
- 11 Set the following information, and select *OK*.

Alias	router4
IP/Netmask	10.12.301.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and the accounting network
Administrative Status	Up

- 12 Edit port4 (ISP) interface.
- 13 Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.102/255.255.255.0
Administrative Access	HTTPS SSH PING

Detect Interface Status for Gateway Load Balancing	enable
Detect Server	172.20.120.5
Detect Protocol	Ping
Detect Interface Status for Gateway Load Balancing	enable
Description	Internet through ISP
Administrative Status	Up

To configure Router2 system information - CLI

```

config system global
    set hostname Router2
end
config router static
    edit 1
        set device "port4"
        set distance 5
        set gateway 172.20.130.5
    end
end
config system interface
    edit port1
        set alias internal
        set ip 10.11.101.102/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal RnD network and Router3"
    next
    edit port2
        set alias router1
        set allowaccess https ssh ping
        set ip 10.11.201.102/255.255.255.0
        set description "Link to Router1"
    next
    edit port3
        set alias router3
        set ip 10.14.202.102/255.255.255.0
        set allowaccess https ssh ping
        set description "Link to Router4"
    next
    edit port4
        set alias ISP
        set ip 172.20.120.102/255.255.255.0
        set allowaccess https ssh ping
        set description "ISP and internet"
    end
end
end

```

To configure Router3 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router3".

- 3 Go to *Router > Static > Static Route*.
- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

- 5 Go to *System > Network > Interface*.
- 6 Edit port1 (internal) interface.
- 7 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.12.101.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router2
Administrative Status	Up

- 8 Edit port2 (router1) interface.
- 9 Set the following information, and select *OK*.

Alias	router1
IP/Netmask	10.13.201.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and Sales network
Administrative Status	Up

- 10 Edit port3 (router4) interface.
- 11 Set the following information, and select *OK*.

Alias	router4
IP/Netmask	10.13.301.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and accounting network
Administrative Status	Up

- 12 Edit port4 (ISP) interface.
- 13 Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect Interface Status for Gateway Load Balancing	enable
Detect Server	172.20.120.5

Detect Protocol	Ping
Description	Internet and ISP
Administrative Status	Up

To configure Router3 system information - CLI

```

config system global
    set hostname Router3
end

config router static
    edit 1
        set device "port4"
        set distance 5
        set gateway 172.20.130.5
    end
end

config system interface
    edit port1
        set alias internal
        set ip 10.12.101.103/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal RnD network and Router2"
    next
    edit port2
        set alias ISP
        set allowaccess https ssh ping
        set ip 10.11.201.103/255.255.255.0
        set description "Link to Router1"
    next
    edit port3
        set alias router3
        set ip 10.14.202.103/255.255.255.0
        set allowaccess https ssh ping
        set description "Link to Router4"
    next
    edit port4
        set alias ISP
        set ip 172.20.120.103/255.255.255.0
        set allowaccess https ssh ping
        set description "ISP and internet"
    end
end

```

To configure Router4 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router4".
- 3 Go to *Router > Static > Static Route*.

- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 5 Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 6 Go to *System > Network > Interface*.
 7 Edit port 1 (internal) interface.
 8 Set the following information, and select **OK**.

Alias	internal
IP/Netmask	10.14.101.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal accounting network
Administrative Status	Up

- 9 Edit port 2 (router2) interface.
 10 Set the following information, and select **OK**.

Alias	router2
IP/Netmask	10.14.201.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network & internet through Router2
Administrative Status	Up

- 11 Edit port 3 (router3) interface.
 12 Set the following information, and select **OK**.

Alias	router3
IP/Netmask	10.14.301.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

To configure Router4 system information - CLI

```
config system global
  set hostname Router4
```

```
end

config router static
  edit 1
    set device "port2"
    set distance 45
    set gateway 10.14.201.102
  next
  edit 2
    set device "port3"
    set distance 45
    set gateway 10.14.202.103
  end
end

config system interface
  edit port1
    set alias internal
    set ip 10.14.101.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Internal sales network"
  next
  edit port2
    set alias router2
    set allowaccess https ssh ping
    set ip 10.14.201.104/255.255.255.0
    set description "Link to R&D network & internet through Router2"
  next
  edit port3
    set alias router3
    set ip 10.14.202.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to R&D network & internet through Router2"
  end
end
```

Configuring FortiGate unit RIP router information

With the interfaces configured, RIP can now be configured on the FortiGate units.

This includes the following steps:

- configure RIP version used
- redistribute static networks
- add networks serviced by RIP
- add interfaces that support RIP on the Fortigate unit

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures — repeat the procedures once for each FortiGate unit.

Configure RIP settings on Router1 and Router4 - web-based manager

- 1 Go to *Router > Dynamic > RIP*.
- 2 Select 2 for *RIP Version*.
- 3 In *Advanced Options*, under *Redistribute* enable *Static*.
- 4 Leave the other *Advanced Options* at default values.
- 5 Enter the following networks, and select *Add* after each:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0
- 6 For interface, select *Create New* and set the following information.

Interface	port1 (internal)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

- 7 For interface, select *Create New* and set the following information.

Interface	port2 (router2)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

- 8 For interface, select *Create New* and set the following information.

Interface	port3 (router3)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

Configure RIP settings on Router1 and Router4 - CLI

```

config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2

```

```

        set send-version 1 2
    next
    edit "port3"
        set receive-version 1 2
        set send-version 1 2
    end
    config network
    edit 1
        set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
        set prefix 10.12.0.0 255.255.0.0
    next
    edit 3
        set prefix 10.14.0.0 255.255.0.0
    next
    edit 4
        set prefix 172.20.120.0 255.255.255.0
    end
    config redistribute "static"
        set status enable
    end
end
end

```

Configure RIP settings on Router2 and Router3- web-based manager

- 1 Go to *Router > Dynamic > RIP*.
- 2 Select 2 for *RIP Version*.
- 3 In *Advanced Options*, under *Redistribute* enable *Static*.
- 4 Leave the other *Advanced Options* at default values.
- 5 Enter the following networks, and select *Add* after each:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0
- 6 For interface, select *Create New* and set the following information.

Interface	port1 (internal)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

- 7 For interface, select *Create New* and set the following information.

Interface	port2 (router1)
Send Version	Both
Receive Version	Both

Authentication	None
Passive Interface	disabled

- 8 For interface, select *Create New* and set the following information.

Interface	port3 (router4)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

- 9 For interface, select *Create New* and set the following information.

Interface	port4 (ISP)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

Configure RIP settings on Router2 and Router3- web-based manager

```

config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port3"
      set receive-version 1 2
      set send-version 1 2
    end
    edit "port4"
      set receive-version 1 2
      set send-version 1 2
    end
  config network
    edit 1
      set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
      set prefix 10.12.0.0 255.255.0.0
    next
    edit 3
      set prefix 10.14.0.0 255.255.0.0
    next
    edit 4

```

```
        set prefix 172.20.120.0 255.255.255.0
    end
    config redistribute "static"
        set status enable
    end
end
```

Configuring other networking devices

In this example there are two groups of other devices on the the network — internal devices, and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers and other network devices. Once the FortiGate units are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise there should be no configuration required.

The second group of devices is the ISP. This consists of the RIP router the FortiGate routers 2 and 3 connect to. You need to contact your ISP and ensure they have your information for your network such as the IP addresses of the connecting RIP routers, what version of RIP your network supports, and what authentication (if any) is used.

Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult [“Troubleshooting RIP” on page 1726](#), and the general [“” on page 1714](#).

RIPng — RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units system information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.


All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.

Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows.

Table 102: Rip example network topology

Network	Router	Interface & Alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

Figure 166: Network topology for the IPV6 RIPng example

Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 4.0 MR3 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 and port2 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices are support IPv6 and are running RIPng.

General configuration steps

This example is very straight forward. The only steps involved are:

- [Configuring the FortiGate units system information](#)
- [Configuring RIPng on FortiGate units](#)
- [Configuring other network devices](#)
- [Testing the configuration](#)

Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.

To configure system information on Router1 - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 For *Host name*, select *Change*.
- 3 Enter "Router1".

- 4 Go to *System > Admin > Settings*.
- 5 In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
- 6 Go to *System > Network > Interface*.
- 7 Edit port1 (internal) interface.
- 8 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	2002:A0B:6565::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

- 9 Edit port2 (ISP) interface.
- 10 Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	2002:AC14:7865::/0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router1 - CLI

```

config system global
    set hostname Router1
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6565::/0
        end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and internet"
        config ipv6
            set ip6-address 2002:AC14:7865::
        end
    end
end

```

To configure system information on Router2 - web-based manager

- 1 Go to *System > Dashboard > Status*.

- 2 For *Host name*, select *Change*.
- 3 Enter "Router2".
- 4 Go to *System > Admin > Settings*.
- 5 In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
- 6 Go to *System > Network > Interface*.
- 7 Edit port1 (internal) interface.
- 8 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	2002:A0B:6566::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

- 9 Edit port2 (ISP) interface.
- 10 Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	2002:AC14:7866::/0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router2 - CLI

```

config system global
    set hostname Router2
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6566::/0
        end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and internet"
        config ipv6
            set ip6-address 2002:AC14:7866::
        end
    end
end

```

Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include — the internal network, and the ISP network. There is no redistribution, and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only confirmation, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

To configure RIPng on Router1 - CLI

```
config router ripng
config interface
edit port1
next
edit port2
end
config neighbor
edit 1
set interface port1
set ipv6 2002:a0b:6566::/0
next
edit 2
set interface port2
set ipv6 2002:AC14:7805::/0
end
```

To configure RIPng on Router2 - CLI

```
config router ripng
config interface
edit port1
next
edit port2
end
config neighbor
edit 1
set interface port1
set ipv6 2002:a0b:6565::/0
next
edit 2
set interface port2
set ipv6 2002:AC14:7805::/0
end
```

Configuring other network devices

The other devices on the internal network all support IPv6, and are running RIPng where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information such as IPv6 addresses.

Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the Troubleshooting chapter.

For troubleshooting problems with RIP, see [“Troubleshooting RIP” on page 1726](#).

Use the following section for testing and troubleshooting RIPng.

Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems.

```
FGT# diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit.

```
FGT# diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table.

```
FGT# get router info6 routing-table
```

View the routing table. This information is almost the same as the previous command (diagnose ipv6 route list) however it is presented in an easier to read format.

```
FGT# get router info6 rip interface external
```

View brief output on the RIP information for the interface listed. The information includes if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
FGT# get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.



Border Gateway Protocol (BGP)

This section describes Border Gateway Protocol (BGP).

The following topics are included in this section:

- [BGP background and concepts](#)
- [Troubleshooting BGP](#)
- [BGP routing examples](#)

BGP background and concepts

The border gateway protocol contains two distinct subsets — internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together, and is the main routing protocol for the Internet backbone. FortiGate units support iBGP, and eBGP only for communities.

The following topics are included in this section:

- [Background](#)
- [Parts and terminology of BGP](#)
- [How BGP works](#)

Background

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by the more recent RFC 4271. The main benefits of BGP-4 are classless inter-domain routing, and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545. BGP-4+ also supports

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. In doing so, BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP are not explained here as they are common to other dynamic routing protocols as well. For more information on parts of BGP that are not listed here, see [“Dynamic routing terminology” on page 1708](#).

The following topics are included in this section:

- [BGP and IPv6](#)
- [Roles of routers in BGP networks](#)
- [Network Layer Reachability Information \(NLRI\)](#)
- [BGP attributes](#)
- [Confederations](#)

BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4, but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as with `config netowrk6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the [FortiGate CLI Reference](#).

IPv6 BGP commands include:

```
config bgp
  set allowas-in6 <max_num_AS_integer>
  set allowas-in-enable6 {enable | disable}
  set attribute-unchanged6 [as-path] [med] [next-hop]
  set capability-default-originate6 {enable | disable}
  set capability-graceful-restart6 {enable | disable}
  set capability-orf6 {both | none | receive | send}
  set default-originate-route-map6 <routemap_str>
  set distribute-list-in6 <access-list-name_str>
  set distribute-list-out6 <access-list-name_str>
  set filter-list-in6 <aspath-list-name_str>
  set filter-list-out6 <aspath-list-name_str>
  set maximum-prefix6 <prefix_integer>
  set maximum-prefix-threshold6 <percentage_integer>
  set maximum-prefix-warning-only6 {enable | disable}
  set next-hop-self6 {enable | disable}
  set prefix-list-in6 <prefix-list-name_str>
  set prefix-list-out6 <prefix-list-name_str>
  set remove-private-as6 {enable | disable}
  set route-map-in6 <routemap-name_str>
  set route-map-out6 <routemap-name_str>
  set route-reflector-client6 {enable | disable}
  set route-server-client6 {enable | disable}
  set send-community6 {both | disable | extended | standard}
  set soft-reconfiguration6 {enable | disable}
  set unsuppress-map6 <route-map-name_str>
config network6
config redistribute6
end
```

Roles of routers in BGP networks

Dynamic routing has a number of different roles routers can fill such as those covered in [“Dynamic routing terminology” on page 1708](#). BGP has a number of custom roles that routers can fill. These include:

- [Speaker routers](#)
- [Peer routers or neighbors](#)
- [Route reflectors \(RR\)](#)

Speaker routers

Any router configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that are not speaker routers, are not treated as BGP routers.

Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to your FortiGate unit. Your FortiGate unit learns about all other routers through these peers.

You need to manually configure BGP peers on your FortiGate unit as neighbors. Otherwise these routers will not be seen as peers, but instead as simply other routers on the network that don't support BGP. You can optionally use MD5 authentication to password protect BGP sessions with those neighbors. (see RFC 2385).

You can configure up to 1000 BGP neighbors on your FortiGate unit. You can clear all or some BGP neighbor connections (sessions) using the `exec router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
FGT# exec router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:


```
FGT# exec router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
FGT# exec router clear bgp dampening 10.10.0.0/16
```

In Figure 1, Router A is directly connected to five other routers in a network that contains 12 routers overall. These routers, the ones in the blue circle, are Router A's peers or neighbors.

Figure 167: Router A and its 5 peer routers



Router A's peer routers

As a minimum, when configuring BGP neighbors you must enter their IP address, and the AS number (remote-as). This is all the information the web-based manager interface allows you to enter for a neighbor.

The BGP commands related to neighbors are quite extensive and include:

```
config router bgp
  config neighbor
    edit <neighbor_address_ipv4>
      set activate {enable | disable}
      set advertisement-interval <seconds_integer>
      set allowas-in <max_num_AS_integer>
      set allowas-in-enable {enable | disable}
      set attribute-unchanged [as-path] [med] [next-hop]
      set bfd {enable | disable}
      set capability-default-originate {enable | disable}
      set capability-dynamic {enable | disable}
      set capability-graceful-restart {enable | disable}
      set capability-orf {both | none | receive | send}
      set capability-route-refresh {enable | disable}
      set connect-timer <seconds_integer>
      set description <text_str>
      set distribute-list-in <access-list-name_str>
      set distribute-list-out <access-list-name_str>
      set dont-capability-negotiate {enable | disable}
      set ebgp-enforce-multihop {enable | disable}
      set ebgp-multihop {enable | disable}
      set ebgp-multihop-ttl <seconds_integer>
      set filter-list-in <aspath-list-name_str>
      set filter-list-out <aspath-list-name_str>
      set holdtime-timer <seconds_integer>
      set interface <interface-name_str>
      set keep-alive-timer <seconds_integer>
      set maximum-prefix <prefix_integer>
      set maximum-prefix-threshold <percentage_integer>
      set maximum-prefix-warning-only {enable | disable}
      set next-hop-self {enable | disable}
      set override-capability {enable | disable}
      set passive {enable | disable}
      set password <string>
      set prefix-list-in <prefix-list-name_str>
      set prefix-list-out <prefix-list-name_str>
      set remote-as <id_integer>
      set remove-private-as {enable | disable}
      set retain-stale-time <seconds_integer>
      set route-map-in <routemap-name_str>
      set route-map-out <routemap-name_str>
      set route-reflector-client {enable | disable}
      set route-server-client {enable | disable}
      set send-community {both | disable | extended | standard}
      set shutdown {enable | disable}
      set soft-reconfiguration {enable | disable}
      set strict-capability-match {enable | disable}
      set unsuppress-map <route-map-name_str>
      set update-source <interface-name_str>
```



```

        set weight <weight_integer>
    end
end
end

```

Route reflectors (RR)

Route reflectors in BGP concentrate route updates so other routers need only talk to the route reflectors to get all the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. BGP route reflectors are defined in RFC 1966.


In a BGP route reflector configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other route reflectors and border routers. Only the reflectors need to be configured, not the clients — the clients will find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. FortiGate units can be configured as either reflectors or clients.

Since route reflectors are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running BGP typically don't require route reflectors (RR). However, RR is a useful feature for large companies, where their AS may include 100 routers or more. For example, for a full mesh 20 router configuration within an AS there would have to be 190 unique BGP sessions — just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. From these numbers, it's plain that updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how route reflectors can improve the situation when only six routers are involved. The AS without route reflectors requires 15 sessions between the routers. In the AS with route reflectors, the two route reflectors receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster as well as other route reflectors and pass them on to the border router. The RR configuration only requires six sessions. This example shows a reduction of 60% in the number of required sessions.

Figure 168: Required sessions within an AS with and without route reflectors



The BGP commands related to route reflectors includes:

```
config router bgp
  config neighbor
    set route-reflector-client {enable | disable}
    set route-server-client {enable | disable}
  end
end
```

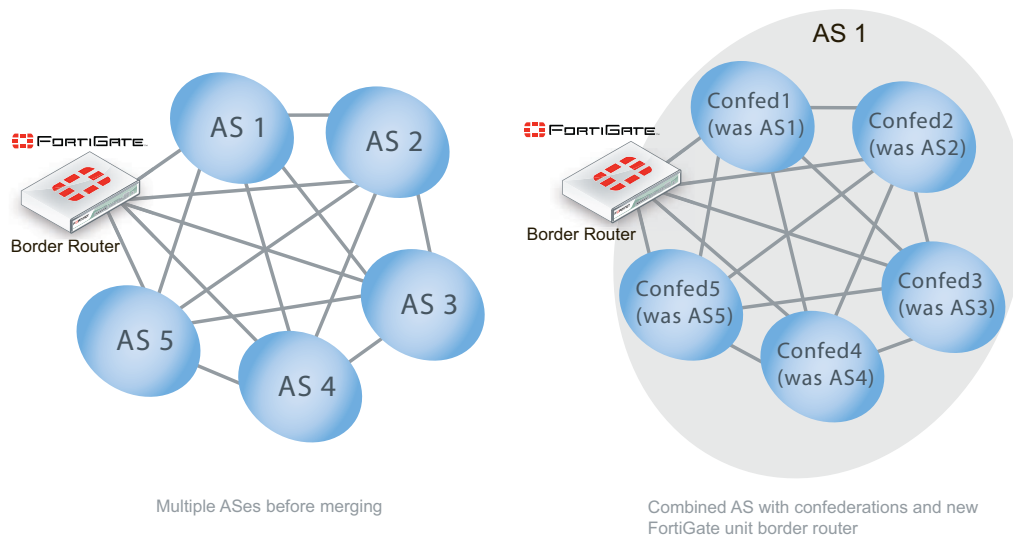
Confederations

Confederations were introduced to reduce the number of BGP advertisements on a segment of the network, and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units. Confederations are defined in RFC 3065 and RFC 1965.

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications in that many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging ASs. Each AS being merged can easily become a confederation, requiring few changes. Any additional permanent changes can then be implemented over time as required. The figure below shows the group of ASs before merging, and the corresponding confederations afterward as part of the single AS with the addition of a new border router. It should be noted that after merging if the border router becomes a route reflector, then each confederation only needs to communicate with one other router, instead of five others.

Figure 169: AS merging using confederations



Confederations and route reflectors perform similar functions — they both sub-divide large ASes for more efficient operation. They differ in that route reflector clusters can include routers that are not members of a cluster, where routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS_PATH attribute making it easier to trace.

It is important to note that while confederations essentially create sub-ASs, all the confederations within an AS appear as a single AS to external ASs.

Confederation related BGP commands include:

```
config router bgp
  set confederation-identifier <peerid_integer>
end
```

Network Layer Reachability Information (NLRI)

Network Layer Reachability Information (NLRI) is unique to BGP-4. It is sent as part of the update messages sent between BGP routers, and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that when combined are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route, and are modified as required along the route.

BGP can work well with mostly default settings, but if you are going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include:

AS_PATH	A list of ASes a route has passed through. See “AS_PATH” on page 1757 .
MULTI_EXIT_DESC (MED)	Which router to use to exit an AS with more than one external connection. See “MULTI_EXIT_DESC (MED)” on page 1758 .
COMMUNITY	Used to apply attributes to a group of routes. See “COMMUNITY” on page 1758 .
NEXT_HOP	Where the IP packets should be forwarded to, like a gateway in static routing. See “NEXT_HOP” on page 1759 .
ATOMIC_AGGREGATE	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. See “ATOMIC_AGGREGATE” on page 1759 .
ORIGIN	Used to determine if the route is from the local AS or not. See “ORIGIN” on page 1759 .
LOCAL_PREF	Used only within an AS to select the best route to a location (like MED)

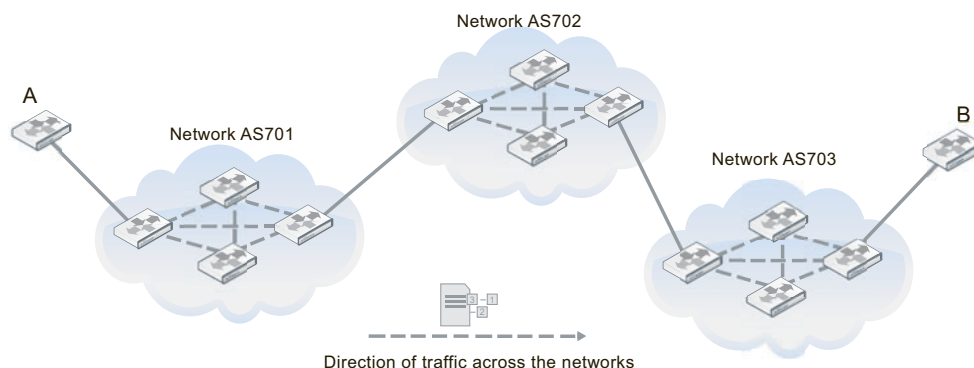


Inbound policies on FortiGate units can change the NEXT-HOP, LOCAL-PREF, MED and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the unit cannot affect these attributes.

AS_PATH

AS_PATH is the BGP attribute that keeps track of each AS a route advertisement has passed through. AS_PATH is used by confederations and by exterior BGP (EBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS_PATH with that routers AS in it. The figure below shows the route between router A and router B. The AS_PATH from A to B would read 701,702,703 for each AS the route passes through.

As of the start of 2010, the industry is upgrading from 2-byte to 4-byte AS_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS_PATH numbers.

Figure 170: AS_PATH of 701,702, 703 between routers A and B

The BGP commands related to AS_PATH include:

```
config router bgp
  set bestpath-as-path-ignore {enable | disable}
end
```

MULTI_EXIT_DESC (MED)

BGP AS systems can have one or more routers that connect them to other ASes. For ASes with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes such as delay. It is a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When the FortiGate unit receives a BGP update, the FortiGate unit examines the Multi-Exit Discriminator (MED) attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiGate unit routing table.

FortiGate units have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information which can be suspicious — possibly a hacking attempt or an attack on the network. At best it is an unreliable route to select.

The BGP commands related to MED include:

```
config router bgp
  set always-compare-med {enable | disable}
  set bestpath-med-confed {enable | disable}
  set bestpath-med-missing-as-worst {enable | disable}
  set deterministic-med {enable | disable}
config neighbor
  set attribute-unchanged [as-path] [med] [next-hop]
end
end
```

COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

The FortiGate unit can set the COMMUNITY attribute of a route to assign the route to predefined paths (see RFC 1997). The FortiGate unit can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include:

```
config router bgp
  set send-community {both | disable | extended | standard}
end
```

NEXT_HOP

The NEXT_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT_HOP attribute is much like a gateway in static routing.

FortiGate units allow you to change the advertising of the FortiGate unit's IP address (instead of the neighbor's IP address) in the NEXT_HOP information that is sent to IBGP peers. This is changed with the config neighbor, set next-hop-self command.

The BGP commands related to NEXT_HOP include:

```
config router bgp
  config neighbor
    set attribute-unchanged [as-path] [med] [next-hop]
    set next-hop-self {enable | disable}
  end
end
```

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that is easier to send in updates. When it reaches its destination, the summarized routes are split back up into the individual routes.

Your FortiGate unit doesn't specifically set this attribute in the BGP router command, but it is used in the route map command.

The commands related to ATOMIC_AGGREGATE include:

```
config router route-map
  edit <route_map_name>
    config rule
      edit <route_map_rule_id>
        set set-aggregator-as <id_integer>
        set set-aggregator-ip <address_ipv4>
        set set-atomic-aggregate {enable | disable}
      end
    end
  end
```

ORIGIN

The ORIGIN attribute records where the route came from. The options can be IBGP, EBGP, or incomplete. This information is important because internal routes (IBGP) are higher priority than external routes (EBGP). However incomplete ORIGINS are the lowest priority of the three.

The commands related to ORIGIN include:

```
config router route-map
  edit <route_map_name>
    set comments <string>
```

```
config rule
edit <route_map_rule_id>
    set match-origin {egp | igp | incomplete | none}
end
end
end
```

How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other, and establish a connection they go from the idle state, through the various states until they reach the established state. An error can cause the connection to be dropped and the state of the router to be reset to either active or idle. These errors can be caused by: TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used such as multiprotocol extensions that can include IPv6 and VPNs.

IBGP versus EBGP

When you read about BGP, often you see EBGP or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASes) where interior BGP (IBGP) involves packets that stay within a single AS. For example the AS_PATH attribute is only useful for EBGP where routes pass through multiple ASes.

These two modes are important because some features of BGP are only used for one of EBGP or IBGP. For example confederations are used in EBGP, and route reflectors are only used in IBGP. Also routes learned from IBGP have priority over EBGP learned routes.

FortiGate units have some commands specific to EBGP. These include:

- automatically resetting the session information to external peers if the connection goes down — `set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (must also configure local and internal distances if this is set) — `set distance-external <distance_integer>`
- enforcing EBGP multihops and their TTL (number of hops) — `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

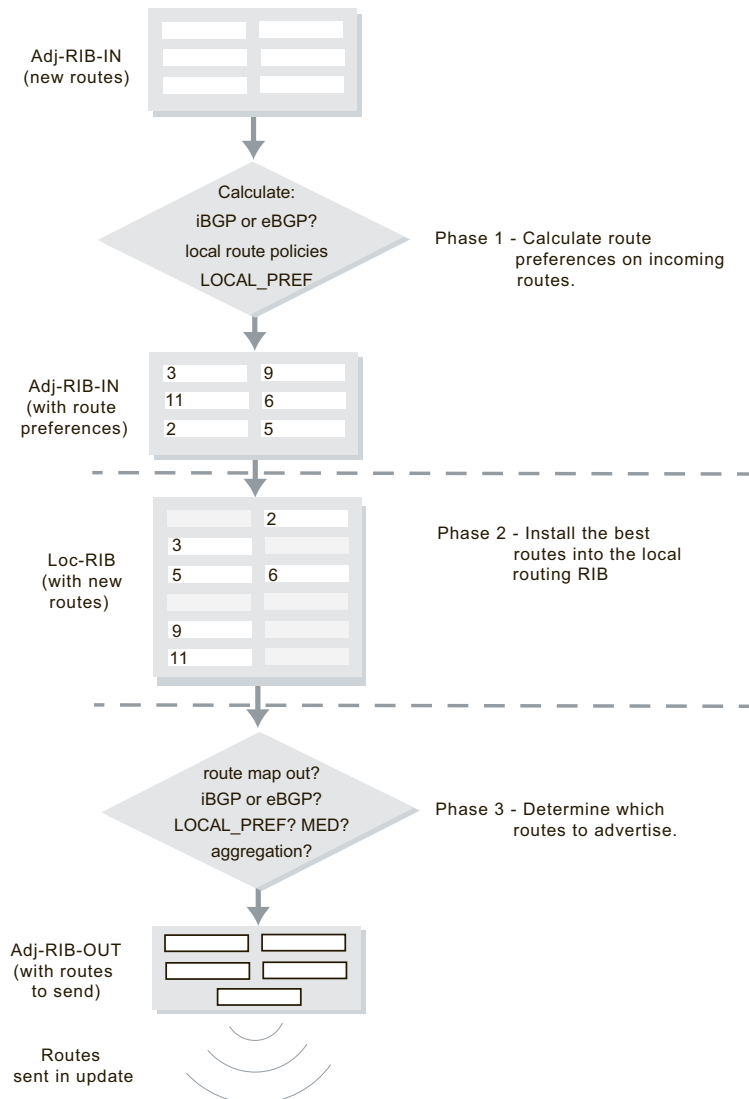
BGP path determination — which route to use

All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination do not change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute to enable an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiGate unit receives BGP updates, or when the FortiGate unit sends out BGP updates.

Figure 171: Three phases of BGP routing decision



Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (IBGP), policy information or LOCAL_PREF is used. For external peer learned routes, it is based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the master routing table. Each route from Phase 1 has their NEXT_HOP checked to ensure the destination is reachable. If it is reachable, the AS_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there is only one route to a location, it is installed.
- If multiple routes to the same location, use the most preferred route from Level 1.
- If there is a tie, break the tie based on the following in descending order of importance: shortest AS_PATH, smallest ORIGIN number, smallest MED, EBGp over IBGP, smallest metric or cost for reaching the NEXT_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed the Loc-RIB will consist of the best of both the new and older routes.

Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there is any route aggregation or summarizing, it happens here. Also any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

Aggregate routes and addresses

BGP4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing enables the configuration of aggregate routes by stating the address bits the aggregated addresses have in common. For more information, see [“Aggregated routes and addresses” on page 1709](#).

In BGP there is an ATOMIC_AGGREGATE attribute that when set informs routers that the route has been aggregated, and should not be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are:

```
config router bgp
  config aggregate-address
    edit <aggr_addr_id>
    set as-set {enable | disable}
    set prefix <address_ipv4mask>
    set summary-only {enable | disable}
  end
  config aggregate-address6
    edit <aggr_addr_id>
    set as-set {enable | disable}
    set prefix6 <address_ipv6mask>
    set summary-only {enable | disable}
  end
end
```


Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically the problems with a BGP network that has been configured, involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

This section includes:

- [Clearing routing table entries](#)
- [Route flap](#)

Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `exec router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
FGT# exec router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:

```
FGT# exec router clear bgp as 650001
```

Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables which creates a lot of administration traffic on the network. And the same traffic happens again when that router comes back online. If the problem is something like a faulty network cable that wobbles on and offline every 10 seconds, there could easily be overwhelming amounts of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiGate units in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline resulting in route flap. While this doesn't occur often, or more than once at a time, it can still result in an interruption in traffic which is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they do not expire during the failover process. Also configuring graceful restart on the HA cluster will help with a smooth failover.

The first method of dealing with router flap should be to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either don't use that interface or swap in a good router. If the power source is bad on a router either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include:

- [Holddown timer](#)
- [Dampening](#)
- [Graceful restart](#)
- [Bi-directional forwarding detection \(BFD\)](#)

Holddown timer

The first line of defence to a flapping route is the hold down timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

Once activated, the holddown timer won't allow the FortiGate unit to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage will be recognized by the FortiGate unit — for the duration of the other outages there will be no changes because the Fortigate unit is essentially treating this router as down. After the timer expires, if the route is still flapping it will happen all over again.

Even if the route isn't flapping — if it goes down, comes up, and stays back up — the timer still counts down and the route is ignored for the duration of the timer. In this situation the route will be seen as down longer than it really is, but there will be only the one set of route updates. This is not a problem in normal operation because updates are not frequent.

Also the potential for a route to be treated as down when it is really up can be viewed as a robustness feature. Typically you do not want most of your traffic being routed over an unreliable route. So if there is route flap going on, it is best to avoid that route if you can. This is enforced by the holddown timer.

How to configure the holddown timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the holddown timer.

For example, your network has two routes that you want to set the holddown timer for. One is your main route (to 10.12.101.4) that all your Internet traffic goes through, and it can't be down for long if its down. The second is a low speed connection to a custom network that is used infrequently (to 10.13.101.4). The holddown timer for the main route should be fairly short, lets say 60 seconds instead of the default 180 seconds. The second route timer can be left at the default or even longer since it is rarely used. In your BGP configuration this looks like:

```
config router bgp
  config neighbor
    edit 10.12.101.4
      set holddown-timer 60
    next
    edit 10.13.101.4
      set holddown-timer 180
    next
  end
end
```

Dampening

Dampening is a method used to limit the amount of network problems due to flapping routes. With dampening the flapping still occurs, but the peer routers pay less and less attention to that route as it flaps more often. One flap doesn't start dampening, but the second starts a timer where the router will not use that route — it is considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There is a period of time called the reachability half-life after which a route flap will only be suppressed for half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiGate units cache by using one of the `execute router clear bgp` commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
```

or

```
execute router clear bgp flap-statistics {<ip_address> |  
<ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
FGT# exec router clear bgp dampening 10.10.0.0/16
```

The BGP commands related to route dampening are:

```
config router bgp  
  set dampening {enable | disable}  
  set dampening-max-suppress-time <minutes_integer>  
  set dampening-reachability-half-life <minutes_integer>  
  set dampening-reuse <reuse_integer>  
  set dampening-route-map <routemap-name_str>  
  set dampening-suppress <limit_integer>  
  set dampening-unreachability-half-life <minutes_integer>end  
end
```

Graceful restart

BGP4 has the capability to gracefully restart.

In some situations, route flap is caused by routers that appear to be offline but the hardware portion of the router (control plane) can continue to function normally. One example of this is when some software is restarting or being upgraded, but the hardware can still function normally.

Graceful restart is best used for these situations where routing will not be interrupted, but the router is unresponsive to routing update advertisements. Graceful restart does not have to be supported by all routers in a network, but the network will benefit when more routers support it.



FortiGate HA clusters can benefit from graceful restart. When a failover takes place, the HA cluster will advertise it is going offline, and will not appear as a route flap. It will also enable the new HA main unit to come online with an updated and usable routing table — if there is a flap the HA cluster routing table will be out of date.

Scheduled time offline

Graceful restart is a means for a router to advertise it is going to have a scheduled shutdown for a very short period of time. When neighboring routers receive this notice, they will not remove that router from their routing table until after a set time elapses. During that time if the router comes back online, everything continues to function as normal. If that router remains offline longer than expected, then the neighboring routers will update their routing tables as they assume that router will be offline for a long time.

FortiGate units support both graceful restart of their own BGP routing software, and also neighboring BGP routers.

For example, if a neighbor of your FortiGate unit, with an IP address of 172.20.120.120, supports graceful restart, enter the command:

```
config router bgp
  config neighbor
    edit 172.20.120.120
      set capability-graceful-restart enable
    end
  end
```

If you want to configure graceful restart on your FortiGate unit where you expect the FortiGate unit to be offline for no more than 2 minutes, and after 3 minutes the BGP network should consider the FortiGate unit offline, enter the command:

```
config router bgp
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
end
```



You can configure graceful restarting and other advanced settings only through CLI commands. For more information on advanced BGP settings, see the “router” chapter of the [FortiGate CLI Reference](#).

The BGP commands related to BGP graceful restart are:

```
config router bgp
  set graceful-restart { disable| enable}
  set graceful-restart-time <seconds_integer>
  set graceful-stalepath-time <seconds_integer>
  config neighbor
    set capability-graceful-restart {enable | disable}
  end
end

execute router restart
```

Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

While BGP can detect route failures, BFD can be configured to detect these failures more quickly allowing faster responses and improved convergence. This can be balanced with the bandwidth BFD uses in its frequent route checking.

Configurable granularity

BFD can run on the entire FortiGate unit, selected interfaces, or on BGP for all configured interfaces. The hierarchy allows each lower level to override the upper level's BFD setting. For example, if BFD was enabled for the FortiGate unit, it could be disabled only for a single interface or for BGP. For information about FortiGate-wide BFD options, see config system settings in the [FortiGate CLI Reference](#).

BFD support was added in FortiOS v3.0 MR4, and can only be configured through the CLI.

The BGP commands related to BFD are:

```
config router bgp
  config neighbor
    edit <neighbor_address_ipv4>
      set bfd {enable | disable}
    end
  end

execute router clear bfd session <src_ipv4> <dst_ipv4>
<interface>
```

BGP routing examples

BGP is a complex dynamic routing protocol. There are many BGP configurations and features that can benefit from in-depth examples.

This section includes:

- [Dual-homed BGP example](#)
- [Redistributing and blocking routes in BGP](#)

Dual-homed BGP example

This is an example of a small network that uses BGP routing connections to two ISPs. This is a common configuration for companies that need redundant connections to the Internet for their business.

This configuration is for a small company connected to two ISPs. The company has one main office, the Head Office, and uses static routing for internal routing on that network.

Both ISPs use BGP routing, and connect to the Internet directly. They want the company to connect to the ISP networks using BGP. They also use graceful restart to prevent unneeded updates, and use smaller timer values to detect network failures faster.

As can be expected, the company wants to keep their BGP configuration relatively simple and easy to manage. The current configuration has only 3 routers to worry about — the 2 ISP border routers, and the FortiGate unit. This means the FortiGate unit will only have two neighbour routers to configure.

This configuration has the added benefit of being easy to expand if the Company wants to add a remote office in the future.

To keep the configuration simple, the Company is allowing only HTTP, HTTPS, FTP, and DNS traffic out of the local network. This will allow employees access to the Internet and their web-mail.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate unit](#)
- [Configuring other networking devices](#)
- [Testing this configuration](#)

Why dual home?

Dual homing means having two separate independent connections to the Internet. Servers in this configuration have also been called bastion hosts and can include DNS servers which require multiple connections.

Benefits of dual homing can include:

- redundant Internet connection that essentially never fails
- faster connections through one ISP or the other for some destinations, such as other clients of those ISPs
- load balancing traffic to your Company network
- easier to enable more traffic through two connections than upgrading one connection to bigger bandwidth
- easier to create protection policies for different traffic through a specific ISP

Some companies require reliable internet access at all times as part of their business. Consider a doctor operating remotely who has their Internet connection fail — the consequences could easily be life or death.

Dual homing is extra expense for the second ISP connection, and more work to configure and maintain the more complex network topology.

Potential dual homing issues

BGP comes with load balancing issues, and dual homing is the same category. BGP does not inherently deal well with load balancing, or getting default routes through BGP. Ideally one connect may be best for certain destinations, but it may not have that traffic routed to it making the load balancing less than perfect. This kind of fine tuning can be very time consuming, and usually results in a best effort situation.

When dual coming is not configured properly, your network may become a link between your ISPs and result in very high traffic between the ISPs that does not originate from your network. The problems with this situation are that your traffic may not have the bandwidth it needs, and you will be paying for a large volume of traffic that is not yours. This problem can be solved by not broadcasting or redistributing BGP routes between the ISPs.

If you learn your default routes from the ISPs in this example, you may run into an asymmetric routing problem where your traffic loops out one ISP and back to you through the other ISP. If you think this may be happening you can turn on asymmetric routing on the FortiGate unit (config system settings, set asymmetric enable) to verify that really is the problem. Turn this feature off once this is established since it disables many features on the FortiGate by disabling stateful inspection. Solutions for this problem can include using static routes for default routes instead of learning them through BGP, or configuring VDOMs on your FortiGate unit to provide a slightly different path back that is not a true loop.

Network layout and assumptions

This section includes:


- [Network layout](#)
- [Assumptions](#)

Network layout

The network layout for the basic BGP example involves the company network being connected to both ISPs as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, ISP1's AS, and ISP2's AS.

The components of the layout include:

- The Company AS (AS number 1) is connected to ISP1 and ISP2 through the FortiGate unit.
- The Company has one internal network — the Head Office network at 10.11.101.0/24.
- The FortiGate unit internal interface is on the the Company internal network with an IP address of 10.11.101.110.
- The FortiGate unit external1 interface is connected to ISP1's network with an IP address of 172.21.111.5, an address supplied by the ISP.
- ISP1 AS has an AS number of 6501, and ISP2 has an AS number of 6502
- Both ISPs are connected to the Internet.
- The ISP1 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.21.111.4.
- The ISP2 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.22.222.4.
- Apart from graceful restart, and shorter timers (holdtimer, and keepalive) default settings are to be used whenever possible.

Figure 172: Basic BGP network topology

Assumptions

The basic BGP configuration procedure follows these assumptions:

- ISP1 is the preferred route, and ISP2 is the secondary route
- all basic configuration can be completed in both GUI and CLI
- only one AS is used for the Company

For these reasons this example configuration does not include:

- [Bi-directional forwarding detection \(BFD\)](#)
- [Route maps](#)
- [Access lists](#)
- changing redistribution defaults — make link when example is set up
- IPv6

For more information on these features, see the corresponding section.

General configuration steps

In this basic example, only two routers need to be configured — the FortiGate unit, and the ISP BGP router. After they are configured, the network configuration should be tested to ensure its working as expected.

To configure a simple BGP network

- 1 [Configuring the FortiGate unit](#)
- 2 [Configuring other networking devices](#)
- 3 [Testing this configuration](#)

Configuring the FortiGate unit

In this topology, the FortiGate unit is the link between the Company Network and the ISP network. The FortiGate unit is the only BGP router on the Company Network, but there is at least one other BGP router on the ISP Network — there may be more but we don't have that information.

As mentioned in the general configuration steps, the ISP must be notified of the Company's BGP router configuration when complete as it will need to add the FortiGate BGP router as a neighbor router on its domain. This step is required for the FortiGate unit to receive BGP routing updates from the ISP network and outside networks.

If the ISP has any special BGP features enabled such as graceful restart, or route dampening that should be determined up front so those features can be enabled on the FortiGate unit.

To configure the FortiGate unit as a BGP router

- 1 [Configure interfaces and default routes](#)
- 2 [Configure firewall services, addresses, and policies](#)
- 3 [Set the FortiGate BGP information](#)
- 4 [Add the internal network to the AS](#)
- 5 [Additional FortiGate BGP configuration](#)

Configure interfaces and default routes

The FortiGate unit is connected to three networks — Company Network on the internal interface, ISP1 Network on external1 interface, and ISP2 on external2 interface.

This example uses basic interface settings. Check with your ISP to determine if additional settings are required such as setting the maximum MTU size, or if gateway detection is supported.

High end FortiGate units do not have interfaces labeled Internal, or External. Instead, for clarity's sake, we are using the alias feature to name interfaces for these roles.

Default routes to both external interfaces are configured here as well. Both are needed in case one goes offline. ISP1 is the primary connection and has a smaller administrative distance so it will be preferred over ISP2. Both distances are set low so they will be preferred over any learned routes.

To configure the FortiGate interfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Edit port 1 (internal) interface.
- 3 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Company internal network
Administrative Status	Up

- 4 Edit port 2 (external1) interface.

- 5 Set the following information, and select **OK**.

Alias	external1
IP/Netmask	172.21.111.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP1 External BGP network
Administrative Status	Up

- 6 Edit port 3 (external2) interface.

- 7 Set the following information, and select **OK**.

Alias	external2
IP/Netmask	172.22.222.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP2 External BGP network
Administrative Status	Up

To configure the FortiGate interfaces - CLI

```
config system interface
  edit port1
    set alias internal
    set ip 10.11.101.110 255.255.255.0
    set allowaccess http https ssh
    set description "Company internal network"
    set status up
  next
  edit port2
    set alias external1
    set ip 172.21.111.5 255.255.255.0
    set allowaccess https ssh
    set description "ISP1 External BGP network"
    set status up
  next
  edit port3
    set alias external2
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "ISP2 External BGP network"
    set status up
  next
end
```

To configure default routes for both ISPs - web-based manager

- 1 Go to *Router > Static > Static Route*.
- 2 Delete any existing routes with a IP/Mask of address of 0.0.0.0/0.0.0.0
- 3 Select *Create New*, and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2

Gateway	172.21.111.5
Distance	10

- 4 Select OK.
- 5 Select *Create New*, and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3
Gateway	172.22.222.5
Distance	15

- 6 Select OK.

To configure default routes for both ISPs - CLI

```
config router static
  edit 1
    set device "port2"
    set distance 10
    set gateway 172.21.111.5
  next
  edit 2
    set device "port3"
    set distance 15
    set gateway 172.22.222.5
  next
end
```

Configure firewall services, addresses, and policies

To create the security policies, first you must create the firewall services group that will include all the services that will be allowed, then you must define the addresses that will be used in the security policies, and lastly you configure the security policies themselves.

To keep the configuration simple, the Company is allowing only HTTP traffic out of the local network. This will allow employees access to the Internet and their web-mail. DNS services will also be allowed through the firewall.

The security policies will allow HTTP traffic (port 80 and port 8080), HTTPS traffic (port 443), FTP traffic (port 21), and DNS traffic (port 53 and port 953) in both directions. Also BGP (port 179) may need access through the firewall.



For added security, you may want to define a smaller range of addresses for the internal network. For example if only 20 addresses are used, only allow those addresses in the range.

In the interest of keeping things simple, a zone will be used to group the two ISP interfaces together. This will allow using one security policy to apply to both ISPs at the same time. Remember to block intra-zone traffic as this will help prevent one ISP sending traffic to the other ISP through your FortiGate unit using your bandwidth. The zone keeps configuration simple, and in the future if there is a need for separate policies for each ISP, they can be created and the zone can be deleted.

The addresses that will be used are the addresses of the FortiGate unit internal and external ports, and the internal network.

More policies or services can be added in the future as applications are added to the network. For more information on security policies, see the firewall chapter of the [FortiGate Administration Guide](#).



When configuring security policies always enable logging to help you track and debug your traffic flow.

To create a firewall services group - web-based manager

- 1 Go to *Firewall Objects > Service > Group*, and select *Create New*.
- 2 For *Group Name*, enter "Basic_Services".
- 3 From *Available Services*, move the following six services over to the *Member* list — BGP, FTP, FTP_GET, FTP_PUT, DNS, HTTP, and HTTPS.
- 4 Select *OK*.

To create a firewall services group - CLI

```
config firewall service group
  edit "Basic_Services"
    set member "BGP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP"
    "HTTPS"
  next
end
```

To create a zone for the ISP interfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select the caret to the right of *Create New* and then select *Zone*.
- 3 Enter the following information.

Zone Name	ISPs
Block Intra-zone traffic	enable
interface members	port2 port3

- 4 Select *OK*.

To create a zone for the ISP interfaces - CLI

```
config system zone
  edit "ISPs"
    set interface "dmz1" "dmz2"
    set intrazone block
  next
end
```

To add the firewall addresses - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, and set the following information.

Address Name	Internal_network
Type	Subnet / IP Range

Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	port1

- 3 Select OK.

To add the firewall addresses - CLI

```
config firewall address
  edit "Internal_network"
    set associated-interface "port1"
    set subnet 10.11.101.0 255.255.255.0
  next
end
```

To add the HTTP and DNS security policies - web-based manager

- 1 Go to *Policy > Policy > Policy*, and select *Create New*.
- 2 Set the following information.

Source Interface/Zone	port1(internal)
Source Address	Internal_network
Destination Interface/Zone	ISPs
Destination Address	All
Schedule	always
Service	Basic_services
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	Enable
Comments	ISP1 basic services out policy

- 3 Select OK.
- 4 Select *Create New*, and set the following information.

Source Interface/Zone	ISPs
Source Address	all
Destination Interface/Zone	port1(internal)
Destination Address	Internal_network
Schedule	always
Service	Basic_services
Action	ACCEPT
Log Allowed Traffic	enable
NAT	Enable
Comments	ISP1 basic services in policy

To add the security policies - CLI

```
config firewall policy
  edit 1
    set srcintf "port1"
    set srcaddr "Internal_network"
    set dstintf "ISPs"
    set dstaddr "all"
    set schedule "always"
    set service "Basic_services"
    set action accept
    set nat enable
    set profile-status enable
    set logtraffic enable
    set comments "ISP1 basic services out policy"
  next
  edit 2
    set srcintf "ISPs"
    set srcaddr "all"
    set dstintf "port1"
    set dstaddr "Internal_network"
    set schedule "always"
    set service "Basic_services"
    set action accept
    set nat enable
    set profile-status enable
    set logtraffic enable
    set comments "ISP1 basic services in policy"
  next
end
```

Set the FortiGate BGP information

When using the default information, there are only two fields to set to configure the FortiGate unit as a BGP router.

For this configuration the FortiGate unit will be in a stub area with one route out — the ISP BGP router. Until you configure the ISP router as a neighbour, even that route out is not available. So while after this part of the configuration is complete your FortiGate unit will be running BGP, it won't know about any other routers running BGP until the next part of the configuration is complete.

To set the BGP router information - web-based manager

- 1 Go to *Router > Dynamic > BGP*.
- 2 Set the following information, and select OK.

Local AS	1
Router ID	10.11.101.110

To set the BGP router information - CLI

```
config router BGP
  set as 1
  set router-id 10.11.101.110
end
```

Add the internal network to the AS

The Company is one AS with the FortiGate unit configured as the BGP border router connecting that AS to the two ISPs ASes. The internal network in the Company's AS must be defined. If there were other networks in the company such as regional offices, they would be added here as well.

To set the networks in the AS - web-based manager

- 1 Go to *Router > Dynamic > BGP*.
- 2 In *Networks*, set the following information and select *OK*.

IP/Netmask	10.11.101.0/255.255.255.0
------------	---------------------------

To set the networks in the AS (CLI)

```
config router bgp
  config network
  edit 1
    set prefix 10.11.101.0 255.255.255.0
  next
end
end
```

Additional FortiGate BGP configuration

At this point that is all the settings that can be done in both the web-based manager and the CLI. The remaining configuration must be completed in the CLI.

These additional settings are mainly determined by your ISP requirements. They will determine your timers such as keep alive timers, if extended features like BFD and graceful restart are being used, and so on. For this example, some common simple features are being used to promote faster detections of network failures which will result in better service for the Company's internal network users.

The ISPs do not require authentication between peer routers.

These commands will enable or modify the following features on the FortiGate unit, and where possible on neighboring routers as well:

- `bestpath-med-missing-as-worst` — treats a route without an MED as the worst possible available route due to expected unreliability
- `fast-external-failover` — immediately reset the session information associated with BGP external peers if the link used to reach them goes down
- `graceful-restart*` — advertise reboots to neighbors so they do not see the router as offline, wait before declaring them offline, and how long to wait when they reboot before advertising updates. These commands apply to neighbors and are part of the BGP capabilities. This prevents unneeded routing updates.
- `holdtime-timer` — how long the router will wait for a keepalive message before declaring a router offline. A shorter time will find an offline router faster.
- `keepalive-timer` — how often the router sends out keepalive messages to neighbor routers to maintain those sessions.
- `log-neighbor-changes` — log changes to neighbor routers' status. This can be useful for troubleshooting from both internal and external networks.
- `connect-timer` — how long in seconds the FortiGate unit will try to reach this neighbor before declaring it offline.

- **weight** — used to prefer routes from one neighbor over the other. In this example ISP1 is the primary connection so it is weighted higher than ISP2

To configure additional BGP options - CLI

```
config router bgp
  set bestpath-med-missing-as-worst enable
  set fast-external-failover enable
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
  set graceful-update-delay 180
  set holdtime-timer 120
  set keepalive-timer 45
  set log-neighbor-changes enable
config neighbor
  edit 172.21.111.4
    set connect-timer 60
    set description "ISP1"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 250
  next
  edit 172.22.222.4
    set connect-timer 60
    set description "ISP2"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 100
  next
end
end
```

Configuring other networking devices

There are two other networking devices that need to be configured both ISPs' BGP routers.

The ISPs' routers must add the FortiGate unit as a neighbor so route updates can be sent in both directions. Note that ISP1 is not directly connected to ISP2 that we are aware of.

Inform both of your ISPs of your FortiGate unit's BGP information. Once they have configured their router, you can test your BGP connection to the internet.

They will require your FortiGate unit's:

- IP address of the connected interface
- Router ID
- your Company's AS number

Testing this configuration

With the dual-homed BGP configuration in place, you should be able to send and receive traffic, send and receive routes, and not have any routing loops. Testing the networks will confirm things are working as expected.

In general for routing you need to look at the routing table on different routers to see what routes are being installed. You also need to sniff packets to see how traffic is being routed in real time. These two sources of information will normally tell you what you need to know.

Basic networking tools and methods can be found in “” on page 1714.

Testing of this example’s network configuration should be completed in two parts:

- [Testing network connectivity](#)
- [Verifying the FortiGate unit’s routing tables](#)
- [Verifying traffic routing](#)
- [Verifying the dual-homed side of the configuration](#)

Testing network connectivity

A common first step in testing a new network topology is to test if you can reach the internet and other locations as you expect you should. If not, you may be prevented by cabling issues, software or other issues.

The easiest way to test connections is to use ping, once you ensure that all the FortiGate unit’s interfaces and ISP routers have ping support enabled. Also ensure that the security policies allow ping through the firewall.

Connections to test in this example are the internal network to ISP1’s router or the internet, and the same for ISP2. If you can connect on the external side of the Fortinet unit, try to ping the internal network. Those three tests should prove your basic network connections are working.



Once you have completed testing the network connectivity, turn off ping support on the external interfaces for additional security.

Verifying the FortiGate unit’s routing tables

The FortiGate routing table contains the routes stored for future use. If you are expecting certain routes to be there and they are not, that is a good indicator that your configuration is not what you expected.

The CLI command `get router info routing-table details` will provide you with every route’s routing protocol, destination address, gateway address, interface, weighting, and if the address is directly connected or not.

If you want to limit the display to BGP routes only, use the CLI command `get router info routing-table bgp`. If there are no BGP routes in the routing table, nothing will be displayed. In the CLI command you can replace BGP with static, or other routing protocols to only display those routes.

If you want to see the contents of the routing information database (RIB), use the CLI command `get router info routing-table database`. This will display the incoming routes that may or may not make it into the routing table.

Verifying traffic routing

Traffic may be reaching the internal network, but it may be using a different route than you think to get there.

Use a browser to try and access the Internet.

If needed, allow traceroute and other diag ports to be opened until things are working properly. Then remove access for them again.

Look for slow hops on the traceroute, or pings to a location, as they may indicate network loops that need to be fixed.

Any locations that have an unresolved traceroute or ping must be examined and fixed.

Use network packet sniffing to ensure traffic is being routed as you expect.

Verifying the dual-homed side of the configuration

Since there are two connections to the internet in this example, theoretically you can pull the plug on one of the ISP connections, and all traffic will go through the other connection. Alternately, you may choose to remove a default route to one ISP, remove that ISP's neighbor settings, or change the weightings to prefer other other ISP. These alternate ways to test dual-homing do not change physical cabling, which may be preferred in some situations.

If this does not work as expected, things to check include:

- default static routes — if these are wrong or don't exist, the traffic can't get out.
- BGP neighbor information — If the ISP router information is incorrect, the FortiGate unit won't be able to talk to it.

Redistributing and blocking routes in BGP

During normal BGP operation, peer routers redistribute routes from each other. However, in some specific situations it may be best to not advertise routes from one peer for various reasons. Some reasons may be the peer is redundant with another peer (they share the same routes exactly), it might be unreliable in some way, or some other reason.

The FortiGate can also take routes it learns from other protocols and advertise them in BGP, for example OSPF or RIP. If your Company hosts its own web or email servers, external locations will require routes to your networks to reach those services.

In this example the Company has a internal networks in an OSPF area, and is connected to a BGP AS and two BGP peers. Company goes through these two peers to reach the Internet. However, Peer 1 routes will not be advertised to Peer 2. The Company internal user and server networks are running OSPF, and will redistribute those routes to BGP so external locations can reach the web and email servers.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate unit](#)
- [Configuring other networking devices](#)
- [Testing this configuration](#)

Network layout and assumptions

This section includes:

- [Network layout](#)
- [Assumptions](#)


Network layout

The network layout for the BGP redistributing routes example involves the company network being connected to two BGP peers as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, and the peer routers.

The components of the layout include:

- There is only one BGP AS in this example — AS 65001, shared by the FortiGate unit and both peers.
- The Company's FortiGate unit connects to the Internet through two BGP peers.
- The Company internal networks on the dmz interface of the FortiGate unit with an IP of 10.11.201.0/24.
- The FortiGate units' interfaces are connected as follows:
 - port1 (dmz) has IP 10.11.201.110 and is the internal user and server network
 - port2 (external1) has IP 172.21.111.4 and is connected to Peer 1's network
 - port3 (external2) has IP 172.22.222.4 and is connected to Peer 2's network
- Peer 1 has IP 172.21.111.5, and Peer 2 has IP 172.22.222.5.
- OSPF Area 1 is configured on the dmz interface of the FortiGate unit, and is the routing protocol used by the internal users and servers.

Figure 173: BGP network topology



Assumptions

The the BGP redistributing routes configuration procedure follows these assumptions:

- the FortiGate unit has been configured following the Install Guide
- interfaces port1, port2, and port 3 exist on the FortiGate unit
- we don't know the router manufacturers of Peer 1 and Peer 2
- we don't know what other devices are on the BGP AS or OSPF Area
- all basic configuration can be completed in both GUI and CLI
- access lists and route maps will only be configured in CLI
- VDOMs are not enabled on the FortiGate unit

General Configuration Steps

- 1 [Configuring the FortiGate unit — networks and firewalls](#)

- 2 [Configuring the FortiGate unit - BGP](#)
- 3 [Configuring the FortiGate unit - OSPF](#)
- 4 [Configuring other networking devices](#)
- 5 [Testing network configuration](#)

Configuring the FortiGate unit — networks and firewalls

The FortiGate unit has three interfaces connected to networks — two external and one dmz.

Security policies must be in place to allow traffic to flow between these networks.

Firewall services will change depending on which routing protocol is being used on that network — either BGP or OSPF. Beyond that, all services that are allowed will be allowed in both directions due to the internal servers. The services allowed are web-server services (DNS, HTTP, HTTPS, SSH, NTP, FTP*, SYSLOG, and MYSQL), email services (POP3, IMAP, and SMTP), and general troubleshooting services (PING, TRACEROUTE). Those last two can be removed once the network is up and working properly to increase security. Other services can be added later as needed.

To configure the interfaces - GUI

- 1 Go to *System > Network > Interface*.
- 2 Edit port1 (dmz) interface.
- 3 Set the following information, and select *OK*.

Alias	dmz
IP/Netmask	10.11.201.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	OSPF internal networks
Administrative Status	Up

- 4 Edit port2 (external1) interface.
- 5 Set the following information, and select *OK*.

Alias	external1
IP/Netmask	172.21.111.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external Peer 1
Administrative Status	Up

- 6 Edit port3 (external2) interface.
- 7 Set the following information, and select *OK*.

Alias	external2
IP/Netmask	172.22.222.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external2 Peer2
Administrative Status	Up

To configure the FortiGate interfaces (CLI)

```
config system interface
  edit port1
    set alias dmz
    set ip 10.11.101.110 255.255.255.0
    set allowaccess https ssh ping
    set description "OSPF internal networks"
    set status up
  next
  edit port2
    set alias external1
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "external1 Peer 1"
    set status up
  next
  edit port3
    set alias external2
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "external2 Peer 2"
    set status up
  next
end
```

To configure the firewall addresses - GUI

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, and set the following information.

Address Name	Internal_networks
Type	Subnet / IP Range
Subnet / IP Range	10.11.201.0 255.255.255.0
Interface	port1

- 3 Select *OK*.
- 4 Select *Create New*, and enter the following information:
- 5 Select *OK*.

To configure the firewall addresses - CLI

```
config firewall address
  edit "BGP_services"
```

To configure firewall service groups - GUI

- 1 Go to *Firewall Objects > Service > Group*.
- 2 Select *Create New*.
- 3 Name the group *OSPF_Services*.
- 4 Move the following services to the right list: DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, OSPF, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.

- 5 Select *OK*.
- 6 Select *Create New*.
- 7 Name the group BGP_Services.
- 8 Move the following services to the right list: BGP, DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
- 9 Select *OK*.

To configure firewall service groups - CLI

```
config firewall service group
edit "BGP_services"
set member "BGP", "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT"
"HTTP" "HTTPS" "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP"
"SSH" "TRACEROUTE" "SYSLOG"
next
edit "OSPF_services"
set member "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP"
"HTTPS" "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH"
"TRACEROUTE" "SYSLOG" "OSPF"
next
end
```

Configuring the FortiGate unit - BGP

The only change from the standard BGP configuration for this example is configuring the blocking Peer 1's routes from being advertised to Peer 2. From the network topology you can guess that both of these peers likely share many routes in common and it makes no sense to advertise unneeded routes.

Blocking Peer 1's routes to Peer 2 is done with distribute-list-out keyword. They allow you to select which routes you will advertise to a neighbor using an access list. In this case we will block all incoming routes from Peer 1 when we send updates to Peer 2. Otherwise Peer 1 and Peer 2 are regular neighbors.

The FortiGate unit will redistribute routes learned from OSPF into BGP.

This is advanced configuration and the commands are only available in the CLI.

To create access list to block Peer 1 - CLI

```
config access-list
edit "block_peer1"
config rule
edit 1
set prefix 172.21.111.0 255.255.255.0
set action deny
set exact-match enable
end
end
end
```

To configure BGP on the FortiGate unit - CLI

```
config router bgp
set as 65001
```

```

set router-id 10.11.201.110
config redistribute ospf
    set status enable
end
config neighbor
    edit 172.22.222.5
        set remote-as 65001
        set distribute-list-out "block_peer1"
    next
    edit 172.21.111.5
        set remote-as 65001
    end
end
end

```

Configuring the FortiGate unit - OSPF

This configuration involves only one OSPF Area, so all traffic will be intra-area. If there were two or more areas with traffic going between them it would be inter-area traffic. These two types are comparable to BGP's traffic within one AS (iBGP) or between multiple ASes (eBGP). Redistributing routes from OSPF to BGP is considered external because either the start or end point is a different routing protocol.

The OSPF configuration is basic apart from redistributing BGP routes learned.

To configure OSPF on the FortiGate unit - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 For Router ID enter 10.11.201.110 and then select *Apply*.
- 3 Under *Advanced Options* and *Redistribute*, select *BGP* and set BGP metric to 1.
- 4 For *Areas*, select *Create New*, enter the following information and then select *OK*.

Area (IP)	0.0.0.0
Type	Regular
Authentication	None

- 5 For *Networks*, select *Create New*.
- 6 Enter 10.11.201.0/255.255.255.0 for *IP/Netmask*, and select *OK*.
- 7 For *Interfaces*, select *Create New*.
- 8 Enter OSPF_dmz_network for *Name*.
- 9 Select port1 (dmz) for *Interface*, and then select *OK*.

To configure OSPF on the FortiGate unit - CLI

```

config router ospf
    set router-id 10.11.201.110
    config area
        edit 0.0.0.0
            set type regular
            set authentication none
        end
    config network
        edit 1
            set area 0.0.0.0
            set prefix 10.11.201.0 255.255.255.0
        end
    end
end

```

```
end
config interface
  edit "OSPF_dmz_network"
    set interface port1(dmz)
    set status enable
  end
config redistribute bgp
  set status enable
  set metric 1
end
end
```

Configuring other networking devices

As with all BGP configurations, the peer routers will need to be updated with the FortiGate unit's BGP information including IP address, AS number, and what capabilities are being used such as IPv6, graceful restart, BFD, and so on.

Testing network configuration

Testing this configuration involves the standard connectivity checks, but also ensuring that routes are being passed between protocols as expected.

Check the routing table on the FortiGate unit to ensure that routes from both OSPF and BGP are present.

Check the routing table on devices on the OSPF network for routes redistributed from BGP. Also check those devices for connectivity to the Internet.

Check the routing table on Peer 2 to ensure no routes from Peer 1 are present, but routes from the internal OSPF network are present.

For help with troubleshooting, see [“Troubleshooting BGP”](#) on page 1763.



Open Shortest Path First (OSPF)

This section describes OSPF routing.

The following topics are included in this section:

- [OSPF Background and concepts](#)
- [Troubleshooting OSPF](#)
- [OSPF routing examples](#)

OSPF Background and concepts

This section includes:

- [Background](#)
- [The parts and terminology of OSPF](#)
- [How OSPF works](#)

Background

OSPF is a link-state interior routing protocol, that is widely used in large enterprise organizations. It only routes packets within a single autonomous system (AS). This is different from BGP as BGP can communicate between ASes.

The main benefit of OSPF is that it detects link failures in the network quickly and within seconds has converged network traffic successfully without any networking loops. Also OSPF has many features to control which routes are propagated and which are not, maintaining smaller routing tables. OSPF can also provide better load-balancing on external links than other interior routing protocols.

OSPF version 2 was defined in 1998 in RFC 2328. OSPF was designed to support classless IP addressing, and variable subnet masks. This was a shortcoming of the earlier RIP protocols.

Updates to OSPF version 2 are included in OSPF version 3 defined in 2008 in RFC 5340. OSPF3 includes support for IPv6 addressing where previously OSPF2 only supports IPv4 addressing.

The parts and terminology of OSPF

Parts and terminology of OSPF includes:

- [OSPF and IPv6](#)
- [Router ID](#)
- [Adjacency](#)
- [Designated router \(DR\) and backup router \(BDR\)](#)
- [Area](#)
- [Authentication](#)
- [Hello and dead intervals](#)

OSPF and IPv6

OSPF version 3 includes support for IPv6. Generally all IP addresses are in IPv6 format instead of IPv4.

OSPF3 area numbers use the same 32-bit numbering system as OSPF2.

Router ID

In OSPF, each router has a unique 32-bit number called its Router ID. Often this 32-bit number is written the same as a 32-bit IPv4 address would be written in dotted decimal notation. However some brands of routers, such as Cisco routers, support a router ID entered as an integer instead of an IP address.

It is a good idea to not use IP address in use on the router for the router ID number. The router ID does not have to be a particular IP address on the router. By choosing a different number, it will be harder to get confused which number you are looking at. A good idea can be to use the as much of the area's number as possible. For example if you have 15 routers in area 0.0.0.0 they could be numbered from 0.0.0.1 to 0.0.0.15. If you have an area 1.1.1.1, then routers in that area could start at 1.1.1.10 for example.

You can manually set the router ID on your FortiGate unit.

To manually set an OSPF router ID of 0.0.1.1 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 For *Router ID*, enter 0.0.1.1.
- 3 Select *OK*.

To manually set an OSPF router ID of 0.0.1.1 - CLI

```
config router ospf
  set router-id 0.0.1.1
end
```

Adjacency

In an OSPF routing network, when an OSPF router boots up it sends out OSPF Hello packets to find any neighbors, or routers that have access to the same network as the router booting up. Once neighbors are discovered and Hello packets are exchanged, updates are sent, and the Link State databases of both neighbors are synchronized. At this point these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met.

- The subnet mask used on both routers must be the same subnet.
- The subnet number derived using the subnet mask and each router's interface IP address must match.
- The Hello interval & The Dead interval must match.
- The routers must have the same OSPF area ID. If they are in different areas, they are not neighbors.
- If authentication is used, they must pass authentication checks.

If any of these parameters are different between the two routers, the routers do not become OSPF neighbors and cannot be adjacent. If the routers become neighbors, they are adjacent.

Adjacency and neighbors

Neighbor routers can be in a Two-Way state, and not be adjacent. Adjacent routers normally have a neighbour state of FULL. Neighbors only exchange Hello packets, and do not exchange routing updates. Adjacent routers exchange LSAs (LSDB information) as well as Hello packets. A good example of an adjacent pair of routers is the DR and BDR.

You can check on the state of an OSPF neighbor using the CLI command `get router info ospf neighbor all`. See [“Checking the state of OSPF neighbors” on page 1800](#).

Why adjacency is important

It is important to have adjacent pairs of routers in the OSPF routing domain because routing protocol packets are only passed between adjacent routers. This means adjacency is required for two OSPF routers to exchange routes. If there is no adjacency between two routers, such as one on the 172.20.120.0 network and another on the 10.11.101.0 network, the routers do not exchange routes. This makes sense because if all OSPF routers on the OSPF domain exchanged updates it would flood the network. Also its better for updates to progress through adjacent routers to ensure there are no outages along the way. Otherwise updates could skip over routers that are potentially offline, causing longer routing outages and delays while the OSPF domain learns of this outage later on.

If the OSPF network has multiple border routers and multiple connections to external networks, the designated router (DR) determines which router pairs become adjacent. The DR can accomplish this because it maintains the complete topology of the OSPF domain, including which router pairs are adjacent. The BDR also has this information in case the DR goes offline.

Designated router (DR) and backup router (BDR)

In OSPF a router can have a number of different roles to play.

A designated router (DR) is the designated broadcasting router interface for an AS. It looks after all the initial contact and other routing administration traffic. Having only one router do all this greatly reduces the network traffic and collisions.

If something happens and the designated router goes offline, the backup designated router (BDR) takes over. An OSPF FortiGate unit interface can become either a DR or BDR. Both the DR and the BDR cover the same area, and are elected at the same time. The election process doesn't have many rules, but the exceptions can become complex.

Benefits

The OSPF concept of the designated router is a big step above RIP. With all RIP routers doing their own updates all the time, RIP suffers from frequent and sometimes unnecessary updates that can slow down your network. With OSPF, not only do routing changes only happen when a link-state changes instead of any tiny change to the routing table, but the designated router reduces this overhead traffic even more.

However, smaller network topologies may only have a couple routers besides the designated router. This may seem excessive, but it maintains the proper OSPF form and it will still reduce the administration traffic but to a lesser extent than on a large network. Also your network topology is ready for when you expand your network.

DR and BDR election

An election chooses the DR and BDR from all the available routers. The election is primarily based on the priority setting of the routers—the highest priority becomes the DR, and the second highest becomes BDR. To resolve any ties, the router with the highest router ID wins. For example 192.168.0.1 would win over 10.1.1.2.

The router priority can vary from 0 to 255, but at 0 a router will never become a DR or BDR. If a router with a higher priority comes on line after the election, it must wait until after the DR and BDR go offline before it would become the DR.

If the original DR goes offline, but then is available when the BDR goes offline later on, the original DR will be promoted back to DR without an election leaving the new BDR as it is.

With your FortiGate unit, to configure the port1 interface to be a potential OSPF designated router or backup designated router called `ospf_DR` on the network, you need to raise the priority of the router to a very high number such as 250 out of 255. This will ensure the interface has a chance to be a DR, but will not guarantee that it will be one. Give the interface a low numbered IP address—such as 10.1.1.1 instead of 192.168.1.1—to help ensure it becomes a DR, but that is not part of this example. Enter the following command:

```
config router ospf
  config ospf-interface
    edit "ospf_DR"
      set priority 250
    end
  end
```

Area

An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.

Within an area if there are two or more routers that are viable, there will always be a designated router (DR) and a backup DR (BDR). For more on these router roles, see [“Designated router \(DR\) and backup router \(BDR\)” on page 1789](#).

Defining a private OSPF area, involves:

- assigning a 32-bit number to the area that is unique on your network
- defining the characteristics of one or more OSPF areas
- creating associations between the OSPF areas that you defined and the local networks to include in the OSPF area
- if required, adjusting the settings of OSPF-enabled interfaces.



IPv6 OSPF area numbers use the same 32-bit number notation as IPv4 OSPF

If you are using the web-based manager to perform these tasks, follow the procedures summarized below.

FortiGate units support the four main types of OSPF area:

- [Backbone area](#)
- [NSSA](#)
- [Stub area](#)

- Regular area

Backbone area

Every OSPF network has at least one AS, and every OSPF network has a backbone area. The backbone is the main area, or possibly the only area. All other OSPF areas are connected to a backbone area. This means if two areas want to pass routing information back and forth, that routing information will go through the backbone on its way between those areas. For this reason the backbone not only has to connect to all other areas in the network, but also be uninterrupted to be able to pass traffic to all points of the network. The backbone area is referred to as area 0 because it has an IP address of 0.0.0.0.

Stub area

A stub area is an OSPF area that receives no outside routes advertised into it, and all routing in it is based on a default route. This essentially isolates it from outside areas.

Stub areas are useful for small networks that are part of a larger organization, especially if the networking equipment can't handle routing large amounts of traffic passing through, or there are other reasons to prevent outside traffic, such as security. For example most organizations don't want their accounting department to be the center of their network with everyone's traffic passing through there. It would increase the security risks, slow down their network, and it generally doesn't make sense.

A variation on the stub area is the totally stubby area. It is a stub area that does not allow summarized routes.

NSSA

A not-so-stubby-area (NSSA) is a stub area that allows for external routes to be injected into it. While it still does not allow routes from external areas, it is not limited to only using the default route for internal routing.

Regular area

A regular area is what all the other ASes are, all the non-backbone, non-stub, non-NSSA areas. A regular area generally has a connection to the backbone, does receive advertisements of outside routes, and does not have an area number of 0.0.0.0.

Authentication

In the OSPF packet header are two authentication related fields —AuType, and Authentication.

All OSPF packet traffic is authenticated. Multiple types of authentication are supported in OSPFv2. However in OSPFv3, there is no authentication built-in but it is assumed that IPsec will be used for authentication instead.

Packets that fail authentication are discarded.

Null authentication

Null authentication indicates there is no authentication being used. In this case the 16-byte Authentication field is not checked, and can be any value. However checksumming is still used to locate errors. On your FortiGate this is the `none` option for authentication.

Simple Password authentication

Simple password refers to a standard plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication.

Cryptographic authentication

Cryptographic authentication involves the use of a shared secret key to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Your FortiGate unit supports all three levels of authentication through the authentication keyword associated with creating an OSPF interface .

For example to create an OSPF interface called `Accounting` on the `port1` interface that is a broadcast interface, has a hello interval of 10 seconds, has a dead interval of 40 seconds, uses text authentication (simple password) with a password of “`ospf_test`”, enter the command:

```
config router ospf
  config ospf-interface
    edit Accounting
      set interface port1
      set network-type broadcast
      set hello-interval 10
      set dead-interval 40
      set authentication text
      set authentication-key "ospf_test"
    end
  end
```

Hello and dead intervals

The OSPF Hello protocol is used to discover and maintain communications with neighboring routers.

Hello packets are sent out at a regular interval for this purpose. The DR sends out the Hello packets. In a broadcast network, the multicast address of 224.0.0.5 is used to send out Hello packets. New routers on the network listen for and reply to these packets to join the OSPF area. If a new router never receives a Hello packet, other routers will not know it is there and will not communicate with it. However, once a new router is discovered the DR adds it to the list of routers in that area and it is integrated into the routing calculations.

Dead interval is the time it takes when a router doesn't respond before it is declared dead, or offline. If this interval is too short routers will be declared offline when they aren't and the link-state updates will happen more than they need to. If the dead interval is too long, it will slow down network traffic while that router it attempted to be contacted when it is already offline.

Access Lists

Access lists are filters used by FortiGate unit OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also its easier to troubleshoot since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of 10.10.10.10 and 11.11.11.11, enter the command:

```
config router access-list
  edit test_list
  config rule
    edit 1
      set prefix 10.10.10.10 255.255.255.255
      set action allow
      set exact-match enable
    next
    edit 2
      set prefix 11.11.11.11 255.255.255.255
      set action allow
      set exact-match enable
    end
  end
end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the command access-list6 as follows:

```
config router access-list6
  edit test_list_ip6
  config rule
    edit 1
      set prefix6 2002:A0A:A0A:0:0:0:0:0/48
      set action deny
    next
    edit 2
      set prefix6 2002:B0B:B0B:0:0:0:0:0/48
      set action deny
    end
  end
```

To use an access_list, you must call it from a routing protocol such as RIP. The following example uses the access_list from the earlier example called test_list to match routes coming in on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially increase. Enter the following command:

```
config router ospf
  config distribute-list
    edit 5
      set access-list test_list
      set protocol connected
    end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose

How OSPF works

An OSPF network one or more areas. An OSPF area is typically divided into logical areas linked by Area Border Routers. A group of contiguous networks form an area. An Area Border Router (ABR) links one or more areas to the OSPF network backbone (area ID 0). See [“Area border router \(ABR\)” on page 1711](#).

OSPF is an interior routing protocol. It includes a backbone AS, and possibly additional ASes. The DR and BDR are elected from potential routers with the highest priorities. The DR handles much of the administration to lower the network traffic required. New routers are discovered through hello packets sent from the DR using the multicast address of 224.0.0.5. If the DR goes offline at any time, the BDR has a complete table of routes that it uses when it takes over as the DR router.

OSPF does not use UDP or TCP, but is encapsulated directly in IP datagrams as protocol 89. This is in contrast to RIP, or BGP. OSPF handles its own error detection and correction functions.

The OSPF protocol, when running on IPv4, can operate securely between routers, optionally using a variety of authentication methods to allow only trusted routers to participate in routing. OSPFv3, running on IPv6, no longer supports protocol-internal authentication. Instead, it relies on IPv6 protocol security (IPsec).

Other important parts of how OSPF works includes:

- [OSPF router discovery](#)
- [How OSPF works on FortiGate units](#)
- [External routes](#)
- [Link-state Database \(LSDB\) and route updates](#)
- [OSPF packets](#)

OSPF router discovery

OSPF-enabled routers generate Link-State Advertisements (LSA) and send them to their neighbors whenever the status of a neighbor changes or a new neighbor comes online. As long as the OSPF network is stable, LSAs between OSPF neighbors do not occur. An LSA identifies the interfaces of all OSPF-enabled routers in an area, and provides information that enables OSPF-enabled routers to select the shortest path to a destination. All LSA exchanges between OSPF-enabled routers are authenticated.

When a network of OSPF routers comes online, the follow steps occur.

- 1 When OSPF routers come online, they send out Hello packets to find other OSPF routers on their network segment.

- 2 When they discover other routers on their network segment, generally they become adjacent. Adjacent routers can exchange routing updates. See [“Adjacency” on page 1788](#).
- 3 A DR and BDR are elected from the available routers using priority settings, and router ID. See [“Designated router \(DR\) and backup router \(BDR\)” on page 1789](#), and [“DR and BDR election issues” on page 1801](#).
- 4 Link state updates are sent between adjacent routers to map the topology of the OSPF area.
- 5 Once complete, the DR floods the network with the updates to ensure all OSPF routers in the area have the same OSPF route database. After the initial update, there are very few required updates if the network is stable.

How OSPF works on FortiGate units

When a FortiGate unit interface is connected to an OSPF area, that unit can participate in OSPF communications. FortiGate units use the OSPF Hello protocol to acquire neighbors in an area. A neighbor is any router that is directly connected to the same area as the FortiGate unit, and ideally is adjacent with a state of Full. After initial contact, the FortiGate unit exchanges Hello packets with its OSPF neighbors regularly to confirm that the neighbors can be reached.

The number of routes that a FortiGate unit can learn through OSPF depends on the network topology. A single unit can support tens of thousands of routes if the OSPF network is configured properly.

External routes

OSPF is an internal routing protocol. OSPF external routes are routes with the destination of the connection using a routing protocol other than OSPF. OSPF handles external routes by adjusting the cost of the route to include the cost of the other routing protocol. There are two methods of calculating this cost, used for OSPF E1 and OSPF E2.

OSPF external1 (E1)

In OSPF E1 the destination is outside of the OSPF domain. This requires a different metric to be used beyond the normal OSPF metrics. The new metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.

OSPF external2 (E2)

OSPF E2 is the default external type when routes are redistributed outside of OSPF. OSPF E2 is similar to E1, except in this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. Dropping the OSPF portion can be useful in a number of situations, on border routers that have no OSPF portion for example or where the OSPF routing cost is negligible compared to the external routing cost.

Comparing E1 and E2

The best way to understand OSPF E1 and E2 routes is to check routing tables on OSPF routers. If you look at the routes on an OSPF border router, the redistributed routes will have an associated cost that represents only the external route, as there is no OSPF cost to the route due to it already being on the edge of the OSPF domain. However, if you look at that same route on a different OSPF router inside the OSPF routing domain, it will have a higher associated cost - essentially the external cost plus the cost over the OSPF domain to that border router. The border router uses OSPF E2, where the internal OSPF router uses OSPF E2 for the same route.

Viewing external routes

When you are trying to determine the costs for routes in your network to predict how traffic will be routed, you need to see the external OSPF routes and their associated costs. On your FortiGate unit, you find this information through your CLI.

To view external routes - CLI

You can view the whole routing table using `get router info routing-table all` to see all the routes including the OSPF external routes, or for a shorter list you can use the command `get router info routing-table ospf`. The letter at the left will be either E1 or E2 for external OSPF routes. The output will look similar to the following, depending on what routes are in your routing table.

```
FGT620B# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
    inter area
* - candidate default

O*E2   0.0.0.0/0 [110/10] via 10.1.1.3, tunnel_wan2, 00:02:11
O      10.0.0.1/32 [110/300] via 10.1.1.3, tunnel_wan2, 00:02:11
S      0.0.0.0/0 [10/0] via 192.168.183.254, port2
S      1.0.0.0/8 [10/0] via 192.168.183.254, port2
```

Link-state Database (LSDB) and route updates

OSPF is based on links. The links between adjacent neighbor routers allow updates to be passed along the network. Network links allow the DR to flood the area with Link-state database (LSDB) updates. External links allow the OSPF area to connect to destinations outside the OSPF autonomous system. Information about these links is passed throughout the OSPF network as link-state updates.

The LSDB contains the information that defines the complete OSPF area, but the LSDB is not the routing table. It contains the information from all the link-state updates passed along the network. When there are no more changes required, and the network is stable then the LSDB on each router in the network will be the same. The DR will flood the LSDB to the area to ensure each router has the same LSDB.

To calculate the best route (shortest path) to a destination, the FortiGate unit applies the Shortest Path First (SPF) algorithm — based on Dijkstra's algorithm — to the accumulated link-state information. OSPF uses relative path cost metric for choosing the best route. The path cost can be any metric, but is typically the bandwidth of the path — how fast traffic will get from one point to another.

The path cost, similar to “distance” for RIP, imposes a penalty on the outgoing direction of a FortiGate unit interface. The path cost of a route is calculated by adding together all of the costs associated with the outgoing interfaces along the path to the destination. The lowest overall path cost indicates the best route, and generally the fastest route. Some brands of OSPF routers, such as Cisco, implement cost as a direct result of bandwidth between the routers. Generally this is a good cost metric because larger bandwidth means more traffic can travel without slowing down. To achieve this type of cost metric on FortiGate units, you need to set the cost for each interface manually in the CLI.



The inter-area routes may not be calculated when a Cisco type ABR has no fully adjacent neighbor in the backbone area. In this situation, the router considers summary-LSAs from all Actively summary-LSAs from all Actively Attached areas (RFC 3509).

The FortiGate unit dynamically updates its routing table based on the results of the SPF calculation to ensure that an OSPF packet will be routed using the shortest path to its destination. Depending on the network topology, the entries in the FortiGate unit routing table may include:

- the addresses of networks in the local OSPF area (to which packets are sent directly)
- routes to OSPF area border routers (to which packets destined for another area are sent)
- if the network contains OSPF areas and non-OSPF domains, routes to area boundary routers, which reside on the OSPF network backbone and are configured to forward packets to destinations outside the OSPF AS.

OSPF Route updates

Once the OSPF domain is established, there should be few updates required on a stable network. When updates occur and a decision is required concerning a new route, this is the general procedure.

- 1 Our router gets a new route, and needs to decide if it should go in the routing table.
- 2 The router has an up to date LSDB of the entire area, containing information about each router, the next hop to it, and most importantly the cost to get there.
- 3 Our router, turns the LSDB into a shortest path first (SPF) tree using Dijkstra's algorithm. It doesn't matter if there is more than one path to a router on the network, the SPF tree only cares about the shortest path to that router.
- 4 Once the SPF tree has been created, and shows the shortest paths to all the OSPF routers on the network, the work is done. If the new route is the best route, it will be part of that tree. If it is not the shortest route, it will not be included in the LSDB.
- 5 If there has been a change from the initial LSDB to the new SPF tree, a link state update will be sent out to let the other routers know about the change so they can update their LSDBs as well. This is vital since all routers on the OSPF area must have the same LSDB.
- 6 If there was no change between the LSDB and the SPF tree, no action is taken.

OSPF packets

Every OSPF packet starts with a standard 24-byte header, and another 24 bytes of information or more. The header contains all the information necessary to determine whether the packet should be accepted for further processing.

Table 103: OSPF packet

1-byte Version field	1-byte Type field	2-byte Packet length	3-byte Router ID
4-byte Area ID	2-byte Checksum	2-byte Auth Type	8-byte Authentication
4-byte Network Mask	2-byte Hello interval	1-byte Options field	1-byte Router Priority
4-byte Dead Router interval	4-byte DR field	4-byte BDR field	4-byte Neighbor ID

The following descriptions summarize the OSPF packet header fields.

Version field— The OSPF version number. This specification documents version 2 of the protocol.

Type field— There are 5 OSPF packet types. From one to five, respectively, they are Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment.

Packet length— The length of the OSPF protocol packet in bytes. This length includes the standard OSPF 24-byte header, so all OSPF packets are at 24-bytes long.

Router ID— The Router ID of the packet's source.

Area ID— A 32-bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.

Checksum— The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field. If the packet's length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. The checksum is considered to be part of the packet authentication procedure; for some authentication types the checksum calculation is omitted.

Auth Type— Identifies the authentication procedure to be used for the packet. Authentication types include Null authentication (0), Simple password (1), Cryptographic authentication (2), and all others are reserved for future use.

Authentication— A 64-bit field for use by the authentication scheme. When AuType indicates no authentication is being used, the Authentication fields is not checked and can be any value. When AuType is set to 2 (Cryptographic authentication), the 64-bit authentication field is split into the following four fields: Zero field, Key ID field, Authentication data length field, and Cryptographic sequence field.

The Key ID field indicates the key and algorithm used to create the message digest appended to the packet. The authentication data length field indicates how many bytes long the message digest is, and the cryptographic sequence number is at non-decreasing number that is set when the packet is received and authenticated to prevent replay attacks.

Network Mask—The subnet where this packet is valid.

Hello interval—The period of time between sending out Hello packets. See [“Hello and dead intervals” on page 1792](#).

Options field— The OSPF protocol defines several optional capabilities. A router indicates the optional capabilities that it supports in its OSPF Hello packets, Database Description packets and in its LSAs. This enables routers supporting a mix of optional capabilities to coexist in a single Autonomous System.

Router priority—The priority between 0 and 255 that determines which routers become the DR and BDR. See [“Designated router \(DR\) and backup router \(BDR\)” on page 1789](#).

Dead router interval—The period of time when there is no response from a router before it is declared dead. See [“Hello and dead intervals” on page 1792](#).

DR and BDR fields—The DR and BDR fields each list the router that fills that role on this network, generally the routers with the highest priorities. See [“Designated router \(DR\) and backup router \(BDR\)” on page 1789](#).

Neighbor ID—The ID number of a neighboring router. This ID is used to discover new routers and respond to them.

Troubleshooting OSPF

As with other dynamic routing protocols, OSPF has some issues that may need troubleshooting from time to time. For basic troubleshooting, see the Troubleshooting chapter.

The more common issues include:

- [Clearing OSPF routes from the routing table](#)
- [Checking the state of OSPF neighbors](#)
- [Passive interface problems](#)
- [Timer problems](#)
- [Authentication issues](#)
- [DR and BDR election issues](#)

Clearing OSPF routes from the routing table

If you think the wrong route has been added to your routing table and you want to check it out, you first have to remove that route from your table before seeing if it is added back in or not. You can clear all or some OSPF neighbor connections (sessions) using the `exec router clear OSPF` command. The `exec router clear` command is much more limiting for OSPF than it is for BGP. See [“Clearing routing table entries” on page 1763](#).

For example, if you have routes in the OSPF routing table and you want to clear the specific route to IP address 10.10.10.1, you will have to clear all the OSPF entries. Enter the command:

```
FGT# exec router clear ospf process
```

Checking the state of OSPF neighbors

In OSPF each router sends out link state advertisements to find other routers on its network segment, and to create adjacencies with some of those routers. This is important because routing updates are only passed between adjacent routers. If two routers you believe to be adjacent are not, that can be the source of routing failures.

To identify this problem, you need to check the state of the OSPF neighbors of your FortiGate unit. Use the CLI command `get router info ospf neighbor all` to see all the neighbors for your FortiGate unit. You will see output in the form of:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID  Pri  State      Dead Time   Address      Interface
10.0.0.2     1    Full/ -   00:00:39   10.1.1.2     tunnel_wan1
10.0.0.2     1    Full/ -   00:00:34   10.1.1.4     tunnel_wan2
```

The important information here is the `State` column. Any neighbors that are not adjacent to your FortiGate unit will be reported in this column as something other than `Full`. If the state is `Down`, that router is offline.

Passive interface problems

A passive OSPF interface doesn't send out any updates. This means it can't be a DR, BDR, or an area border router among other things. It will depend on other neighbor routers to update its link-state table.

Passive interfaces can cause problems when they aren't receiving the routing updates you expect from their neighbors. This will result in the passive OSPF FortiGate unit interface having an incomplete or out of date link-state database, and it will not be able to properly route its traffic. It is possible that the passive interface is causing a hole in the network where no routers are passing updates to each other, however this is a rare situation.

If a passive interface is causing problems, there are some easy methods to determine it is the cause. The easiest method is to make it an active interface, and if the issues disappear that was the cause. Another method is to examine the OSPF routing table and related information to see if it is incomplete compared to other neighbor routers. If this is the case.

If you cannot make the interface active for some reason, you will have to change your network to fix the "hole" by adding more routers, or changing the relationship between the passive router's neighbors to provide better coverage.

Timer problems

A timer mismatch is when two routers have different values set for the same timer. For example if one router declares a router dead after 45 seconds and another waits for 4 minutes that difference in time will result in those two routers being out of synch for that period of time—one will still see that offline router as being online.

The easiest method to check the timers is to check the configuration on each router. Another method is to sniff some packets, and read the timer values in the packets themselves from different routers. Each packet contains the hello interval, and dead interval periods, so you can compare them easily enough.

Bi-directional Forwarding Detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

Authentication issues

OSPF has a number of authentication methods you can choose from. You may encounter problems with routers not authenticating as you expect. This will likely appear simply as one or more routers that have a blind spot in their routing - they won't acknowledge a router. This can be a problem if that router connects areas to the backbone as it will appear to be offline and unusable.

To confirm this is the issue, the easiest method is to turn off authentication on the neighboring routers. With no authentication between any routers, everything should flow normally.

Another method to confirm that authentication is the problem is to sniff packets, and look at their contents. The authentication type and password are right in the packets which makes it easy to confirm they are what you expect during real time. Its possible one or more routers is not configured as you expect and may be using the wrong authentication. This method is especially useful if there are a group of routers with these problems—it may only be one router causing the problem that is seen in multiple routers.

Once you have confirmed the problem is authentication related, you can decide how to handle it. You can turn off authentication and take your time to determine how to get your preferred authentication type back online. You can try another type of authentication, text instead of md5 for example, which may have more success and still provide some level of protection. The important part is that once you confirm the problem, you can decide how to fix it properly.

DR and BDR election issues

You can force a particular router to become the DR and BDR by setting their priorities higher than any other OSPF routers in the area. This is a good idea when those routers have more resources to handle the traffic and extra work of the DR and BDR roles, since not all routers may be able to handle all that traffic.

However, if you set all the other routers to not have a chance at being elected, a priority of zero, you can run into problems if the DR and BDR go offline. The good part is that you will have some warning generally as the DR goes offline and the BDR is promoted to the DR position. But if the network segment with both the DR and BDR goes down, your network will have no way to send hello packets, send updates, or the other tasks the DR performs.

The solution to this is to always allow routers to have a chance at being promoted, even if you set their priority to one. In that case they would be the last choice, but if there are no other candidates you want that router to become the DR. Most networks would have already alerted you to the equipment problems, so this would be a temporary measure to keep the network traffic moving until you can find and fix the problem to get the real DR back online.

OSPF routing examples

This section includes:

- [Basic OSPF example](#)

- [Advanced inter-area OSPF example](#)

Basic OSPF example

This example sets up an OSPF network at a small office. There are 3 routers, all running OSPF v2. The border router connects to a BGP network.

All three routers in this example are FortiGate units. Router1 will be the designated router (DR) and router2 will be the backup DR (BDR) due to their priorities. Router3 will not be considered for either the DR or BDR elections. Instead, Router3 is the area border router (ASBR) routing all traffic to the ISP's BGP router on its way to the Internet.

Router2 has a modem connected that provides dialup access to the Internet as well, at a reduced bandwidth. This is a PPPoE connection to a DSL modem. This provides an alternate route to the Internet if the other route goes down. The DSL connection is slow, and is charged by the amount of traffic. For these reasons OSPF will highly favor Router3's Internet access.

The DSL connection connects to an OSPF network with the ISP, so no redistribution of routes is required. The ISP network does have to be added to that router's configuration however.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

This section includes:

- [Network layout](#)
- [Assumptions](#)

Network layout

There are three FortiGate units acting as OSPF v2 routers on the network—Router1, Router2, and Router3. Router1 will be the designated router (DR), and Router 2 the BDR. Router3 is the area border router (ASBR) that connects to the external ISP router running BGP. Router2 has a PPPoE DSL connection that can access the Internet.

The Head Office network is connected to Router1 and Router2 on the 10.11.101.0 subnet.

Router1 and Router3 are connected over the 10.11.103.0 subnet.

Router2 and Router3 are connected over the 10.11.102.0 subnet.


The following table lists the router, interface, address, and role it is assigned.

Table 104: Routers, interfaces, and IP addresses for basic OSPF example network

Router name	Interface	IP address	Interface is connected to:
Router1 (DR)	Internal (port1)	10.11.101.1	Head office network, and Router2
	External (port2)	10.11.102.1	Router3

Table 104: Routers, interfaces, and IP addresses for basic OSPF example network

Router2 (BDR)	Internal (port1)	10.11.101.2	Head office network, and Router1
	External (port2)	10.11.103.2	Router3
	DSL (port3)	10.12.101.2	PPPoE DSL access
Router3 (ASBR)	Internal1 (port1)	10.11.102.3	Router1
	Internal2 (port2)	10.11.103.3	Router2
	External (port3)	172.20.120.3	ISP's BGP network

Figure 174: Basic OSPF network topology

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed, and are in NAT/Route operation mode.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF networks.
- Both Internet connections are always available.
- The modem connection is very slow and expensive.
- Other devices may be on the network, but do not affect this basic configuration.
- Router3 is responsible for redistributing all routes into and out of the OSPF AS.

General configuration steps

The general configuration steps involved are:

- 1 [Configuring the FortiGate units](#)
 - basic interface configuration
 - general system configuration
- 2 [Configuring OSPF on the FortiGate units](#)
 - configure OSPF for each interface
 - configure general OSPF settings for each router
 - Configure each router as one of DR, BDR, or ASBR
 - Configure route redistribution between BGP and OSPF
- 3 [Configuring other networking devices](#)
- 4 [Testing network configuration](#)

Configuring the FortiGate units

Each FortiGate unit needs the interfaces, and basic system information such as hostname configured.

This section includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)

Configuring Router1

Router1 has two interfaces connected to the network—internal (port1) and external (port2). Its host name must be changed to Router1.

To configure Router1 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Beside the host name, select *Change*.
- 3 Enter a hostname of `Router1`, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1, set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router2
Administrative Status	Up

- 5 Edit port2, set the following information, and then select *OK*.

Alias	External
IP/Netmask	10.11.102.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

Configuring Router2

Router2 configuration is the same as Router1, except Router2 also has the DSL interface to configure.

The DSL interface is configured with a username of “user1” and a password of “ospf_example”. The default gateway will be retrieved from the ISP, and the defaults will be used for the rest of the PPPoE settings.

To configure Router2 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Beside the host name, select *Change*.
- 3 Enter a hostname of `Router2`, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1, set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.101.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router1
Administrative Status	Up

- 5 Edit port2, set the following information, and then select *OK*.

Alias	External
IP/Netmask	10.11.103.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

- 6 Edit DSL (port3), set the following information, and then select *OK*.

Alias	DSL
Addressing Mode	PPPoE
Username	user1
Password	ospf_example
Unnumbered IP address	10.12.101.2/255.255.255.0
Retrieve default gateway from server	Enable
Administrative Access	HTTPS SSH PING
Description	DSL
Administrative Status	Up

Configuring Router3

Router3 is similar to Router1 and Router2 configurations. The main difference is the External (port3) interface connected to the ISP BGP network which has no administration access enabled for security reasons.

To configure Router3 interfaces - web-based manager

- 1 Go to *System > Status > Dashboard*.
- 2 Next to hostname, select *Change*.
- 3 Enter a hostname of *Router3*, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1, set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.102.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router1
Administrative Status	Up

- 5 Edit port2, set the following information, and then select *OK*.

Alias	Internal2
IP/Netmask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2
Administrative Status	Up

- 6 Edit port3, set the following information, and then select *OK*.

Alias	External
IP/Netmask	172.20.120.3/255.255.255.0
Administrative Access	
Description	ISP BGP
Administrative Status	Up

Configuring OSPF on the FortiGate units

With the interfaces configured, now the FortiGate units can be configured for OSPF on those interfaces. All routers are part of the backbone 0.0.0.0 area, so there is no inter-area communications needed.

For a simple configuration there will be no authentication, no graceful restart or other advanced features, and timers will be left at their defaults. Also the costs for all interfaces will be left at 10, except for the modem and ISP interfaces where cost will be used to load balance traffic. Nearly all advanced features of OSPF are only available from the CLI.

The network that is defined covers all the subnets used in this example - 10.11.101.0, 10.11.102.0, and 10.11.103.0. All routes for these subnets will be advertised. If there are other interfaces on the FortiGate units that you do not want included in the OSPF routes, ensure those interfaces use a different subnet outside of the 10.11.0.0 network. If you want all interfaces to be advertised you can use an OSPF network of 0.0.0.0 .

Each router will configure:

- router ID
- area
- network
- two or three interfaces depending on the router
- priority for DR (Router1) and BDR (Router2)
- redistribute for ASBR (Router3)

This section includes:

- [Configuring OSPF on Router1](#)
- [Configuring OSPF on Router2](#)
- [Configuring OSPF on Router3](#)

Configuring OSPF on Router1

Router1 has a very high priority to ensure it becomes the DR for this area. Also Router1 has the lowest IP address to help ensure it will win in case there is a tie at some point.

Otherwise it is a standard OSPF configuration.

Setting the priority can only be done in the CLI, and it is for a specific OSPF interface.

To configure OSPF on Router1 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Set *Router ID* to `10.11.101.1` and select *Apply*.
- 3 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

- 4 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

- 5 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Router1-Internal-DR
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

- 6 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Router1-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

- 7 Using the CLI, enter the following commands to set the priority for the Router1-Internal OSPF interface to maximum, ensuring this interface becomes the DR.

```
config router ospf
  config ospf-interface
    edit Router1-Internal-DR
      set priority 255
    end
```

To configure OSPF on Router1 - CLI

```
config router ospf
  set router-id 10.11.101.1
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router1-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router1-External"
      set interface "port2"
    next
  end
end
```

Configuring OSPF on Router2

Router2 has a high priority to ensure it becomes the BDR for this area, and configures the DSL interface slightly differently—assume this will be a slower connection resulting in the need for longer timers, and a higher cost for this route.

Otherwise it is a standard OSPF configuration.

To configure OSPF on Router2 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Set *Router ID* to 10.11.101.2 and select *Apply*.

- 3 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

- 4 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

- 5 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name		Router2-Internal
Interface		port1 (Internal)
IP		0.0.0.0
Authentication		none
Timers (seconds)		
	Hello Interval	10
	Dead Interval	40

- 6 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name		Router2-External
Interface		port2 (External)
IP		0.0.0.0
Authentication		none
Timers (seconds)		
	Hello Interval	10
	Dead Interval	40

- 7 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name		Router2-DSL
Interface		port3 (DSL)
IP		0.0.0.0
Authentication		none
Timers (seconds)		
	Hello Interval	20
	Dead Interval	80

- 8 Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
  config ospf-interface
    edit Router2-Internal
      set priority 250
```

```
    next
end
```

- 9 Using the CLI, enter the following commands to set the cost of the DSL interface higher than the other routes to reflect its higher monetary cost, and slower speed.

```
config router ospf
  config ospf-interface
    edit DSL
      set cost 50
    next
  end
```

To configure OSPF on Router2 - CLI

```
config router ospf
  set router-id 10.11.101.2
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.0.0
    next
  end
  config ospf-interface
    edit "Router2-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router2-External"
      set interface "port2"
    next
    edit "Router2-DSL"
      set interface "port3"
      set cost 50
    next
  end
end
```

Configuring OSPF on Router3

Router3 is more complex than the other two routers. The interfaces are straightforward, but this router has to import and export routes between OSPF and BGP. That requirement makes Router3 a border router or ASBR. Also Router3 needs a lower cost on its route to encourage all traffic to the Internet to route through it.

In the advanced OSPF options, Redistribute is enabled for Router3. It allows different types of routes, learned outside of OSPF, to be used in OSPF. Different metrics are assigned to these other types of routes to make them more or less preferred to regular OSPF routes.

To configure OSPF on Router3 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Set *Router ID* to 10.11.101.2 and select *Apply*.

- 3 Expand *Advanced Options*.
- 4 In *Redistribute*, set the following information, and select *OK*.

Route type	Redistribute	Metric
Connected	Enable	15
Static	Enable	15
RIP	Disable	n/a
BGP	Enable	5

- 5 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

- 6 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

- 7 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Router3-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

- 8 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Router3-Internal2
Interface	port2 (Internal2)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

- 9 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Router3-ISP-BGP
Interface	port3 (ISP-BGP)
IP	0.0.0.0

Authentication	none
Timers (seconds)	
Hello Interval	20
Dead Interval	80

- 10** Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
  config ospf-interface
    edit Router3-Internal
      set priority 250
    next
  end
```

- 11** Using the CLI, enter the following commands to set the cost of the DSL interface higher than the other routes to reflect its higher monetary cost, and slower speed.

```
config router ospf
  config ospf-interface
    edit ISP_BGP
      set cost 2
    next
  end
```

To configure OSPF on Router3 - CLI

```
config router ospf
  set router-id 10.11.102.3
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
    edit 2
      set prefix 172.20.120.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router3-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router3-External"
      set interface "port2"
    next
    edit "Router3-ISP-BGP"
      set interface "port3"
      set cost 2
    next
  end
end
```

Configuring other networking devices

The other networking devices required in this configuration are on the two ISP networks - the BGP network for the main Internet connection, and the DSL backup connection.

In both cases, the ISPs need to be notified of the OSPF network settings including router IP addresses, timer settings, and so on. The ISP will use this information to configure its routers that connect to this OSPF network.

Testing network configuration

Testing the network configuration involves two parts — testing the network connectivity, and testing the OSPF routing.

To test the network connectivity use ping, traceroute, and other network tools.

To test the OSPF routing in this example, refer to the troubleshooting outlined in [“Troubleshooting OSPF” on page 1799](#).

Advanced inter-area OSPF example

This example sets up an OSPF network at a large office. There are 3 areas, each with 2 routers. Typically OSPF areas would not be this small, and if they were the areas would be combined into one bigger area. However, the stub area services the accounting department which is very sensitive about their network and do not want any of their network information broadcast through the rest of the company. The backbone area contains the bulk of the company network devices. The regular area was established by IT for various reasons such as hosting the company servers on a separate area with extra security

One area is a small stub area that has no independent Internet connection, and only one connection to the backbone area. That connection between the stub area and the backbone area is only through a default route - no routes outside the stub area are advertised into that area.

Another area is the backbone, which is connected to the other two areas. The third area has the Internet connection, and all traffic to and from the Internet must use that area's connection. If that traffic comes from the stub area, then that traffic is treating the backbone like a transit area - an area it only uses to get to another area.

In the stub area, a subnet of computers is running the RIP routing protocol and those routes must be redistributed into the OSPF areas.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

This section includes:

- [Network layout](#)
- [Assumptions](#)

Network layout

There are four FortiGate units in this network topology acting as OSPF routers.

Area 1.1.1.1 is a stub area with one FortiGate unit OSPF router called Router1 (DR). Its only access outside of that area is a default route to the backbone area, which is how it accesses the Internet—traffic must go from the stub area, through the backbone, to the third area to reach the Internet. The backbone area in this configuration is called a transit area. Also in area 1.1.1.1 there is a RIP router that will be providing routes to the OSPF area through redistribution.

Area 0.0.0.0 is the backbone area, and has two FortiGate unit routers named Router2 (BDR) and Router3 (DR).

Area 2.2.2.2 is a regular area that has an Internet connection accessed by both the other two OSPF areas. There is only one FortiGate unit router in this area called Router4 (DR). This area is more secure and requires MD5 authentication by routers.

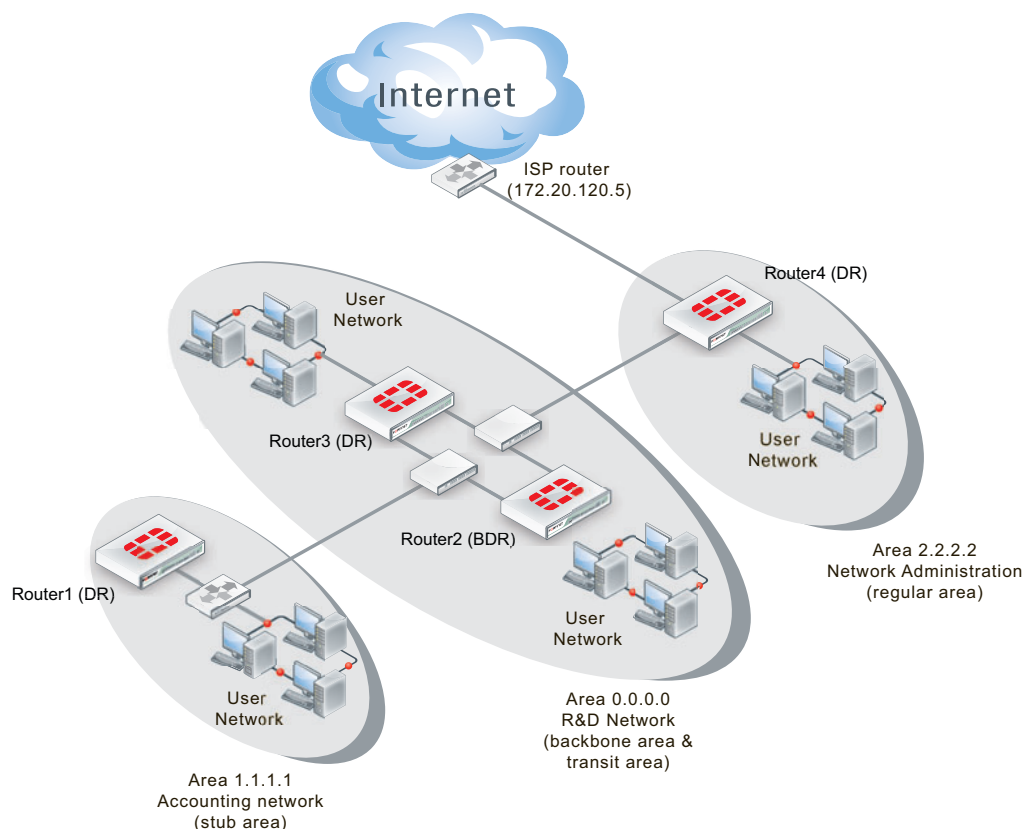
All areas have user networks attached, but they are not important for configuring the network layout for this example.

Internal interfaces are connected to internal user networks only. External1 interfaces are connected to the 10.11.110.0 network, joining Area 1.1.1.1 and Area 0.0.0.0.

External2 interfaces are connected to the 10.11.111.0 network, joining Area 0.0.0.0 and Area 2.2.2.2. The ISP interface is called ISP.

Table 105: Routers, areas, interfaces, IP addresses for advanced OSPF network

Router name	Area number and type	Interface	IP address
Router1 (DR)	1.1.1.1 - stub area (Accounting)	port1 (internal)	10.11.101.1
		port2 (external1)	10.11.110.1
Router2 (BDR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.102.2
		port2 (external1)	10.11.110.2
		port3 (external2)	10.11.111.2
Router3 (DR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.103.3
		port2 (external1)	10.11.110.3
		port3 (external2)	10.11.111.3
Router4 (DR)	2.2.2.2 - regular area (Network Admin)	port1 (internal)	10.11.104.4
		port2 (external2)	10.11.111.4
		port3 (ISP)	172.20.120.4

Figure 175: Advanced inter-area OSPF network topology

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed, and are in NAT/Route operation mode.
- During configuration, if settings are not directly referred to they will be left at default settings.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF areas outside of this example.
- The Internet connection is always available.
- Other devices may be on the network, but do not affect this configuration.

General configuration steps

The general configuration steps involved are:

- 1 **Configuring the FortiGate units**
 - basic interface configuration
 - general system configuration

- 2 [Configuring OSPF on the FortiGate units](#)
 - configure OSPF for each interface
 - configure general OSPF settings for each router
 - Configure each router as one of DR, BDR, or ASBR
 - Configure route redistribution between BGP and OSPF
- 3 [Configuring other networking devices](#)
- 4 [Testing network configuration](#)

Configuring the FortiGate units

This section configures the basic settings on the FortiGate units to be OSPF routers in this example. These configurations include multiple interface settings, and hostname.

There are four FortiGate units in this example. The two units in the backbone area can be configured exactly the same except for IP addresses, so only router3 (the DR) configuration will be given with notes indicating router2 (the BDR) IP addresses. These addresses can also be obtained from the “[Network layout](#)” on page 1814.

Configuring the FortiGate units includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)
- [Configuring Router4](#)

Configuring Router1

Router1 is part of the Accounting network stub area (1.1.1.1).

This section configures interfaces and hostname.

To configure Router1 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Next to hostname, select *Change*.
- 3 Enter a hostname of `Router1`, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1, set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

- 5 Edit port2, set the following information, and then select *OK*.

Alias	External1
IP/Netmask	10.11.110.1/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Backbone network and internet
Administrative Status	Up

Configuring Router2

Router2 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

Router2 has three interfaces configured—one to the internal network, and two to Router3 for redundancy.

This section configures interfaces and hostname.

To configure Router2 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Next to hostname, select *Change*.
- 3 Enter a hostname of `Router2`, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1 (internal), set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.102.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

- 5 Edit port2 (external1), set the following information, and then select *OK*.

Alias	external1
IP/Netmask	10.11.110.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 first connection
Administrative Status	Up

- 6 Edit port3 (external2), set the following information, and then select *OK*.

Alias	external2
IP/Netmask	10.11.111.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 second connection
Administrative Status	Up

Configuring Router3

Router3 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

This section configures interfaces and hostname.

To configure Router3 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Next to hostname, select *Change*.
- 3 Enter a hostname of `Router3`, and select *OK*.
- 4 Go to *System > Network > Interface*, edit port1 (internal), set the following information, and then select *OK*.

Alias	internal
IP/Netmask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

- 5 Edit port2 (external1), set the following information, and then select *OK*.

Alias	external1
IP/Netmask	10.11.110.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 first connection
Administrative Status	Up

- 6 Edit port3 (external2), set the following information, and then select *OK*.

Alias	external2
IP/Netmask	10.11.111.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 second connection
Administrative Status	Up

Configuring Router4

Router4 is part of the Network Administration regular area (2.2.2.2). This area provides internet access for both area 1.1.1.1 and the backbone area.

This section configures interfaces and hostname.

To configure Router4 interfaces - web-based manager

- 1 Go to *System > Dashboard > Status*.
- 2 Next to hostname, select *Change*.
- 3 Enter a hostname of `Router4`, and select *OK*.
- 1 Go to *System > Network > Interface*.
- 2 Edit port1 (internal).
- 3 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.4/255.255.255.0

Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

- 4 Edit port2 (external2).
- 5 Set the following information, and select *OK*.

Alias	external2
IP/Netmask	10.11.110.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone and Accounting network
Administrative Status	Up

- 6 Edit port3 (ISP).
- 7 Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

Configuring OSPF on the FortiGate units

Three of the routers are designated routers (DR) and one is a backup DR (BDR). This is achieved through the lowest router ID numbers, or OSPF priority settings.

Also each area needs to be configured as each respective type of area - stub, backbone, or regular. This affects how routes are advertised into the area.

To configure OSPF on Router1 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Enter 10.11.101.1 for the *Router ID* and then select *Apply*.
- 3 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	1.1.1.1
Type	Stub
Authentication	None

- 4 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.101.0/255.255.255.0
Area	1.1.1.1

- 5 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Accounting
Interface	port1 (internal)

IP	10.11.101.1
Authentication	None

- 6 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.1
Authentication	None

To configure OSPF on Router2 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Enter 10.11.102.2 for the *Router ID* and then select *Apply*.
- 3 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	None

- 4 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.102.2/255.255.255.0
Area	0.0.0.0

- 5 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.110.2/255.255.255.0
Area	0.0.0.0

- 6 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.111.2/255.255.255.0
Area	0.0.0.0

- 7 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	RnD network
Interface	port1 (internal)
IP	10.11.102.2
Authentication	None

- 8 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.2
Authentication	None

- 9 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.2
Authentication	None

To configure OSPF on Router3 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Enter 10.11.103.3 for the *Router ID* and then select *Apply*.
- 3 In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	None

- 4 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.102.3/255.255.255.0
Area	0.0.0.0

- 5 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.110.3/255.255.255.0
Area	0.0.0.0

- 6 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.111.3/255.255.255.0
Area	0.0.0.0

- 7 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	RnD network
Interface	port1 (internal)
IP	10.11.103.3
Authentication	None

- 8 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.3
Authentication	None

- 9 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.3
Authentication	None

To configure OSPF on Router4 - web-based manager

- 1 Go to *Router > Dynamic > OSPF*.
- 2 Enter 10.11.104.4 for the *Router ID* and then select *Apply*.
- 3 In *Areas*, select *Create New*.
- 4 Set the following information, and select *OK*.

Area	2.2.2.2
Type	Regular
Authentication	None

- 5 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.104.0/255.255.255.0
Area	0.0.0.0

- 6 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	10.11.111.0/255.255.255.0
Area	0.0.0.0

- 7 In *Networks*, select *Create New*, set the following information, and then select *OK*.

IP/Netmask	172.20.120.0/255.255.255.0
Area	0.0.0.0

- 8 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Network Admin network
Interface	port1 (internal)
IP	10.11.104.4
Authentication	None

- 9 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	Backbone2
Interface	port2 (external2)
IP	10.11.111.4
Authentication	None

10 In *Interfaces*, select *Create New*, set the following information, and then select *OK*.

Name	ISP
Interface	port3 (ISP)
IP	172.20.120.4
Authentication	None

Configuring other networking devices

All network devices on this network are running OSPF routing. The user networks (Accounting, R&D, and Network Administration) are part of one of the three areas.

The ISP needs to be notified of your network configuration for area 2.2.2.2. Your ISP will not advertise your areas externally as they are intended as internal areas. External areas have assigned unique numbers. The area numbers used in this example are similar to the 10.0.0.0 and 192.168.0.0 subnets used in internal networking.

Testing network configuration

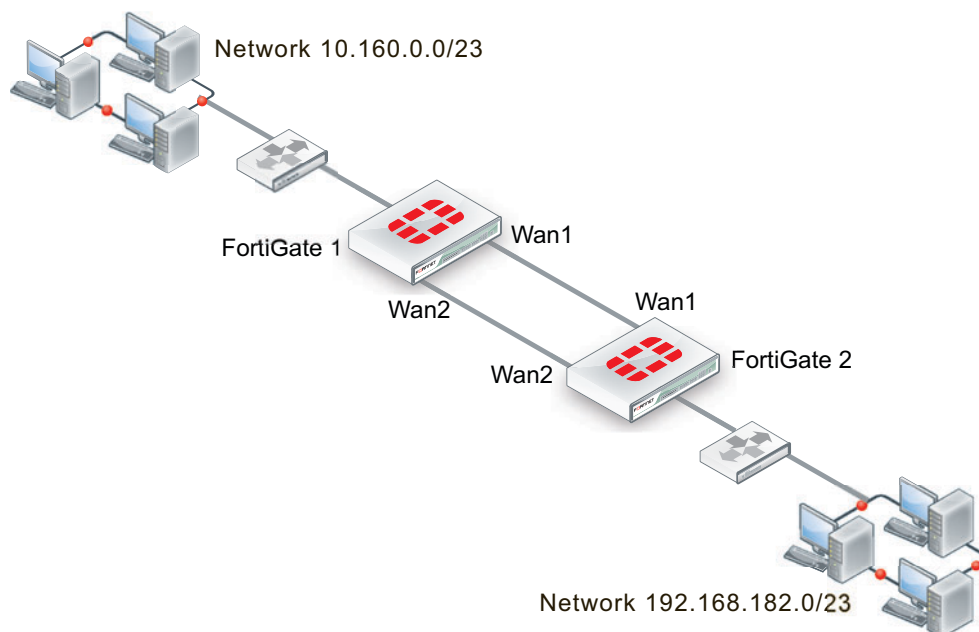
There are two main areas to test in this network configuration — network connectivity, and OSPF routing.

To test the network connectivity, see if computers on the Accounting or R&D networks can access the internet. If you need to troubleshoot network connectivity, see the Troubleshooting chapter.

To test the OSPF routing, check the routing tables on the FortiGate units to ensure the expected OSPF routes are present. If you need help troubleshooting OSPF routing, see [“Troubleshooting OSPF” on page 1799](#).

Controlling redundant links by cost

In this scenario, two FortiGate units have redundant links: one link between their WAN1 interfaces and another between their WAN2 interfaces.



FortiGate 1 should learn the route to network 192.168.182.0 and FortiGate 2 should learn the route to network 10.160.0.0. Under normal conditions, they should learn these routes through the WAN1 link. The WAN2 link should be used only as a backup.

With the default settings, each FortiGate unit learns these routes from both WAN1 and WAN2.

FortiGate 1:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 1 Full/Backup 00:00:33 10.182.0.187 wan1
10.2.2.2 1 Full/Backup 00:00:31 10.183.0.187 wan2

FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:01
[110/10] via 10.182.0.187, wan1, 00:00:01
O 192.168.182.0/23 [110/20] via 10.183.0.187, wan2, 00:02:04
[110/20] via 10.182.0.187, wan1, 00:02:04
```

FortiGate 2:

```
FGT2 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.1.1.1 1 Full/DR 00:00:38 10.182.0.57 wan1
10.1.1.1 1 Full/DR 00:00:38 10.183.0.57 wan2

FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.183.0.57, wan2, 00:00:39
[110/20] via 10.182.0.57, wan1, 00:00:39
```

Adjusting the route costs

On both FortiGate units, the cost of the route through WAN2 is adjusted higher so that this route will only be used if the route through WAN1 is unavailable. The default cost is 10. The WAN2 route will be changed to a cost of 200.

On both FortiGate units:

```
config router ospf
  config ospf-interface
    edit "WAN2_higher_cost"
      set cost 200
      set interface "wan2"
    end
```

Now both FortiGate units use only the WAN1 route:

FortiGate 1:

```
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.182.0.187, wan1, 00:00:40
O 192.168.182.0/23 [110/20] via 10.182.0.187, wan1, 00:00:40
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.182.0.57, wan1, 00:09:37
```

LSDB check on FortiGate 1:

```
FGT1 # get router info ospf database router lsa
Router Link States (Area 0.0.0.0)
```

LS age: 81
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x0
LS Type: router-LSA
Link State ID: 10.1.1.1
Advertising Router: 10.1.1.1
LS Seq Number: 8000000b
Checksum: 0xe637
Length: 60
Number of Links: 3

Link connected to: Stub Network
(Link ID) Network/subnet number: 10.160.0.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.57
Number of TOS metrics: 0
TOS 0 Metric: 200

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.57
Number of TOS metrics: 0
TOS 0 Metric: 10

LS age: 83
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x2 : ASBR
LS Type: router-LSA
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 8000000e
Checksum: 0xfc9b
Length: 60
Number of Links: 3

Link connected to: Stub Network
(Link ID) Network/subnet number: 192.168.182.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.187
Number of TOS metrics: 0
TOS 0 Metric: 200

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57

```
(Link Data) Router Interface address: 10.182.0.187
Number of TOS metrics: 0
TOS 0 Metric: 10
```

Verifying route redundancy

Bring down WAN1 and then check the routes on the two FortiGate units.

FortiGate 1:

```
FGT1 # get router info routing-table ospf
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:06
O 192.168.182.0/23 [110/210] via 10.183.0.187, wan2, 00:00:06
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/210] via 10.183.0.57, wan2, 00:00:14
```

The WAN2 interface is now in use on both units.



Intermediate System To Intermediate System Protocol (IS-IS)

This section describes the Intermediate System To Intermediate System Protocol (IS-IS). The following topics are included in this section:

- [IS-IS background and concepts](#)
- [Troubleshooting IS-IS](#)
- [Simple IS-IS example](#)

IS-IS background and concepts

This section contains:

- [Background](#)
- [How IS-IS works](#)

Background

Intermediate System To Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS).

IS-IS is an Interior Gateway Protocol (IGP). It is not intended to be used between Autonomous Systems (ASes). IS-IS is a link state protocol and uses Dijkstra's algorithm to find the best path, as does OSPF.

IS-IS was developed by Digital Equipment Corporation and later standardized by ISO in 1992 as ISO 19589 (see RFC 1142 — note this RFC is different from the ISO version). At roughly the same time the Internet Engineering Task Force developed OSPF. After the initial version, IP support was added to IS-IS and this version is called Integrated IS-IS (see RFC 1195).

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures.

How IS-IS works

As one of the original modern dynamic routing protocols, IS-IS is straight forward. Its routing algorithm is not complex, there are some options to allow fine tuning, and it's straight forward to configure IS-IS on FortiGate units.

From RFC 1058:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

IS-IS versus static routing

IS-IS was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, IS-IS is a big step forward from simple static routing.

While IS-IS may be slow in response to network outages, static routing has zero response. The same is true for convergence — static routing has zero convergence. Both IS-IS and static routing have the limited hop count, so its not a strength or a weakness. Count to infinity can be a problem, but typically can be fixed as it happens or is the result of a network outage that would cause even worse problems on static routing network.

Overall, IS-IS is a large step forward when compared to static routing.

IS-IS packet structure

It is hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how will help you better understand the IS-IS protocol, and better configure your network for it.

This section provides information on the contents of IS-IS 1 and IS-IS 2 packets.

IS-IS version 1

IS-IS version 1 packets are 24 bytes long. The zero fields were left for future expansion.

Table 106: IS-IS IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

The following descriptions summarize the IS-IS version 1 packet fields.

Command — Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple IS-IS packets are used to convey information from large routing tables.

Version — Specifies the IS-IS version used. This field can signal different potentially incompatible versions.

Zero field — This field defaults to zero, and is not used by RFC 1058 IS-IS.

Address-family identifier (AFI) — Specifies the address family used. IS-IS is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.

IP Address — Specifies the IP address for the entry.

Metric — This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable the metric is 16.

IS-IS version 2

IS-IS version 2 has more features than IS-IS 1 and this is reflected in its packets. IS-IS 2 packets are similar in format to IS-IS 1, but carry more information. Only one zero field remains.

Table 107: IS-IS 2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag	4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric
-------------------	-------------------	------------------	---------------	------------------------	----------------------	------------------	-----------------------	------------------

The following descriptions summarize the fields IS-IS 2 adds to the IS-IS IP header. The other fields have been described above for IS-IS 1.

Unused — Has a value set to zero, and is intended for future use

Route tag — Provides a method for distinguishing between internal routes learned by IS-IS and external routes learned from other protocols.

Subnet mask — Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.

Next hop — Indicates the IP address of the next hop to which packets for the entry should be forwarded.

Troubleshooting IS-IS

This section is about troubleshooting IS-IS. For general troubleshooting information, see the Troubleshooting chapter.

This section includes:

- [Routing Loops](#)
- [Split horizon and Poison reverse updates](#)

Routing Loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

Routing loops' effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the IS-IS routers connected to it. In a network loop, it's possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- [Check your logs](#)
- [Use SNMP network monitoring](#)
- [Use dead gateway detection and e-mail alerts](#)
- [Look at the packet flow](#)

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it.

Check your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report > Log & Archive Access*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

- 1 Go to *System > Config > SNMP*.
- 2 Enable *SNMP Agent*.
- 3 Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.
- 4 In either *SNMP v1/v2c* section or *SNMP v3* section, as appropriate, select *Create New*.
- 5 Enter the *Community Name* that you want to use.
- 6 In *Hosts*, select *Add* to add an IP address where you will be monitoring the FortiGate unit. You can add up to 8 different addresses.
- 7 Ensure that ports 161 and 162 (SNMP queries and traps) are allowed through your security policies.

- 8 In *SNMP Event*, select the events you want to be notified of. For routing loops this should include *CPU Overusage*, *Memory Low*, and possibly *Log disk space low*. If there are problems, the log will be filling up quickly, and the FortiGate unit's resources will be overused.
- 9 Select *OK*.
- 10 Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Use dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with dead gateway detection and e-mail alerts

- 1 To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
- 2 Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.
- 3 To configure interface status change notification, go to *Log&Report > Log Config > Alert E-mail*.
- 4 After you enter your email details, select the events you want to be alerted about — in our case *Configuration changes*. You may also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email to the outage.

Look at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable. Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

- 1 Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.

- 2 Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will not work. However, if B tries to use it this begins an endless loop.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This "poisoned" route is marked as unreachable for routers that cannot use it. In IS-IS this means that route is marked with a distance of 16.

Simple IS-IS example

This is an example of a typical medium sized network configuration using IS-IS routing.

Your company has 3 small local networks, one for each department. These networks are connected by IS-IS, and then connected to the Internet. Each subnet has more than one route, for redundancy. There are two central routers that are both connected to the internet, and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running IS-IS, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running those will need to be redistributed through the IS-IS network.

To keep the example simple, there will be no authentication of router traffic.

With IS-IS properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. IS-IS is good for a smaller network due to its lack of complex configurations.


Basic network layout

Table 108: IS-IS example network topology

Network	Router	Interface & Alias	IP address
Sales	Router1	port1 (internal)	10.11.101.101
		port2 (router2)	10.11.201.101
		port3 (router3)	10.11.202.101

Table 108: IS-IS example network topology

R&D	Router2	port1 (internal)	10.12.101.102
		port2 (router1)	10.11.201.102
		port3 (router4)	10.14.201.102
		port4 (ISP)	172.20.120.102
	Router3	port1 (internal)	10.12.101.103
		port2 (router1)	10.11.201.103
		port3 (router4)	10.14.202.103
		port4 (ISP)	172.20.120.103
Accounting	Router4	port1 (internal)	10.14.101.104
		port2 (router2)	10.14.201.104
		port3 (router3)	10.14.202.104

Figure 176: Network topology for the simple IS-IS example

Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 4.0 MR3 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 through port4 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- Only FortiGate units are running IS-IS on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.

- Router2 and Router3 each have their own connection to the Internet, indicated in black in Figure 176.

General configuration steps

This example is very straight forward. The only steps involved are:

- [Configuring FortiGate hostnames, interfaces, and default routes](#)
- [Configuring FortiGate unit IS-IS router information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Configuring FortiGate hostnames, interfaces, and default routes

On each FortiGate unit, configure its hostname, interfaces, and default routes as follows:

To configure Router1 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router1".
- 3 Go to *Router > Static > Static Route*.
- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 5 Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 6 Go to *System > Network > Interface*.
- 7 Edit port1 (internal) interface.
- 8 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal sales network
Administrative Status	Up

- 9 Edit port2 (router2) interface.
- 10 Set the following information, and select *OK*.

Alias	router2
IP/Netmask	10.11.201.101/255.255.255.0

Administrative Access	HTTPS SSH PING
Description	Link to R&D network & internet through Router2
Administrative Status	Up

11 Edit port3 (router3) interface.

12 Set the following information, and select *OK*.

Alias	router3
IP/Netmask	10.11.202.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

To configure Router2 system information - web-based manager

1 Go to *System > Dashboard > Status > System Information*.

2 Next to *Host Name* select *Change*, and enter "Router2".

3 Go to *Router > Static > Static Route*.

4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

5 Go to *System > Network > Interface*.

6 Edit port1 (internal) interface.

7 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.12.101.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router3
Administrative Status	Up

8 Edit port2 (router1) interface.

9 Set the following information, and select *OK*.

Alias	router1
IP/Netmask	10.12.201.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and the Sales network
Administrative Status	Up

10 Edit port3 (router4) interface.

11 Set the following information, and select **OK**.

Alias	router4
IP/Netmask	10.12.301.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and the accounting network
Administrative Status	Up

12 Edit port4 (ISP) interface.

13 Set the following information, and select **OK**.

Alias	ISP
IP/Netmask	172.20.120.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internet through ISP
Administrative Status	Up

14 Go to *Router > Static > Settings*, select *Create New* and enter the following:

Interface	port4
Ping Server	172.20.120.5
Detect Protocol	ICMP Ping
Leave other settings at default values.	

15 Select **OK**.

To configure Router3 system information - web-based manager

1 Go to *System > Dashboard > Status > System Information*.

2 Next to *Host Name* select *Change*, and enter "Router3".

3 Go to *Router > Static*.

4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

5 Go to *System > Network > Interface*.

6 Edit port1 (internal) interface.

7 Set the following information, and select **OK**.

Alias	internal
IP/Netmask	10.12.101.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router2
Administrative Status	Up

- 8 Edit port2 (router1) interface.
- 9 Set the following information, and select **OK**.

Alias	router1
IP/Netmask	10.13.201.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and Sales network
Administrative Status	Up

- 10 Edit port3 (router4) interface.
- 11 Set the following information, and select **OK**.

Alias	router4
IP/Netmask	10.13.301.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and accounting network
Administrative Status	Up

- 12 Edit port4 (ISP) interface.
- 13 Set the following information, and select **OK**.

Alias	ISP
IP/Netmask	172.20.120.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Administrative Status	Up

- 14 Go to *Router > Static > Settings*, select *Create New* and enter the following:

Interface	port4
Ping Server	172.20.120.5
Detect Protocol	ICMP Ping
Leave other settings at default values.	

- 15 Select **OK**.

To configure Router4 system information - web-based manager

- 1 Go to *System > Dashboard > Status > System Information*.
- 2 Next to *Host Name* select *Change*, and enter "Router4".
- 3 Go to *Router > Static*.
- 4 Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 5 Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

- 6 Go to *System > Network > Interface*.
 7 Edit port 1 (internal) interface.
 8 Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.14.101.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal accounting network
Administrative Status	Up

- 9 Edit port 2 (router2) interface.
 10 Set the following information, and select *OK*.

Alias	router2
IP/Netmask	10.14.201.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network & internet through Router2
Administrative Status	Up

- 11 Edit port 3 (router3) interface.
 12 Set the following information, and select *OK*.

Alias	router3
IP/Netmask	10.14.301.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

Configuring FortiGate unit IS-IS router information

With the interfaces configured, IS-IS can now be configured on the FortiGate units.

This includes the following steps:

- add interfaces that support IS-IS on the FortiGate unit
- add networks serviced by IS-IS
- redistribute static networks

This configuration is available only in the CLI.

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures — repeat the procedures once for each FortiGate unit.

Configure IS-IS settings on Router1 and Router4

```
config router isis
  config isis-interface
    edit port1
      set status enable
    next
    edit port2
      set status enable
    next
    edit port3
      set status enable
    next
    edit port4
      set status enable
  end
  config isis-net
    edit 1
      set net 10.11.0.0 255.255.0.0
    next
    edit 2
      set net 10.12.0.0 255.255.0.0
    next
    edit 3
      set net 10.14.0.0 255.255.0.0
    next
    edit 4
      set net 172.20.120.0 255.255.255.0
  end
  config redistribute static
    set status enable
  end
```

Configure IS-IS settings on Router2 and Router3

```
config router isis
  config isis-interface
    edit "port1"
      set status enable
    next
    edit "port2"
      set status enable
    next
    edit "port3"
      set status enable
    end
    edit "port4"
      set status enable
    end
  config isis-net
    edit 1
```

```
        set net 10.11.0.0 255.255.0.0
    next
    edit 2
        set net 10.12.0.0 255.255.0.0
    next
    edit 3
        set net 10.14.0.0 255.255.0.0
    next
    edit 4
        set net 172.20.120.0 255.255.255.0
    end
    config redistribute "static"
        set status enable
    end
```

Configuring other networking devices

In this example there are two groups of other devices on the the network — internal devices, and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers and other network devices. Once the FortiGate units are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise there should be no configuration required.

The second group of devices is the ISP. This consists of the IS-IS router the FortiGate routers 2 and 3 connect to. You need to contact your ISP and ensure they have your information for your network such as the IP addresses of the connecting IS-IS routers, what version of IS-IS your network supports, and what authentication (if any) is used.

Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult [“Troubleshooting IS-IS” on page 1829](#).



Router Reference

This section introduces you to the web-based manager Router menu.

This section contains the following topics:

- [Static](#)
- [Dynamic](#)
- [Monitor](#)



The word “unit” refers to the FortiGate unit. The words “FortiGate unit” are used when talking about different Fortinet products in one sentence. For example, “The Central Management menu provides the option of remotely managing your FortiGate unit by a FortiManager unit.”

Static

The Static menu provides settings for configuring both static and policy routes. A static route causes packets to be forwarded to a destination other than the factory configured default gateway. A policy route allows you to redirect traffic away from a static route, and can be useful when you want to route certain type of network traffic differently.

The factory configured static default route provides you with a starting point to configure the default gateway. You must either edit the factory configured static default route to specify a different default gateway for the unit, or delete the factory configured route and specify your own static default route that points to the default gateway for the unit.

Static routes are defined manually and control traffic exiting from the unit—you can specify through which interface the packet will leave and to which device the packet should be routed.

Route policies or policy routes, specify additional criteria for examining the properties of incoming packets. By using route policies, you can configure the unit to route packets based on the IP source and destination addresses in packet headers and other criteria such as on which interface the packet was received and which protocol (service) and port are being used to transport the packet.

The following topics are included in this section:

- [Static Route](#)
- [Policy Route](#)
- [Settings](#)

Static Route

Static routes are configured by defining the destination IP address and netmask of packets that you intend the unit to intercept, and by specifying a (gateway) IP address for those packets. The gateway address specifies the next-hop router to which traffic will be routed.

The Static Route page displays a list of routes that the unit compares to packet headers in order to route packets. Initially, the list contains the factory configured static default route. For more information, see [“Default route and default gateway” on page 1844](#). You can add new entries manually.

When you add a static route to the Static Route list, the unit performs a check to determine whether a matching route and destination already exist in the FortiGate routing table. If no match is found, the unit adds the route to the routing table.

When IPv6 is enabled in the web-based manager, IPv6 routes are visible on the Static Route list and you can select IPv6 when creating a new static route. Otherwise, IPv6 routes are not displayed.

Static routing configuration settings

The following are static route configuration settings in *Router > Static > Static Route*.

Static Route page	
Lists all the static routes that you created, including the default static route. On this page, you can edit, delete or create a new static route.	
Create New	<p>Creates a new static route. When you select <i>Create New</i>, you are automatically redirected to the New Static Route page.</p> <p>If you are configuring IPv6 addresses, select the down arrow beside <i>Create New</i> to create an IPv6 static Route.</p> <p>Tip: You can create a new static route by selecting the down arrow beside <i>Create New</i>, and then selecting <i>Route</i>.</p>
Edit	Modifies settings within a static route. When you select <i>Edit</i> , you are automatically redirected to the Edit Static Route page.
Delete	<p>Removes a static route from within the list on the Static Route page.</p> <p>To remove multiple static routes in the list, on the Static Route page, in each of the rows of the routes you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all static routes in the list, on the Static Route page, select the check box in the check box column, and then select <i>Delete</i>.</p>
Column Settings	Select to add, remove, or change the order of information columns. By default, the <i>Priority</i> and <i>Distance</i> columns are not displayed.
IP/Mask	The destination IP addresses and network masks of packets that the FortiGate unit intercepts.
Gateway	The IP addresses of the next-hop routers to which intercepted packets are forwarded.
Device	The names of the FortiGate interfaces through which intercepted packets are received and sent.
Comment	Optionally, add a description for this route.
Distance	The administrative distances associated with each route. The values represent distances to next-hop routers.
Weight	If ECMP Route Failover & Load Balance Method is set to <i>weighted</i> , add weights for each route. Add higher weights to routes that you want to assign more sessions to when load balancing.
New Static Route page	
Provides settings for defining the destination IP address and netmask of packets that you intend the unit to intercept, and by specifying a (gateway) IP address for those packets.	
Tip: You can display the priority and distance information on the Static Route page using <i>Column Settings</i> .	
Destination IP/Mask	Enter the destination IP address and netmask of the packets that you intend the unit to intercept.
Device	Select the interface through which intercepted packets are received and sent.

Gateway	Enter the gateway IP address for those packets that you intend the unit to intercept.
Comments	Enter a description about the static route.
Advanced	Select to show the options <i>Distance</i> and <i>Priority</i> .
Distance	Enter the number that represents the distances to the next-hop routers. The administrative distance allows you to weight one route to be preferred over another. This is useful when one route is unreliable. For example, if route A has an administrative distance of 10 and route B has an administrative distance of 30, the preferred route is route A with the smaller administrative distance of 10. If you discover that route A is unreliable, you can change the administrative distance for route A from 10 to 40, which will make the route B the preferred route.
Priority	Enter the number for the priority of the static route.



Unless otherwise specified, static route examples and procedures are for IPv4 static routes. You can use the `config router static6` CLI command to add, edit, or delete static routes for IPv6 traffic.

Default route and default gateway

In the factory default configuration, entry number 1 in the Static Route list is associated with a destination address of 0.0.0.0/0.0.0.0, which means any/all destinations. This route is called the static default route. If no other routes are present in the routing table and a packet needs to be forwarded beyond the unit, the factory configured static default route causes the unit to forward the packet to the default gateway.

To prevent this, you must either edit the factory configured static default route to specify a different default gateway for the unit, or delete the factory configured route and specify your own static default route that points to the default gateway for the unit.

Changing the gateway for the default route

The default gateway determines where packets matching the default route will be forwarded.

If you are using DHCP or PPPoE over a modem interface on your unit, you may have problems configuring a static route on this interface. After trying to either renew your DHCP license, or reconnect the PPPoE connection, go to the CLI and enable `dynamic-gateway` under `config system interface` for the modem interface. This will remove the need to specify a gateway for this interface's route.



For network traffic to pass, even with the correct routes configured, you must have the appropriate security policies.

To change the gateway for the default route

- 1 Go to *Router > Static > Static Route*.
- 2 Select the default route and then select *Edit*.

- 3 If the unit reaches the next-hop router through an interface other than the interface that is currently selected in the *Device* field, select the name of the interface from the *Device* field.
- 4 In the *Gateway* field, type the IP address of the next-hop router to which outbound traffic may be directed.
- 5 In the *Distance* field, optionally adjust the administrative distance value.
The default route distance should be set high enough to allow other routes to be configured at lower distances so they will be preferred over the default route.
- 6 Select *OK*.

Adding a static route to the routing table

Static routes are defined manually. They control traffic exiting the unit—you can specify through which interface the packet will leave and to which device the packet should be routed. When adding a static route to the routing table, use the “[New Static Route page](#)” on page 1843 table.

When you add a static route through the web-based manager, the unit adds the entry to the Static Route list.

However, if multiple routes to the same destination have the same priority but different distances, the route with the lowest distance is used. If multiple routes to the same destination have the same distance but different priorities, the route with the lowest priority is used. Distance takes precedence over priority. If multiple routes to the same destination have the different distances and different priorities, the route with the lowest distance is always used even if it has the highest priority.

Policy Route

Policy route options define which attributes of an incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the unit routes the packet through the specified interface to the specified gateway.

The following are policy route configuration settings in *Router > Static > Policy Route*.

Policy Route page	
Lists all policy routes that you have created. On this page, you can edit, delete or create a new policy route.	
Create New	Creates a new policy route. When you select <i>Create New</i> , you are automatically redirected to the New Routing Policy page.
Edit	Modifies settings within a static route. When you select <i>Edit</i> , you are automatically redirected to the Edit Routing Policy page.
Delete	<p>Removes a policy route from within the list on the Policy Route page.</p> <p>To remove multiple policy routes in the list, on the Policy Route page, in each of the rows of the routes you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all policy routes in the list, on the Policy Route page, select the check box in the check box column, and then select <i>Delete</i>.</p>

Move To	<p>Moves a policy route to any position within the list on the Policy Route page. You must select the check box in the row of the policy route that you want moved so that it will be moved in the list.</p> <p>When you select <i>Move To</i>, the Move Policy Route window appears. Enter the policy route ID number for the route's new position in the (<i>Policy route ID</i>) field. Select either <i>Before</i> or <i>After</i> to indicate the new position of the policy route. Select <i>OK</i>.</p>
#	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.
Incoming	The interfaces on which packets subjected to route policies are received.
Outgoing	The interfaces through which policy routed packets are routed.
Source	The IP source addresses and network masks that cause policy routing to occur.
Destination	The IP destination addresses and network masks that cause policy routing to occur.
New Routing Policy page Provides settings for configuring how to redirect traffic away from the static route.	
If incoming traffic matches:	
Protocol	<p>To perform policy routing based on the value in the protocol field of the packet, enter the protocol number to match. The Internet Protocol Number is found in the IP packet header. RFC 5237 describes protocol numbers and you can find a list of the assigned protocol numbers. The range is from 0 to 255. A value of 0 disables the feature.</p> <p>Tip: Commonly used <i>Protocol</i> settings include 6 to route TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.</p> <p>For protocols other than 6 and 17, the port number is ignored.</p>
Incoming interfaces	Select the name of the interface through which incoming packets subjected to the policy are received.
Source address/mask	To perform policy routing based on the IP source address of the packet, type the source address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination address/mask	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination ports	<p>To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature.</p> <p>The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.</p>

Type of Service	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see “Type of Service” on page 1695 .
Force traffic to:	
Outgoing interfaces	Select the name of the interface through which packets affected by the policy will be routed.
Gateway Address	Type the IP address of the next-hop router that the unit can access through the specified interface. A value of 0.0.0.0 is not valid.

Settings

In the Settings menu, you can configure dead gateway detection as well as the ECMP load balancing method.

Dead gateway detection is a way to detect the response between the connection to a server and the number of times the test fails before the unit assumes that the interface cannot connect to the server. ECMP load balancing is a load balancing method that helps to distribute traffic to the same destination, such as the Internet or network. For more information, see [“ECMP route failover and load balancing” on page 1685](#).

The following are configuration settings in *Router > Static > Settings*.

Settings page	
Lists the dead gateway detection settings as well as the chosen ECMP load balancing method.	
ECMP Load Balancing Method	
Select from one of the following ECMP route failover and load balancing methods that will be used for static routes:	
<ul style="list-style-type: none"> • <i>Source IP based</i> – balances sessions among ECMP routes based on the source IP address of the sessions to be load balanced. There are no configuration changes required to support source IP load balancing. • <i>Weighted Load Balance</i> – balances sessions among ECMP routes based on weights that are added to ECMP routes. More traffic is directed to routes with higher weights. If you select this option, you must also add weights to the static routes. • <i>Spillover</i> – distributes sessions among ECMP routes based on how busy the FortiGate interfaces are that have been added to the routes. After selecting Spillover, you must then add spillover thresholds to the interfaces added to ECMP routes. 	
Dead Gateway Detection	
Lists the dead gateways that you have created.	
Create New	Creates a new dead gateway. When you select <i>Create New</i> , you are automatically redirected to the New Dead Gateway Detection page.
Edit	Modifies settings within a dead gateway. When you select <i>Edit</i> , you are automatically redirected to the Edit Dead Gateway Detection page.

Delete	<p>Removes a dead gateway from the list.</p> <p>To remove multiple dead gateways from the list, within the Dead Gateway Detection section, in each of the rows of the of gateways you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all dead gateways in the list, within the Dead Gateway Detection section, select the check box in the check box column, and then select <i>Delete</i>.</p>
Interface	The FortiGate interface that the dead gateway will be detected on.
Ping Server	The IP address of the ping server.
Detect Protocol	The protocol that is used to detect the dead gateway.
Interval	The ping interval in seconds.
Failover	The failover threshold. This is the number of consecutive pings that do not get a response.
<i>New Dead Gateway Detection page</i> Provides settings for configuring a dead gateway. You are automatically redirected to this page when you select Create New.	
Interface	Select the interface that will be detecting dead gateways.
Ping Server	Enter the IP address of the ping server.
Detect Protocol	Select a protocol that will be used in detecting the dead gateway.
Ping Interval (seconds)	Enter the number of seconds for the ping interval.
Failover Threshold (Ping lost consecutively)	Enter the number for the threshold.
HA Priority	Set the HA remote IP monitoring priority for this interface if HA remote IP monitoring is configured for this interface from the CLI.

Dynamic

The Dynamic menu provides settings for configuring dynamic routing. Dynamic routing protocols allow the unit to automatically share information about routes with neighboring routers and learn about routes and networks advertised by them.

The unit supports these dynamic routing protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP).

This topic contains the following:

- [RIP](#)
- [OSPF](#)
- [BGP](#)
- [Multicast](#)



A unit can operate as a Protocol Independent Multicast (PIM) version 2 router in the root virtual domain. FortiGate units support PIM sparse mode and dense mode and can service multicast servers or receivers on the network segment to which a FortiGate interface is connected. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations.

RIP

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. The FortiGate implementation of RIP supports RIP version 1 (see RFC 1058), RIP version 2 (see RFC 2453), as well as the IPv6 version RIPng (see RFC 2080).

For detailed information about RIP, see [“RIP background and concepts” on page 1715](#).

RIP configuration settings

The following are RIP configuration settings in *Routing > Dynamic > RIP*.

RIP page Lists all the networks and interfaces that you have created. This page also allows you to configure basic RIP settings, including creating interfaces and networks.	
RIP Version	Select the level of RIP compatibility needed at the unit. You can enable global RIP settings on all FortiGate interfaces connected to RIP-enabled networks: <ul style="list-style-type: none"> 1 – send and receive RIP version 1 packets. 2 – send and receive RIP version 2 packets. You can override the global settings for a specific FortiGate interface if required. For more information, see “RIP-enabled interface” on page 1853 .
Advanced Options	Select the Expand Arrow to view or hide advanced RIP options. For more information, see “Advanced RIP options” on page 1851 .
Networks section of the RIP page The IP addresses and network masks of the major networks (connected to the unit) that run RIP. When you add a network to the Networks list, the FortiGate interfaces that are part of the network are advertised in RIP updates. You can enable RIP on all FortiGate interfaces whose IP addresses match the RIP network address space.	
IP/Netmask	Enter the IP address and netmask that defines the RIP-enabled network.
Add	Adds a new network IP address and netmask to the Networks section of the RIP page. When you select Add, the IP address and netmask is automatically added to the list.
Delete	Removes an IP address and netmask from within the Networks section of the RIP page. To remove multiple IP addresses and netmasks in the list, within the Networks section, in each of the rows of the networks you want removed, select the check box and then select <i>Delete</i> . To remove all IP addresses and netmasks in the list, within the Networks section, select the check box in the check box column, and then select <i>Delete</i> .
Interfaces section of the RIP page Any additional settings needed to adjust RIP operation on a FortiGate interface.	
Create New	Creates a new RIP interface. These parameters will override the global RIP settings for that interface. When you select <i>Create New</i> in the Interfaces section of the RIP page, you are automatically redirected to the New/Edit RIP Interface. See “RIP-enabled interface” on page 1853 .
Edit	Modifies settings within a RIP interface. When you select <i>Edit</i> , you are automatically redirected to the New/Edit RIP Interfaces.

	Delete	Removes a RIP interface from within the Interfaces section on the RIP page. To remove multiple interfaces in the list, within the Interfaces section, in each of the rows of the routes you want removed, select the check box and then select <i>Delete</i> . To remove all interfaces in the list, within the Interfaces section, select the check box in the check box column, and then select <i>Delete</i> .
	Interface	The name of the unit RIP interface.
	Send Version	The version of RIP used to send updates through each interface: 1, 2, or both.
	Receive Version	The versions of RIP used to listen for updates on each interface: 1, 2, or both.
	Authentication	The type of authentication used on this interface: <i>None</i> , <i>Text</i> or <i>MD5</i> .
	Passive	Permissions for RIP broadcasts on this interface. A green checkmark means the RIP broadcasts are blocked.



The `get router info rip` CLI command provides detailed information about configured RIP settings. You can view the entries in the RIP routing database or get status information about FortiGate interfaces.

Advanced RIP options

With advanced RIP options, you can specify settings for RIP timers and define metrics for redistributing routes that the unit learns through some means other than RIP updates. For example, if the unit is connected to an OSPF or BGP network or you add a static route to the FortiGate routing table manually, you can configure the unit to advertise those routes on RIP-enabled interfaces.

You can configure additional advanced options through customizable widgets, and the CLI. For example, you can filter incoming or outgoing updates by using a route map, an access list, or a prefix list. The unit also supports offset lists, which add the specified offset to the metric of a route.

Advanced RIP options are configured in *Router > Dynamic > RIP*, in the *Advanced Options* area of the page. You must expand Advanced Options to reveal the hidden settings so that you can configure these advanced options. Use the following table when configuring advanced RIP options.

Advanced Options section of the RIP page	
Advanced Options	Select the Expand Arrow to view or hide advanced options.
Default Metric	Enter the default hop count that the unit should assign to routes that are added to the FortiGate routing table. The range is from 1 to 16. This metric is the hop count, with 1 being best or shortest. This value also applies to Redistribute unless otherwise specified.
Enable Default-information-originate	Select to generate and advertise a default route into the unit's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both.

RIP Timers		<p>Enter new values to override the default RIP timer settings. The default settings are effective in most configurations — if you change these settings, ensure that the new settings are compatible with local routers and access servers.</p> <p>If the update timer is smaller than <i>Timeout</i> or <i>Garbage</i> timers, you will get an error.</p>
	Update	<p>Enter the amount of time (in seconds) that the unit will wait between sending RIP updates.</p> <p>The update timer determines the interval between routing updates. The update time should be at least three times smaller than the timeout timer; otherwise you will experience an error.</p>
	Timeout	<p>Enter the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the unit will keep a reachable route in the routing table while no updates for that route are received. If the unit receives an update for the route before the timeout period expires, the timer is restarted.</p> <p>The <i>Timeout</i> period should be at least three times longer than the <i>Update</i> period.</p> <p>Tip: If you are experiencing problems with routers not responding in time to updates, increase this time; however, remember that longer timeout intervals result in longer overall update periods.</p>
	Garbage	<p>Enter the amount of time (in seconds) that the unit will advertise a route as being unreachable before deleting the route from the routing table. The value determines how long an unreachable route is kept in the routing table.</p> <p>The garbage timer is the amount of time that the unit will advertise a route as being unreachable before deleting the route from the routing table.</p>
Redistribute		Select one or more of the options to redistribute RIP updates about routes that were not learned through RIP. The unit can use RIP to redistribute routes learned from directly connected networks, static routes, OSPF, and BGP.
	Connected	Select to redistribute routes learned from directly connected networks. To specify a hop count for those routes, select <i>Metric</i> , and enter the hop count in the <i>Metric</i> field. The valid hop count range is from 1 to 16.
	Static	Select to redistribute routes learned from static routes. To specify a hop count for those routes, select <i>Metric</i> , and enter the hop count in the <i>Metric</i> field. The range is from 1 to 16.
	OSPF	Select to redistribute routes learned through OSPF. To specify a hop count for those routes, select <i>Metric</i> , and enter the hop count in the <i>Metric</i> field. The range is from 1 to 16.
	BGP	Select to redistribute routes learned through BGP. To specify a hop count for those routes, select <i>Metric</i> , and enter the hop count in the <i>Metric</i> field. The range is from 1 to 16.

RIP-enabled interface

You can use RIP interface options to override the global RIP settings that apply to all unit interfaces connected to RIP-enabled networks. For example, if you want to suppress RIP advertising on an interface that is connected to a subnet of a RIP-enabled network, you can set the interface to operate passively. Passive interfaces listen for RIP updates but do not respond to RIP requests.

If RIP version 2 is enabled on the interface, you can optionally choose password authentication to ensure that the unit authenticates a neighboring router before accepting updates from that router. The unit and the neighboring router must both be configured with the same password. Authentication guarantees the authenticity of the update packet, not the confidentiality of the routing information in the packet.

RIP-enabled interface configuration settings

RIP-enabled interfaces are configured in *Router > Dynamic > RIP*. The following are RIP interface configuration settings.



Additional options such as split-horizon and key-chains can be configured per interface through the CLI.

New/Edit RIP Interface page

Provides settings for configuring a RIP Interface. When you select *Create New* in the Interfaces section of the RIP page, you are automatically redirected to the New/Edit RIP Interface page.

Interface	Select the name of the FortiGate interface to which these settings apply. The interface must be connected to a RIP-enabled network. The interface can be a virtual IPsec or GRE interface.
Send Version	Select to override the default RIP-compatibility setting for sending updates through the interface. You can choose either RIP version 1, 2, or both RIP versions.
Receive Version	Select to override the default RIP-compatibility setting for receiving updates through the interface. You can choose either RIP version 1, 2 or both RIP versions.
Authentication	Select an authentication method for RIP exchanges on the specified interface: <ul style="list-style-type: none"> <i>None</i> — Disable authentication. <i>Text</i> — Select if the interface is connected to a network that runs RIP version 2. Type a password (up to 35 characters) in the <i>Password</i> field. The unit and the RIP updates router must both be configured with the same password. The password is sent in clear text over the network.
Password	Enter the password for authentication.
Passive Interface	Select to suppress the advertising of unit routing information over the specified interface. Clear the check box to allow the interface to respond normally to RIP requests.

OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol that is most often used in large heterogeneous networks to share routing information among routers in the same Autonomous System (AS). FortiGate units support OSPF version 2 (see RFC 2328). For detailed information about OSPF, see [“OSPF Background and concepts” on page 1787](#).

The main benefit of OSPF is that it advertises routes only when neighbors change state instead of at timed intervals, so routing overhead is reduced.

When you configure OSPF settings, you have to define the AS in which OSPF is enabled and specify which of the FortiGate interfaces participate in the AS. As part of the AS definition, you specify the AS areas and specify which networks to include those areas. You may optionally adjust the settings associated with OSPF operation on the FortiGate interfaces.

Basic OSPF configuration settings

The following are OSPF configuration settings in *Router > Dynamic > OSPF*.

OSPF page	
Lists all areas, networks and interfaces that you created for OSPF.	
Router ID	<p>Enter a unique router ID to identify the unit to other OSPF routers. By convention, the router ID is the numerically highest IP address assigned to any of the FortiGate interfaces in the OSPF AS.</p> <p>If you change the router ID while OSPF is configured on an interface, all connections to OSPF neighbors will be broken temporarily. The connections will re-establish themselves.</p> <p>If <i>Router ID</i> is not explicitly set, the highest IP address of the VDOM or unit will be used.</p> <p>Note: Make sure to select <i>Apply</i> to apply the settings.</p>
Advanced Options	Expand to view or hide advanced OSPF settings. See “Advanced OSPF options” on page 1856 .
Areas section of the OSPF page	
Information about the areas making up an OSPF AS. The header of an OSPF packet contains an area ID, which helps to identify the origination of a packet inside the AS.	
Create New	Creates a new OSPF area. When you select <i>Create New</i> , you are automatically redirected to the New/Edit OSPF Area page.
Edit	Modifies settings within the OSPF area. When you select <i>Edit</i> , you are automatically redirected to the New/Edit OSPF Area page.
Delete	<p>Removes an OSPF area from within the Areas section on the OSPF page.</p> <p>To remove multiple areas in the list, within the Areas section, in each of the rows of the OSPF areas you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all areas in the list, within the Areas section, select the check box in the check box column, and then select <i>Delete</i>.</p>
Area	The unique 32-bit identifiers of areas in the AS, in dotted-decimal notation. Area ID 0.0.0.0 references the backbone of the AS and cannot be changed or deleted.

	Type	<p>The types of areas in the AS:</p> <ul style="list-style-type: none"> • <i>Regular</i> - a normal OSPF area • <i>NSSA</i> - a not so stubby area • <i>Stub</i> - a stub area. <p>For more information, see “Defining OSPF areas” on page 1857.</p>
	Authentication	<p>The methods for authenticating OSPF packets sent and received through all FortiGate interfaces linked to each area:</p> <ul style="list-style-type: none"> • <i>None</i> — authentication is disabled • <i>Text</i> — text-based authentication is enabled • <i>MD5</i> — MD5 authentication is enabled. <p>A different authentication setting may apply to some of the interfaces in an area, as displayed under Interfaces. For example, if an area employs simple passwords for authentication, you can configure a different password for one or more of the networks in that area.</p>
<p>Networks section of the OSPF page</p> <p>The networks in the OSPF AS and their area IDs. When you add a network to the Networks list, all FortiGate interfaces that are part of the network are advertised in OSPF link-state advertisements. You can enable OSPF on all FortiGate interfaces whose IP addresses match the OSPF network address space. For more information, see “OSPF networks” on page 1858.</p>		
	Create New	Creates a new OSPF network. When you select <i>Create New</i> , you are automatically redirected to the New/Edit OSPF Network page.
	Edit	Modifies settings within that OSPF network. When you select <i>Edit</i> , you are automatically redirected to the New/Edit OSPF Network page.
	Delete	<p>Removes an OSPF network from within the Networks section on the OSPF page.</p> <p>To remove multiple networks in the list, within the Networks section, in each of the rows of the OSPF networks you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all areas in the list, within the Networks section, select the check box in the check box column, and then select <i>Delete</i>.</p>
	Network	The IP addresses and network masks of networks in the AS on which OSPF runs. The FortiGate unit may have physical or VLAN interfaces connected to the network.
	Area	The area IDs that have been assigned to the OSPF network address space.
<p>Interfaces section of the OSPF page</p> <p>Any additional settings needed to adjust OSPF operation on a FortiGate interface. For more information, see “Operating parameters for an OSPF interface” on page 1859.</p>		
	Create New	Create additional/different OSPF operating parameters for a unit's interface and add the configuration to the Interfaces list. When you select <i>Create New</i> , you are automatically redirected to the New/Edit OSPF Interface page.

Edit	Modifies settings within that OSPF interface. When you select <i>Edit</i> , you are automatically redirected to the New/Edit OSPF Interface page.
Delete	Removes an OSPF interfaces from within the Interfaces section on the OSPF page. To remove multiple networks in the list, within the Interface section, in each of the rows of the OSPF interfaces you want removed, select the check box and then select <i>Delete</i> . To remove all areas in the list, within the Interfaces section, select the check box in the check box column, and then select <i>Delete</i> .
Name	The names of OSPF interface definitions.
Interface	The names of FortiGate physical or VLAN interfaces having OSPF settings that differ from the default values assigned to all other interfaces in the same area.
IP	The IP addresses of the OSPF-enabled interfaces having additional/different settings.
Authentication	The methods for authenticating LSA exchanges sent and received on specific OSPF-enabled interfaces. These settings override the area Authentication settings.

Advanced OSPF options

By selecting advanced OSPF options, you can specify metrics for redistributing routes that the unit learns through some means other than OSPF link-state advertisements. For example, if the unit is connected to a RIP or BGP network or you add a static route to the FortiGate routing table manually, you can configure the unit to advertise those routes on OSPF-enabled interfaces.

You can configure additional advanced options through customizable GUI widgets, and the CLI. For example, you can filter incoming or outgoing updates by using a route map, an access list, or a prefix list. The unit also supports offset lists, which add the specified offset to the metric of a route.

Advanced OSPF options are located in *Router > Dynamic > OSPF*. You must expand the Advanced Options on the page to access these options.

Advanced Options on the OSPF page	
Router ID	Enter a unique router ID to identify the unit to other OSPF routers.
Expand Arrow	Select to view or hide <i>Advanced Options</i> .
Default Information	Generate and advertise a default (external) route to the OSPF AS. You may base the generated route on routes learned through a dynamic routing protocol, routes in the routing table, or both.
None	Prevent the generation of a default route.
Regular	Generate a default route into the OSPF AS and advertise the route to neighboring autonomous systems only if the route is stored in the FortiGate routing table.
Always	Generate a default route into the OSPF AS and advertise the route to neighboring autonomous systems unconditionally, even if the route is not stored in the FortiGate routing table.

Redistribute		Select one or more of the options listed to redistribute OSPF link-state advertisements about routes that were not learned through OSPF. The unit can use OSPF to redistribute routes learned from directly connected networks, static routes, RIP, and BGP.
	Connected	Select to redistribute routes learned from directly connected networks. Enter a cost for those routes in the Metric field. The range is from 1 to 16 777 214.
	Static	Select to redistribute routes learned from static routes. Enter a cost for those routes in the Metric field. The range is from 1 to 16 777 214.
	RIP	Select to redistribute routes learned through RIP. Enter a cost for those routes in the Metric field. The range is from 1 to 16 777 214.
	BGP	Select to redistribute routes learned through BGP. Enter a cost for those routes in the Metric field. The range is from 1 to 16 777 214.

Defining OSPF areas

An area logically defines part of the OSPF AS. Each area is identified by a 32-bit area ID expressed in dotted-decimal notation, for example 192.168.0.1. Area ID 0.0.0.0 is reserved for the OSPF network backbone. You can classify the remaining areas of an AS as regular, stub, or NSSA.

A regular area contains more than one router, each having at least one OSPF-enabled interface to the area.

To reach the OSPF backbone, the routers in a stub area must send packets to an area border router. Routes leading to non-OSPF domains are not advertised to the routers in stub areas. The area border router advertises to the OSPF AS a single default route (destination 0.0.0.0) into the stub area, which ensures that any OSPF packet that cannot be matched to a specific route will match the default route. Any router connected to a stub area is considered part of the stub area.

In a Not-So-Stubby Area (NSSA), routes that lead out of the area into a non-OSPF domain are made known to OSPF AS. However, the area itself continues to be treated like a stub area by the rest of the AS.

Regular areas and stub areas (including not-so-stubby areas) are connected to the OSPF backbone through area border routers.

The following are configuration settings for defining OSPF areas on the OSPF page in *Router > Dynamic > OSPF*.



If required, you can define a virtual link to an area that has lost its physical connection to the OSPF backbone. Virtual links can be set up only between two units that act as area border routers.

New/Edit OSPF Area page

Provides settings for defining an OSPF area. When you select *Create New* in the *Areas* section of the OSPF page, you are automatically redirected to the New/Edit OSPF Area page.

Area	Type a 32-bit identifier for the area. The value must resemble an IP address in dotted-decimal notation. Once you have created the OSPF area, the area IP value cannot be changed; you must delete the area and restart.
Type	<p>Select an area type to classify the characteristics of the network that will be assigned to the area:</p> <ul style="list-style-type: none"> • <i>Regular</i> — If the area contains more than one router, each having at least one OSPF-enabled interface to the area. • <i>NSSA</i> — If you want routes to external non-OSPF domains made known to OSPF AS and you want the area to be treated like a stub area by the rest of the AS. • <i>STUB</i> — If the routers in the area must send packets to an area border router in order to reach the backbone and you do not want routes to non-OSPF domains to be advertised to the routers in the area.
Authentication	<p>Select the method for authenticating OSPF packets sent and received through all interfaces in the area:</p> <ul style="list-style-type: none"> • <i>None</i> — Disable authentication. • <i>Text</i> — Enables text-based password authentication. to authenticate LSA exchanges using a plain-text password. The password is sent in clear text over the network. • <i>MD5</i> — Enable MD5-based authentication using an MD5 cryptographic hash (RFC 1321). <p>If required, you can override this setting for one or more of the interfaces in the area. For more information, see “Operating parameters for an OSPF interface” on page 1859.</p>

OSPF networks

OSPF areas group a number of contiguous networks together. When you assign an area ID to a network address space, the attributes of the area are associated with the network.

Assigning an OSPF area ID to a network is configured in *Router > Dynamic > OSPF*. You must be in the Network section of the page to assign an OSPF area ID to a network. The following are configuration settings for OSPF networks.

New/Edit OSPF Network page

Provides settings for configuring networks that are assigned to an area ID. When you select *Create New* in the Network section of the OSPF page, you are automatically redirected to the New/Edit OSPF Network page.

IP/Netmask	Enter the IP address and network mask of the local network that you want to assign to an OSPF area.
Area	Select an area ID for the network. The attributes of the area must match the characteristics and topology of the specified network. You must define the area before you can select the area ID. For more information, see “Defining OSPF areas” on page 1857 .

Operating parameters for an OSPF interface

An OSPF interface definition contains specific operating parameters for a FortiGate OSPF-enabled interface. The definition includes the name of the interface (for example, external or VLAN_1), the IP address assigned to the interface, the method for authenticating LSA exchanges through the interface, and timer settings for sending and receiving OSPF Hello and dead-interval packets.

You can enable OSPF on all FortiGate interfaces whose IP addresses match the OSPF-enabled network space. For example, define an area of 0.0.0.0 and the OSPF network as 10.0.0.0/16. Then define vlan1 as 10.0.1.1/24, vlan2 as 10.0.2.1/24 and vlan3 as 10.0.3.1/24. All three VLANs can run OSPF in area 0.0.0.0. To enable all interfaces, you would create an OSPF network 0.0.0.0/0

When entering the operating parameters for MD5 keys for the interface, the following special characters are not supported:

- < >
- ()
- ‘ ’
- #
- “ ”
- .

You can configure different OSPF parameters for the same FortiGate interface when more than one IP address has been assigned to the interface. For example, the same FortiGate interface could be connected to two neighbors through different subnets. You could configure an OSPF interface definition containing one set of Hello and dead-interval parameters for compatibility with one neighbor's settings, and a second OSPF interface definition for the same interface to ensure compatibility with the second neighbor's settings.

The following are configuration settings for OSPF operating parameters in *Router > Dynamic > OSPF*, on the Interfaces section of the page.

New/Edit OSPF Interface page	
Provides settings for configuring an OSPF interface. When you select Create New in the Interface section of the OSPF page, you are automatically redirected to the New/Edit OSPF Interface page.	
Name	Enter a name to identify the OSPF interface definition. For example, the name could indicate to which OSPF area the interface will be linked.
Interface	Select the name of the FortiGate interface to associate with this OSPF interface definition (for example, port1, external, or VLAN_1). The unit can have physical, VLAN, virtual IPsec or GRE interfaces connected to the OSPF-enabled network.

IP	<p>Enter the IP address that has been assigned to the OSPF-enabled interface. The interface becomes OSPF-enabled because its IP address matches the OSPF network address space.</p> <p>For example, if you defined an OSPF network of 172.20.120.0/24 and port1 has been assigned the IP address 172.20.120.140, type 172.20.120.140.</p>
Authentication	<p>Select an authentication method for LSA exchanges on the specified interface:</p> <ul style="list-style-type: none"> • <i>None</i> — Disable authentication. • <i>Text</i> — Authenticate LSA exchanges using a plain-text password. The password can be up to 35 characters, and is sent in clear text over the network. • <i>MD5</i> — Use one or more keys to generate an MD5 cryptographic hash.
Password	<p>Enter the plain-text password. Enter an alphanumeric value of up to 15 characters. The OSPF neighbors that send link-state advertisements to this FortiGate interface must be configured with an identical password. This field is available only if you selected plain-text authentication.</p>
MD5 Keys	<p>Enter the key identifier for the (first) password in the ID field (the range is from 1 to 255) and then type the associated password in the Key field. The password is a 128-bit hash, represented by an alphanumeric string of up to 16 characters. When entering the characters, do not use < >, (), #, ", and ' because they are not supported.</p> <p>The OSPF neighbors that send link-state advertisements to this FortiGate interface must be configured with an identical MD5 key. If the OSPF neighbor uses more than one password to generate MD5 hash, select the Add icon to add additional MD5 keys to the list.</p> <p>This field is available only if you selected MD5 authentication.</p>
Timer(seconds) section of the page	
Hello Interval	<p>Optionally, set the Hello Interval to be compatible with Hello Interval settings on all OSPF neighbors.</p> <p>This setting defines the period of time (in seconds) that the unit waits between sending Hello packets through this interface.</p>
Dead Interval	<p>Optionally, set the Dead Interval to be compatible with Dead Interval settings on all OSPF neighbors.</p> <p>This setting defines the period of time (in seconds) that the unit waits to receive a Hello packet from an OSPF neighbor through the interface. If the unit does not receive a Hello packet within the specified amount of time, the unit declares the neighbor inaccessible.</p> <p>By convention, the Dead Interval value is usually four times greater than the Hello Interval value.</p>

BGP

Border Gateway Protocol (BGP) is an Internet routing protocol typically used by ISPs to exchange routing information between different ISP networks. For example, BGP enables the sharing of network paths between the ISP network and an autonomous system (AS) that uses RIP, OSPF, or both to route packets within the AS. The FortiGate implementation of BGP supports BGP-4 and complies with RFC 1771 and RFC 2385.

For more information, see [“BGP background and concepts” on page 1751](#).



You can configure graceful restarting and other advanced settings only through CLI commands.

When you configure BGP settings, you need to specify the AS to which the unit belongs and enter a router ID to identify this unit to other BGP routers. You must also identify the unit's BGP neighbors and specify which of the networks local to the unit should be advertised to BGP neighbors.

BGP configuration settings

The following are BGP configuration settings in *Router > Dynamic > BGP*. You can also configure many advanced BGP options through the CLI.

BGP page Lists all neighbors and networks that you have created. This page also allows you to configure neighbors, networks and a local AS. You can also configure four-byte AS paths as well. If you want additional information about configuring four-byte AS paths, see RFC 4893.	
Local AS	Enter the number of the local AS to which the FortiGate unit belongs. Select <i>Apply</i> to apply the local AS setting.
Router ID	Enter a unique router ID to identify the unit to other BGP routers. The router ID is an IP address written in dotted-decimal format, for example 192.168.0.1. If you change the router ID while BGP is configured on an interface, all connections to BGP peers will be broken temporarily. The connections will re-establish themselves. If Router ID is not explicitly set, the highest IP address of the VDOM will be used. Note: You must select <i>Apply</i> after entering the router ID to save the setting.
Neighbors section of the BGP page The IP addresses and AS numbers of BGP peers in neighboring autonomous systems.	
IP	Enter the IP address of the neighbor interface to the BGP-enabled network.
Remote AS	Enter the number of the AS that the neighbor belongs to.
Add/Edit	Add the neighbor information to the Neighbors list, or edit an entry in the list.
Neighbor	The IP addresses of BGP peers.
Remote AS	The numbers of the autonomous systems associated with the BGP peers.
Delete	Removes a BGP neighbor from within the Neighbors section on the BGP page. To remove multiple neighbors in the list, within the Neighbors section, in each of the rows of the BGPneighbors you want removed, select the check box and then select <i>Delete</i> . To remove all neighbors in the list, within the Neighbors section, select the check box in the check box column, and then select <i>Delete</i> .
Networks section of the BGP page The IP addresses and network masks of networks to advertise to BGP peers. The unit may have a physical or VLAN interface connected to those networks.	
IP/Netmask	Enter the IP address and netmask of the network to be advertised.
Add	Add the network information to the Networks list.

	Network	The IP addresses and network masks of major networks that are advertised to BGP peers.
	Delete	<p>Removes a BGP network from within the Networks section on the BGP page.</p> <p>To remove multiple networks in the list, within the Networks section, in each of the rows of the BGP networks you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all networks in the list, within the Networks section, select the check box in the check box column, and then select <i>Delete</i>.</p>



The `get router info bgp` CLI command provides detailed information about configured BGP settings.

Multicast

A FortiGate unit can operate as a Protocol Independent Multicast (PIM) version 2 router in the root virtual domain. FortiGate units support PIM sparse mode (RFC 2362) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate interface is connected.



You can configure basic options through the web-based manager. Many additional options are available, but only through the CLI.

When multicast (PIM) routing is enabled, you can configure sparse mode or dense mode operation on any FortiGate interface.

Multicast (PIM) interface options are used to set operating parameters for FortiGate interfaces connected to PIM domains. For example, you can enable dense mode on an interface that is connected to a PIM-enabled network segment. When sparse mode is enabled, you can adjust the priority number that is used to advertise Rendezvous Point (RP) and/or Designated Router (DR) candidacy on the interface.



Multicast offloading is available only on accelerated modules, and only in NAT Route mode. Multicast offloading is available only in the CLI.

Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast group(s) that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.



When a FortiGate interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP's distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that is associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly connected network segments. The PIM routers exchange state refresh messages to update their distribution trees. FortiGate units store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that do not have downstream receivers—PIM routers that do not manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

Multicast configuration settings

The following are multicast configuration settings in *Router > Dynamic > Multicast*. Advanced PIM settings are configured in the CLI.

Multicast page Lists each individual multicast route that you created. This page also allows you to configure each multicast route and add RP addresses. Note: Multicast offload is available only on accelerated modules, and only in NAT Route mode. Multicast offload is available only in the CLI.	
Enable Multicast Routing	Select to enable PIM version 2 routing. A security policy must be created on PIM-enabled interfaces to pass encapsulated packets and decapsulated data between the source and destination.
Static Rendezvous Points (RPs)	<p>If required for sparse mode operation, enter the IP address of a Rendezvous Point (RP) that may be used as the root of a packet distribution tree for a multicast group. Join messages from the multicast group are sent to the RP, and data from the source is sent to the RP.</p> <p>If an RP for the specified IP's multicast group is already known to the Bootstrap Router (BSR), the RP known to the BSR is used and the static RP address that you specify is ignored.</p>
Apply	Save the specified static RP addresses.
Create New	<p>Creates a new multicast entry for an interface. When you select <i>Create New</i>, you are automatically redirected to the New page.</p> <p>You can use the new entry to fine-tune PIM operation on a specific FortiGate interface or override the global PIM settings on a particular interface.</p>
Edit	Modifies settings within the multicast route. When you select <i>Edit</i> , you are automatically redirected to the Edit page.

Delete	Removes a multicast route from within the Multicast page. To remove multiple routes in the list, on the Multicast Route page, in each of the rows of the routes you want removed, select the check box and then select <i>Delete</i> . To remove all neighbors in the list, on the Multicast Route page, select the check box in the check box column, and then select <i>Delete</i> .
Interface	The names of FortiGate interfaces having specific PIM settings.
Mode	The mode of PIM operation (<i>Sparse</i> or <i>Dense</i>) on that interface.
Status	The status of parse-mode RP candidacy on the interface. To change the status of RP candidacy on an interface, select the Edit icon in the row that corresponds to the interface.
Priority	The priority number assigned to RP candidacy on that interface. Available only when RP candidacy is enabled.
DR Priority	The priority number assigned to Designated Router (DR) candidacy on the interface. Available only when sparse mode is enabled.
New page Provides settings for configuring a new multicast interface. When you select <i>Create New</i> on the Multicast page, you are automatically redirected to the New page.	
Interface	Select the name of the root VDOM FortiGate interface to which these settings apply. The interface must be connected to a PIM version 2 enabled network segment.
PIM Mode	Select the mode of operation: Sparse Mode or Dense Mode. All PIM routers connected to the same network segment must be running the same mode of operation. If you select Sparse Mode, adjust the remaining options as described below.
DR Priority	Enter the priority number for advertising DR candidacy on the unit's interface. The range is from 1 to 4 294 967 295. The unit compares this value to the DR interfaces of all other PIM routers on the same network segment, and selects the router having the highest DR priority to be the DR.
RP Candidate	Enable RP candidacy on the interface.
RP Candidate Priority	Enter the priority number for advertising RP candidacy on the FortiGate interface. The range is from 1 to 255.

Multicast destination NAT

Multicast destination NAT (DNAT) allows you translate externally received multicast destination addresses to addresses that conform to an organization's internal addressing policy.

By using this feature that is available only in the CLI, you can avoid redistributing routes at the translation boundary into their network infrastructure for Reverse Path Forwarding (RPF) to work properly. They can also receive identical feeds from two ingress points in the network and route them independently.

Configure multicast DNAT in the CLI using the `firewall multicast policy` command.

Bi-directional Forwarding Detection (BFD)

The Bi-directional Forwarding Detection (BFD) protocol is designed to deal with dynamic routing protocols' lack of a fine granularity for detecting device failures on the network and re-routing around those failures. BFD can more quickly react to these failures, since it detects them on a millisecond timer, where other dynamic routing protocols can only detect them on a second timer.

Your unit supports BFD as part of OSPF and BGP dynamic networking. BFD is configured only in the CLI.

Configuring BFD

BFD is intended for networks that use BGP or OSPF routing protocols. This generally excludes smaller networks.

BFD configuration on your unit is very flexible. You can enable BFD for the whole unit, and turn it off for one or two interfaces. Alternatively you can specifically enable BFD for each neighbor router, or interface. Which method you choose will be determined by the amount of configuring required for your network

The timeout period determines how long the unit waits before labeling a connection as down. The length of the timeout period is important—if it is too short connections will be labeled down prematurely, and if it is too long time will be wasted waiting for a reply from a connection that is down. There is no easy number, as it varies for each network and unit. High end models will respond very quickly unless loaded down with traffic. Also the size of the network will slow down the response time—packets need to make more hops than on a smaller network. Those two factors (CPU load and network traversal time) affect how long the timeout you select should be. With too short a timeout period, BFD will not connect to the network device but it will keep trying. This state generates unnecessary network traffic, and leaves the device unmonitored. If this happens, you should try setting a longer timeout period to allow BFD more time to discover the device on the network.

For this example, BFD is enabled on the unit using the default values. This means that once a connection is established, your unit will wait for up to 150 milliseconds for a reply from a BFD router before declaring that router down and rerouting traffic—a 50 millisecond minimum transmit interval multiplied by a detection multiplier of 3. The port that BFD traffic originates from will be checked for security purposes as indicated by disabling `bfd-dont-enforce-src-port`.

```
config system settings
  set bfd enable
  set bfd-desired-min-tx 50
  set bfd-required-min-rx 50
  set bfd-detect-mult 3
  set bfd-dont-enforce-src-port disable
end
```



The minimum receive interval (`bfd-required-min-rx`) and the detection multiplier (`bfd-detect-mult`) combine to determine how long a period your unit will wait for a reply before declaring the neighbor down. The correct value for your situation will vary based on the size of your network and the speed of your unit's CPU. The numbers used in this example may not work for your network.

Disabling BFD for a specific interface

The previous example enables BFD for your entire unit. If an interface is not connected to any BFD enabled routers, you can reduce network traffic by disabling BFD for that interface. For this example, BFD is disabled for the internal interface using CLI commands.

```
config system interface
  edit <interface>
    set bfd disable
  end
```

Access List

Access lists are filters used by the unit's routing processes to limit access to the network based on IP addresses. For an access list to take effect, it must be called by a unit routing process (for example, a process that supports RIP or OSPF). The offset list is part of the RIP and OSPF routing protocols. Access lists are configured only in the CLI.

For more information about access lists, see (for RIP) [“Access Lists” on page 1719](#) or (for OSPF) [“Access Lists” on page 1793](#).

Monitor

The Monitor menu provides a way to view the activity of configured routes. The Routing Monitor page allows you to view specific information as well as all routes. The list displays the entries in the FortiGate routing table.

This topic includes the following:

- [Viewing routing information](#)
- [Searching the routing monitor table](#)

Viewing routing information

By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of “any/all” packets.

View the list in *Router > Monitor > Routing Monitor*.

Routing Monitor page Lists all routes that are being monitored, including the default static route. On this page, you can also filter the information that is displayed on the page by applying a filter.	
IP version	Select IPv4 or IPv6 routes. Fields displayed vary depending on which IP version is selected. Displays only if IPv6 display is enabled on the web-based manager.
Type	Select one of the following route types to search the routing table and display routes of the selected type only: <ul style="list-style-type: none"> • <i>All</i> – all routes recorded in the routing table. • <i>Connected</i> – all routes associated with direct connections to FortiGate interfaces. • <i>Static</i> – the static routes that have been added to the routing table manually. • <i>RIP</i> – all routes learned through RIP. • <i>OSPF</i> – all routes learned through OSPF. • <i>BGP</i> – all routes learned through BGP. • <i>ISIS</i> – all routes learned through ISIS. • <i>HA</i> – RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you are viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster. Not displayed when IP version IPv6 is selected. For more information about HA routing synchronization, see the HA chapter.
Network	Enter an IP address and netmask (for example, 172.16.14.0/24) to search the routing table and display routes that match the specified network. Not displayed when IP version IPv6 is selected.

Gateway	<p>Enter an IP address and netmask (for example, 192.168.12.1/32) to search the routing table and display routes that match the specified gateway.</p> <p>Not displayed when IP version IPv6 is selected.</p>
Apply Filter	<p>Select to search the entries in the routing table based on the specified search criteria and display any matching routes.</p> <p>Not displayed when IP version IPv6 is selected.</p>
Type	<p>The type values assigned to FortiGate routes (Static, Connected, RIP, OSPF, or BGP).</p> <p>Not displayed when IP version IPv6 is selected.</p>
Subtype	<p>If applicable, the subtype classification assigned to OSPF routes.</p> <ul style="list-style-type: none"> • An empty string implies an intra-area route. The destination is in an area to which the unit is connected. • <i>OSPF inter area</i> — the destination is in the OSPF AS, but the unit is not connected to that area. • <i>External 1</i> — the destination is outside the OSPF AS. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together. • <i>External 2</i> — the destination is outside the OSPF AS. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. • <i>OSPF NSSA 1</i> — same as External 1, but the route was received through a not-so-stubby area (NSSA). • <i>OSPF NSSA 2</i> — same as External 2, but the route was received through a not-so-stubby area. <p>Not displayed when IP version IPv6 is selected.</p>
Network	<p>The IP addresses and network masks of destination networks that the unit can reach.</p>
Distance	<p>The administrative distance associated with the route. A value of 0 means the route is preferable compared to routes to the same destination.</p> <p>To modify the administrative distance assigned to static routes, see “Adding a static route to the routing table” on page 1845.</p>
Metric	<p>The metric associated with the route type. The metric of a route influences how the unit dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to.</p> <ul style="list-style-type: none"> • <i>Hop count</i> — routes learned through RIP. • <i>Relative cost</i> — routes learned through OSPF. • <i>Multi-Exit Discriminator (MED)</i> — routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network.
Gateway	<p>The IP addresses of gateways to the destination networks.</p>

Interface	The interface through which packets are forwarded to the gateway of the destination network.
Up Time (d h:m:s)	The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable. This column does not display when IPv6 is selected.

Searching the routing monitor table

You can apply a filter to search the routing table and display certain routes only. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed (an implicit AND condition is applied to all of the search parameters you specify).

For example, if the unit is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select *Connected* from the Type list, type *172.16.14.0/24* in the Network field, and then select *Apply Filter* to display the associated routing table entry or entries. Any entry that contains the word “Connected” in its Type field and the specified value in the Gateway field will be displayed.



All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.



Chapter 11 Virtual Domains

This FortiOS Handbook chapter contains the following sections:

[Virtual Domains](#) provides an overview of the VDOM technologies, and the basic concepts and rules for using them. We recommend that you begin with this chapter before attempting to configuring VDOMs on your FortiGate unit.

[Virtual Domains in NAT/Route mode](#) provides detailed explanations and examples for configuring VDOM features in your FortiGate unit using the NAT/Route mode.

[Virtual Domains in Transparent mode](#) provides detailed explanations, as well as basic and advanced examples for configuring these features in your FortiGate unit using Transparent mode.

[Inter-VDOM routing](#) describes inter-VDOM routing concepts and scenarios, and gives examples that illustrate them.

[Troubleshooting Virtual Domains](#) provides diagnostic and troubleshooting information for some potential VDOM issues.



Virtual Domains

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs can provide separate firewall policies and, in NAT/Route mode, completely separate configurations for routing and VPN services for each connected network or organization.

This chapter will cover the basics of VDOMs, how they change your FortiGate unit, and how to work with VDOMs.

VDOMs let you split your physical FortiGate unit into multiple virtual units. The resulting benefits range from limiting Transparent mode ports to simplified administration, and reduced space and power requirements.

When VDOMs are disabled on any FortiGate unit, there is still one VDOM active: the root VDOM. It is always there in the background. When VDOMs are disabled, the root VDOM is not visible but it is still there.

The root VDOM must be there because the FortiGate unit needs a management VDOM for management traffic among other things. It is also why when you enable VDOMs, all your configuration is preserved in the root VDOM-because that is where you originally configured it.

This section includes:

- [Benefits of Virtual Domains](#)
- [Enabling and accessing Virtual Domains](#)
- [Configuring Virtual Domains](#)

Benefits of Virtual Domains

VDOMs provide the following benefits:

- [Improving Transparent mode configuration](#)
- [Easier administration](#)
- [Continued security](#)
- [Savings in physical space and power](#)
- [More flexible MSSP configurations](#)

Improving Transparent mode configuration

When VDOMs are not enabled, and you put your FortiGate unit into Transparent mode all the interfaces on your unit become broadcast interfaces. The problem is there are no interfaces free to do anything else.

With multiple VDOMs you can have one of them configured in Transparent mode, and the rest in NAT/Route mode. In this configuration, you have an available transparent mode FortiGate unit you can drop into your network for troubleshooting, and you also have the standard.

Easier administration

VDOMs provide separate security domains that allow separate zones, user authentication, firewall policies, routing, and VPN configurations. VDOMs separate security domains and simplify administration of complex configurations—you do not have to manage as many settings at one time. For more information, see [“Global and per-VDOM settings” on page 1881](#).

By default, each FortiGate unit has a VDOM named root. This VDOM includes all of the unit’s physical interfaces, modem, VLAN subinterfaces, zones, firewall policies, routing settings, and VPN settings.

Also, you can optionally assign an administrator account restricted to one VDOM. If the VDOM is created to serve an organization, this feature enables the organization to manage its own configuration. For more information, see [“Administrators in Virtual Domains” on page 1898](#).

Each physical FortiGate unit requires a FortiGuard license to access security updates. VDOMs do not require any additional FortiGuard licenses, or updating — all the security updates for all the VDOMs are performed once per update at the global level. Combined this can be a potentially large money and time saving feature in your network.

Management systems such as SNMP, logging, alert email, FDN-based updates, and NTP-based time setting use addresses and routing in the management VDOM to communicate with the network. They can connect only to network resources that communicate with the management VDOM. Using a separate VDOM for management traffic enables easier management of the FortiGate unit global settings, and VDOM administrators can also manage their VDOMs more easily. For more information, see [“Changing the management virtual domain” on page 1903](#).

Continued security

When a packet enters a VDOM, it is confined to that VDOM and is subject to any firewall policies for connections between VLAN subinterfaces or zones in that VDOM, just like those interfaces on a FortiGate unit without VDOMs enabled.

To travel between VDOMs, a packet must first pass through a firewall policy on a physical interface. The packet then arrives at another VDOM on that same FortiGate unit, but on a different interface, where it must pass through another firewall before entering. It doesn’t matter if the interface is physical or virtual — inter-VDOM packets still require the same security measures as when passing through physical interfaces.

VDOMs provide an additional level of security because regular administrator accounts are specific to one VDOM — an administrator restricted to one VDOM cannot change information on other VDOMs. Any configuration changes and potential errors will apply only to that VDOM and limit any potential down time. Using this concept, you can farther split settings so that the management domain is only accessible by the super_admin and does not share any settings with the other VDOMs.

Savings in physical space and power

To increase the number of physical FortiGate units, you need more rack space, cables, and power to install the new units. You also need to change your network configuration to accommodate the new physical units. In the future, if you need fewer physical units you are left with expensive hardware that is idle.

Increasing VDOMs involves no additional hardware, no additional cabling, and very few changes to existing networking configurations. VDOMs save physical space and power. You are limited only by the size of the VDOM license you buy and the physical resources on the FortiGate unit.

For example if you are using one FortiGate 620B with 10 VDOMs instead of 10 of those units, over a year you will save an estimated 18,000 kWh. You could potentially save ten times that amount with a 100 VDOM license.

By default, FortiGate units support a maximum of 10 VDOMs in any combination of NAT/Route and Transparent modes. For FortiGate models numbered 3000 and higher, you can purchase a license key to increase the maximum number of VDOMs to 25, 50, 100, or 250. For more information on VDOM licences, see [“Virtual Domain Licensing” on page 1892](#).

More flexible MSSP configurations

If you are a managed security and service provider (MSSP), VDOMs are fundamental to your business. As a service provider you have multiple customers, each with their own needs and service plans. VDOMs allow you to have a separate configuration for each customer, or group of customers; you can have up to 250 VDOMs configured on a FortiGate unit on high end models. See [“Virtual Domain Licensing” on page 1892](#).

Not only does this provide the exact level of service needed by each customer, but administration of the FortiGate unit is easier as well - you can provide uninterrupted service generally with immediate changes as required. Most importantly, it allows you to only use the resources that each customer needs. Inter-VDOM links allow you to customize the level of interaction you need between each of your customers and your administrators. See [“Inter-VDOM routing” on page 1941](#).

Enabling and accessing Virtual Domains

While Virtual Domains are essentially the same as your regular FortiGate unit for menu configuration, CLI command structure, and general task flow, there are some small differences.

After first enabling VDOMs on your FortiGate unit, you should take the time to familiarize yourself with the interface. This section will help walk you through virtual domains.

This section includes:

- [Enabling Virtual Domains](#)
- [Viewing the VDOM list](#)
- [Global and per-VDOM settings](#)
- [Resource settings](#)
- [Virtual Domain Licensing](#)
- [Logging in to VDOMs](#)

Enabling Virtual Domains

Using the default admin administration account, you can enable or disable VDOM operation on the FortiGate unit.

To enable VDOM configuration - web-based manager

- 1 Log in with a super_admin account.
- 2 Go to *System > Dashboard > Status*.
- 3 Under *System Information > Virtual Domain*, select *Enable* and confirm your selection.

The FortiGate unit logs off all sessions. You can now log in again as admin. For more information, see [“Administrators in Virtual Domains” on page 1898](#).

Figure 177: System Information

▼ System Information	
Host Name	FG200B3909600899 [Change]
Serial Number	FG200B3909600899
HA Status	Standalone [Configure]
System Time	Wed May 4 11:10:43 2011 [Change]
Firmware Version	v4.0,build0441,110420 (MR3) [Update]
System Configuration	Last Backup: N/A [Backup] [Restore]
Current Administrator	admin [Change Password] /1 in Total [Details]
Uptime	10 day(s) 1 hour(s) 14 min(s)
Virtual Domain	Enabled [Disable]

VDOMs are enabled

To enable VDOM configuration - CLI

```
config system global
    set vdom-admin enable
end
```

Changes to the web-based manager and CLI

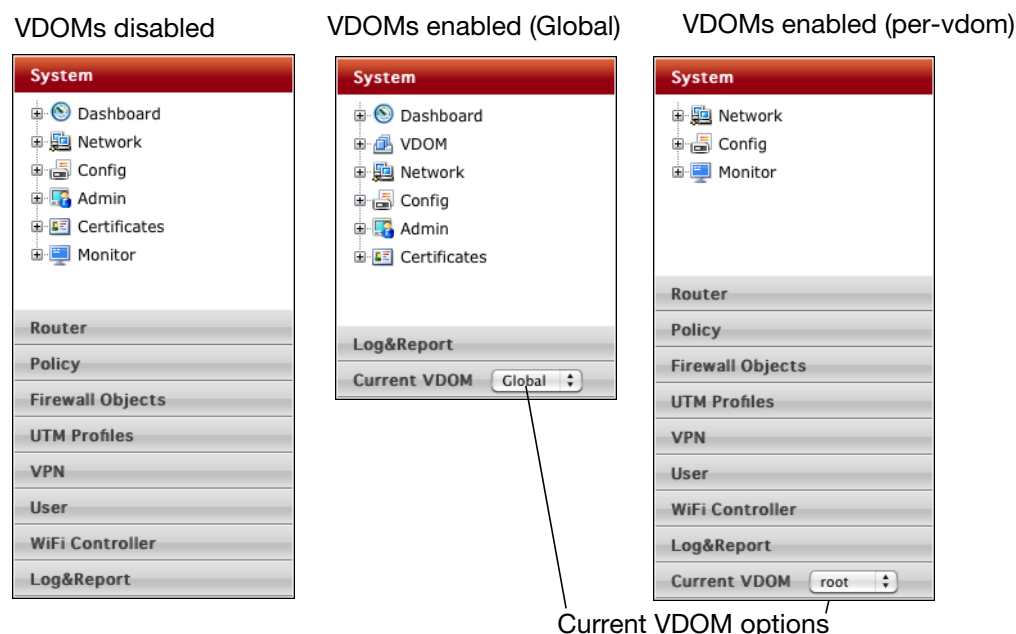
When Virtual Domains are enabled, your FortiGate unit will change. The changes will be visible in both the web-based manager and CLI, just the web-based manager, or just the CLI.

When enabling VDOMs, the web-based manager and the CLI are changed as follows:

- Global and per-VDOM configurations are separated. This is indicated in the Online Help by Global and VDOM icons. See [“Global and per-VDOM settings” on page 1881](#).
- Only admin accounts using the super_admin profiles can view or configure global options. See [“Administrators in Virtual Domains” on page 1898](#).
- Admin accounts using the super_admin profile can configure all VDOM configurations.
- All other administrator accounts can configure only the VDOM to which they are assigned.

The following changes are specific to the web-based manager:

- The *System > Dashboard > Status* view is different for VDOMs.
- In the Global view, the System menu includes a VDOM sub-menu.
- For admin accounts using the super_admin profile, a new control called Current VDOM is added at the bottom of the left menu. It indicates which VDOM you are in, and allows you to easily select either another VDOM or Global settings to configure. See [Figure 178 on page 1879](#).

Figure 178: Menu with VDOMs disabled, at the global level, and VDOM level

In the CLI, admin accounts using the super_admin profile must specify either the global or a VDOM-specific shell before entering commands:

- To change FortiGate unit system settings, from the top level you must first enter `config global` before entering commands.
- To change VDOM settings, from the top level you must first enter `config vdom` `edit <vdom_name>` before entering your commands for that VDOM. For information on which commands are global and which are per-VDOM, see [“Global and per-VDOM settings” on page 1881](#).

Changes to FortiGate unit settings

Settings configured outside of a VDOM are called global settings. These settings affect the entire FortiGate unit and include areas such as interfaces, HA, maintenance, some antivirus, and some logging. In general, any unit settings that should only be changed by the top level administrator are global settings.

Settings configured within a VDOM are called VDOM settings. These settings affect only that specific VDOM and include areas such as operating mode, routing, firewall, VPN, some antivirus, some logging, and reporting.

For more information, see [“Global and per-VDOM settings” on page 1881](#).

Viewing the VDOM list

The VDOM list shows all virtual domains, their status, and which VDOM is the management VDOM. It is accessible if you are logged in on an administrator account with the super_admin profile such as the “admin” administrator account.

In the VDOM list you can create or delete VDOMs, edit VDOMs, change the management VDOM, and enable or disable VDOMs.



The root domain cannot be disabled, even if it is not the management VDOM.

To view the VDOM list

- 1 For *Current VDOM*, select *Global*.
- 2 Go to *System > VDOM > VDOM*.

Figure 179: List of VDOMs

Name	Operation Mode	Interfaces	Enable	Comments	Ref.
root	NAT	example_wan , modem , ssl.root , wan1 , wan2 , wlan , wlan_employee	✓		0
vdom1	NAT	dmz , internal , ssl.vdom1	✓		0
vdom2	NAT	ssl.vdom2	✗		0

Create New	Select to add a new VDOM. See “Creating a Virtual Domain” on page 1895 .
Edit	Select to change an existing selected VDOM.
Delete	Select to delete the selected VDOM. See “Deleting a VDOM” on page 1897 .
Switch Management	Select to switch the management VDOM. Also shows the current management VDOM. You must select an active non-management VDOM before this option becomes available. See “Changing the management virtual domain” on page 1903 .
Selected	When checked, this checkbox indicates this VDOM has been selected. Nearly all operations such as Edit, Delete, and Switch Management require a VDOM to first be selected.
Name	The name of the VDOM. VDOMs are listed in alphabetical order. When the VDOM is active, you can select the VDOM name to enter that VDOM. See “Enabling and accessing Virtual Domains” on page 1877 .
Operation Mode	Indicates the operation mode as either NAT (for NAT/Route mode) or TP (for Transparent mode).
Interfaces	The interfaces associated with this VDOM. Each VDOM also includes an interface that starts with “ssl.” that is created by default.
Enable	A green checkmark indicates this VDOM is active. See “Disabling a Virtual Domain” on page 1896 . A grey X indicated this VDOM is disabled. See “Disabling a Virtual Domain” on page 1896 .

Comments	Comments entered when the VDOM was created are displayed here.
Ref.	The number of references to this VDOM in the configuration.

Global and per-VDOM settings

Settings configured outside of a VDOM are called global settings. These settings affect the entire FortiGate unit and include areas such as interfaces, HA, maintenance, some antivirus, and some logging. In general, any unit settings that should only be changed by the top level administrator are global settings.

Settings configured within a VDOM are called VDOM settings. These settings affect only that specific VDOM and include areas such as operating mode, routing, firewall, VPN, some antivirus, some logging, and reporting.

When Virtual Domains are not enabled, the entire FortiGate unit is effectively a single VDOM. Per-VDOM limits apply. For some resource types, the global limit cannot be reached with only one VDOM.

Some FortiGate unit documentation indicates which parts of the web-based manager, or the CLI are global and which are per-VDOM using the icons shown below. These icons are also present in the Online Help, available on your FortiGate unit.

Figure 180: Global and VDOM icons



For more information on CLI commands, see the [FortiGate CLI Reference](#).

This section includes:

- [Global settings - web-based manager](#)
- [Per-VDOM settings - web-based manager](#)
- [Global settings - CLI](#)
- [Per-VDOM settings - CLI](#)

Global settings - web-based manager

The following table lists commands in the web-based manager that are considered global settings when VDOMs are enabled.

The following configuration settings affect all virtual domains. When virtual domains are enabled, only accounts with the default super_admin profile can access global settings.

Table 109: Global configuration settings

System	Dashboard > Status - Host name Dashboard > Status - HA Status Dashboard > Status - System Time Dashboard > Status - Firmware version Dashboard > Status - Configuration backup and restore VDOM > VDOM - list VDOM > VDOM - edit VDOM (mode and resources) VDOM > Global Resources
---------------	---

Table 109: Global configuration settings (Continued)

	Network > Interfaces Network > DNS - DNS and DDNS settings Config > HA Config > SNMP Config > Replacement Message - messages and images Config > Firmware Config > FortiGuard - configuration Config > Advanced - scripts, USB Auto-install, debug log download Admin > Administrators Admin > Admin Profile Admin > Central Management - configuration Admin > Settings - web administration ports, password policy, display settings, timeouts, LCD panel Certificates - local, remote, and CA certificates, CRLs
Log&Report	Log Config - Log Setting and Alert E-mail

Per-VDOM settings - web-based manager

The following table lists commands in the web-based manager that are considered per-VDOM settings when VDOMs are enabled.

Table 110: VDOM configuration settings

System	Dashboard > Status - read-only except for administrator password Network > Interface (and zones) Network > DNS Server Network > DHCP Server Network > Explicit Proxy Network > Routing Table (Transparent mode only) Network > Modem Config > Replacement Message (messages and images) Config > Replacement Message Group Config > Tag Management Monitor > DHCP Monitor Monitor > Modem Monitor
Router	All settings, including dead gateway detection
Policy	All settings
Firewall Objects	All settings
UTM Profiles	All settings

Table 110: VDOM configuration settings (Continued)

VPN	All settings
User	All settings
WiFi Controller	All settings
Log&Report	Log & Archive Access for Events, UTM, Traffic Log & Archive Access - Vulnerability Scan Log Log Config > Log Setting Log Config > Alert E-mail
Monitor	Logging Monitor

Global settings - CLI

The following table lists commands in the web-based manager that are considered global settings when VDOMs are enabled.

From a super_admin profile account, use this command to configure features that apply to the complete FortiGate unit including all virtual domains. Virtual domain configuration (vdom-admin) must be enabled first.

This command syntax shows how you access the commands within config global. For information on these commands, refer to the relevant sections in this Reference. If there are multiple versions of the same command with a “2” or “3” added, the additional commands are not listed but fall under the unnumbered command of the same name.

```
config global
  config antivirus heuristic
  config antivirus quarfilepattern
  config antivirus service
  config application name
  config dlp settings
  config endpoint-control app-detect
  config firewall ssl
  config gui console
  config ips decoder
  config ips global
  config ips rule
  config log fortianalyzer setting
  config log fortiguard setting
  config log memory global-setting
  config log syslogd filter
  config log syslogd setting
  config log webtrends ...
  config spamfilter fortishield
  config spamfilter options
  config system accprofile
  config system admin
  config system alertemail
  config system amc
  config system auto-install
  config system autoupdate ...
  config system aux
  config system bug-report
  config system central-management
```

```
config system chassis-loadbalance
config system console
config system ddns
config system dialinsvr
config system dns
config system dynamic-profile
config system fips-cc
config system fortiguard
config system fortiguard-log
config system global
config system ha
config system interface
config system npu
config system ntp
config system password-policy
config system replacemsg ...
config system replacemsg-image
config system resource-limits
config system session-helper
config system session-sync
config system sflow
config system snmp ...
config system switch-interface
config system tos-based-priority
config system vdom-link
config system vdom-property
config vpn certificate ...
config wanopt storage
config webfilter fortiguard
config wireless-controller global
config wireless-controller timers
config wireless-controller vap
execute backup
execute batch
execute central-mgmt
execute cfg reload
execute cfg save
execute cli check-template-status
execute cli status-msg-only
execute date
execute disconnect-admin-session
execute disk
execute enter
execute factoryreset
execute firmware-list
execute formatlogdisk
execute forticlient
execute fortiguard-log
execute ha disconnect
execute ha manage
execute ha synchronize
execute log ...
execute log-report
execute reboot
```

```
execute report-config
execute restore
execute revision
execute router ... (except clear)
execute scsi-dev
execute send-fds-statistics
execute set-next-reboot
execute sfp-mode-sgmii
execute shutdown
execute tac
execute time
execute update-ase
execute update-av
execute update-ips
execute update-netscan
execute update-now
execute upload
execute usb-disk
execute vpn certificate ...
execute wireless-controller ... (except reset-wtp)
get firewall vip ...
end
```

Per-VDOM settings - CLI

The following table lists commands in the web-based manager that are considered global settings when VDOMs are enabled.

From the super admin account, use this command to add and configure virtual domains. The number of virtual domains you can add is dependent on the FortiGate model. Virtual domain configuration (vdom-admin) must be enabled.

Once you add a virtual domain you can configure it by adding zones, firewall policies, routing settings, and VPN settings. You can also move physical interfaces from the root virtual domain to other virtual domains and move VLAN subinterfaces from one virtual domain to another.

By default all physical interfaces are in the root virtual domain. You cannot remove an interface from a virtual domain if the interface is part of any of the following configurations:

- routing
- proxy arp
- DHCP server
- zone
- firewall policy
- redundant pair
- link aggregate (802.3ad) group

Delete these objects, or modify them, to be able to remove the interface.

This command syntax shows how you access the commands within a VDOM. Refer to the relevant sections in this Reference for information on these commands.

```
config vdom
  edit <vdom_name>
    config antivirus profile
```

```
config antivirus quarantine
config antivirus settings
config application list
config application rule-settings
config dlp ... (except settings)
config endpoint-control app-detect
config endpoint-control profile
config endpoint-control settings
config firewall ... (except ssl)
config ftp-proxy
config icap
config imp2p
config ips DoS
config ips custom
config ips rule-settings
config ips sensor
config ips settings
config log custom-field
config log disk
config log eventfilter
config log fortianalyzer
config log gui
config log memory
config log syslogd
config log trafficfilter
config log visibility
config netscan
config router
config spamfilter ... (except fortishield and options)
config system 3g-modem
config system admin
config system arp-table
config system carrier-endpoint-translation
config system dhcp ...
config system dhcp6 ...
config system dns-database
config system dns-server
config system gre-tunnel
config system interface
config system ipv6-tunnel
config system modem
config system monitors
config system object-tag
config system proxy-arp
config system replacemsg-group
config system session-ttl
config system settings
config system sit-tunnel
config system switch-interface
config system wccp
config system zone
config user ...
config voip
config vpn ...
```

```
config wanopt
config web-proxy
config webfilter (except fortiguard)
config wireless-controller (except global and timers)
execute backup
execute clear system arp table
execute cli check-template-status
execute cli status-msg-only
execute dhcp lease-clear
execute dhcp lease-list
execute dhcp6 lease-clear
execute dhcp6 lease-list
execute enter
execute fortitoken ...
execute fsso refresh
execute interface dhcpclient-renew
execute interface pppoe-reconnect
execute log ...
execute log-report ...
execute modem dial
execute modem hangup
execute modem trigger
execute mrouter clear
execute netscan ...
execute ping, ping6
execute ping-options, ping6-options
execute restore
execute revision
execute router clear bgp
execute router clear ospf process
execute router restart
execute sfp-mode-sgmii
execute ssh
execute tac
execute telnet
execute traceroute
execute tracert6
execute upload
execute usb-disk
execute vpn ipsec tunnel
execute vpn sslvpn ...
execute wireless-controller reset-wtp
next
edit <another_vdom>
  config ...
  execute ...
end
end
```

For more information, see “Global and per-VDOM settings” on page 1881.

Resource settings

Your FortiGate unit has a limited amount of hardware resources such as memory, disk storage, CPU operations. When Virtual Domains are disabled, this limit is not a major concern because all sessions, users, and other processes share all the resources equally.

When using Virtual Domains, hardware resources can be divided differently between Virtual Domains as they are needed. Also minimum levels of resources can be set so that no Virtual Domain will suffer a complete lack of resources.

For example if one VDOM has only a web server and logging server connected, and a second VDOM has an internal network of 20 users these two VDOMs will require different levels of resources. The first VDOM will require many sessions but no user accounts. This compares to the second VDOM where user accounts and management resources are required, but fewer sessions.

Using the global and per-VDOM resource settings, you can customize the resources allocated to each VDOM to ensure the proper level of service is maintained on each VDOM.

This section includes:

- [Global resource settings](#)
- [Per-VDOM resource settings](#)

Global resource settings

Global Resources apply to the whole FortiGate unit. They represent all of the hardware capabilities of your unit. By default the values are set to their maximum values. These values vary by your model due to each model having differing hardware capabilities.

It can be useful to change the maximum values for some resources to ensure there is enough memory available for other resources that may be more important to your configuration.

To use the earlier example, if your FortiGate unit is protecting a number of web servers and other publicly accessible servers you would want to maximize the available sessions and proxies while minimizing other settings that are unused such as user settings, VPNs, and dial-up tunnels.

Global Resources are only configurable at the global level, and only the admin account has access to these settings.

Note that global resources, such as the log disk quota resource, will only be visible if your FortiGate unit hardware supports those resources, such as having a hard disk to support the log disk resource.

Figure 181: Global Resources- web-based manager

Edit Reset to default value				
<input type="checkbox"/>	Resource	Configured Maximum	Default Maximum	Current Usage
<input type="checkbox"/>	Sessions	0	0	26
<input type="checkbox"/>	VPN IPsec Phase1 Tunnels	10000	10000	0
<input type="checkbox"/>	VPN IPsec Phase2 Tunnels	10000	10000	0
<input type="checkbox"/>	Dial-up Tunnels	0	0	0
<input type="checkbox"/>	Firewall Policies	100000	100000	3
<input type="checkbox"/>	Firewall Addresses	20000	20000	11
<input type="checkbox"/>	Firewall Address Groups	10000	10000	0
<input type="checkbox"/>	Firewall Custom Services	0	0	0
<input type="checkbox"/>	Firewall Service Groups	0	0	0
<input type="checkbox"/>	Firewall One-time Schedules	0	0	0
<input type="checkbox"/>	Firewall Recurring Schedules	0	0	5
<input type="checkbox"/>	Local Users	0	0	0
<input type="checkbox"/>	User Groups	0	0	0
<input type="checkbox"/>	SSL VPN	0	0	0
<input type="checkbox"/>	Concurrent web proxy users	2000	2000	0
<input type="checkbox"/>	log disk quota	0	0	0

To view global resource settings - web-based manager

- 1 For *Current VDOM*, select *Global*.
- 2 Select *System > VDOM > Global Resources*.

The following information is displayed:

Edit	Select to edit the <i>Configured Maximum</i> value for a single selected <i>Resource</i> . If multiple <i>Resources</i> are selected, <i>Edit</i> is not available.
Reset to default value	Select to return one or more selected <i>Resources</i> to factory default settings.
Checkbox	Select a <i>Resource</i> for editing or resetting to default values.
Resource	The name of the available global resources.
Configured Maximum	The currently configured maximum for this resource. This value can be changed by selecting the <i>Resource</i> and editing it.
Default Maximum	The factory configured maximum value for this resource. You cannot set the <i>Configured Maximum</i> higher than the <i>Default Maximum</i> .
Current Usage	The amount of this resource that is currently being used. This value is useful for determining when and if you may need to adjust <i>Configured Maximum</i> values for some resources on your FortiGate unit.

To view global resource settings - CLI

```
config global
  config system resource-limits
  get
```

When viewing the global resource limits in the CLI, the output appears similar to:

```
FGT1000A (global) # config system resource-limits
FGT1000A (resource-limits) # get

session                : 0
ipsec-phase1           : 10000
ipsec-phase2           : 10000
dialup-tunnel          : 0
firewall-policy        : 100000
firewall-address       : 20000
firewall-addrgroup    : 10000
custom-service         : 0
service-group          : 0
onetime-schedule       : 0
recurring-schedule    : 0
user                   : 0
user-group             : 0
sslvpn                 : 0
webproxy               : 2000
```



For explicit proxies when configuring limits on the number of concurrent users, you need to allow for the number of users based on their authentication method. Otherwise you may run out of user resources prematurely.

- Each session-based authenticated user is counted as a single user using their authentication membership (RADIUS, LDAP, FSAE, local database etc.) to match users in other sessions. So one authenticated user in multiple sessions is still one user.
- For all other situations, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

Per-VDOM resource settings

Global resources apply to resources shared by the whole FortiGate unit. Per-VDOM resources are specific to only one Virtual Domain.

By default all the per-VDOM resource settings are set to no limits. This means that any single VDOM can use up all the resources of the entire FortiGate unit if it needs to do so. This would starve the other VDOMs for resources to the point where they would be unable to function. For this reason, it is recommended that you set some maximums on resources that are most vital to your customers.

Each Virtual Domain has its own resource settings. These settings include both maximum, and minimum levels. The maximum level is the highest amount of that resource that this VDOM can use if it is available on the FortiGate unit. Minimum levels are a guaranteed level that this minimum level of the resource will always be available no matter what the other VDOMs may be using.

Figure 182: per-VDOM resources - web-based manager

Resource Usage			
Resource	Maximum	Guaranteed	Current
Sessions	0	0	24
VPN IPsec Phase1 Tunnels	0	0	0
VPN IPsec Phase2 Tunnels	0	0	0
Dial-up Tunnels	0	0	0
Firewall Policies	0	0	3
Firewall Addresses	0	0	3
Firewall Address Groups	0	0	0
Firewall Custom Services	0	0	0
Firewall Service Groups	0	0	0
Firewall One-time Schedules	0	0	0
Firewall Recurring Schedules	0	0	1
Local Users	0	0	0
User Groups	0	0	0
SSL VPN	0	0	0
Concurrent web proxy users	0	0	0
log disk quota	0	0	0

For example your FortiGate unit has ten VDOMs configure. vdom1 has a maximum of 5000 sessions and a minimum of 1000 sessions. If the FortiGate unit has a global maximum of 20,000 sessions, it is possible that vdom1 will not be able to reach its 5000 session upper limit. However, at all times vdom1 is guaranteed to have 1000 sessions available that it can use. On the other hand, if the remaining nine VDOMs use only 1000 sessions each, vdom1 will be able to reach its maximum of 5000.

To view per-VDOM resource settings - web-based manager

- 1 For *Current VDOM*, select *Global*.
- 2 Select *System > VDOM > VDOM*.
- 3 Select the `root` VDOM, and select *Edit*. Adjust the settings in the *Resource Usage* section of the page.

Resource	Name of the resource. Includes dynamic and static resources.
Maximum	Override the global limit to reduce the amount of each resource available for this VDOM. The maximum must be the same as or lower than the global limit. The default value is 0, which means the maximum is the same as the global limit. Note: If you set the maximum resource usage for a VDOM you cannot reduce the default maximum global limit for all VDOMs below this maximum.
Guaranteed	Enter the minimum amount of the resource available to this VDOM regardless of usage by other VDOMs. The default value is 0, which means that an amount of this resource is not guaranteed for this VDOM.
Current	The amount of the resource that this VDOM currently uses.

- 4 Select *OK*.

To view per-VDOM resource settings - CLI

```
config global
  config system vdom-property
    edit root
  get
```

When viewing the per-VDOM resource limits in the CLI, the output appears similar to the following. Note that the first two lines are not part of the resource limits. In the CLI, the first number is the maximum value, and the second number is the guaranteed minimum.

```
FGT1KA3607500810 (vdom-property) # edit root
FGT1KA3607500810 (root) # get

name                : root
description         : property limits for vdom root
session             : 0 0
ipsec-phase1        : 0 0
ipsec-phase2        : 0 0
dialup-tunnel       : 0 0
firewall-policy     : 0 0
firewall-address    : 0 0
firewall-addrgrp    : 0 0
custom-service      : 0 0
service-group       : 0 0
onetime-schedule    : 0 0
recurring-schedule  : 0 0
user                : 0 0
user-group          : 0 0
sslvpn              : 0 0
webproxy            : 0 0
```

Virtual Domain Licensing

All FortiGate models except the FortiGate-30B and FortiWiFi-30B models support VDOMs. By default 10 VDOMs are available.

For FortiGate models numbered 1240 and higher, you can purchase a license key to increase the maximum number of VDOMs. Model 1240B supports up to 25 VDOMs. Most Enterprise and Large Enterprise models can support 250 VDOMs. Chassis-based models can support up to 3000 VDOMs. For specific information, see the product data sheet.

Configuring 250 or more VDOMs will result in reduced system performance. See [“FortiGate unit running very slowly” on page 1977](#).



Your FortiGate unit has limited resources that are divided among all configured VDOMs. These resources include system memory and CPU. You cannot run Unified Threat Management (UTM) features when running 250 or more VDOMs. UTM features include proxies, web filtering, and antivirus—your FortiGate unit can provide only basic firewall functionality.



It is important to backup your configuration before upgrading the VDOM license on your FortiGate unit or units, especially with FortiGate units in HA mode.

To obtain a VDOM license key

- 1 Log in with a super_admin account.
- 2 Go to *System > Dashboard > Status*.
- 3 Record your FortiGate unit serial number as shown in “System Information” on page 1878.
- 4 Under *License Information > Virtual Domain*, select *Purchase More*.



If you do not see the *Purchase More* option on the System Dashboard, your FortiGate model does not support more than 10 VDOMs.

Figure 183: VDOM License Information

License Information		
Support Contract		
Registration	Unreachable	ⓘ
FortiGuard Services		
AntiVirus	Unreachable [Configure]	ⓘ
AV Definitions	9.00795 (Updated 2008-12-08) [Update]	
Extended set	0.00000 (Updated 2003-01-01)	
Intrusion Protection	Unreachable [Configure]	ⓘ
IPS Definitions	2.00720 (Updated 2009-12-01) [Update]	
Vulnerability Compliance and Management	Unreachable [Configure]	ⓘ
VCM Plugin	1.00098 (Updated 2010-02-11) [Update]	
Web Filtering	Unreachable [Configure]	ⓘ
Email Filtering	Unreachable [Configure]	ⓘ
Analysis & Management Service	Unreachable	ⓘ
Services Account ID	[Change]	
Virtual Domain		
VDOMs Allowed	10 [Purchase More]	
Endpoint Security		
FortiClient Software	Unreachable	
Application Signature Package	1.131 (Updated 2010-02-16)	

Purchase a larger VDOM license

- 5 You will be taken to the Fortinet customer support web site where you can log in and purchase a license key for 25, 50, 100, 250, or 500 VDOMs.
- 6 When you receive your license key, go to the Dashboard and select *Upload License* under *License Information, Virtual Domains*.
- 7 In the *Input License Key* field, enter the 32-character license key you received from Fortinet customer support.
- 8 Select *Apply*.

To verify the new VDOM license, in global configuration go to *System > Dashboard*. Under *License Information, Virtual Domains* the maximum number of VDOMs allowed is shown.



VDOMs created on a registered FortiGate unit are recognized as real devices by any connected FortiAnalyzer unit. The FortiAnalyzer unit includes VDOMs in its total number of registered devices. For example, if three FortiGate units are registered on the FortiAnalyzer unit and they contain a total of four VDOMs, the total number of registered FortiGate units on the FortiAnalyzer unit is seven. For more information, see the [FortiAnalyzer Administration Guide](#).

Logging in to VDOMs

Only super_admin administrator accounts can access all global settings on the FortiGate unit and all of the VDOMs as well. Other administrator accounts can access and configure only their single VDOM and they must connect to an interface that is part of that VDOM. For example, administratorB is the admin for vdomB. If he tries to log into vdomA, or an interface that is part of vdomA he will not be able to log on. For more information on administrators in VDOMs, see [“Administrators in Virtual Domains” on page 1898](#).

Management services communicate using the management VDOM, which is the root VDOM by default. For more information, see [“Changing the management virtual domain” on page 1903](#).



Management traffic requires an interface that has access to the Internet. If there is no interface assigned to the VDOM containing the management traffic, services including updates will not function. For more information, see [“Changing the management virtual domain” on page 1903](#).

To access a VDOM with a super_admin account - web-based manager

- 1 Log in with a super_admin account.
- 2 In *Current VDOM*, select the VDOM to configure.
The system network page for that VDOM opens.
- 3 When you have finished configuring the VDOM, you can
 - in *Current VDOM*, select *Global* to return to global configuration
 - log out.

To access a VDOM with a super_admin account - CLI

With the super_admin, logging into the CLI involves also logging into the specific VDOM. If you need a reminder, use `edit ?` to see a list of existing VDOMs before you editing a VDOM.



If you misspell a VDOM you are trying to switch to, you will create a new VDOM by that name. Any changes you make will be part of the new VDOM, and not the intended VDOM. If you are having problems where your changes aren't visible, back up to the top level and use `edit ?` to see a list of VDOMs to ensure this has not happened. If it has happened, see [“Deleting a VDOM” on page 1897](#).

```
config vdom
edit ?
edit <chosen_vdom>
..
<enter vdom related commands>
..
end
exit
```

To access a VDOM with a non super_admin account - web-based manager

- 1 Connect to the FortiGate unit using an interface that belongs to the VDOM to be configured.

- 2 Log in using an administrator account that has access to the VDOM.

The main web-based manager page opens. From here you can access VDOM-specific settings.

To access a VDOM with a non-super_admin account - CLI

A non-super_admin account has access to only one VDOM and must log in through an interface that belongs to the same VDOM.

```
Login: regular_admin
Password: <password>
..
<enter vdom related commands>
..
exit
```

Configuring Virtual Domains

Only a super_admin administrator account such as the default “admin” account can create, disable, or delete VDOMs. That account can create additional administrators for each VDOM.

This section includes:

- [Creating a Virtual Domain](#)
- [Disabling a Virtual Domain](#)
- [Deleting a VDOM](#)
- [Administrators in Virtual Domains](#)

Creating a Virtual Domain

Once you have enabled Virtual Domains on your FortiGate unit, you can create additional Virtual Domains beyond the default root Virtual Domain.

By default new Virtual Domains are set to NAT/Route operation mode. If you want a Virtual Domain to be in Transparent operation mode, you must manually change it. See [“Virtual Domains in Transparent mode” on page 1921](#).

You can name new Virtual Domains as you like with the following restrictions:

- only letters, numbers, “-”, and “_” are allowed
- no more than 11 characters are allowed
- no spaces are allowed
- VDOMs cannot have the same names as interfaces, zones, switch interfaces, or other VDOMs.



When creating large numbers of VDOMs (up to 250), you cannot enable advanced features such as proxies, web filtering, and antivirus due to limited FortiGate unit resources. Also when creating large numbers of VDOMs, you may experience reduced performance for the same reason.

To create a VDOM - web-based manager

- 1 Log in with a super_admin account.
- 2 Go to *System > Dashboard > Status* and ensure that Virtual Domains are enabled. If not, see [“Enabling and accessing Virtual Domains” on page 1877](#).

- 3 Select *System > VDOM > VDOM*.
- 4 Select *Create New*.
- 5 Enter a unique name for your new VDOM.
- 6 Enter a short and descriptive comment to identify this VDOM.
- 7 Select *OK*.

Repeat Steps 4 through 7 to add additional VDOMs.

To create a VDOM - CLI

```
config vdom
  edit <new_vdom_name>
end
```



If you want to edit an existing Virtual Domain in the CLI, and mistype the name a new Virtual Domain will be created with this new misspelled name. If you notice expected configuration changes are not visible, this may be the reason. You should periodically check your VDOM list to ensure there are none of these misspelled VDOMs present.

Disabling a Virtual Domain

The status of a VDOM can be Enabled, or Disabled.

Active status VDOMs can be configured. Active is the default status when a VDOM is created. The management VDOM must be an Active VDOM. For more information on the management VDOM, see [“Changing the management virtual domain” on page 1903](#).

Disabled status VDOMs are considered “offline”. The configuration remains, but you cannot use the VDOM, and only the super_admin administrator can view it. You cannot delete a disabled VDOM without first enabling it, and removing references to it like usual—there is no *Delete* icon for disabled status VDOMs. You can assign interfaces to a disabled VDOM. See [“Deleting a VDOM” on page 1897](#).

The following procedures show how to disable a VDOM called “test-vdom”.

To disable a VDOM - web-based manager

- 1 In *Current VDOM*, select *Global*.
 - 2 Go to *System > VDOM > VDOM*.
 - 3 Open the VDOM for editing.
 - 4 Ensure *Enable* is not selected and then select *OK*.
- The VDOM’s Enable icon in the VDOM list is a grey X.

To disable a VDOM - CLI

```
config vdom
  edit test-vdom
    config system settings
      set status disable
    end
end
```

To enable a VDOM - web-based manager

- 1 For *Current VDOM*, select *Global*.
- 2 Go to *System > VDOM > VDOM*.
- 3 Open the VDOM for editing.

- 4 Ensure *Enable* is selected and then select *OK*.

The VDOM's Enable icon in the VDOM list is a green checkmark.

To enable a VDOM - CLI

```
config vdom
  edit test-vdom
    config system settings
      set status enable
    end
  end
```

Deleting a VDOM

Deleting a VDOM removes it from the FortiGate unit configuration.

Before you can delete a VDOM, all references to it must be removed. This includes any objects listed in “[Per-VDOM settings - web-based manager](#)” on page 1882. If there are any references to the VDOM remaining, you will see an error message and not be able to delete the VDOM.

The VDOM must also be enabled. A disabled VDOM cannot be deleted. You cannot delete the root VDOM or the management VDOM.



Before deleting a VDOM, a good practice is to reset any interface referencing that VDOM to its default configuration, with “root” selected as the Virtual Domain.

The following procedures show how to delete the `test-vdom` VDOM.

To delete a VDOM - web-based manager

- 1 For *Current VDOM*, select *Global*.
- 2 Go to *System > VDOM > VDOM*.
- 3 Select the check box for the VDOM and then select the *Delete* icon.
If the *Delete* icon is not active, there are still references to the VDOM that must first be removed. The *Delete* icon is available when all the references to this VDOM are removed.
- 4 Confirm the deletion.

To delete a VDOM - CLI

```
config vdom
  delete test-vdom
end
```

Removing references to a VDOM

When you are doing to delete a VDOM, all references to that VDOM must first be removed. It can be difficult to find all the references to the VDOM. This section provides a list of common objects that must be removed before a VDOM can be deleted, and a CLI command to help list the dependencies.

Interfaces are an important part of VDOMs. If you can move all the interfaces out of a VDOM, generally you will be able to delete that VDOM.

Common objects that refer to VDOMs

When you are getting ready to delete a VDOM check for, and remove the following objects that refer to that VDOM or its components:

- Routing - both static and dynamic routes
- Firewall addresses, policies, groups, or other settings
- UTM
- VPN configuration
- Users or user groups
- Logging
- DHCP servers
- Network interfaces, zones, custom DNS servers
- VDOM Administrators

Administrators in Virtual Domains

When Virtual Domains are enabled, permissions change for administrators. Administrators are now divided into per-VDOM administrators, and `super_admin` administrators. Only `super_admin` administrator accounts can create other administrator accounts and assign them to a VDOM.

This section includes:

- [Administrator VDOM permissions](#)
- [Creating administrators for Virtual Domains](#)
- [Virtual Domain administrator dashboard display](#)

Administrator VDOM permissions

Different types of administrator accounts have different permissions within VDOMs. For example, if you are using a `super_admin` profile account, you can perform all tasks. However, if you are using a regular admin account, the tasks available to you depend on whether you have read only or read/write permissions. The following table shows what tasks can be performed by which administrators.

Table 111: Administrator VDOM permissions

Tasks	Regular administrator account		Super_admin profile administrator account
	Read only permission	Read/write permission	
View global settings	yes	yes	yes
Configure global settings	no	no	yes
Create or delete VDOMs	no	no	yes
Configure multiple VDOMs	no	no	yes
Assign interfaces to a VDOM	no	no	yes
Revision Control Backup and Restore	no	no	yes

Table 111: Administrator VDOM permissions

Create VLANs	no	yes - for 1 VDOM	yes - for all VDOMs
Assign an administrator to a VDOM	no	no	yes
Create additional admin accounts	no	yes - for 1 VDOM	yes - for all VDOMs
Create and edit protection profiles	no	yes - for 1 VDOM	yes - for all VDOMs

The only difference in admin accounts when VDOMs are enabled is selecting which VDOM the admin account belongs to. Otherwise, by default the administration accounts are the same as when VDOMs are disabled and closely resemble the `super_admin` account in their privileges.

Creating administrators for Virtual Domains

Using the admin administrator account, you can create additional administrator accounts and assign them to VDOMs.



The newly-created administrator can access the FortiGate unit only through network interfaces that belong to their assigned VDOM or through the console interface. The network interface must be configured to allow management access, such as HTTPS and SSH. Without these in place, the new administrator will not be able to access the FortiGate unit and will have to contact the `super_admin` administrator for access.

The following procedure creates a new Local administrator account called `admin_sales` with a password of `fortinet` in the `sales` VDOM using the `admin_prof` default profile.

To create an administrator for a VDOM - web-based manager

- 1 Log in with a `super_admin` account.
- 2 Go to *System > Admin > Administrators*.
- 3 Select *Create New*.
- 4 Select *Regular* for Type, as you are creating a Local administrator account.
- 5 If this admin will be accessing the VDOM from a particular IP address or subnet, enter it in *Trusted Host #1*. See [“Using trusted hosts” on page 1900](#).
- 6 Select `prof_admin` for the *Admin Profile*.
- 7 Select `sales` from the list of *Virtual Domains*.
- 8 Select *OK*.

To create administrators for VDOMs - CLI

```
config global
  config system admin
    edit <new_admin_name>
      set vdom <vdom_for_this_account>
      set password <pwd>
      set accprofile <an_admin_profile>
      ...
    end
```

Using trusted hosts

Setting trusted hosts for all of your administrators increases the security of your network by further restricting administrative access. In addition to knowing the password, an administrator must connect only through the subnet or subnets you specify. You can even restrict an administrator to a single IP address if you define only one trusted host IP address with a netmask of 255.255.255.255.

When you set trusted hosts for all administrators, the FortiGate unit does not respond to administrative access attempts from any other hosts. This provides the highest security. If you leave even one administrator unrestricted, the unit accepts administrative access attempts on any interface that has administrative access enabled, potentially exposing the unit to attempts to gain unauthorized access.

The trusted hosts you define apply both to the web-based manager and to the CLI when accessed through Telnet or SSH. CLI access through the console is not affected.

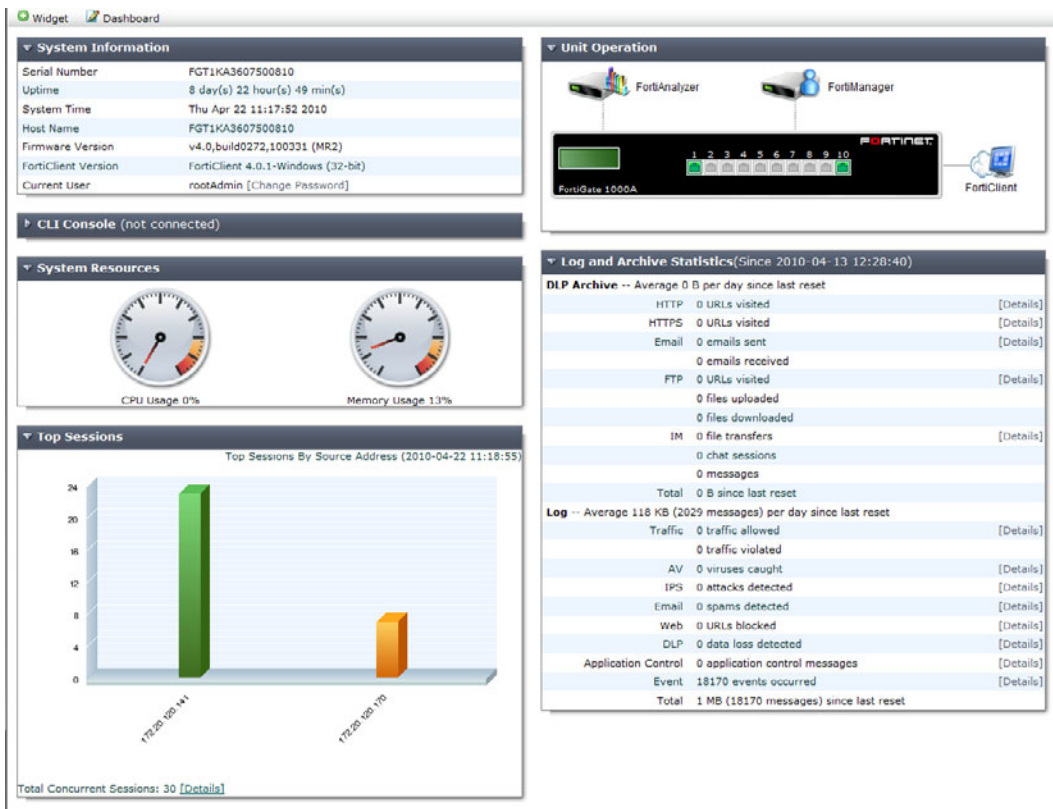
The trusted host addresses all default to 0.0.0.0/0.0.0.0 for IPv4, or ::/0 for IPv6. If you set one of the zero addresses to a non-zero address, the other zero addresses will be ignored. The only way to use a wildcard entry is to leave the trusted hosts at 0.0.0.0/0.0.0.0 or ::0. However, this configuration is less secure.

Virtual Domain administrator dashboard display

When administrators logs into their virtual domain, they see a different dashboard than the global administrator will see. The VDOM dashboard displays information only relevant to that VDOM — no global or other VDOM information is displayed.

Information	per-VDOM	Global
System Information	read-only	yes
License Information	no	yes
CLI console	yes	yes
Unit Operation	read-only	yes
Alert Message Console	no	yes
Top Sessions	limited to VDOM sessions	yes
Traffic	limited to VDOM interfaces	yes
Statistics	yes	yes

Figure 184: VDOM administrator dashboard





Virtual Domains in NAT/Route mode

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that each function as independent units. Each virtual domain has separate routing and security policies. A single FortiGate unit with virtual domains is flexible enough to serve multiple departments of an organization, separate organizations, or be the basis for a service provider's managed security service.



The examples in this chapter are intended to be followed in order as procedures build on previous procedures. If you do not complete the previous procedures, the procedure you are working on may not work properly. If this happens, consult previous procedures or FortiGate documentation.

This chapter contains the following sections:

- [Virtual domains in NAT/Route mode](#)
- [Example NAT/Route VDOM configuration](#)

Virtual domains in NAT/Route mode

Once you have enabled virtual domains and created one or more VDOMs, you need to configure them. Configuring VDOMs on your FortiGate unit includes tasks such as the ones listed here; while you may not require all for your network topology, it is recommended that you perform them in the order given:

- [Changing the management virtual domain](#)
- [Configuring interfaces in a NAT/Route VDOM](#)
- [Configuring VDOM routing](#)
- [Configuring security policies for NAT/Route VDOMs](#)
- [Configuring UTM profiles for NAT/Route VDOMs](#)

Changing the management virtual domain

The management virtual domain is the virtual domain where all the management traffic for the FortiGate unit originates. This management traffic needs access to remote servers, such as FortiGuard services and NTP, to perform its duties. It needs access to the Internet to send and receive this traffic.

Management traffic includes, but is not limited to:

- DNS lookups
- logging to FortiAnalyzer or syslog
- FortiGuard service
- sending alert emails
- Network time protocol traffic (NTP)
- Sending SNMP traps
- Quarantining suspicious files and email.

By default the management VDOM is the root domain. When other VDOMs are configured on your FortiGate unit, management traffic can be moved to one of these other VDOMs.

Reasons to move the management VDOM include selecting a non-root VDOM to be your administration VDOM, or the root VDOM not having an interface with a connection to the Internet.



You cannot change the management VDOM if any administrators are using RADIUS authentication.

The following procedure will change the management VDOM from the default `root` to a VDOM named `mgmt_vdom`. It is assumed that `mgmt_vdom` has already been created and has an interface that can access the Internet.

To change the management VDOM - web-based manager

- 1 In *Current VDOM*, select *Global*.
- 2 Select *System > VDOM > VDOM*.
- 3 Select the checkbox next to the required VDOM.
- 4 Select *Switch Management*

The current management VDOM is shown in square brackets, “[root]” for example.

To change the management VDOM - CLI

```
config global
  config system global
    set management-vdom mgmt_vdom
  end
```

Management traffic will now originate from `mgmt_vdom`.

Configuring interfaces in a NAT/Route VDOM

A VDOM must contain at least two interfaces to be useful. These can be physical interfaces or VLAN interfaces. By default, all physical interfaces are in the root VDOM. When you create a new VLAN, it is in the root VDOM by default.

When there are VDOMs on the FortiGate unit in both NAT and Transparent operation modes, some interface fields will be displayed as “-” on *System > Network > Interface*. Only someone with a `super_admin` account can view all the VDOMs.



When moving an interface to a different VDOM, firewall IP pools and virtual IPs for this interface are deleted. You should manually delete any routes that refer to this interface. Once the interface has been moved to the new VDOM, you can add these services to the interface again.



When configuring VDOMs on FortiGate units with accelerated interfaces, such as NP2 or NP4 interfaces, you must assign both interfaces in the pair to the same VDOM for those interfaces to retain their acceleration. Otherwise they will become normal interfaces.

This section includes the following topics:

- [Adding a VLAN to a NAT/Route VDOM](#)
- [Moving an interface to a VDOM](#)
- [Deleting an interface](#)
- [Adding a zone to a VDOM](#)

Adding a VLAN to a NAT/Route VDOM

The following example shows one way that multiple companies can maintain their security when they are using one FortiGate unit with VLANs that share interfaces on the unit.

This procedure will add a VLAN interface called `client1-v100` with a VLAN ID of 100 to an existing VDOM called `client1` using the physical interface called `port2`.



The physical interface does not need to belong to the VDOM that the VLAN belongs to.

To add a VLAN subinterface to a VDOM - web-based manager

- 1 In *Current VDOM*, select *Global*.
- 2 Go to *System > Network > Interface*.
- 3 Select *Create New*.
- 4 Enter the following information and select *OK*:

Name	client1-v100
Interface	port2
VLAN ID	100
Virtual Domain	Client1
Addressing mode	Manual
IP/Netmask	172.20.120.110/255.255.255.0
Administrative Access	HTTPS, SSH

You will see an expand arrow added to the port2 interface. When the arrow is expanded, the interface shows the `client1-v100` VLAN subinterface.

To add a VLAN subinterface to a VDOM - CLI

```
config global
  config system interface
    edit client1-v100
      set type vlan
      set vlanid 100
      set vdom Client1
      set interface port2
      set ip 172.20.120.110 255.255.255.0
      set allowaccess https ssh
    end
```

Moving an interface to a VDOM

Interfaces belong to the root VDOM by default. Moving an interface is the same procedure no matter if its moving from the root VDOM or a any other VDOM.

If you have an accelerated pair of physical interfaces, such as NP2 interfaces, both interfaces must be in the same VDOM or you will loose their acceleration.

The following procedure will move the port3 interface to the Client2 VDOM. This is a common action when configuring a VDOM. It is assumed that the Client2 VDOM has already been created. It is also assumed that your FortiGate unit has a port3 interface. If you are using a different model, your physical interfaces may not be named `port2`, `external` or `port3`.

To move an existing interface to a different VDOM - web-based manager

- 1 For Current VDOM, select Global.
- 2 Go to *System > Network > Interface*.
- 3 Select *Edit* for the port3 interface.
- 4 Select `Client2` as the new *Virtual Domain*.
- 5 Select *OK*.

To move an existing interface to a different VDOM - CLI

```
config global
  config system interface
    edit port3
      set vdom Client2
    end
```

Deleting an interface

Before you can delete a virtual interface, or move an interface from one VDOM to another, all references to that interface must be removed. For a list of objects that can refer to an interface see [“Per-VDOM settings - web-based manager” on page 1882](#).

The easiest way to be sure an interface can be deleted is when the Delete icon is no longer greyed out. If it remains greyed out when an interface is selected, that interface still has objects referring to it, or it is a physical interface that cannot be deleted.

To delete a virtual interface - web-based manager

- 1 Ensure all objects referring to this interface have been removed.
- 2 In *Current VDOM*, select Global.
- 3 Select *System > Network > Interface*.
- 4 Select the interface to delete.
- 5 Select the delete icon.

Adding a zone to a VDOM

Grouping interfaces and VLAN subinterfaces into zones simplifies policy creation. You can configure policies for connections to and from a zone, but not between interfaces in a zone.

Zones are VDOM-specific. A zone cannot be moved to a different VDOM. Any interfaces in a zone cannot be used in another zone. To move a zone to a new VDOM requires deleting the current zone and re-creating a zone in the new VDOM.

The following procedure will create a zone called `accounting` in the `client2` VDOM. It will not allow intra-zone traffic, and both `port3` and `port2` interfaces belong to this zone. This is a method of grouping and isolating traffic over particular interfaces—it is useful for added security and control within a larger network.

To add a zone to a VDOM - web-based manager

- 1 In *Current VDOM*, select the `client2` VDOM.
- 2 Go to *System > Network > Interface*.
- 3 Select *Create New > Zone*.
- 4 Enter the following information and select *OK*:

Zone Name	accounting
Block intra-zone traffic	Select
Interface Members	port3, port2

To add a zone to a VDOM - CLI

```
config vdom
  edit client2
    config system zone
      edit accounting
        set interface port3 port2
        set intrazone deny
      end
    end
  end
```

Configuring VDOM routing

Routing is VDOM-specific. Each VDOM should have a default static route configured as a minimum. Within a VDOM, routing is the same as routing on your FortiGate unit without VDOMs enabled.

When configuring dynamic routing on a VDOM, other VDOMs on the FortiGate unit can be neighbors. The following topics give a brief introduction to the routing protocols, and show specific examples of how to configure dynamic routing for VDOMs. Figures are included to show the FortiGate unit configuration after the successful completion of the routing example.

This section includes:

- [Default static route for a VDOM](#)
- [Dynamic Routing in VDOMs](#)

Default static route for a VDOM

The routing you define applies only to network traffic entering non-ssl interfaces belonging to this VDOM. Set the administrative distance high enough, typically 20, so that automatically configured routes will be preferred to the default.

In the following procedure, it is assumed that a VDOM called “Client2” exists. The procedure will create a default static route for this VDOM. The route has a destination IP of 0.0.0.0, on the `port3` interface. It has a gateway of 10.10.10.1, and an administrative distance of 20.

The values used in this procedure are very standard, and this procedure should be part of configuring all VDOMs.

To add a default static route for a VDOM - web-based manager

- 1 For Current VDOM, select Global.
- 2 Go to *System > VDOM > VDOM*.
- 3 Select the *Client2* VDOM and select *Enter*.
- 4 Go to *Router > Static > Static Route*.
- 5 Select *Create New*.
- 6 Enter the following information and select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2
Gateway	10.10.10.1
Distance	20

To add a default static route for a VDOM - CLI

```

config vdom
  edit Client2
    config router static
      edit 4
        set device port2
        set dst 0.0.0.0 0.0.0.0
        set gateway 10.10.10.1
        set distance 20
      end
    end
  end
end

```

Dynamic Routing in VDOMs

Dynamic routing is VDOM-specific, like all other routing. Dynamic routing configuration is the same with VDOMs as with your FortiGate unit without VDOMs enabled, once you are at the routing menu. If you have multiple VDOMs configured, the dynamic routing configuration between them can become quite complex.

VDOMs provide some interesting changes to dynamic routing. Each VDOM can be a neighbor to the other VDOMs. This is useful in simulating a dynamic routing area or AS or network using only your FortiGate unit.

You can separate different types of routing to different VDOMs if required. This allows for easier troubleshooting. This is very useful if your FortiGate unit is on the border of a number of different routing domains.

For more information on dynamic routing in FortiOS, see [“Dynamic Routing Overview” on page 1699](#).

Inter-VDOM links must have IP addresses assigned to them if they are part of a dynamic routing configuration. Inter-VDOM links may or may not have IP addresses assigned to them. Without IP addresses, you need to be careful how you configure routing. While the default static route can be assigned an address of 0.0.0.0 and rely instead on the interface, dynamic routing almost always requires an IP address.

RIP

The RIP dynamic routing protocol uses hop count to determine the best route, with a hop count of 1 being directly attached to the interface and a hop count of 16 being unreachable. For example if two VDOMs on the same FortiGate unit are RIP neighbors, they have a hop count of 1.

OSPF

OSPF communicates the status of its network links to adjacent neighbor routers instead of the complete routing table. When compared to RIP, OSPF is more suitable for large networks, it is not limited by hop count, and is more complex to configure. For smaller OSPF configurations its easiest to just use the backbone area, instead of multiple areas.

BGP

BGP is an Internet gateway protocol (IGP) used to connect autonomous systems (ASes) and is used by Internet service providers (ISPs). BGP stores the full path, or path vector, to a destination and its attributes which aid in proper routing.

Configuring security policies for NAT/Route VDOMs

Security policies are VDOM-specific. This means that all firewall settings for a VDOM, such as firewall addresses and security policies, are configured within the VDOM.

In VDOMs, all firewall related objects are configured per-VDOM including addresses, service groups, UTM profiles, schedules, traffic shaping, and so on. If you want firewall addresses, you will have to create them on each VDOM separately. If you have many addresses, and VDOMs this can be tedious and time consuming. Consider using a FortiManager unit to manage your VDOM configuration — it can get firewall objects from a configured VDOM or FortiGate unit, and push those objects to many other VDOMs or FortiGate units. See the [FortiManager Administration Guide](#).



You can customize the *Policy* display by including some or all columns, and customize the column order onscreen. Due to this feature, security policy screenshots may not appear the same as on your screen.

Configuring a security policy for a VDOM

Your security policies can involve only the interfaces, zones, and firewall addresses that are part of the current VDOM, and they are only visible when you are viewing the current VDOM. The security policies of this VDOM filter the network traffic on the interfaces and VLAN subinterfaces in this VDOM.

A firewall service group can be configured to group multiple services into one service group. When a descriptive name is used, service groups make it easier for an administrator to quickly determine what services are allowed by a security policy.

In the following procedure, it is assumed that a VDOM called `Client2` exists. The procedure will configure an outgoing security policy. The security policy will allow all HTTPS and SSH traffic for the `SalesLocal` address group on VLAN_200 going to all addresses on port3. This traffic will be scanned and logged.

To configure a security policy for a VDOM - web-based manager

- 1 Go to *System > VDOM > VDOM*.
- 2 Select the *Client2* VDOM and select *Enter*.
- 3 Go to *Policy > Policy*.
- 4 Select *Create New*.

5 Enter the following information and select OK:

Source Interface/Zone	VLAN_200
Source Address	SalesLocal
Destination Interface/Zone	port3
Destination Address	any
Schedule	always
Service	Multiple - HTTPS, SSH
Action	ACCEPT
Log Allowed Traffic	enable

To configure a security policy for a VDOM - CLI

```
config vdom
  edit Client2
    config firewall policy
      edit 12
        set srcintf VLAN_200
        set srcaddr SalesLocal
        set dstintf port3(dmz)
        set dstaddr any
        set schedule always
        set service HTTPS SSH
        set action accept
        set status enable
        set logtraffic enable
      end
    end
  end
```

Configuring UTM profiles for NAT/Route VDOMs

In NAT/Route VDOMs, UTM profiles are exactly like regular FortiGate unit operation with one exception. In VDOMs, there are no default UTM profiles.

If you want UTM profiles in VDOMs, you must create them yourself. If you have many UTM profiles to create in each VDOM, you should consider using a FortiManager unit. It can get existing profiles from a VDOM or FortiGate unit, and push those profiles down to multiple other VDOMs or FortiGate units. See [FortiManager Administration Guide](#).

When VDOMs are enabled, you only need one FortiGuard license for the physical unit, and download FortiGuard updates once for the physical unit. This can result in a large time and money savings over multiple physical units if you have many VDOMs.

Configuring VPNs for a VDOM

Virtual Private Networking (VPN) settings are VDOM-specific, and must be configured within each VDOM. Configurations for IPsec Tunnel, IPsec Interface, PPTP and SSL are VDOM-specific. However, certificates are shared by all VDOMs and are added and configured globally to the FortiGate unit.

Example NAT/Route VDOM configuration

Company A and Company B each have their own internal networks and their own ISPs. They share a FortiGate unit that is configured with two separate VDOMs, with each VDOM running in NAT/Route mode enabling separate configuration of network protection profiles. Each ISP is connected to a different interface on the FortiGate unit.

This network example was chosen to illustrate one of the most typical VDOM configurations.

This example has the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Creating the VDOMs](#)
- [Configuring the FortiGate interfaces](#)
- [Configuring the vdomA VDOM](#)
- [Configuring the vdomB VDOM](#)
- [Testing the configuration](#)

Network topology and assumptions

Both companies have their own ISPs and their own internal interface, external interface, and VDOM on the FortiGate unit.

For easier configuration, the following IP addressing is used:

- all IP addresses on the FortiGate unit end in “.2” such as 10.11.101.2.
- all IP addresses for ISPs end in “.7”, such as 172.20.201.7.
- all internal networks are 10.*.* networks, and sample internal addresses end in “.55”.

The IP address matrix for this example is as follows.

Address	Company A	Company B
ISP	172.20.201.7	192.168.201.7
Internal network	10.11.101.0	10.012.101.0
FortiGate / VDOM	172.20.201.2 (port1)	192.168.201.2 (port3)
	10.11.101.2 (port4)	10.012.101.2 (port2)

The Company A internal network is on the 10.11.101.0/255.255.255.0 subnet. The Company B internal network is on the 10.12.101.0/255.255.255.0 subnet.


There are no switches or routers required for this configuration.

There are no VLANs in this network topology.

The interfaces used in this example are port1 through port4. Different FortiGate models may have different interface labels. port1 and port3 are used as external interfaces. port2 and port4 are internal interfaces.

The administrator is a super_admin account. If you are using a non-super_admin account, refer to [“Global and per-VDOM settings” on page 1881](#) to see which parts a non-super_admin account can also configure.

When configuring security policies in the CLI always choose a policy number that is higher than any existing policy numbers, select `services` before `profile-status`, and `profile-status` before `profile`. If these commands are not entered in that order, they will not be available to enter.

Figure 185: Example VDOM configuration

General configuration steps

For best results in this configuration, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 [Creating the VDOMs](#)
- 2 [Configuring the FortiGate interfaces](#)
- 3 [Configuring the vdomA VDOM, and Configuring the vdomB VDOM](#)
- 4 [Testing the configuration](#)

Creating the VDOMs

In this example, two new VDOMs are created — vdomA for Company A and vdomB for Company B. These VDOMs will keep the traffic for these two companies separate while enabling each company to access its own ISP.

To create two VDOMs - web-based manager

- 1 Log in with a super_admin account.
- 2 For *Current VDOM*, select Global.
- 3 Go to *System > VDOM > VDOM*, and select *Create New*.
- 4 Enter vdomA and select *OK*.
- 5 Select *OK* again to return to the VDOM list.
- 6 Select *Create New*.
- 7 Enter vdomB and select *OK*.

To create two VDOMs - CLI

```
config vdom
  edit vdomA
  next
```

```
edit vdomB
end
```

Configuring the FortiGate interfaces

This section configures the interfaces that connect to the companies' internal networks, and to the companies' ISPs.

All interfaces on the FortiGate unit will be configured with an IP address ending in ".2" such as 10.11.101.2. This will simplify network administration both for the companies, and for the FortiGate unit global administrator. Also the internal addresses for each company differ in the second octet of their IP address - Company A is 10.11.*, and Company B is 10.12.*.

This section includes the following topics:

- [Configuring the vdomA interfaces](#)
- [Configuring the vdomB interfaces](#)



If you cannot change the VDOM of a network interface it is because something is referring to that interface that needs to be deleted. Once all the references are deleted the interface will be available to switch to a different VDOM. For example a common reference to the external interface is the default static route entry. See "[Configuring interfaces in a NAT/Route VDOM](#)" on page 1904.

Configuring the vdomA interfaces

The `vdomA` VDOM includes two FortiGate unit interfaces: `port1` and `external`.

The `port4` interface connects the Company A internal network to the FortiGate unit, and shares the internal network subnet of 10.11.101.0/255.255.255.0.

The `external` interface connects the FortiGate unit to ISP A and the Internet. It shares the ISP A subnet of 172.20.201.0/255.255.255.0.

To configure the vdomA interfaces - web-based manager

- 1 For Current VDOM, select `Global`.
- 2 Go to *System > Network > Interface*.
- 3 Select *Edit* on the `port1` interface.
- 4 Enter the following information and select *OK*:

Virtual Domain	vdomA
Addressing mode	Manual
IP/Netmask	172.20.201.2/255.255.255.0

- 5 Select *Edit* on the `port4` interface.
- 6 Enter the following information and select *OK*:

Virtual Domain	vdomA
Addressing mode	Manual
IP/Netmask	10.11.101.2/255.255.255.0

To configure the vdomA interfaces - CLI

```

config global
  config system interface
    edit port1
      set vdom vdomA
      set mode static
      set ip 172.20.201.2 255.255.255.0
    next
    edit port4
      set vdom ABCdomain
      set mode static
      set ip 10.11.101.2 255.255.255.0
    end
  end
end

```

Configuring the vdomB interfaces

The vdomB VDOM uses two FortiGate unit interfaces: port2 and port3.

The port2 interface connects the Company B internal network to the FortiGate unit, and shares the internal network subnet of 10.12.101.0/255.255.255.0.

The port3 interface connects the FortiGate unit to ISP B and the Internet. It shares the ISP B subnet of 192.168.201.0/255.255.255.0.

To configure the DEFdomain interfaces - web-based manager

- 1 For Current VDOM, select Global.
- 2 Go to *System > Network > Interface*.
- 3 Select *Edit* on the port3 interface.
- 4 Enter the following information and select *OK*:

Virtual domain	vdomB
Addressing mode	Manual
IP/Netmask	192.168.201.2/255.255.255.0

- 5 Select *Edit* on the port2 interface.
- 6 Enter the following information and select *OK*:

Virtual domain	vdomB
Addressing mode	Manual
IP/Netmask	10.12.101.2/255.255.255.0

To configure the vdomB interfaces - CLI

```

config global
  config system interface
    edit port3
      set vdom vdomB
      set mode static
      set ip 192.168.201.2 255.255.255.0
    next
    edit port2
      set vdom vdomB

```



```
set mode static
set ip 10.12.101.2 255.255.255.0
end
```

Configuring the vdomA VDOM

With the VDOMs created and the ISPs connected, the next step is to configure the vdomA VDOM.

Configuring the vdomA includes the following:

- [Adding vdomA firewall addresses](#)
- [Adding the vdomA security policy](#)
- [Adding the vdomA default route](#)

Adding vdomA firewall addresses

You need to define the addresses used by Company A's internal network for use in security policies. This internal network is the 10.11.101.0/255.255.255.0 subnet.

The FortiGate unit provides one default address, "all", that you can use when a security policy applies to all addresses as the source or destination of a packet.

To add the vdomA firewall addresses - web-based manager

- 1 For Current VDOM, select *vdomA*.
- 2 Go to *Firewall Objects > Address > Address*.
- 3 Select *Create New*.
- 4 Enter the following information and select *OK*:

Address Name	Ainternal
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0/255.255.255.0
Interface	port4

To add the ABCdomain VDOM firewall addresses - CLI

```
config vdom
edit vdomA
config firewall address
edit Ainternal
set type ipmask
set subnet 10.11.101.0 255.255.255.0
end
end
```

Adding the vdomA security policy

You need to add the vdomA security policy to allow traffic from the internal network to reach the external network, and from the external network to internal as well. You need two policies for this domain.

To add the vdomA security policy - web-based manager

- 1 In Current VDOM, select *vdomA*.
- 2 Go to *Policy > Policy*.

- 3 Select *Create New*.
- 4 Enter the following information and select *OK*:

Source Interface/Zone	port4
Source Address	Ainternal
Destination Interface/Zone	port1
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

- 5 Select *Create New*.
- 6 Enter the following information and select *OK*:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port4
Destination Address	Ainternal
Schedule	Always
Service	ANY
Action	ACCEPT

To add the vdomA security policy - CLI

```

config vdom
  edit vdomA
    config firewall policy
      edit 1
        set srcintf port4
        set srcaddr Ainternal
        set dstintf port1
        set dstaddr all
        set schedule always
        set service ANY
        set action accept
        set status enable
      next
      edit 2
        set srcintf port1
        set srcaddr all
        set dstintf port4
        set dstaddr Ainternal
        set schedule always
        set service ANY
        set action accept
        set status enable
      end
    end
  end

```

Adding the vdomA default route

You also need to define a default route to direct packets from the Company A internal network to ISP A. Every VDOM needs a default static route, as a minimum, to handle traffic addressed to external networks such as the Internet.

The administrative distance should be set slightly higher than other routes. Lower admin distances will get checked first, and this default route will only be used as a last resort.

To add a default route to the vdomA - web-based manager

- 1 For Current VDOM, select *vdomA*.
- 2 Go to *Router > Static > Static Route*.
- 3 Select *Create New*.
- 4 Enter the following information and select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port1
Gateway	172.20.201.7
Distance	20

To add a default route to the vdomA - CLI

```
config vdom
  edit vdomA
    config router static
      edit 1
        set device port1
        set gateway 172.20.201.7
      end
```

Configuring the vdomB VDOM

In this example, the vdomB VDOM is used for Company B. Firewall and routing settings are specific to a single VDOM.

vdomB includes the FortiGate port2 interface to connect to the Company B internal network, and the FortiGate port3 interface to connect to ISP B. Security policies are needed to allow traffic from port2 to external and from external to port2 interfaces.

This section includes the following topics:

- [Adding the vdomB firewall address](#)
- [Adding the vdomB security policy](#)
- [Adding a default route to the vdomB VDOM](#)

Adding the vdomB firewall address

You need to define addresses for use in security policies. In this example, the vdomB VDOM needs an address for the port2 interface and the “all” address.

To add the vdomB firewall address - web-based manager

- 1 In *Current VDOM*, select *vdomB*.
- 2 Go to *Firewall Objects > Address > Address*.
- 3 Select *Create New*.

- 4 Enter the following information and select OK:

Address Name	Binternal
Type	Subnet / IP Range
Subnet / IP Range	10.12.101.0/255.255.255.0
Interface	port2

To add the vdomB firewall address - CLI

```
config vdom
  edit vdomB
    config firewall address
      edit Binternal
        set type ipmask
        set subnet 10.12.101.0 255.255.255.0
      end
    end
  end
```

Adding the vdomB security policy

You also need a security policy for the Company B domain. In this example, the security policy allows all traffic.

To add the vdomB security policy - web-based manager

- 1 Log in with a super_admin account.
- 2 In *Current VDOM*, select vdomB.
- 3 Go to *Policy > Policy*.
- 4 Select *Create New*.
- 5 Enter the following information and select OK:

Source Interface/Zone	port2
Source Address	Binternal
Destination Interface/Zone	port3
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

- 6 Select *Create New*.
- 7 Enter the following information and select OK:

Source Interface/Zone	port3
Source Address	all
Destination Interface/Zone	port2
Destination Address	Binternal
Schedule	Always

Service	ANY
Action	ACCEPT

To add the vdomB security policy - CLI

```
config vdom
  edit vdomB
    config firewall policy
      edit 1
        set srcintf port2
        set dstintf port3
        set srcaddr Binternal
        set dstaddr all
        set schedule always
        set service ANY
        set action accept
        set status enable
      edit 1
        set srcintf port3
        set dstintf port2
        set srcaddr all
        set dstaddr Binternal
        set schedule always
        set service ANY
        set action accept
        set status enable
    end
  end
end
```

Adding a default route to the vdomB VDOM

You need to define a default route to direct packets to ISP B.

To add a default route to the vdomB VDOM - web-based manager

- 1 Log in as the super_admin administrator.
- 2 In *Current VDOM*, select vdomB.
- 3 Go to *Router > Static > Static Route*.
- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3
Gateway	192.168.201.7
Distance	20

To add a default route to the vdomB VDOM - CLI

```
config vdom
  edit vdomB
    config router static
      edit 1
        set dst 0.0.0.0/0
```

```

        set device external
        set gateway 192.168.201.7
    end
end

```

Testing the configuration

Once you have completed configuration for both company VDOMs, you can use diagnostic commands, such as `tracert` in Windows, to test traffic routed through the FortiGate unit. Alternately, you can use the `tracert` command on a Linux system with similar output.

Possible errors during the traceroute test are:

- “***Request timed out” - the trace was not able to make the next connection towards the destination fast enough
- “Destination host unreachable” - after a number of timed-out responses the trace will give up

Possible reasons for these errors are bad connections or configuration errors.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 1977](#).

Testing traffic from the internal network to the ISP

In this example, a route is traced from the Company A internal network to ISP A. The test was run on a Windows PC with an IP address of 10.11.101.55.

The output here indicates three hops between the source and destination, the IP address of each hop, and that the trace was successful.

From the Company A internal network, access a command prompt and enter this command:

```

C:\>tracert 172.20.201.7
Tracing route to 172.20.201.7 over a maximum of 30 hops:
  1  <10 ms  <10 ms  <10 ms  10.11.101.2
  2  <10 ms  <10 ms  <10 ms  172.20.201.2
  3  <10 ms  <10 ms  <10 ms  172.20.201.7

```

Trace complete.



You can customize the Firewall Policy display by including some or all columns, and customize the column order onscreen. Due to this feature, firewall policy screenshots may not appear the same as on your screen.



To complete the setup, configure devices on the VLANs with default gateways. The default gateway for VLAN 10 is the FortiGate VLAN 10 subinterface. Configure the rest of the devices, similarly matching the default gateway and FortiGate VLAN subinterface numbers.



Virtual Domains in Transparent mode

In Transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering and intrusion protection to traffic. There are some limitations in Transparent mode in that you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in Transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

VDOMs can each be configured to operate either in Transparent or NAT/Route operation mode, with each VDOM behaving like a separate FortiGate unit operating in the respective mode. VLANs configured on a VDOM in Transparent mode are the same as VLANs configured on the FortiGate unit when VDOMs are disabled.

This chapter includes the following sections:

- [Before you begin](#)
- [Transparent operation mode](#)
- [Configuring VDOMs in Transparent mode](#)
- [Example of VDOMs in Transparent mode](#)

Before you begin

Before you begin using this chapter, take a moment to note the following:

- The information in this chapter applies to all FortiGate units. All FortiGate models except the FortiGate-30B model support VDOMs, and all FortiGate models support VLANs.
- By default, your FortiGate unit supports a maximum of 10 VDOMs in any combination of NAT/Route and Transparent operating modes. For FortiGate models numbered 1240 and higher, you can purchase a license key to increase the maximum number of VDOMs. Model 1240B supports up to 25 VDOMs. Most Enterprise and Large Enterprise models can support 250 VDOMs. Chassis-based models can support up to 3000 VDOMs. For specific information, see the product data sheet.
- This chapter uses port1 through port4 for interfaces in examples, where possible aliases have been assigned to the interfaces for extra clarity. The interface names on some models will vary. For example, some models do not have interfaces labeled external or internal.
- A super_admin administrator account is assumed for the procedures and examples; however, if you are an administrator restricted to a VDOM, you may be able to perform some procedures. For more information, see [“Administrators in Virtual Domains” on page 1898](#).

Transparent operation mode

In transparent mode, the FortiGate unit becomes a layer-2 IP forwarding bridge. This means that Ethernet frames are forwarded based on destination MAC address, and no other routing is performed. All incoming traffic that is accepted by the firewall, is broadcast out on all interfaces.

In transparent mode the FortiGate unit is a forwarding bridge, not a switch. A switch can develop a port table and associated MAC addresses, so that it can bridge two ports to deliver the traffic instead of broadcasting to all ports. In transparent mode, the FortiGate unit does not follow this switch behavior, but instead is the forwarding bridge that broadcasts all packets out over all interfaces, subject to security policies.

Features such as broadcast domains, forwarding domains, and STP apply to both FortiGate units and VDOMs in Transparent mode.

Broadcast domains

A broadcast domain is a network segment in which any network equipment can transmit data directly to another device without going through a routing device. All the devices share the same subnet. The subnets are separated by layer-3 devices, such as routers, that can forward traffic from one broadcast domain to the next.

Broadcast domains are important to transparent mode FortiGate units because the broadcast domain is the limit of where the FortiGate unit can forward packets when it is in transparent mode.

Forwarding domains

Address Resolution Protocol (ARP) packets are vital to communication on a network, and ARP support is enabled on FortiGate unit interfaces by default. Normally you want ARP packets to pass through the FortiGate unit. However, in Transparent mode ARP packets arriving on one interface are sent to all other interfaces including VLANs giving the appearance of duplicates of the same MAC address on different interfaces. Some layer-2 switches become unstable when they detect these duplicate MAC addresses. Unstable switches may become unreliable or reset and cause network traffic to slow down considerably.

When you are using VLANs in Transparent mode, the solution to the duplicate MAC address issue is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic—it is like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0.

To assign VLAN 200 to collision group 2, VLAN 300 to collision group 3, and all other interfaces to stay in the default collision group 0 enter the following CLI commands:

```
config system interface
  edit vlan200
    set vlanid 200
    set forward_domain 2
  next
  edit vlan300
    set vlanid 300
    set forward_domain 3
  next
end
```


When using forwarding domains, you may experience connection issues with layer-2 traffic, such as ping, if your network configuration has

- packets going through the FortiGate unit in Transparent mode multiple times,
- more than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled.

Spanning Tree Protocol

VDOMs and FortiGate units do not participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, quickly reducing available bandwidth to zero.

If you use your VDOM or FortiGate unit in a network topology that relies on STP for network loop protection, you need to make changes to your FortiGate configuration. Otherwise, STP recognizes your FortiGate unit as a blocked link and forwards the data to another path. By default, your FortiGate unit blocks STP as well as other non-IP protocol traffic. Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the port2 interface:

```
config global
  config system interface
    edit port2
      set l2forward enable
      set stpforward enable
    next
  end
```

There are different CLI commands to allow other common layer-2 protocols such as IPX, PPTP or L2TP on the network. For more information, see the [FortiOS CLI Reference](#).

Differences between NAT/Route and Transparent mode

The differences between NAT/Route mode and Transparent mode include:

Table 112: Differences between NAT/Route and Transparent modes

Features	NAT/Route mode	Transparent mode
Specific Management IP address required	No	Yes
Perform Network Address Translation (NAT)	Yes	Yes
Stateful packet inspection	Yes	Yes
Layer-2 forwarding	Yes	Yes
Layer-3 routing	Yes	No
Unicast Routing / Policy Based routing	Yes	No
DHCP server	Yes	No
IPsec VPN	Yes	Yes
PPTP/L2TP VPN	Yes	No
SSL VPN	Yes	No

Table 112: Differences between NAT/Route and Transparent modes

UTM features	Yes	Yes
VLAN support	Yes	Yes - limited to VLAN trunks.
Ping servers (dead gateway detection)	Yes	No

To provide administrative access to a FortiGate unit or VDOM in Transparent mode, you must define a management IP address and a gateway. This step is not required in NAT/Route mode where you can access the FortiGate unit through the assigned IP address of any interface where administrative access is permitted.

If you incorrectly set the Transparent mode management IP address for your FortiGate unit, you will be unable to access your unit through the web-based manager. In this situation, you will need to connect to the FortiGate unit using the console cable and change the settings so you can access the unit. Alternately, if your unit has an LCD panel, you can change the operation mode and interface information through the LCD panel.

Operation mode differences in VDOMs

A VDOM, such as root, can have a maximum of 255 interfaces in Network Address Translation (NAT) mode or Transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. To have more than a total of 255 interfaces configured, you need multiple VDOMs with multiple interfaces on each.

In Transparent mode without VDOMs enabled, all interfaces on the FortiGate unit act as a bridge — all traffic coming in on one interface is sent back out on all the other interfaces. This effectively turns the FortiGate unit into a two interface unit no matter how many physical interfaces it has. When VDOMs are enabled, this allows you to determine how many interfaces to assign to a VDOM running in Transparent mode. If there are reasons for assigning more than two interfaces based on your network topology, you are able to. However, the benefit of VDOMs in this case is that you have the functionality of Transparent mode, but you can use interfaces for NAT/Route traffic as well.

You can add more VDOMs to separate groups of VLAN subinterfaces. When using a FortiGate unit to serve multiple organizations, this configuration simplifies administration because you see only the security policies and settings for the VDOM you are configuring. For information on adding and configuring virtual domains, see [“Benefits of Virtual Domains” on page 1875](#).

One essential application of VDOMs is to prevent problems caused when a FortiGate unit is connected to a layer-2 switch that has a global MAC table. FortiGate units normally forward ARP requests to all interfaces, including VLAN subinterfaces. It is then possible for the switch to receive duplicate ARP packets on different VLANs. Some layer-2 switches reset when this happens. As ARP requests are only forwarded to interfaces in the same VDOM, you can solve this problem by creating a VDOM for each VLAN. For a configuration example, see [“Example of VDOMs in Transparent mode” on page 1926](#).

Configuring VDOMs in Transparent mode

In Transparent mode, your FortiGate unit becomes a layer-2 bridge — any traffic coming in on one port is broadcast out on all the other ports. If your FortiGate unit has many interfaces, this is not the best use of those interfaces. VDOMs can limit Transparent mode to only a few interfaces while allowing the rest of the FortiGate unit to remain in NAT/Route mode.

The essential steps to configure your FortiGate unit to work with VLANs in Transparent mode are:

- [Switching to Transparent mode](#)
- [Adding VLAN subinterfaces](#)
- [Creating security policies.](#)

You can also configure the UTM profiles that manage antivirus scanning, web filtering and spam filtering. For more information, see [“UTM overview” on page 879](#).

In Transparent mode, you can access the FortiGate web-based manager by connecting to an interface configured for administrative access and using HTTPS to access the management IP address. On the FortiGate unit used for examples in this guide, administrative access is enabled by default on the internal interface and the default management IP address is 10.11.0.1.

Switching to Transparent mode

A VDOM is in NAT/Route mode by default when it is created. You must switch it to Transparent mode, and add a management IP address so you can access the VDOM from your management computer.



Before applying the change to Transparent mode, ensure the VDOM has administrative access on the selected interface, and that the selected management IP address is reachable on your network.

To switch the `tpVDOM` VDOM to Transparent mode - web-based manager

- 1 Go to *Current VDOM* menu and select *Global*.
- 2 Go to *System > VDOM > VDOM*.
- 3 Edit the *tpVDOM*.
- 4 Select *Transparent for Operation mode*.
- 5 Enter the management IP/Netmask.

The IP address must be accessible to the subnet where the management computer is located. For example 10.11.0.99/255.255.255.0 will be able to access the 10.11.0.0 subnet.

- 6 Select *Apply*.

When you select *Apply*, the FortiGate unit will log you out. When you log back in, the VDOM will be in Transparent mode.

To switch the `tpVDOM` VDOM to Transparent mode - CLI

```
config vdom
  edit tpVDOM
    config system settings
      set opmode transparent
      set manageip 10.11.0.99 255.255.255.0
    end
  end
```

Adding VLAN subinterfaces

There are a few differences when adding VLANs in Transparent mode compared to NAT/Route mode.

In Transparent mode, VLAN traffic is trunked across the VDOM. That means VLAN traffic cannot be routed, changed, or inspected. For this reason when you assign a VLAN to a Transparent mode VDOM, you will see the *Addressing Mode* section of the interface configuration disappear in from the web-based manager. It is because with no routing, inspection, or any activities able to be performed on VLAN traffic the VDOM simply re-broadcasts the VLAN traffic. This requires no addressing.

Also any routing related features such as dynamic routing or Virtual Router Redundancy Protocol (VRRP) are not available in Transparent mode for any interfaces.

Creating security policies

Security policies permit communication between the FortiGate unit's network interfaces based on source and destination IP addresses. Typically you will also limit communication to desired times and services for additional security.

In Transparent mode, the FortiGate unit performs antivirus and antispam scanning on each packet as it passes through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another. For more information, see the [FortiGate Administration Guide](#), or [FortiGate Fundamentals Guide](#).

Example of VDOMs in Transparent mode

In this example, the FortiGate unit provides network protection to two organizations — Company A and Company B. Each company has different policies for incoming and outgoing traffic, requiring three different security policies and protection profiles.

VDOMs are not required for this configuration, but by using VDOMs the profiles and policies can be more easily managed on a per-VDOM basis either by one central administrator or separate administrators for each company. Also future expansion is simply a matter of adding additional VDOMs, whilst not disrupt the existing VDOMs.

For this example, firewalls are only included to deal with web traffic. This is to provide an example without making configuration unnecessarily complicated.

This example includes the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Configuring common items](#)
- [Creating virtual domains](#)
- [Configuring the Company_A VDOM](#)
- [Configuring the Company_B VDOM](#)
- [Configuring the VLAN switch and router](#)
- [Testing the configuration](#)

Network topology and assumptions

Each organization's internal network consists of a different range of IP addresses:

- 10.11.0.0/255.255.0.0 for Company A.


- 10.12.0.0/255.255.0.0 for Company B.

For the procedures in this section, it is assumed that you have enabled VDOM configuration on your FortiGate unit. For more information, see [“Enabling and accessing Virtual Domains” on page 1877](#).

The VDOM names are similar to the company names for easy recognition. The root VDOM cannot be renamed and is not used in this example.

Interfaces used in this example are port1 and port2. Some FortiGate models may not have interfaces with these names. port1 is an external interface. port2 is an internal interface.

Figure 186: VLAN and VDOM Transparent example network topology



General configuration steps

The following steps summarize the configuration for this example. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 [Configuring common items](#)
- 2 [Creating virtual domains](#)
- 3 [Configuring the Company_A VDOM](#)
- 4 [Configuring the Company_B VDOM](#)
- 5 [Configuring the VLAN switch and router](#)
- 6 [Testing the configuration](#)

Configuring common items

Both VDOMs require you configure UTM profiles. These will be configured the same way, but need to be configured in both VDOMs.

The relaxed profile allows users to surf websites they are not allowed to visit during normal business hours. Also a quota is in place to restrict users to one hour of access to these websites to ensure employees do not take long and unproductive lunches.

To create a strict web filtering profile - web-based manager

- 1 Go to the proper VDOM, and select *UTM Profiles > Web Filter > Profile*.
- 2 Select *Create New*.
- 3 Enter `strict` for the *Name*.
- 4 Expand FortiGuard Web Filtering, and select block for all Categories except Business Oriented, and Other.
- 5 Block all Classifications except Cached Content, and Image Search.
- 6 Ensure *FortiGuard Quota* for all Categories and Classifications is Disabled.
- 7 Select *OK*.

To create a strict web filtering profile - CLI

```
config vdom
  edit <vdom_name>
    config webfilter profile
      edit strict
        config ftgd-wf
          set allow g07 g08 g21 g22 c01 c03
          set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
        end
        set web-ftgd-err-log enable
      end
    end
  end
```

To create a relaxed web filtering profile - web-based manager

- 1 Go to the proper VDOM, and select *UTM Profiles > Web Filter > Profile*.
- 2 Select *Create New*.
- 3 Enter `relaxed` for the *Name*.
- 4 Expand FortiGuard Web Filtering, and select block for Potentially Security Violating Category, and Spam URL Classification.
- 5 Enable FortiGuard Quotas to allow 1 hour for all allowed Categories and Classifications.

Creating virtual domains

The FortiGate unit supports 10 virtual domains. Root is the default VDOM. It cannot be deleted or renamed. The root VDOM is not used in this example. New VDOMs are created for Company A and Company B

To create the virtual domains - web-based manager

- 1 With VDOMs enabled, select *System > VDOM > VDOM*.
- 2 Select *Create New*.
- 3 Enter `Company_A` for Name, and select *OK*.
- 4 Select *Create New*.
- 5 Enter `Company_B` for Name, and select *OK*.

To create the virtual domains - CLI

```
config system vdom
  edit Company_A
  next
  edit Company_B
end
```

Configuring the Company_A VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company_A VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating the Lunch schedule](#)
- [Configuring Company_A firewall addresses](#)
- [Creating Company_A security policies](#)

Adding VLAN subinterfaces

You need to create a VLAN subinterface on the port2 interface and another one on the port1 interface, both with the same VLAN ID.

To add VLAN subinterfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Name	VLAN_100_int
Interface	port2
VLAN ID	100
Virtual Domain	Company_A

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Name	VLAN_100_ext
Interface	port1
VLAN ID	100
Virtual Domain	Company_A

To add the VLAN subinterfaces - CLI

```
config system interface
  edit VLAN_100_int
    set interface port2
    set vlanid 100
    set vdom Company_A
  next
  edit VLAN_100_ext
    set interface port1
```

```
set vlanid 100
set vdom Company_A
end
```

Creating the Lunch schedule

Both organizations have the same lunch schedule, but only Company A has relaxed its security policy to allow employees more freedom in accessing the Internet during lunch. Lunch schedule will be Monday to Friday from 11:45am to 2:00pm (14:00).

To create a recurring schedule for lunchtime - web-based manager

- 1 In Company_A VDOM, go to *Firewall Objects > Schedule > Recurring*.
- 2 Select *Create New*.
- 3 Enter *Lunch* as the name for the schedule.
- 4 Select *Mon, Tues, Wed, Thu, and Fri*.
- 5 Set the *Start* time as 11:45 and set the *Stop* time as 14:00.
- 6 Select *OK*.

To create a recurring schedule for lunchtime - CLI

```
config vdom
edit Company_A
config firewall schedule recurring
edit Lunch
set day monday tuesday wednesday thursday friday
set start 11:45
set end 14:00
end
```

Configuring Company_A firewall addresses

For Company A, its networks are all on the 10.11.0.0 network, so restricting addresses to that domain provides added security.

To configure Company_A firewall addresses - web-based manager

- 1 In the Company_A VDOM, go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter *CompanyA* in the *Address Name* field.
- 4 Type 10.11.0.0/255.255.0.0 in the *Subnet / IP Range* field.
- 5 Select *OK*.

To configure vdomA firewall addresses - CLI

```
config firewall address
edit CompanyA
set type ipmask
set subnet 10.11.0.0 255.255.0.0
end
```


Creating Company_A security policies

A security policy can include varying levels of UTM protection. This example only deals with web filtering. The following security policies use the custom UTM `strict` and `relaxed` profiles configured earlier. See [“Configuring common items” on page 1927](#).

For these security policies, we assume that all protocols will be on their standard ports, such as port 80 for http traffic. If the ports are changed, such as using port 8080 for http traffic, you will have to create custom services for protocols with non-standard ports, and assign them different names.

The firewalls configured in this section are:

- internal to external — always deny all
- external to internal — always deny all
- internal to external — always allow all, UTM - web filtering: strict
- internal to external — Lunch allow all, UTM - web filtering:relaxed

Security policies allow packets to travel between the internal VLAN_100 interface to the external interface subject to the restrictions of the protection profile. Entering the policies in this order means the last one configured is at the top of the policy list, and will be checked first. This is important because the policies are arranged so if one does not apply the next is checked until the end of the list.

To configure Company_A security policies - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_int
Source Address	CompanyA
Destination Interface/Zone	VLAN_100_ext
Destination Address	all
Schedule	always
Service	all
Action	DENY

This policy is a catch all for outgoing traffic to ensure that if it doesn't match any of the other policies, it will not be allowed. This is standard procedure.

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_ext
Source Address	all
Destination Interface/Zone	VLAN_100_int
Destination Address	CompanyA
Schedule	always

Service	all
Action	DENY

This policy is a catch all for incoming traffic to ensure that if it doesn't match any of the other policies, it will not be allowed. This is standard procedure.

6 Select *Create New*.

7 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_int
Source Address	CompanyA
Destination Interface/Zone	VLAN_100_ext
Destination Address	all
Schedule	always
Service	all
Action	ACCEPT
UTM	Enable
Web Filtering	strict

This policy enforces strict scanning at all times, while allowing all traffic. It ensures company policies are met for network security.

8 Select *Create New*.

9 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_100_int
Source Address	CompanyA
Destination Interface/Zone	VLAN_100_ext
Destination Address	all
Schedule	Lunch
Service	all
Action	ACCEPT
UTM	enable
Web Filtering	relaxed

This policy provides relaxed protection during lunch hours — going from strict down to scan for protocol options and web filtering. AntiVirus and Email Filtering remain at strict for security — relaxing them would not provide employees additional access to the Internet and it would make the company vulnerable.

10 Verify that the policies entered appear in the list with the last policy (lunch) at the top, and the first policy (deny all) at the bottom. Otherwise traffic will not flow as expected.

To configure Company_A security policies - CLI

```
config vdom
  edit Company_A
    config firewall policy
      edit 1
        set srcintf VLAN_100_int
```

```

        set dstintf VLAN_100_ext
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule Lunch
        set UTM enabled
        set webfiltering relaxed
    next
    edit 3
        set srcintf VLAN_100_int
        set dstintf VLAN_100_ext
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule BusinessDay
        set service HTTP
        set profile_status enable
        set profile BusinessOnly
    end

```

Configuring the Company_B VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company B VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating Company_B service groups](#)
- [Configuring Company_B firewall addresses](#)
- [Configuring Company_B security policies](#)

Adding VLAN subinterfaces

You need to create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID.

To add VLAN subinterfaces - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Name	VLAN_200_int
Interface	port2
VLAN ID	200
Virtual Domain	Company_B

- 4 Select *Create New*.
- 5 Enter the following information and select *OK*:

Name	VLAN_200_ext
Interface	port1

VLAN ID	200
Virtual Domain	Company_B

To add the VLAN subinterfaces - CLI

```
config system interface
  edit VLAN_200_int
    set interface internal
    set vlanid 200
    set vdom Company_B
  next
  edit VLAN_200_ext
    set interface external
    set vlanid 200
    set vdom Company_B
end
```

Creating Company_B service groups

Company_B does not want its employees to use online gaming software or any online chat software except NetMeeting, which the company uses for net conferencing. To simplify the creation of a security policy for this purpose, you create a service group that contains all of the services you want to restrict. A security policy can manage only one service or one group. The administrator decided to simply name this group “Games” although it also restricts chat software.

To create a games service group - web-based manager

- 1 Go to *Firewall Objects > Service > Group*.
- 2 Select *Create New*.
- 3 Enter *Games* in the *Group Name* field.
- 4 For each of AOL, IRC, QUAKE, SIP-MSNmessenger and TALK, select the service in the *Available Services* list and select the right arrow to add it to the *Members* list.
- 5 Select *OK*.

To create a games and chat service group - CLI

```
config firewall service group
  edit Games
    set member IRC QUAKE AOL TALK
end
```

Configuring Company_B firewall addresses

Company B’s network is all in the 10.12.0.0 network. Security can be improved by only allowing traffic from IP addresses on that network.

To configure Company_B firewall address - web-based manager

- 1 In the Company_B VDOM, go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*.
- 3 Enter *new* in the *Address Name* field.
- 4 Type 10.12.0.0/255.255.0.0 in the *Subnet / IP Range* field.
- 5 Select *OK*.

To configure DEFdomain firewall addresses - CLI

```

config vdom
  edit Company_B
    config firewall address
      edit all
        set type ipmask
        set subnet 10.12.0.0 255.255.0.0
      end
    end
  end

```

Configuring Company_B security policies

Security policies allow packets to travel between the internal and external VLAN_200 interfaces subject to the restrictions of the protection profile.

To configure Company_B security policies - web-based manager

- 1 Go to *Policy > Policy*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_200_int
Source Address	all
Destination Interface/Zone	VLAN_200_ext
Destination Address	all
Schedule	BusinessDay
Service	games-chat
Action	DENY

This policy prevents the use of network games or chat programs (except NetMeeting) during business hours.

- 4 Enter the following information and select *OK*:

Source Interface/Zone	VLAN_200_int
Source Address	all
Destination Interface/Zone	VLAN_200_ext
Destination Address	all
Schedule	Lunch
Service	HTTP
Action	ACCEPT
Protection Profile	Relaxed

This policy relaxes the web category filtering during lunch hour.

- 5 Select *Create New*.

- 6 Enter the following information and select **OK**:

Source Interface/Zone	VLAN_200_int
Source Address	all
Destination Interface/Zone	VLAN_200_ext
Destination Address	all
Schedule	BusinessDay
Service	HTTP
Action	ACCEPT
Protection Profile	BusinessOnly

This policy provides rather strict web category filtering during business hours.

- 7 Select *Create New*.
- 8 Enter the following information and select **OK**:

Source Interface/Zone	VLAN_200_int
Source Address	all
Destination Interface/Zone	VLAN_200_ext
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
Protection Profile	Relaxed

Because it is last in the list, this policy applies to the times and services not covered in preceding policies. This means that outside of regular business hours, the Relaxed protection profile applies to email and web browsing, and online chat and games are permitted. Company B needs this policy because its employees sometimes work overtime. The other companies in this example maintain fixed hours and do not want any after-hours Internet access.

To configure Company_B security policies - CLI

```
config firewall policy
  edit 1
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
    set dstaddr all
    set schedule BusinessDay
    set service Games
    set action deny
  next
  edit 2
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
```

```
set dstaddr all
set action accept
set schedule Lunch
set service HTTP
set profile_status enable
set profile Relaxed
next
edit 3
set srcintf VLAN_200_int
set srcaddr all
set dstintf VLAN_200_ext
set dstaddr all
set action accept
set schedule BusinessDay
set service HTTP
set profile_status enable
set profile BusinessOnly
next
edit 4
set srcintf VLAN_200_int
set srcaddr all
set dstintf VLAN_200_ext
set dstaddr all
set action accept
set schedule always
set service ANY
set profile_status enable
set profile Relaxed
end
```

Configuring the VLAN switch and router

The Cisco switch is the first VLAN device internal passes through, and the Cisco router is the last device before the Internet or ISP.

This section includes the following topics:

- [Configuring the Cisco switch](#)
- [Configuring the Cisco router](#)

Configuring the Cisco switch

On the Cisco Catalyst 2900 ethernet switch, you need to define the VLANs 100, 200 and 300 in the VLAN database, and then add configuration files to define the VLAN subinterfaces and the 802.1Q trunk interface.

Add this file to Cisco VLAN switch:

```
!
interface FastEthernet0/1
  switchport access vlan 100
!
interface FastEthernet0/5
  switchport access vlan 300
!
interface FastEthernet0/6
  switchport trunk encapsulation dot1q
```

```
switchport mode trunk
!
```

Switch 1 has the following configuration:

Port 0/1	VLAN ID 100
Port 0/3	VLAN ID 200
Port 0/6	802.1Q trunk

Configuring the Cisco router

The configuration for the Cisco router in this example is the same as in the basic example, except we add VLAN_300. Each of the three companies has its own subnet assigned to it.

The IP addresses assigned to each VLAN on the router are the gateway addresses for the VLANs. For example, devices on VLAN_100 would have their gateway set to 10.11.0.1/255.255.0.0.

```
!
interface FastEthernet0/0
!
interface FastEthernet0/0.1
 encapsulation dot1Q 100
 ip address 10.11.0.1 255.255.0.0
!
interface FastEthernet0/0.3
 encapsulation dot1Q 200
 ip address 10.12.0.1 255.255.0.0
!
```

The router has the following configuration:

Port 0/0.1	VLAN ID 100
Port 0/0.3	VLAN ID 200
Port 0/0	802.1Q trunk

Testing the configuration

Use diagnostic commands, such as `tracert`, to test traffic routed through the network.

You should test traffic between the internal VLANs as well as from the internal VLANs to the Internet to ensure connectivity.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 1977](#).

This section includes the following topics:

- [Testing traffic from VLAN_100 to the Internet](#)
- [Testing traffic from VLAN_100 to VLAN_200](#)

Testing traffic from VLAN_100 to the Internet

In this example, a route is traced from VLANs to a host on the Internet. The route target is `www.example.com`.

From a host on VLAN_100, access a command prompt and enter this command:

```
C:\>tracert www.example.com
Tracing route to www.example.com [208.77.188.166]
```



```
over a maximum of 30 hops:
 1   <10 ms   <10 ms   <10 ms   10.100.0.1
...
14   172 ms   141 ms   140 ms   208.77.188.166
Trace complete.
```

The number of steps between the first and the last hop, as well as their IP addresses, will vary depending on your location and ISP. However, all successful tracerts to `www.example.com` will start and end with these lines.

Repeat the `tracert` for `VLAN_200`.

The `tracert` for each VLAN will include the gateway for that VLAN as the first step. Otherwise, the `tracert` should be the same for each VLAN.

Testing traffic from VLAN_100 to VLAN_200

In this example, a route is traced between two internal networks. The route target is a host on `VLAN_200`. The Windows `tracert` command `tracert` is used.

From `VLAN_100`, access a Windows command prompt and enter this command:

```
C:\>tracert 10.12.0.2
Tracing route to 10.12.0.2 over a maximum of 30 hops:
 1   <10 ms   <10 ms   <10 ms   10.100.0.1
 2   <10 ms   <10 ms   <10 ms   10.12.0.2
Trace complete.
```

You can repeat this for different routes in the topology. In each case the IP addresses will be the gateway for the starting VLAN, and the end point at the ending VLAN.



Inter-VDOM routing

In the past, virtual domains (VDOMs) were separate from each other—there was no internal communication. Any communication between VDOMs involved traffic leaving on a physical interface belonging to one VDOM and re-entering the FortiGate unit on another physical interface belonging to another VDOM to be inspected by firewall policies in both directions.

Inter-VDOM routing changes this. With VDOM links, VDOMs can communicate internally without using additional physical interfaces.

Inter-VDOM routing is the communication between VDOMs. VDOM links are virtual interfaces that connect VDOMs. A VDOM link contains a pair of interfaces with each one connected to a VDOM, and forming either end of the inter-VDOM connection.

This chapter contains the following sections:

- [Benefits of inter-VDOM routing](#)
- [Getting started with VDOM links](#)
- [Dynamic routing over inter-VDOM links](#)
- [HA virtual clusters and VDOM links](#)
- [Example of inter-VDOM routing](#)

Benefits of inter-VDOM routing

Inter-VDOM routing has a number of advantages over independent VDOM routing. These benefits include:

- [Freed-up physical interfaces](#)
- [More speed than physical interfaces](#)
- [Continued support for secure firewall policies](#)
- [Configuration flexibility](#)

Freed-up physical interfaces

Tying up physical interfaces on the FortiGate unit presents a problem. With a limited number of interfaces available, configuration options for the old style of communication between VDOMs are very limited. VLANs can be an answer to this, but they have some limitations.

For example, the FortiGate-800 has 8 physical ethernet ports. If they are assigned 2 per VDOM (one each for external and internal traffic) there can only be 4 VDOMs at most configured, not the 10 VDOMs the license will allow. Adding even one additional interface per VDOM to be used to communicate between VDOMs leaves only 2 VDOMs for that configuration, since it would required 9 interfaces for 3 VDOMs. Even using one physical interface for both external traffic and inter-VDOM communication would severely lower the available bandwidth for external traffic on that interface.

With the introduction of inter-VDOM routing, traffic can travel between VDOMs internally, freeing up physical interfaces for external traffic. Using the above example we can use the 4 VDOM configuration and all the interfaces will have their full bandwidth.

More speed than physical interfaces

Internal interfaces are faster than physical interfaces. Their speed depends on the FortiGate unit CPU and its load. That means that an inter-VDOM link interface will be faster than a outbound physical interface connected to another inbound physical interface.

Inter-VDOM links are CPU bound, and cannot be part of an accelerated pair of interfaces. However, while one virtual interface with normal traffic would be considerably faster than on a physical interface, the more traffic and more internal interfaces you configure, the slower they will become until they are slower than the physical interfaces. CPU load can come from other sources such as AV or content scanning. This produces the same effect—internal interfaces such as inter-VDOM links will be slower.

Continued support for secure firewall policies

VDOMs help to separate traffic based on your needs. This is an important step in satisfying regulations that require proof of secure data handling. This is especially important to health, law, accounting, and other businesses that handle sensitive data every day.

By keeping things separate, traffic has to leave the FortiGate unit and re-enter to change VDOMs. This forces traffic to go through the firewall when leaving and enter through another firewall, keeping traffic secure.

With inter-VDOM routing, the need for the physical interfaces is greatly reduced. However, firewall policies still need to be in place for traffic to pass through any interface, physical or virtual, and thus provide the same level of security both internally and externally. Configuration of firewall policies is the same for inter-VDOM links as for any other interface, and your data will continue to have the high level of security.

Configuration flexibility

A typical VDOM uses at least two interfaces, typically physical interfaces, one for internal and one for external traffic. Depending on the configuration, more interfaces may be required. The one exception to this is possibly one-armed IPS.

As explained earlier, the maximum number of VDOMs configurable on a FortiGate unit is the number of physical interfaces available divided by two. VLANs can increase the number by providing multiple virtual interfaces over a single physical interface, but VLANs have some limitations.

Using physical interfaces for inter-VDOM communication severely limits the number of possible configurations on your FortiGate unit, but inter-VDOM routing allows these connections to be moved inside the FortiGate unit. Using virtual interfaces, VDOM links, frees up the physical interfaces for external traffic. Using VDOM links on a FortiGate unit with 8 interfaces, you can have 4 VDOMs communicating with each other (meshed configuration) and continue to have 2 physical interfaces each for internal and external connections. This configuration would have required 20 physical interfaces without inter-VDOM routing. With inter-VDOM routing it only requires 8 physical interfaces, with the other 12 interfaces being internal VDOM links.

Inter-VDOM routing allows you to select [Standalone VDOM configuration](#), [Management VDOM configuration](#) and [Meshed VDOM configuration](#) without being limited by the number of physical interfaces on your FortiGate unit.

Getting started with VDOM links

Once VDOMs are configured on your FortiGate unit, configuring inter-VDOM routing and VDOM-links is very much like creating a VLAN interface.

VDOM-links are managed through the web-based manager or CLI. In the web-based manager, VDOM link interfaces are managed in the network interface list.

This section includes the following topics:

- [Viewing VDOM links](#)
- [Creating VDOM links](#)
- [Deleting VDOM links](#)

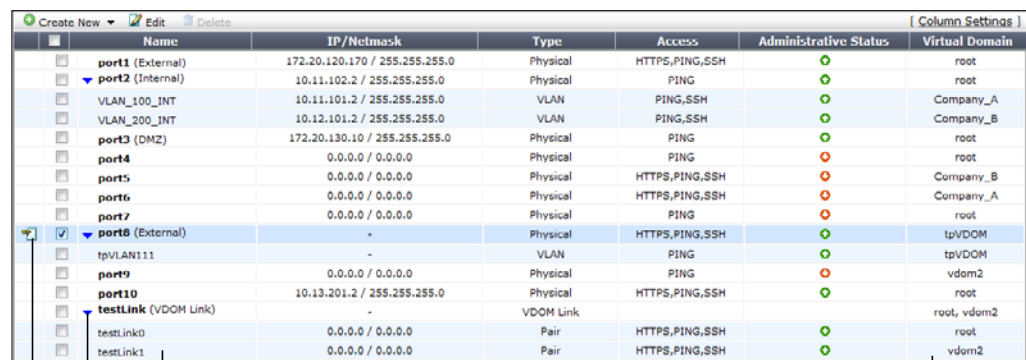
Viewing VDOM links

VDOM links are displayed on the network interface list in the web-based manager.

You can view VDOM links only if you are using a super_admin account and in global configuration.

To view the network interface list, in the Global menu go to *System > Network > Interface*.

Figure 187: Interface list displaying interface names and information



	Name	IP/Netmask	Type	Access	Administrative Status	Virtual Domain
	port1 (External)	172.20.120.170 / 255.255.255.0	Physical	HTTPS,PING,SSH		root
	port2 (Internal)	10.11.102.2 / 255.255.255.0	Physical	PING		root
	VLAN_100_INT	10.11.101.2 / 255.255.255.0	VLAN	PING,SSH		Company_A
	VLAN_200_INT	10.12.101.2 / 255.255.255.0	VLAN	PING,SSH		Company_B
	port3 (DMZ)	172.20.130.10 / 255.255.255.0	Physical	PING		root
	port4	0.0.0.0 / 0.0.0.0	Physical	PING		root
	port5	0.0.0.0 / 0.0.0.0	Physical	HTTPS,PING,SSH		Company_B
	port6	0.0.0.0 / 0.0.0.0	Physical	HTTPS,PING,SSH		Company_A
	port7	0.0.0.0 / 0.0.0.0	Physical	PING		root
	port8 (External)	-	Physical	HTTPS,PING,SSH		tpvDOM
	tpvLAN111	-	VLAN	PING		tpvDOM
	port9	0.0.0.0 / 0.0.0.0	Physical	PING		vdom2
	port10	10.13.201.2 / 255.255.255.0	Physical	HTTPS,PING,SSH		root
	testLink (VDOM Link)	-	VDOM Link			root, vdom2
	testLink0	0.0.0.0 / 0.0.0.0	Pair	HTTPS,PING,SSH		root
	testLink1	0.0.0.0 / 0.0.0.0	Pair	HTTPS,PING,SSH		vdom2

Annotations in the image:

- Create New**: Points to the green plus icon in the top left.
- Description of interface**: Points to the 'Name' column header.
- VDOM link pair**: Points to the 'testLink0' and 'testLink1' rows.
- VDOM link interface**: Points to the 'testLink (VDOM Link)' row.
- VDOM**: Points to the 'Virtual Domain' column header.

Create New	<p>Select the arrow to create a new interface or VDOM link. Interface options include VLAN, Aggregate, Redundant, or loopback interfaces.</p> <p>For more information, see “Creating VDOM links” on page 1944.</p>
Edit	<p>Select to change interface configuration for the selected interface.</p> <p>This option not available if no interfaces or multiple interfaces are selected.</p>
Delete	<p>Select to remove an interface from the list. One or more interfaces must be selected for this option to be available.</p> <p>You cannot delete permanent physical interfaces, or any interfaces that have configuration referring to them. See “Deleting VDOM links” on page 1946 or “Deleting an interface” on page 1906.</p>

Column Settings	Select to change which information is displayed about the interfaces, and in which order the columns appear. Use to display VDOM, VLAN, and other information.
Checkbox	Select the checkbox for an interface to edit or delete that interface. Select multiple interfaces to delete those interfaces. Optionally select the check box at the top of the column to select or unselect all checkboxes.
Name	The name of the interface. The name of the VDOM link (<code>vlink1</code>) has an expand arrow to display or hide the pair of VDOM link interfaces. For more information, see “Viewing VDOM links” on page 1943 .
IP/Netmask	The IP address and netmask assigned to this interface.
Type	The type of interface such as physical, VLAN, or VDOM link pair.
Access	The protocols allowed for administrators to connect to the FortiGate unit.
Administrative Status	The status of this interface, either set to up (active) or down (disabled).
Virtual Domain	The virtual domain this interface belongs to. For more information on VDOMs, see “Virtual Domains in NAT/Route mode” on page 1903 .

Creating VDOM links

VDOM links connect VDOMs together to allow traffic to pass between VDOMs as per firewall policies. Inter-VDOM links are virtual interfaces that are very similar to VPN tunnel interfaces except inter-VDOM links do not require IP addresses. See [“IP addresses are not required for inter-VDOM links” on page 1945](#).

To create a VDOM link, you first create the point-to-point interface, and then bind the two interface objects associated with it to the virtual domains.

In creating the point-to-point interface, you also create two additional interface objects by default. They are called `vlink10` and `vlink11` - the interface name you chose with a 1 or a 0 to designate the two ends of the link.

Once the interface objects are bound, they are treated like normal FortiGate interfaces and need to be configured just like regular interfaces.

The assumptions for this example are as follows:

- Your FortiGate unit has VDOMs enabled and you have 2 VDOMs called `customer1` and `customer2` already configured. For more information on configuring VDOMs see [“Only a super_admin administrator account such as the default “admin” account can create, disable, or delete VDOMs. That account can create additional administrators for each VDOM.” on page 1895](#).
- You are using a super_admin account



Inter-VDOM links cannot include VDOMs in Transparent mode.

To configure an inter-VDOM link - web-based manager

- 1 For Current VDOM, select Global..
- 2 Select *System > Network > Interface*.
- 3 Select *Create New > VDOM link*, enter the following information, and select OK.

Name	vlink1 (The name can be up to 11 characters long. Valid characters are letters, numbers, "-", and "_". No spaces are allowed.)
Interface #0	
Virtual Domain	customer1
IP/Netmask	10.11.12.13/255.255.255.0
Administrative Access	HTTPS, SSL
Interface #1	
Virtual Domain	customer2
IP/Netmask	172.120.100.13/255.255.255.0
Administrative Access	HTTPS, SSL



If your inter-VDOM links have names longer than 8 characters, and you upgrade from FortiOS 3.0 MR3, the names will be truncated to 8 characters and will not function. The solution is to change the names of your inter-VDOM links before you upgrade.

To configure an inter-VDOM link - CLI

```
config global
  config system vdom-link
    edit vlink1
  end
  config system interface
    edit vlink10
      set vdom customer1
    next
    edit vlink11
      set vdom customer2
    end
```

Once you have created and bound the interface ends to VDOMs, configure the appropriate firewall policies and other settings that you require. To confirm the inter-VDOM link was created, find the VDOM link pair and use the expand arrow to view the two VDOM link interfaces. You can select edit to change any information.

IP addresses are not required for inter-VDOM links

Besides being virtual interfaces, here is one main difference between inter-VDOM links and regular interfaces—inter-VDOM links do not require IP addresses. This introduces three possible situations with inter-VDOM links that are:

- **unnumbered** - an inter-VDOM link with no IP addresses for either end of the tunnel

- **half numbered** - an inter-VDOM link with one IP address for one end and none for the other end
- **full numbered** - an inter-VDOM link with two IP addresses, one for each end.

An IP address is not required for inter-VDOM links because it is an internal connection that can be referred to by the interface name in firewall policies, and other system references.

Not using an IP address in the configuration can speed up and simplify configuration for you. Also you will not use up all the IP addresses in your subnets if you have many inter-VDOM links.

Half or full numbered interfaces are required if you are doing NAT, either SNAT or DNAT as you need an IP number on both ends to translate between.

You can use unnumbered interfaces in static routing, by naming the interface and using 0.0.0.0 for the gateway. Running traceroute will not show the interface in the list of hops. However you can see the interface when you are sniffing packets, which is useful for troubleshooting.

Deleting VDOM links

When you delete the VDOM link, the two link objects associated with it will also be deleted. You cannot delete the objects by themselves. The example uses a VDOM routing connection called "vlink1". Removing vlink1 will also remove its two link objects vlink10 and vlink11.



Before deleting the VDOM link, ensure all policies, firewalls, and other configurations that include the VDOM link are deleted, removed, or changed to no longer include the VDOM link.

To remove a VDOM link - web-based manager

- 1 For Current VDOM, select Global..
- 2 Select *System > Network > Interface*.
- 3 Select *Delete* for the VDOM link *vlink1*.

To remove a VDOM link - CLI

```
config global
  config system vdom-link
    delete vlink1
  end
```

For more information, see the [FortiGate CLI Reference](#).



Once the inter-VDOM link is created, you cannot change these IP addresses without deleting the link.

Inter-VDOM configurations

By using fewer physical interfaces to inter-connect VDOMs, inter-VDOM links provide you with more configuration options.

None of these configurations use VLANs to reduce the number of physical interfaces. It is generally assumed that an internal or client network will have its own internal interface and an external interface to connect to its ISP and the Internet.

These inter-VDOM configurations can use any FortiGate model with possible limitations based on the number of physical interfaces. VLANs can be used to work around these limitations.

In the following inter-VDOM diagrams, red indicates the physical FortiGate unit, grey indicate network connections external to the FortiGate unit, and black is used for inter-VDOM links and VDOMs.


This section includes the following topics:

- [Standalone VDOM configuration](#)
- [Independent VDOMs configuration](#)
- [Management VDOM configuration](#)
- [Meshed VDOM configuration](#)

Standalone VDOM configuration

The standalone VDOM configuration uses a single VDOM on your FortiGate unit — the root VDOM that all FortiGate units have by default. This is the VDOM configuration you are likely familiar with. It is the default configuration for FortiGate units before you create additional VDOMs.

Figure 188: Standalone VDOM



The configuration shown in [Figure 188](#) has no VDOM inter-connections and requires no special configurations or settings.


The standalone VDOM configuration can be used for simple network configurations that only have one department or one company administering the connections, firewalls and other VDOM-dependent settings.

However, with this configuration, keeping client networks separate requires many interfaces, considerable firewall design and maintenance, and can quickly become time consuming and complex. Also, configuration errors for one client network can easily affect other client networks, causing unnecessary network downtime.

Independent VDOMs configuration

The independent VDOMs configuration uses multiple VDOMs that are completely separate from each other. This is another common VDOM configuration.

Figure 189: Independent VDOMs



This configuration has no communication between VDOMs and apart from initially setting up each VDOM, it requires no special configurations or settings. Any communication between VDOMs is treated as if communication is between separate physical devices.


The independent inter-VDOM configuration can be used where more than one department or one company is sharing the FortiGate unit. Each can administer the connections, firewalls and other VDOM-dependent settings for only its own VDOM. To each company or department, it appears as if it has its own FortiGate unit. This configuration reduces the amount of firewall configuration and maintenance required by dividing up the work.

However, this configuration lacks a management VDOM for VDOMs 1, 2, and 3. This is illustrated in Figure 50. This management VDOM would enable an extra level of control for the FortiGate unit administrator, while still allowing each company or department to administer its own VDOM.

Management VDOM configuration

In the management VDOM configuration, the root VDOM is the management VDOM. The other VDOMs are connected to the management VDOM with inter-VDOM links. There are no other inter-VDOM connections.

Figure 190: Management VDOM configuration



The inter-VDOM links connect the management VDOM to the other VDOMs. This does not require any physical interfaces, and the bandwidth of inter-VDOM links can be faster than physical interfaces, depending on the CPU workload.

Only the management VDOM is connected to the Internet. The other VDOMs are connected to internal networks. All external traffic is routed through the management VDOM using inter-VDOM links and firewall policies between the management VDOM and each VDOM. This ensures the management VDOM has full control over access to the Internet, including what types of traffic are allowed in both directions. There is no communication directly between the non-root VDOMs. Security is greatly increased with only one point of entry and exit. Only the management VDOM needs to be fully managed to ensure network security in this case. Each client network can manage its own configuration without compromising security or bringing down another client network.

The management VDOM configuration is ideally suited for a service provider business. The service provider administers the management VDOM with the other VDOMs as customers. These customers do not require a dedicated IT person to manage their network. The service provider controls the traffic and can prevent the customers from using banned services and prevent Internet connections from initiating those same banned services. One example of a banned service might be Instant Messaging (IM) at a company concerned about intellectual property. Another example could be to limit bandwidth used by file-sharing applications without banning that application completely. Firewall policies control the traffic between the customer VDOM and the management VDOM and can be customized for each customer.


The management VDOM configuration is limited in that the customer VDOMs have no inter-connections. In many situations this limitation is ideal because it maintains proper security. However, some configurations may require customers to communicate with each other, which would be easier if the customer VDOMs were inter-connected.

Meshed VDOM configuration

The meshed VDOMs configuration, including partial and full mesh, has VDOMs inter-connected with other VDOMs. There is no special feature to accomplish this—they are just complex VDOM configurations.

Partial mesh means only some VDOMs are inter-connected. In a full mesh configuration, all VDOMs are inter-connected to all other VDOMs. This can be useful when you want to provide full access between VDOMs but handle traffic differently depending on which VDOM it originates from or is going to.

Figure 191: Meshed VDOMs



With full access between all VDOMs being possible, it is extra important to ensure proper security. You can achieve this level of security by establishing extensive firewall policies and ensuring secure account access for all administrators and users.

Meshed VDOM configurations can become complex very quickly, with full mesh VDOMs being the most complex. Ensure this is the proper solution for your situation before using this configuration. Generally, these configurations are seen as theoretical and are rarely deployed in the field.

Dynamic routing over inter-VDOM links

BGP is supported over inter-VDOM links. Unless otherwise indicated, routing works as expected over inter-VDOM links.

If an inter-VDOM link has no assigned IP addresses to it, it may be difficult to use that interface in dynamic routing configurations. For example BGP requires an IP address to define any BGP router added to the network.

In OSPF, you can configure a router using a router ID and not its IP address. In fact, having no IP address avoids possible confusing between which value is the router ID and which is the IP address. However for that router to become adjacent with another OSPF router it will have to share the same subnet, which is technically impossible without an IP address. For this reason, while you can configure an OSPF router using an IP-less inter-VDOM link, it will likely be of limited value to you.

In RIP the metric used is hop count. If the inter-VDOM link can reach other nodes on the network, such as through a default route, then it may be possible to configure a RIP router on an inter-VDOM link. However, once again it may be of limited value due to limitations.

As stated earlier, BGP requires an IP address to define a router — an IP-less inter-VDOM link will not work with BGP.

In Multicast, you can configure an interface without using an IP address. However that interface will be unable to become an RP candidate. This limits the roles available to such an interface.

HA virtual clusters and VDOM links

FortiGate HA is implemented by configuring two or more FortiGate units to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate unit, processing network traffic and providing normal security services such as firewall, VPN, IPS, virus scanning, web filtering, and spam filtering.

Virtual clustering extends HA features to provide failover protection and load balancing for a FortiGate unit operating with virtual domains. A virtual cluster consists of a cluster of two FortiGate units operating with virtual domains. Traffic on different virtual domains can be load balanced between the cluster units.

With virtual clusters (vclusters) configured, inter-VDOM links must be entirely within one vcluster. You cannot create links between vclusters, and you cannot move a VDOM that is linked into another virtual cluster. If your FortiGate units are operating in HA mode, with multiple vclusters when you create the vdom-link, the CLI command `config system vdom-link` includes an option to set which vcluster the link will be in. For more information, see the [FortiGate HA Guide](#).

What is virtual clustering?

Virtual clustering is an extension of the FGCP for FortiGate units operating with multiple VDOMs enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

Virtual clustering and failover protection

Virtual clustering operates on a cluster of two (and only two) FortiGate units with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGate units in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM partitioning from the web-based manager by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI config system `ha` command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.

Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

Note: If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.

In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved. From the web-based manager you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.

From the CLI you configure VDOM partitioning by setting the HA mode to a-p. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1. You do the same for virtual cluster 2 by entering the `config secondary-vcluster` command.

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit. If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

Example of inter-VDOM routing

This example shows how to configure a FortiGate unit to use inter-VDOM routing.

This section contains the follow topics:

- [Network topology and assumptions](#)
- [Creating the VDOMs](#)
- [Configuring the physical interfaces](#)
- [Configuring the VDOM links](#)
- [Configuring the firewall and UTM settings](#)
- [Testing the configuration](#)

Network topology and assumptions

Two departments of a company, Accounting and Sales, are connected to one FortiGate-800 unit. To do its work, the Sales department receives a lot of email from advertising companies that would appear to be spam if the Accounting department received it. For this reason, each department has its own VDOM to keep firewall policies and other configurations separate. A management VDOM makes sense to ensure company policies are followed for traffic content.

The traffic between Accounting and Sales will be email and HTTPS only. It could use a VDOM link for a meshed configuration, but we will keep from getting too complex. With the configuration, inter-VDOM traffic will have a slightly longer path to follow than normal—from one department VDOM, through the management VDOM, and back to the other department VDOM. Since inter-VDOM links are faster than physical interfaces, this longer path should not be noticed.

Firewall policies will be in place. For added security, firewall policies will allow only valid office services such as email, web browsing, and FTP between either department and the Internet. Any additional services that are required can be added in the future.

The company uses a single ISP to connect to the Internet. The ISP uses DHCP to provide an IP address to the FortiGate unit. Both departments use the same ISP to reach the Internet.

Other assumptions for this example are as follows:


- Your FortiGate unit has interfaces labelled port1 through port4 and VDOMs are not enabled.
- You are using the super_admin account.
- You have the FortiClient application installed.

- You are familiar with configuring interfaces, firewalls, and other common features on your FortiGate unit.



All configuration is available to a super_admin. A non-super_admin account may also perform certain procedures, but only for the VDOM that the account has access to. For more information, see “Administrators in Virtual Domains” on page 1898.

Figure 192: Management VDOM for two departments



General configuration steps

This example includes the following general steps. For best results, follow the steps in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

- 1 [Creating the VDOMs](#)
- 2 [Configuring the physical interfaces](#)
- 3 [Configuring the VDOM links](#)
- 4 [Configuring the firewall and UTM settings](#)
- 5 [Testing the configuration](#)

Creating the VDOMs

This procedure enables VDOMs and creates the Sales and Accounting VDOMs.

To create the VDOMs - web-based manager

- 1 Log in as the super_admin administrator.
- 2 Go to *System > Dashboard > Status > System Information > Virtual Domain*, and select *Enable*.
- 3 Log in again.

- 4 Go to *System > VDOM > VDOM*.
- 5 Select *Create New*, enter *Accounting* for the VDOM Name, and select *OK*.
- 6 Select *Create New*, enter *Sales* for the VDOM Name, and select *OK*.

To create the VDOMs - CLI

```
config system global
    set vdom enable
end

config system vdom
    edit Accounting
    next
    edit Sales
    next
end
```

Configuring the physical interfaces

Next, the physical interfaces must be configured. This example uses three interfaces on the FortiGate unit - port2 (internal), port3(dmz), and port1(external). port2 and port3 interfaces each have a department's network connected. port1 is for all traffic to or from the Internet and will use DHCP to configure its IP address, which is common with many ISPs.

To configure the physical interfaces - web-based manager

- 1 In *Current VDOM*, select *Global*.
- 2 Select *System > Network > Interface*.
- 3 Select *Edit* for the port2 interface, enter the following information, and select *OK*.

Alias	AccountingLocal
Virtual Domain	Accounting
Addressing mode	Manual
IP/Netmask	172.100.1.1/255.255.0.0
Administrative Access	HTTPS, PING, SSH
Description	This is the accounting department internal interface.

- 4 Select *Edit* for the port3 interface, enter the following information, and select *OK*.

Alias	SalesLocal
Virtual Domain	Sales
Addressing mode	Manual
IP/Netmask	192.168.1.1/255.255.0.0
Administrative Access	HTTPS, PING, SSH
Description	This is the sales department internal interface.

- 5 Select *Edit* for the port1 interface, enter the following information, and select *OK*.

Alias	ManagementExternal
Virtual Domain	root
Addressing Mode	DHCP
Distance	5
Retrieve default gateway from server	Enable
Override internal DNS	Enable
Administrative Access	HTTPS, SSH, SNMP
Description	This is the accounting department internal interface.



When the mode is set to DHCP or PPoE on an interface you can set the distance field. This is the administrative distance for any routes learned through the gateway for this interface. The gateway is added to the static route table with these values. A lower distance indicates a preferred route.

To configure the physical interfaces - CLI

```
config global
  config system interface
    edit port2
      set alias AccountingLocal
      set vdom Accounting
      set mode static
      set ip 172.100.1.1 255.255.0.0
      set allowaccess https ping ssh
      set description "The accounting dept internal interface"
    next
  edit port3
    set alias SalesLocal
    set vdom Sales
    set mode static
    set ip 192.168.1.1 255.255.0.0
    set allowaccess https ping ssh
    set description "The sales dept. internal interface"
  next
  edit port1
    set alias ManagementExternal
    set vdom root
    set mode DHCP
    set distance 5
    set gwdetect enable
    set dns-server-override enable
    set allowaccess https ssh snmp
    set description "The systemwide management interface."
  end
```

Configuring the VDOM links

To complete the connection between each VDOM and the management VDOM, you need to add the two VDOM links; one pair is the Accounting - management link and the other is for Sales - management link.

When configuring inter-VDOM links, you do not have to assign IP addresses to the links unless you are using advanced features such as dynamic routing that require them. Not assigning IP addresses results in faster configuration, and more available IP addresses on your networks.

If you require them, or if you simply want to assign IP addresses for clarity can do so.

To configure the Accounting and management VDOM link - web-based manager

- 1 In *Current VDOM*, select *Global*.
- 2 Select *System > Network > Interface*.
- 3 Select the expand arrow to select *Create New > VDOM link*.
- 4 Enter the following information, and select *OK*.

Name	AccountVlnk
Interface #0	
Virtual Domain	Accounting
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Accounting VDOM side of the link.
Interface #1	
Virtual Domain	root
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Management VDOM side of the link.

To configure the Accounting and management VDOM link - CLI

```
config global
  config system vdom-link
    edit AccountVlnk
    next
  end
  config system interface
    edit AccountVlnk0
      set vdom Accounting
      set ip 0.0.0.0 0.0.0.0
      set allowaccess https ping ssh
      set description "Accounting side of the VDOM link"
    next
    edit AccountVlnk1
      set vdom root
      set ip 0.0.0.0 0.0.0.0
      set allowaccess https ping ssh
      set description "Management side of the VDOM link"
    end
```

To configure the Sales and management VDOM link - web-based manager

- 1 In *Current VDOM*, select *Global*.
- 2 Select *System > Network > Interface*.
- 3 Select the expand arrow and select *Create New > VDOM link*.
- 4 Enter the following information, and select *OK*.

Name	SalesVlnk
Interface #0	
Virtual Domain	Sales
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Sales VDOM side of the link.
Interface #1	
Virtual Domain	root
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Management VDOM side of the link.

To configure the Sales and management VDOM link - CLI

```

config global
  config system vdom-link
    edit SalesVlnk
  end
  config system interface
    edit SalesVlnk0
      set vdom Accounting
      set ip 0.0.0.0 0.0.0.0
      set allowaccess https ping ssh
      set description "Sales side of the VDOM link"
    next
    edit SalesVlnk1
      set vdom root
      set ip 0.0.0.0 0.0.0.0
      set allowaccess https ping ssh
      set description "Management side of the VDOM link"
    end
  end
end

```

Configuring the firewall and UTM settings

With the VDOMs, physical interfaces, and VDOM links configured the firewall must now be configured to allow the proper traffic. Firewalls are configured per-VDOM, and firewall objects must be created for each VDOM separately.

For this example, the firewall group of services allowed between the internal networks and the Internet are the basic services for web browsing, file transfer, and email. These include: HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, and SMTP.

The only services allowed between Sales and Accounting are secure web browsing (HTTPS) and email (POP3 and SMTP)



The limited number of services ensures security between departments. The list of services can be expanded in the future if needed.

UTM settings will block all non-essential business websites while logging all web traffic, scan and file filter all web and email protocols, and block game and peer-to-peer applications using application control.

For added security, FortiClient is required on internal computers with AntiVirus scanning configured. This is enforced by *Endpoint NAC* in firewall policies.

Using firewall addresses makes the firewall policies easier to read. Also if any changes need to be made in the future, you can simply update the addresses without changing the firewall policies. The addresses required are:

- `AccountingLocal` - all traffic from the internal accounting network
- `AccountingVlnk` - all traffic from the VDOM link between accounting and management VDOMs
- `SalesLocal` - all traffic from the internal sales network
- `SalesVlnk` - all traffic from the VDOM link between sales and management VDOM.

The Accounting VDOM requires `AccountingLocal`, `AccountingVlnk`, and `SalesLocal`. The Sales VDOM requires `SalesLocal`, `SalesVlnk`, and `AccountingLocal`.

The firewall policies required on the Accounting VDOM are

- `AccountingLocal` to Internet
- Internet to `AccountingLocal`
- `SalesLocal` to `AccountingLocal`
- `AccountingLocal` to `SalesLocal`

The firewall policies required on the Sales VDOM are

- `SalesLocal` to Internet
- Internet to `SalesLocal`
- `SalesLocal` to `AccountingLocal`
- `AccountingLocal` to `SalesLocal`

This section includes the following topics:

- [Configuring firewall service groups](#)
- [Configuring UTM settings for the Accounting VDOM](#)
- [Configuring firewall settings for the Accounting VDOM](#)
- [Configuring UTM settings for the Sales VDOM](#)
- [Configuring firewall settings for the Sales VDOM](#)
- [Configuring firewall settings between the Accounting and Sales VDOMs](#)

Configuring firewall service groups

Service groups are an easy way to manage multiple services, especially if the same services are used on different networks.

The two service groups used here are intended for normal office traffic to the Internet, and for restricted traffic between departments. In both cases network traffic will be limited to the services listed to prevent any potential security risks or bandwidth-robbing applications.

These service groups can be changed as needed to either include additional valid services that are being used on the network, or to exclude services that are not required. Also, custom services can be created as needed for applications that are not listed.

To configure two firewall service groups - web-based manager

- 1 In *Current VDOM*, select Accounting.
- 2 Go to *Firewall Objects > Service > Group*.
- 3 Select *Create New*, enter the following information, and select *OK*.

Group Name	OfficeServices
Members	HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, PING, SMTP

- 4 Select *Create New*, enter the following information, and select *OK*.

Group Name	AccountingSalesServices
Members	HTTPS, POP3, PING, SMTP

To configure two firewall service groups - CLI

```
config vdom
  edit Accounting
    config firewall service group
      edit OfficeServices
        set member HTTP HTTPS SSL FTP DNS NTP POP3 PING SMTP
      next
      edit AccountingSalesServices
        set member HTTPS POP3 PING SMTP
      end
    end
  end
```

Configuring UTM settings for the Accounting VDOM

UTM settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that

- the business environment is free from viruses
- employees do not surf grossly inappropriate websites, and
- employees do not use games or peer-to-peer applications at work.

To configure web filtering for the Accounting VDOM - web-based manager

- 1 In *Current VDOM*, select Accounting.
- 2 Go to *UTM Profiles > Web Filter > Profile*.
- 3 Select *Create New*.
- 4 Enter `webStrict` for the *Name*.

- 5 Select the arrow to expand the *FortiGuard Web Filtering* section.
- 6 Block all *Categories* except Business Oriented, Other, and Unrated.
- 7 Block all *Classifications* except Image Search..
- 8 Log all *Categories* and *Classifications*.
- 9 Select OK.

To configure web filtering for the Accounting VDOM - CLI

```
config vdom
  edit Accounting
    config webfilter profile
      edit webStrict
        config ftgd-wf
          set allow g07 g08 g21 g22 c01 c03
          set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
        end
        set web-ftgd-err-log enable
      end
    end
  end
```

To configure AntiVirus for the Accounting VDOM - web-based manager

- 1 In *Current VDOM*, select Accounting.
- 2 Go to *UTM Profiles > AntiVirus > Profile*.
- 3 Select *Create New*.
- 4 Enter *avStrict* for the *Name*.
- 5 Enable *Scan* for all protocols.
- 6 Enable *File filter* for all protocols, and select *built-in-patterns* for *Option*.
- 7 Enable logging for both *Scan* and *File Filter*.
- 8 Select OK.

To configure AntiVirus for the Accounting VDOM - CLI

```
config vdom
  edit Accounting
    config antivirus profile
      edit avStrict
        config http
          set options scan file-filter
        end
        config ftp
          set options scan file-filter
        end
        config imap
          set options scan file-filter
        end
        config pop3
          set options scan file-filter
        end
        config smtp
          set options scan file-filter
        end
      end
    end
  end
```

```

config nntp
  set options scan file-filter
end
config im
  set options scan file-filter
end
set filepattable 1
set av-virus-log enable
set av-block-log enable
end
end

```

To configure application control for the Accounting VDOM - web-based manager

- 1 In *Current VDOM*, select *Accounting*.
- 2 Go to *UTM Profiles > Application Control > Application Sensor*.
- 3 Select *Create New* (+ button at top right of page).
- 4 Enter `appStrict` for *Name* and select *OK*.
- 5 Select *Create New*.
- 6 In *Filters*, set *Category* to *game*.
- 7 In *Applications/Settings*, enter the following, and select *OK*.

Action	Block
Packet Logging	Enable

- 8 Select *Create New*.
- 9 In *Filters*, set *Category* to *p2p*.
- 10 In *Applications/Settings*, enter the following, and select *OK*.

Action	Block
Packet Logging	Enable

- 11 Select *Apply*.

To configure application control for the Accounting VDOM - CLI

```

config vdom
  edit Accounting
    config application list
      edit appStrict
        config entries
          edit 1
            set category 2
          next
          edit 2
            set category 8
          end
        end
      end
    end
  end
end

```


Configuring firewall settings for the Accounting VDOM

This configuration includes two firewall addresses and two firewall policies for the Accounting VDOM - one for the internal network, and one for the VDOM link with the management VDOM (root).

For added security, all traffic allowed will be scanned. Only valid office traffic will be allowed using the service group `OfficeServices`. The FortiClient application must be used to ensure additional protection for the sensitive accounting information.

All sales and accounting computers have the FortiClient application installed, so the firewall policies check that FortiClient is installed and that antivirus scanning is enabled.

Note the spelling of `AccountVlnk` which is due to the eleven character limit on VDOM link names.

To configure firewall addresses - web-based manager

- 1 For Current VDOM, select *Accounting*.
- 2 Select *Firewall Objects > Address > Address*
- 3 Select *Create New*, enter the following information, and select *OK*.

Address Name	AccountingLocal
Type	Subnet/ IP Range
Subnet / IP Range	172.100.0.0
Interface	port1

- 4 Select *Create New*, enter the following information, and select *OK*.

Address Name	AccountManagement
Type	Subnet/ IP Range
Subnet / IP Range	10.0.1.0
Interface	AccountVlnk

To configure firewall addresses - CLI

```
config vdom
  edit Accounting
    config firewall address
      edit AccountingLocal
        set type iprange
        set subnet 172.100.0.0
        set associated-interface port1
      next
    edit AccountManagement
      set type iprange
      set subnet 10.0.1.0
      set associated-interface AccountVlnk
    end
  end
```

To configure protocol options for Accounting VDOM - web-based manager

- 1 In *Current VDOM*, select *Accounting*.
- 2 Select *Policy > Policy > Protocol Options*.

- 3 Select *Create New*.
- 4 Enter `default` for the *Name*.
- 5 Select *OK*.

To configure the firewall policies from AccountingLocal to the Internet - web-based manager

- 1 In *Current VDOM*, select *Accounting*.
- 2 Go to *Policy > Policy*.
- 3 Select *Create New*, enter the following information, and then select *OK*.

Source Interface/Zone	port2
Source Address	AccountingLocal
Destination Interface/Zone	AccountVlnk
Destination Address	AccountManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
UTM	enabled
Protocol Option	default
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict
Enable Endpoint NAC	Enforce_FortiClient_AV

- 4 In *Current VDOM*, select *root*.
- 5 Go to *Policy > Policy*.
- 6 Select *Create New*, enter the following information, and then select *OK*.

Source Interface/Zone	AccountVlnk
Source Address	AccountManagement
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
UTM	enable
Protocol Option	default
Web Filtering	webStrict
AntiVirus Filtering	avStrict

Application Control	appStrict
Enable Endpoint NAC	disabled

To configure the firewall policies from AccountingLocal to Internet - CLI

```

config vdom
  edit Accounting
    config firewall policy
      edit 1
        set srcintf "port2"
        set dstintf "AccountVlnk"
        set srcaddr "AccountingLocal"
        set dstaddr "AccountManagement"
        set action accept
        set schedule "always"
        set service "OfficeServices"
        set nat enable
        set utm-status enable
        set av-profile avStrict
        set webfilter-profile webStrict
        set application-list appStrict
        set profile-protocol-options default
        set endpoint-check enable
        set endpoint-profile "FortiClient_installed"
      end
    end
  end

config vdom
  edit root
    config firewall policy
      edit 2
        set srcintf AccountVlnk
        set dstintf port1
        set srcaddr AccountManagement
        set dstaddr all
        set action accept
        set schedule always
        set service OfficeServices
        set nat enable
        set utm-status enable
        set av-profile "scan"
        set webfilter-profile "scan"
        set application-list "AppControlList"
        set profile-protocol-options default
        set endpoint-check disable
      end
    end
  end

```

To configure the firewall policies from Internet to AccountingLocal - web-based manager

- 1 In *Current VDOM*, select root.
- 2 Go to *Policy > Policy*.

- 3 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	AccountVlnk
Destination Address	AccountManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
UTM	enable
Protocol Option	default
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict
Enable Endpoint NAC	disabled

- 4 In *Current VDOM*, select *Accounting*.
 5 Go to *Policy > Policy*.
 6 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	AccountVlnk
Source Address	AccountManagement
Destination Interface/Zone	port2
Destination Address	AccountingLocal
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
UTM	enable
Protocol Option	default
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict
Enable Endpoint NAC	disabled

To configure the firewall policies from Internet to AccountingLocal - CLI

```
config vdom
  edit root
    config firewall policy
```

```
edit 3
    set srcintf port1
    set dstintf AccountVlnk
    set srcaddr all
    set dstaddr AccountManagement
    set action accept
    set schedule always
    set service OfficeServices
    set nat enable
    set utm-status enable
    set av-profile avStrict
    set webfilter-profile webStrict
    set application-list appstrict
    set profile-protocol-options default
    set endpoint-check disable
end
end
config vdom
    edit Accounting
        config firewall policy
            edit 4
                set srcintf AccountVlnk
                set dstintf port2
                set srcaddr AccountManagement
                set dstaddr AccountingLocal
                set action accept
                set schedule always
                set service OfficeServices
                set nat enable
                set utm-status enable
                set av-profile avStrict
                set webfilter-profile webStrict
                set application-list appstrict
                set profile-protocol-options default
                set endpoint-check disable
            end
        end
    end
end
```

Configuring UTM settings for the Sales VDOM

UTM settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that

- the business environment is free from viruses
- employees do not surf grossly inappropriate websites, and
- employees do not use games or peer-to-peer applications at work.

Note that Sales web traffic is different from Accounting, and web filtering is different to account for this.

To configure web filtering for the Sales VDOM - web-based manager

- 1 In *Current VDOM*, select Sales.
- 2 Go to *UTM Profiles > Web Filter > Profile*.
- 3 Select *Create New*.

- 4 Enter `webStrict` for the *Name*.
- 5 In *FortiGuard Categories*, select all of the categories except *Bandwidth Consuming*, *General Interest - Business* and *Unrated*.
- 6 In *Change Action for Selected Categories* select *Block*.
- 7 Select *Apply*.

To configure web filtering for the Sales VDOM - CLI

```
config vdom
  edit Sales
    config webfilter profile
      edit webStrict
        config ftgd-wf
          set allow g07 g08 g21 g22 c01 c03
          set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
        end
        set web-ftgd-err-log enable
      end
    end
  end
```

To configure AntiVirus for the Sales VDOM - web-based manager

- 1 In *Current VDOM*, select *Sales*.
- 2 Go to *UTM Profiles > AntiVirus > Profile*.
- 3 Select *Create New*.
- 4 Enter `avStrict` for the *Name*.
- 5 Enable virus scan for all protocols.
- 6 Select *Apply*.

To configure AntiVirus for the Sales VDOM - CLI

```
config vdom
  edit Sales
    config antivirus profile
      edit "avStrict"
        config http
          set options scan file-filter
        end
        config ftp
          set options scan file-filter
        end
        config imap
          set options scan file-filter
        end
        config pop3
          set options scan file-filter
        end
        config smtp
          set options scan file-filter
        end
        config nntp
          set options scan file-filter
        end
      end
    end
  end
```

```

config im
    set options scan file-filter
end
set filepattable 1
set av-virus-log enable
set av-block-log enable
end
end

```

To configure application control for the Sales VDOM - web-based manager

- 1 In *Current VDOM*, select *Accounting*.
- 2 Go to *UTM Profiles > Application Control > Application Sensor*.
- 3 Select *Create New* (+ button at top right of page).
- 4 Enter *appStrict* for *Name* and select *OK*.
- 5 Select *Create New*.
- 6 In *Filters*, set *Category* to *game*.
- 7 In *Applications/Settings*, enter the following, and select *OK*.

Action	Block
Packet Logging	Enable

- 8 Select *Create New*.
- 9 In *Filters*, set *Category* to *p2p*.
- 10 In *Applications/Settings*, enter the following, and select *OK*.

Action	Block
Packet Logging	Enable

- 11 Select *Apply*.

To configure application control for the Sales VDOM - CLI

```

config vdom
    edit Sales
        config application list
            edit "appStrict"
                config entries
                    edit 1
                        set category 2
                    next
                    edit 2
                        set category 8
                    end
                end
            end
        end
    end
end

```

Configuring firewall settings for the Sales VDOM

Like the Accounting firewall settings, this configuration includes two firewall addresses and two firewall policies for the sales VDOM: one for the internal network, and one for the VDOM link with the management VDOM.

When entering the CLI commands, the number of the firewall policies must be high enough to be a new policy. Depending on the number of firewall policies on your FortiGate unit, this may require starting at a higher number than the 6 required for the default configuration. This number is added automatically when you configure firewall policies using the web manager interface.

The FortiClient application must be used on Sales network computers to ensure additional protection for the sensitive information and for protection against spam.

To configure firewall addresses - web-based manager

- 1 In *Current VDOM*, select *Sales*.
- 2 Go to *Firewall Objects > Address > Address*.
- 3 Select *Create New*, enter the following information, and select *OK*.

Address Name	SalesLocal
Type	Subnet / IP Range
Subnet / IP Range	172.100.0.0
Interface	port3

- 4 Go to *Firewall Objects > Addresses*.
- 5 Select *Create New*, enter the following information, and select *OK*.

Address Name	SalesManagement
Type	Subnet / IP Range
Subnet / IP Range	10.0.1.0
Interface	SalesVlnk

To configure the firewall addresses - CLI

```
config vdom
  edit Sales
    config firewall address
      edit SalesLocal
        set type iprange
        set subnet 172.100.0.0
        set associated-interface port2
      next
      edit SalesManagement
        set type iprange
        set subnet 10.0.1.0
        set associated-interface SalesVlnk
      end
    end
  end
```

To configure the firewall policies from SalesLocal to the Internet - web-based manager

- 1 In *Current VDOM*, select *Sales*.
- 2 Go to *Policy > Policy*.

- 3 Select *Create New*, enter the following information, and select OK.

Source Interface/Zone	port3
Source Address	SalesLocal
Destination Interface/Zone	SalesVlnk
Destination Address	SalesManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Log Allowed Traffic	enabled
Enable Endpoint Control Check	disabled
Redirect Non-conforming Clients to Download Portal	enabled

- 4 In *Current VDOM*, select *root*.
- 5 Go to *Policy > Policy*.
- 6 Select *Create New*, enter the following information, and select OK.

Source Interface/Zone	SalesVlnk
Source Address	SalesManagement
Destination Interface/Zone	external
Destination Address	all
Schedule	always
Service	OfficeServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled
Enable Endpoint Control Check	disabled

To configure the firewall policies from SalesLocal to the Internet - CLI

```
config vdom
  edit root
    config firewall policy
      edit 6
        set srcintf port2
        set srcaddr SalesLocal
        set dstintf SalesVlnk
        set dstaddr SalesManagement
        set schedule always
        set service OfficeServices
        set action accept
        set profile-status enable
        set profile scan
        set logtraffic enable
        set endpoint-check enable
```

```

        set endpoint-redir-portal enable
    end
end

config vdom
    edit Sales
        config firewall policy
            edit 7
                set srcintf SalesVlnk
                set srcaddr SalesManagement
                set dstintf external
                set dstaddr all
                set schedule always
                set service OfficeServices
                set action accept
                set profile-status enable
                set profile scan
                set logtraffic enable
                set endpoint-check enable
            end
        end
    end
end

```

To configure the firewall policies from the Internet to SalesLocal - web-based manager

- 1 In Current VDOM, select *root*.
- 2 Go to *Policy > Policy*.
- 3 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	external
Source Address	all
Destination Interface/Zone	SalesVlnk
Destination Address	SalesManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled
Enable Endpoint Control Check	disabled

- 4 In *Current VDOM*, select *Sales*.
- 5 Go to *Policy > Policy*.
- 6 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	SalesVlnk
Source Address	SalesManagement
Destination Interface/Zone	port2
Destination Address	SalesLocal

Schedule	always
Service	OfficeServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled
Enable Endpoint Control Check	disabled
Redirect Non-conforming Clients to Download Portal	enabled

To configure the firewall policies from the Internet to SalesLocal - CLI

```

config vdom
  edit root
    config firewall policy
      edit 8
        set srcintf external
        set srcaddr all
        set dstintf SalesVlnk
        set dstaddr SalesManagement
        set schedule always
        set service OfficeServices
        set action accept
        set profile-status enable
        set profile scan
        set logtraffic enable
        set endpoint-check enable
        set endpoint-redir-portal enable
      end
    end
  end

config vdom
  edit Sales
    config firewall policy
      edit 9
        set srcintf SalesVlnk
        set srcaddr SalesManagement
        set dstintf port2
        set dstaddr SalesLocal
        set schedule always
        set service OfficeServices
        set action accept
        set profile-status enable
        set profile scan
        set logtraffic enable
        set endpoint-check enable
        set endpoint-redir-portal enable
      end
    end
  end

```

Configuring firewall settings between the Accounting and Sales VDOMs

Firewall policies are required for any communication between each internal network and the Internet. Policies are also required for the two internal networks to communicate with each other through the management VDOM.

The more limited AccountingSalesServices group of services will be used between Sales and Accounting to ensure the traffic is necessary business traffic only. These policies will result in a partially meshed VDOM configuration. The FortiClient application must be used to ensure additional protection for the sensitive accounting information.

Two firewall policies are required to allow traffic in both directions between Sales and Accounting.

To configure the firewall policy between Sales and Accounting on the management VDOM - web-based manager

- 1 For Current VDOM, select *root*.
- 2 Go to *Policy > Policy*.
- 3 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	SalesVlnk
Source Address	SalesManagement
Destination Interface/Zone	AccountVlnk
Destination Address	AccountingManagement
Schedule	always
Service	AccountingSalesServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled
Enable Endpoint Control Check	disabled
Redirect Non-conforming Clients to Download Portal	enabled

- 4 Go to *Policy > Policy*.
- 5 Select *Create New*, enter the following information, and select *OK*.

Source Interface/Zone	AccountVlnk
Source Address	AccountingManagement
Destination Interface/Zone	SalesVlnk
Destination Address	SalesManagement
Schedule	always
Service	AccountingSalesServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled

Enable Endpoint Control Check	disabled
Redirect Non-conforming Clients to Download Portal	enabled

To configure the firewall policy between Sales and Accounting on the management VDOM - CLI

```

config vdom
  edit root
    config system firewall policy
      edit 9
        set srcintf SalesVlnk
        set srcaddr SalesManagement
        set dstintf AccountVlnk
        set dstaddr AccountManagement
        set schedule always
        set service AccountingSalesServices
        set action accept
        set profile-status enable
        set profile scan
        set logtraffic enable
        set endpoint-check enable
        set endpoint-redir-portal enable
      next
      edit 10
        set srcintf AccountVlnk
        set srcaddr AccountManagement
        set dstintf SalesVlnk
        set dstaddr SalesManagement
        set schedule always
        set service AccountingSalesServices
        set action accept
        set profile-status enable
        set profile scan
        set logtraffic enable
        set endpoint-check enable
        set endpoint-redir-portal enable
      end
    end
  end
end

```

Testing the configuration

Once the inter-VDOM routing has been configured, tests must be conducted to confirm proper operation. If there are any problems, use the troubleshooting tips to resolve them.

This section includes the following topics:

- [Testing connectivity](#)
- [Troubleshooting Tips](#)

Testing connectivity

Testing connectivity ensures that physical networking connections as well as FortiGate unit interface configurations, including firewall policies, are properly configured.

The easiest way to test connectivity is to use the `ping` and `tracert` commands to confirm the connectivity of different routes on the network. Include testing:

- from AccountingLocal to Internet
- from Internet to AccountingLocal
- from SalesLocal to Internet
- from Internet to SalesLocal
- from AccountingLocal to SalesLocal.

When using the commands on a Windows computer, go to a command line prompt and enter either `ping <IP address>` or `tracert <IP address>`.


When using the commands on a FortiGate unit, go to the CLI and enter either `exec ping <IP address>` or `exec traceroute <IP address>`.

Troubleshooting Tips

When there are problems with connectivity, the following troubleshooting tips will help resolve the issues.

- If a multiple hop test, such as `tracert`, is not successful then reduce it to a single hop to simplify the test. Test each link of the path to see which hop is down. If all hops are up, check the FortiGate unit policies to ensure they allow basic traffic to flow as expected.
- If ping does not work, confirm that the FortiGate unit interfaces have Ping enabled and also ensure Ping is enabled in the firewall policies. Otherwise the Ping traffic will be blocked.
- If one protocol does not work but others do work, check the FortiGate unit firewall policies for that one protocol to ensure it is allowed.
- If there are unexplained connectivity problems, check the local computer to ensure it does not have a software firewall running that may be blocking traffic. MS Windows computers have a firewall running by default that can cause problems.

For additional troubleshooting, see [“Troubleshooting Virtual Domains” on page 1977](#).



Troubleshooting Virtual Domains

When you are configuring VDOMs you may run into some issues. This section provides answers to some common issues with VDOMs.

This section includes:

- [VDOM admin having problems gaining access](#)
- [FortiGate unit running very slowly](#)
- [General VDOM tips and troubleshooting](#)

VDOM admin having problems gaining access

With VDOMs configured, administrators have an extra layer of permissions and may have problems accessing their information.

Confirm the admin's VDOM

Each administrator account, other than the super_admin account, is tied to one specific VDOM. That administrator is not able to access any other VDOM. It may be possible they are trying to access the wrong VDOM.

Confirm the VDOM's interfaces

An administrator can only access their VDOM through interfaces that are assigned to that VDOM. If interfaces on that VDOM are disabled or unavailable there will be no method of accessing that VDOM by its local administrator. The super_admin will be required to either bring up the interfaces, fix the interfaces, or move another interface to that VDOM to restore access.

Confirm the VDOMs admin access

As with all FortiGate units, administration access on the VDOM's interfaces must be enabled for that VDOM's administrators to gain access. For example if SSH is not enabled, that is not available to administrators.

To enable admin access, the super_admin will go to the global *System > Network > Interface* page, and for the interface in question enable the admin access.

FortiGate unit running very slowly

You may experience a number of problems resulting from your FortiGate unit being overloaded. These problems may appear as:

- CPU and memory threshold limits exceeded on a continual basis
- AV failopen happening on a regular basis
- dropped traffic or sessions due to lack of resources

These problems are caused by a lack of system resources. There are a number of possible reasons for this.

Too many VDOMs

If you have configured many VDOMs on your system, past the default ten VDOMs, this could easily be your problem.

Each VDOM you create on your FortiGate unit requires system resources to function - CPU cycles, memory, and disk space. When there are too many VDOMs configured there are not enough resources for operation. This may be a lack of memory in the session table, or no CPU cycles for processing incoming IPS traffic, or even a full disk drive.

Go to *System > VDOM* and see the number of configured VDOMs on your system. If you are running 250 or more VDOMs, you must have a FortiGate 5000 chassis. Otherwise you need to reduce the number of VDOMs on your system to fix the problem. Even if you have the proper hardware, you may encounter noticeably slow throughput if you are using advanced features such as UTM or deep content inspection with many configured VDOMs.

One or more VDOMs are consuming all the resources

If you have sufficient hardware to support the number of VDOMs you are running, check the global resources on your FortiGate unit. At a glance it will tell you if you are running out of a particular resource such as sessions, or users. If this is the case, you can then check your VDOMs to see if one particular VDOM is using more than its share of resources. If that is the case you can change the resource settings to allow that VDOM (or those VDOMs) fewer resources and in turn allow the other VDOMs access to those resources.

Too many UTM features in use

If you are running 250 or more VDOMs and have a FortiGate 5000 chassis, it is still possible that you are running too many features for the FortiGate unit to support all those VDOMs. To support 250 or more VDOMs, FortiGate units cannot run advanced UTM features. Instead they are limited to less processor intensive features that do not require stateful inspection.

It is likely that reducing the UTM features in use even with fewer VDOM configuration will greatly improve overall system performance and should be considered as an option.

Finally it is possible that your FortiGate unit configuration is incorrect in some other area, which is using up all your resources. For example, forgetting that you are running a network sniffer on an interface will create significant amounts of traffic that may prevent normal operation.

General VDOM tips and troubleshooting

Besides ping and traceroute, there are additional tools for troubleshooting your VDOM configurations. These include packet sniffing and debugging the packet flow.

Perform a sniffer trace

When troubleshooting networks, it helps to look inside the headers of packets to determine if they are traveling along the route you expect that they are. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate unit has NP2 interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP2 interfaces, you should disable offloading on those interfaces.

What can sniffing packets tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the Fortigate unit is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing, which prevents an IP packet from being forwarded if its Source IP does not either belong to a locally attached subnet (local interface), or be part of the routing between the FortiGate and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the FortiGate unit and cause many features to be turned off.

Note If you configure virtual IP addresses on your Fortigate unit, it will use those addresses in preference to the physical IP addresses. You will notice this when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

How do you sniff packets

When you are using VDOMs, you must be in a VDOM to access the `diag sniffer` command. At the global level, the command is not available. This is limit the packets only to the ones on your VDOM, and protects the privacy of other VDOM clients.

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose>
<count>
```

To stop the sniffer, type `CTRL+C`.

<interface_name>	The name of the interface to sniff, such as “port1” or “internal”. This can also be “any” to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. “none” indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <code><CTRL C></code> .

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the port1 interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
```

```

filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh
3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh
1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh
4225311614 ack 3314279933

```

For a more advanced example of packet sniffing, the following commands will report packets on any interface travelling between a computer with the host name of PC1 and the computer with the host name of PC2. With verbosity 4 and above, the sniffer trace will display the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type CTRL+C. Note that PC1 and PC2 may be VDOMs.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and
icmp" 4
```

The following sniffer CLI command includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for instance PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or
arp" 4
```

Debug the packet flow

Traffic should come in and leave the VDOM. If you have determined that network traffic is not entering and leaving the VDOM as expected, debug the packet flow.

Debugging can only be performed using CLI commands. Debugging the packet flow requires a number of debug commands to be entered as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate unit has NP2 interfaces that are offloading traffic, this will change the packet flow. Before performing the debug on any NP2 interfaces, you should disable offloading on those interfaces.

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate unit and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```

FGT# diag debug enable
FGT# diag debug flow filter add <PC1>
FGT# diag debug flow show console enable
FGT# diag debug flow trace start 100
FGT# diag debug enable

```

The `start 100` argument in the above list of commands will limit the output to 100 packets from the flow. This is useful for looking at the flow without flooding your log or your display with too much information.

To stop all other debug activities, enter the command:

```
FGT# diag debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching Firewall Policy, and is in turn blocked by the FortiGate unit. The denied message indicates the traffic was blocked. Note that even with VDOMs not enabled, vd-root is still shown.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825
  msg="vd-root received a packet(proto=6,
    192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924
  msg="allocate a new session-013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597
  msg="find a route: gw-192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg="
  Denied by forward policy check"
```




Chapter 12 High Availability

This FortiOS Handbook chapter contains the following sections:

[Solving the High Availability problem](#) describes the high availability problem and introduces the FortiOS solutions described in this document (FGCP, VRRP, and standalone session synchronization).

[An introduction to the FortiGate Clustering Protocol \(FGCP\)](#) introduces the FGCP clustering protocol and many of its features and terminology.

[Configuring and connecting HA clusters](#) describes configuring HA clusters and contains HA clustering configuration examples.

[Configuring and connecting virtual clusters](#) describes configuring HA virtual clusters and contains virtual clustering configuration examples.

[Configuring and operating FortiGate full mesh HA](#) describes configuring FortiGate Full mesh HA and contains a full mesh HA configuration example.

[Operating a cluster](#) describes how to operate a cluster and includes detailed information about how various FortiGate systems operate differently in a cluster.

[HA and failover protection](#) describes in detail how FortiGate HA device failover, link failover, and session failover work.

[HA and load balancing](#) describes in detail how FortiGate HA active-active load balancing load balances sessions.

[HA with third-party products](#) describes how FortiGate units interact with third-party products.

[VRRP](#) describes FortiOS support of the Virtual Router Redundancy Protocol (VRRP) and its use for high availability.

[TCP session synchronization](#) describes the FortiGate standalone session synchronization feature and its use for high availability.





Solving the High Availability problem

The basic high availability (HA) problem for TCP/IP networks and security gateways is keeping network traffic flowing. Uninterrupted traffic flow is a critical component for online systems and media because critical business processes quickly come to a halt when the network is down.

The security gateway is a crucial component of most networks since all traffic passes through it. A standalone network security gateway is a single point of failure that is vulnerable to any number of software or hardware problems that could compromise the device and bring all traffic on the network to a halt.

A common solution to the high availability problem is to eliminate the security gateway as single point of failure by introducing redundancy. With two or more redundant security gateways, if one fails, the remaining one or more gateways keep the traffic flowing. FortiOS provides three redundancy solutions: VRRP, TCP session synchronization, and Fortinet's proprietary FortiGate Cluster Protocol (FGCP) high availability solution.




FGCP HA, TCP session synchronization and VRRP are not compatible. You cannot configure more than one of these high availability methods on the same FortiGate unit.

A strong and flexible High availability solution is required for many mission-critical firewall and UTM applications. Each FortiOS high availability solution can be fine tuned to fit into many different network scenarios.

FortiGate Cluster Protocol (FGCP)

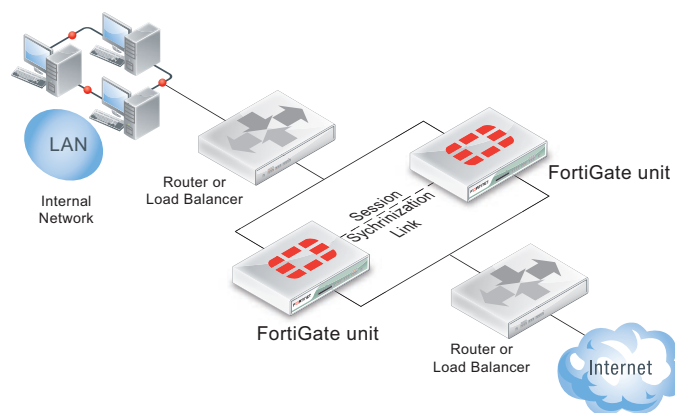
FGCP HA provides a solution for two key requirements of critical enterprise networking components: enhanced reliability and increased performance. Enhanced reliability is achieved through device failover protection, link failover protection and remote link failover protection. Increased performance is achieved through active-active HA load balancing. Extended FGCP features include full mesh HA and virtual clustering. You can also fine tune the performance of the FGCP to change how a cluster forms and shares information among cluster units and how the cluster responds to failures. When configured onto your network an FGCP cluster appears to be a single FortiGate unit operating in NAT/Route or Transparent mode. If a failover occurs, the cluster recovers quickly and automatically and also sends administrator notifications so that the problem that caused the failure can be corrected and any failed equipment restored.



The FGCP is compatible with most network environments and most networking equipment. While initial configuration is relatively quick and easy, a large number of tools and configuration options are available to fine tune the cluster for most situations.

TCP session synchronization

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two or more FortiGate units can be integrated into the load balancing configuration by enabling TCP session synchronization (also called standalone session synchronization) between them. The external load balancers or routers can distribute TCP sessions among the FortiGate units and TCP session synchronization keeps both FortiGate unit's session tables synchronized.




If one of the FortiGate units fails, session failover occurs and active TCP sessions fail over to the unit that is still operating. This failover occurs without any loss of data. As well, the external routers or load balancers detect the failover and re-distribute all sessions to the unit that is still operating.

Unlike the FCGP, TCP session synchronization does not include configuration synchronization. In fact, the configurations of the two FortiGate units (sometimes called peers) are not identical since in most cases the peers would have different IP addresses. Also unlike the FCGP, load balancing and session failover is done by external routers or load balancers instead of by the FGCP. The FortiGate units just perform session synchronization which supports the session failover.

VRRP

FortiGate units can function as master or backup Virtual Router Redundancy Protocol (VRRP) routers and can be quickly and easily integrated into a network that has already deployed VRRP. A FortiGate unit can be integrated into a VRRP group with any third-party VRRP devices and VRRP can provide redundancy between multiple FortiGate units.



In a VRRP configuration, when a FortiGate unit operating as the master unit fails, a backup unit takes its place and continues processing network traffic. If the backup unit is a FortiGate unit, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate unit is back on line. You can include different FortiGate models in the same VRRP group.

FortiOS supports VRRP between two or more FortiGate units and between FortiGate units and third-party routers that support VRRP. Using VRRP you can assign VRRP routers as master or backup routers. The master router processes traffic and the backup routers monitor the master router and can begin forwarding traffic if the master fails. Similar to the FGCP you can configuration VRRP between multiple FortiGate units to provide redundancy. You can also create a VRRP group with a FortiGate units and any routers that support VRRP.

In a VRRP configuration that consists of one FortiGate unit and one router, normally the FortiGate unit would be the master and all traffic would be processed by the FortiGate unit. If the FortiGate unit fails, all traffic switches to the router. Network connectivity is maintained even though FortiGate security features will be unavailable until the FortiGate unit can is back on line.



An introduction to the FortiGate Clustering Protocol (FGCP)

A FortiGate HA cluster consists of two to four FortiGate units configured for HA operation. Each FortiGate unit in a cluster is called a cluster unit. All cluster units must be the same FortiGate model with the same FortiOS firmware build installed. All cluster units must also have the same hardware configuration (for example, the same AMC modules installed in the same slots, the same number of hard disks and so on) and be running in the same operating mode (NAT/Route mode or Transparent mode).



You can create an FGCP cluster of up to four FortiGate units.

On startup, after configuring the cluster units with the same HA configuration, the cluster units use the FortiGate Clustering Protocol (FGCP) to find other FortiGate units configured for HA operation and to negotiate to create a cluster. During cluster operation, the FGCP shares communication and synchronization information among the cluster units. This communication and synchronization is called the FGCP heartbeat or the HA heartbeat. Often, this is shortened to just heartbeat. For a cluster to form, the cluster units must be able to communicate using their configured heartbeat interfaces.

The cluster uses the FGCP to select the primary unit, and to provide device, link and session failover. The FGCP also manages the two HA modes; active-passive (failover HA) and active-active (load balancing HA).


This chapter describes.

- [About the FGCP](#)
- [Configuring a FortiGate unit for FGCP HA operation](#)
- [Active-passive and active-active HA](#)
- [Identifying the cluster and cluster units](#)
- [Device failover, link failover, and session failover](#)
- [Primary unit selection](#)
- [HA override](#)
- [FortiGate HA compatibility with PPPoE and DHCP](#)
- [Hard disk configuration and HA](#)
- [HA Best practices](#)
- [FGCP HA terminology](#)
- [HA web-based manager options](#)

About the FGCP

FortiGate HA is implemented by configuring two or more FortiGate units to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate unit, processing network traffic and providing normal security services such as firewalling, Unified Threat Management (UTM) and VPN services.

Figure 193: HA cluster installed between an internal network and the Internet



Inside the cluster the individual FortiGate units are called cluster units. These cluster units share state and configuration information. If one cluster unit fails, the other units in the cluster automatically replace that unit, taking over the work that the failed unit was doing. After the failure, the cluster continues to process network traffic and provide normal FortiGate services with virtually no interruption.

Every FortiGate cluster contains one primary unit (also called the master unit) and one or more subordinate units (also called slave or backup units). The primary unit controls how the cluster operates. The roles that the primary and subordinate units play in the cluster depend on the mode in which the cluster operates. See [“Active-passive HA \(failover protection\)” on page 1998](#) and [“Active-active HA \(load balancing and failover protection\)” on page 1998](#).

The ability of an HA cluster to continue providing firewall services after a failure is called failover. FortiGate HA failover means that your network does not have to rely on one FortiGate unit to continue functioning. You can install additional units and form an HA cluster. Other units in the cluster will take over if one of the units fails.

A second HA feature, called load balancing, can be used to increase performance. A cluster of FortiGate units can increase overall network performance by sharing the load of processing network traffic and providing Unified Threat Management (UTM) services. The cluster appears to your network to be a single device, adding increased performance without changing your network configuration.

Virtual clustering extends HA features to provide failover protection and load balancing for a FortiGate operating with virtual domains. A virtual cluster consists of a cluster of two FortiGate units operating with virtual domains. Traffic on different virtual domains can be load balanced between the cluster units. For details about virtual clustering, see [“Configuring and connecting virtual clusters” on page 2089](#).

FortiGate models that support redundant interfaces can be configured to support a clustering configuration called full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster. For details about full mesh HA, see [“Configuring and operating FortiGate full mesh HA” on page 2111](#).

FGCP failover protection

The FGCP provides IP/MAC takeover for failover protection by assigning virtual MAC addresses to the primary cluster unit and then sending gratuitous ARP packets from the primary unit interfaces to reprogram the network.

Failover times can be less than a second under optimal conditions. You can fine tune failover performance for your network by adjusting cluster status checking, routing table update, and wait timers.

An HA cluster fails over if the primary unit experiences a device or link failure. The cluster can detect link failures for connections to the primary unit using port monitoring and for connections between downstream network components using remote IP monitoring. To compensate for a link failover, the cluster maintains active links to keep traffic flowing between high-priority networks. Port and remote IP monitoring can be fine tuned without disrupting cluster operation.

Session Failover

FGCP session failover maintains TCP, SIP and IPsec VPN sessions after a failure. Session failover does not failover UDP, multicast, ICMP, or SSL VPN sessions. Session failover may not be required for all networks because many TCP/IP protocols can resume sessions on their own. Supporting session failover adds extra overhead to cluster operations and can be disabled to improve cluster performance if its not required.

Load Balancing

Active-active HA load balances resource-intensive UTM processing among all cluster units to provide better UTM performance than a standalone FortiGate unit. If network traffic consists of mainly TCP sessions, the FGCP can also load balance all TCP sessions to improve TCP performance in some network configurations. You can use accelerated FortiGate interfaces to also accelerate HA load balancing and HA load balancing schedules can be adjusted to optimize performance for the traffic mix on your network. Weighted load balancing can be used to control the relative amount of sessions processed by each cluster unit.

Virtual Clustering

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGate units operating with multiple VDOMS enabled. Not only does virtual clustering provide failover protection for a multiple VDOM configuration, but a virtual cluster can load balance traffic between the cluster units. Load balancing with virtual clustering is quite efficient and load balances all traffic (not just UTM and TCP traffic). Its possible to fine tune virtual clustering load balancing in real time to actively optimize load sharing between the cluster units without affecting the smooth operation of the cluster.

Full Mesh HA

High availability improves the reliability of a network by replacing a single point of failure (a single FortiGate unit) with a cluster that can maintain network traffic if one of the cluster units fails. However, in a cluster configuration single points of failure remain. Full mesh HA removes these single points of failure by allowing you to connect redundant switches to each cluster interface. Full mesh HA is achieved by configuring 802.3ad aggregate or redundant interfaces on the FortiGate unit and connecting redundant switches to these interfaces. Configuration is a relatively simple extension of the normal aggregate/redundant interface and HA configurations.

Cluster Management

FortiOS HA provides a wide range of cluster management features:

- Automatic continuous configuration synchronization. You can get a cluster up and running almost as quickly as a standalone FortiGate unit by performing a few basic steps to configure HA settings and minimal network settings on each cluster unit. When the cluster is operating you can make start configuring FortiGate features such as UTM and IPsec VPN in the same way as for a standalone FortiGate unit. All configuration changes (even complex changes such as switching to multiple VDOM mode or from NAT/Route to Transparent mode) are synchronized among all cluster units.
- Firmware upgrades/downgrades. Upgrading or downgrading cluster firmware is similar to upgrading or downgrading standalone FortiGate firmware. The Firmware is uploaded once to the primary unit and the cluster automatically upgrades or downgrades all cluster units in one operation with minimal or no service interruption.
- Individual cluster unit management. In some cases you may want to manage individual cluster units. You can do so from cluster CLI by navigating to each cluster unit. You can also use the reserved management interface feature to give each cluster unit its own IP address and default route. You can use the reserved management interfaces and IP addresses to connect to the GUI and CLI of each cluster unit and configure an SNMP server to poll each cluster unit.
- Removing and adding cluster units. In one simple step any unit (even the primary unit) can be removed from a cluster and given a new IP address. The cluster keeps operating as it was; the transition happening without interrupting cluster operation. Any unit can also be added to an operating cluster without disrupting network traffic. All you have to do is connect the new unit and change its HA configuration to match the cluster's. The cluster automatically finds and adds the unit and synchronizes its configuration with the cluster.
- Debug and diagnose commands. A full range of debug and diagnose commands can be used to report on HA operation and find and fix problems.
- Logging and reporting. All cluster units can be configured to record all log messages. These messages can be stored on the individual cluster units or sent to a FortiAnalyzer unit. You can view all cluster unit log messages by logging into any cluster unit.
- FortiManager support. FortiManager understands FortiOS HA and automatically recognizes when you add a FortiOS cluster to the FortiManager configuration.

Configuring a FortiGate unit for FGCP HA operation

Each FortiGate unit in the cluster must have the same HA configuration. Once the cluster is connected, you can configure it in the same way as you would configure a standalone FortiGate unit. The following procedures set the HA mode to active-passive and sets the HA password to HA_pass.

To configure a FortiGate unit for HA operation - web-based manager

- 1 Power on the FortiGate unit to be configured.
- 2 Log into the web-based manager.
- 3 On the Dashboard *System Information* dashboard widget, beside *Host Name* select *Change*.
- 4 Enter a new Host Name for this FortiGate unit.
Changing the host name makes it easier to identify individual cluster units when the cluster is operating.
- 5 Go to *System > Config > HA* and change the following settings:

Mode	Active-Passive
Group Name	Example_cluster
Password	HA_pass The password must be the same for all FortiGate units in the cluster.

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

- 6 Select *OK*.
The FortiGate unit negotiates to establish an HA cluster. When you select *OK* you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all ARP table entries). You may be able to delete the ARP table of your management PC from a command prompt using a command similar to `arp -d`.
- 7 Power off the FortiGate unit.
- 8 Repeat this procedure for all of the FortiGate units in the cluster.
Once all of the units are configured, continue with [“Connecting a FortiGate HA cluster” on page 1996](#).

To configure a FortiGate unit for HA operation - CLI

- 1 Power on the FortiGate unit to be configured.
- 2 Log into the CLI.
- 3 Enter the following command to change the FortiGate unit host name.

```
config system global
    set hostname Example1_host
end
```

Changing the host name makes it easier to identify individual cluster units when the cluster is operating.

- 4 Enter the following command to enable HA:

```
config system ha
  set mode active-passive
  set group-name Example_cluster
  set password HA_pass
end
```

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

- 5 Power off the FortiGate unit.
- 6 Repeat this procedure for all of the FortiGate units in the cluster.

Once all of the units are configured, continue with [“Connecting a FortiGate HA cluster”](#).

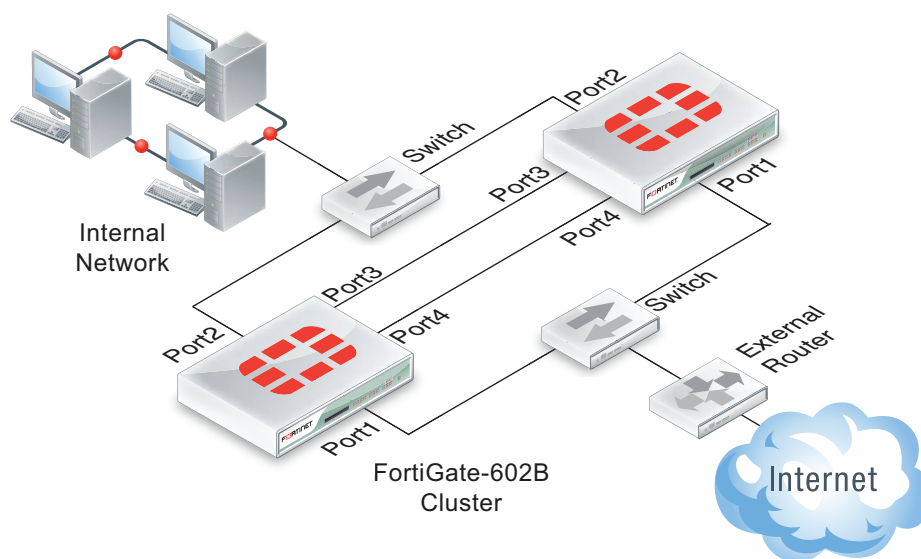
Connecting a FortiGate HA cluster

Use the following procedure to connect a cluster. Connect the cluster units to each other and to your network. You must connect all matching interfaces in the cluster to the same switch, then connect these interfaces to their networks using the same switch.

Although you can use hubs, Fortinet recommends using switches for all cluster connections for the best performance.

Connecting an HA cluster to your network temporarily interrupts communications on the network because new physical connections are being made to route traffic through the cluster. Also, starting the cluster interrupts network traffic until the individual cluster units are functioning and the cluster completes negotiation. Cluster negotiation is automatic and normally takes just a few seconds. During system startup and negotiation all network traffic is dropped.

This section describes how to connect the cluster shown in [Figure 194 on page 1997](#) that consists of two FortiGate-620B units to be connected between the Internet and a head office internal network. The port1 interfaces of the FortiGate unit connect the cluster to the Internet and the port2 interfaces connect the cluster to the internal network. The port3 and port4 interfaces are used for redundant HA heartbeat links.

Figure 194: Example cluster connections**To connect a FortiGate HA cluster**

- 1 Connect the port1 interfaces of each cluster unit to a switch connected to the Internet.
- 2 Connect the port2 interfaces of each cluster unit to a switch connected to the internal network.
- 3 Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on both of the FortiGate units.

As the cluster units start, they negotiate to choose the primary unit and the subordinate units. This negotiation occurs with no user intervention and normally just takes a few seconds.

At least one heartbeat interface should be connected together for the cluster to operate. You can also connect the heartbeat interfaces to a network. If the cluster consists of just two FortiGate units, you can connect the heartbeat interfaces directly using a crossover cable. For more information about heartbeat interfaces, see [“HA heartbeat and communication between cluster units”](#) on page 2170.

You could use one switch to connect all four heartbeat interfaces. However, this is not recommended because if the switch fails both heartbeat interfaces will become disconnected.

You can now configure the cluster as if it is a single FortiGate unit.

Active-passive and active-active HA

The first decision to make when configuring FortiGate HA is whether to choose active-passive or active-active HA mode. To configure the HA mode, go to *System > Config > HA* and set Mode to *Active-Passive* or *Active-Active*.

From the CLI enter the following command to set the HA mode to active-passive:

```
config system ha
  set mode a-p
end
```

To form a cluster, all cluster units must be set to the same mode. You can also change the mode after the cluster is up and running. Changing the mode of a functioning cluster causes a slight delay while the cluster renegotiates to operate in the new mode and possibly select a new primary unit.

Active-passive HA (failover protection)

An active-passive (A-P) HA cluster provides hot standby failover protection.

An active-passive cluster consists of a primary unit that processes communication sessions, and one or more subordinate units. The subordinate units are connected to the network and to the primary unit but do not process communication sessions. Instead, the subordinate units run in a standby state. In this standby state, the configuration of the subordinate units is synchronized with the configuration of the primary unit and the subordinate units monitor the status of the primary unit.

Active-passive HA provides transparent device failover among cluster units. If a cluster unit fails, another immediately take its place. See [“Device failover” on page 2169](#).

Active-passive HA also provides transparent link failover among cluster units. If a cluster unit interface fails or is disconnected, this cluster unit updates the link state database and the cluster negotiates and may select a new primary unit. See [“Link failover” on page 2194](#) for more information.

If session failover (also called session pickup) is enabled, active-passive HA provides session failover for some communication sessions. See [“Session failover \(session pickup\)” on page 2205](#) for information about session failover and its limitations.

The following example shows how to configure a FortiGate unit for active-passive HA operation. You would enter the exact same commands on every FortiGate unit in the cluster.

```
config system ha
  set mode a-p
  set group-name myname
  set password HApas
end
```

Active-active HA (load balancing and failover protection)

Active-active (A-A) HA load balances resource-intensive UTM processing among all cluster units. UTM processing applies protocol recognition, virus scanning, IPS, web filtering, email filtering, data leak prevention (DLP), application control, and VoIP content scanning and protection to HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, NNTP, SIP, SIMPLE, and SCCP sessions accepted by security policies. By load balancing this resource-intensive UTM processing among all cluster units, an active-active HA cluster may provide better UTM performance than a standalone FortiGate unit. Other features enabled in security policies such as Endpoint security, traffic shaping and authentication (identity-based policies) have no effect active-active load balancing.

All non-UTM sessions are not load balanced and are processed by the primary unit. You can also optionally configure active-active HA to load balance all TCP sessions in addition to UTM sessions. For more information see [“Load balancing UTM sessions and TCP sessions” on page 2221](#).

An active-active HA cluster consists of a primary unit that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. In an active-active cluster the subordinate units are also considered active since they also process UTM sessions.

In all other ways active-active HA operates the same as active-passive HA.

The following example shows how to configure a FortiGate unit for active-active HA operation. You would enter the exact same commands on every FortiGate unit in the cluster.

```
config system ha
  set mode a-a
  set group-name myname
  set password HApass
end
```

Identifying the cluster and cluster units

You can use the cluster group name, group id, and password to identify a cluster and distinguish one cluster from another. If you have more than one cluster on the same network, each cluster must have a different group name, group id, and password.

Group name

Use the group name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.

The default group name is *FGT-HA*. The group name appears on the FortiGate dashboard of a functioning cluster as the *Cluster Name*.

To change the group name from the web-based manager go to *Config > System > HA* and change the *Group Name*.

Enter the following CLI command to change the group name to *Cluster_name*:

```
config system ha
  set group-name Cluster_name
end
```

Password

Use the password to identify the cluster. You should always change the password when configuring a cluster. The password must be the same for all FortiGate units before they can form a cluster. The maximum password length is 19 characters. When the cluster is operating you can change the password, if required. Two clusters on the same network cannot have the same password.

To change the password from the web-based manager go to *Config > System > HA* and change the *Password*.

Enter the following CLI command to change the group name to *ha_pwd*:

```
config system ha
  set password ha_pwd
end
```

Group ID

Similar to the group name, the group ID is also used to identify the cluster. In most cases you do not have to change the group ID. However, you should change the group ID if you have more than one cluster on the same network. All members of the HA cluster must have the same group ID. The group ID range is from 0 to 63.

Changing the group ID changes the cluster virtual MAC address. See [“Cluster virtual MAC addresses” on page 2177](#).

Enter the following CLI command to change the group ID to 10:

```
config system ha
    set group-id 10
end
```

Device failover, link failover, and session failover

The FGCP provides transparent device and link failover. You can also enable session pickup to provide session failover. A failover can be caused by a hardware failure, a software failure, or something as simple as a network cable being disconnected. When a failover occurs, the cluster detects and recognizes the failure and takes steps to respond so that the network can continue to operate without interruption. The internal operation of the cluster changes, but network components outside of the cluster notice little or no change.


If a failover occurs, the cluster also records log messages about the event and can be configured to send log messages to a syslog server and to a FortiAnalyzer unit. The cluster can also send SNMP traps and alert email messages. These alerts can notify network administrators of the failover and may contain information that the network administrators can use to find and fix the problem that caused the failure.

For a complete description of device failover, link failover, and session failover, how clusters support these types of failover, and how FortiGate HA clusters compensate for a failure to maintain network traffic flow see [“HA and failover protection” on page 2167](#).

Primary unit selection

Once FortiGate units recognize that they can form a cluster, the cluster units negotiate to select a primary unit. Primary unit selection occurs automatically based on the criteria shown in [Figure 195](#). After the cluster selects the primary unit, all of the remaining cluster units become subordinate units.

Negotiation and primary unit selection also takes place if a primary unit fails (device failover) or if a monitored interface fails or is disconnected (link failover). During a device or link failover, the cluster renegotiates to select a new primary unit also using the criteria shown in [Figure 195](#).

Figure 195: Selecting the primary unit

For many basic HA configurations primary unit selection simply selects the cluster unit with the highest serial number to become the primary unit. A basic HA configuration involves setting the HA mode to active-passive or active-active and configuring the cluster group name and password. Using this configuration, the cluster unit with the highest serial number becomes the primary unit because primary unit selection disregards connected monitored interfaces (because interface monitoring is not configured), the age of the cluster units would usually always be the same, and all units would have the same device priority.

Using the serial number is a convenient way to differentiate cluster units; so basing primary unit selection on the serial number is predictable and easy to understand and interpret. Also the cluster unit with the highest serial number would usually be the newest FortiGate unit with the most recent hardware version. In many cases you may not need active control over primary unit selection, so basic primary unit selection based on serial number is sufficient.

In some situations you may want control over which cluster unit becomes the primary unit. You can control primary unit selection by setting the device priority of one cluster unit to be higher than the device priority of all other cluster units. If you change one or more device priorities, during negotiation, the cluster unit with the highest device priority becomes the primary unit. As shown in [Figure 195](#) the FGCP selects the primary unit based on device priority before serial number. For more information about how to use device priorities, see [“Primary unit selection and device priority” on page 2006](#).

The only other way that you can influence primary unit selection is by configuring interface monitoring (also called port monitoring). Using interface monitoring you can make sure that cluster units with failed or disconnected monitored interfaces cannot become the primary unit. See [“Primary unit selection and monitored interfaces” on page 2002](#).

Finally, the age of a cluster unit is determined by a number of cluster operating factors. Normally the age of all cluster units is the same so normally age has no effect on primary unit selection. Age does affect primary unit selection after a monitored interface failure. For more information about age, see [“Primary unit selection and age” on page 2003](#).

This section describes:

- [Primary unit selection and monitored interfaces](#)
- [Primary unit selection and age](#)
- [Primary unit selection and device priority](#)
- [Primary unit selection and FortiGate unit serial number](#)
- [Points to remember about primary unit selection](#)

Primary unit selection and monitored interfaces

If you have configured interface monitoring the cluster unit with the highest number of monitored interfaces that are connected to networks becomes the primary unit. Put another way, the cluster unit with the highest number of failed or disconnected monitored interfaces cannot become the primary unit.

Normally, when a cluster starts up, all monitored interfaces of all cluster units are connected and functioning normally. So monitored interfaces do not usually affect primary unit selection when the cluster first starts.

A cluster always renegotiates when a monitored interface fails or is disconnected (called link failover). A cluster also always renegotiates when a failed or disconnected monitored interface is restored.

If a primary unit monitored interface fails or is disconnected, the cluster renegotiates and if this is the only failed or disconnected monitored interface the cluster selects a new primary unit.

If a subordinate unit monitored interface fails or is disconnected, the cluster also renegotiates but will not necessarily select a new primary unit. However, the subordinate unit with the failed or disconnected monitored interface cannot become the primary unit.

Multiple monitored interfaces can fail or become disconnected on more than one cluster unit. Each time a monitored interface is disconnected or fails, the cluster negotiates to select the cluster unit with the most connected and operating monitored interfaces to become the primary unit. In fact, the intent of the link failover feature is just this, to make sure that the primary unit is always the cluster unit with the most connected and operating monitored interfaces. For information about monitored interfaces and link failover see [“Link failover” on page 2194](#).

Primary unit selection and age

The cluster unit with the highest age value becomes the primary unit. The age of a cluster unit is the amount of time since a monitored interface failed or is disconnected. Age is also reset when a cluster unit starts (boots up). So, when all cluster units start up at the same time, they all have the same age. Age does not affect primary unit selection when all cluster units start up at the same time. Age also takes precedence over priority for primary unit selection.

If a link failure of a monitored interface occurs, the age value for the cluster unit that experiences the link failure is reset. So, the cluster unit that experienced the link failure also has a lower age value than the other cluster units. This reduced age does not effect primary unit selection because the number of link failures takes precedence over the age.

If the failed monitored interface is restored the cluster unit that had the failed monitored interface cannot become the primary unit because its age is still lower than the age of the other cluster units.

In most cases, the way that age is handled by the cluster reduces the number of times the cluster selects a new primary unit, which results in a more stable cluster since selecting a new primary unit has the potential to disrupt traffic.

Cluster age difference margin (grace period)

In any cluster, some of the cluster units may take longer to start up than others. This startup time difference can happen as a result of a number of issues and does not affect the normal operation of the cluster. To make sure that cluster units that start slower can still become primary units, by default the FGCP ignores age differences of up to 5 minutes (300 seconds).

In most cases, during normal operation this age difference margin or grace period helps cluster function as expected. However, the age difference margin can result in some unexpected behavior in some cases:

- During a cluster firmware upgrade with `uninterruptable-upgrade` enabled (the default configuration) the cluster should not select a new primary unit after the firmware of all cluster units has been updated. But since the age difference of the cluster units is most likely less than 300 seconds, age is not used to affect primary unit selection and the cluster may select a new primary unit. See [“Upgrading cluster firmware” on page 2152](#) for more information.
- During failover testing where cluster units are failed over repeatedly the age difference between the cluster units will most likely be less than 5 minutes. During normal operation, if a failover occurs, when the failed unit rejoins the cluster its age will be very different from the age of the still operating cluster units so the cluster will not select a new primary unit. However, if a unit fails and is restored in a very short time the age difference may be less than 5 minutes. As a result the cluster may select a new primary unit during some failover testing scenarios.

Changing the cluster age difference margin

You can change the cluster age difference margin using the following command:

```
config system ha
  set ha-uptime-diff-margin 60
end
```

This command sets the cluster age difference margin to 60 seconds (1 minute). The age difference margin range 1 to 65535 seconds. The default is 300 seconds.

You may want to reduce the margin if during failover testing you don't want to wait the default age difference margin of 5 minutes. You may also want to reduce the margin to allow uninterruptable upgrades to work. See [“Upgrading cluster firmware” on page 2152](#).

You may want to increase the age margin if cluster unit startup time differences are larger than 5 minutes.

Displaying cluster unit age differences

You can use the CLI command `diagnose sys ha dump 1` to display the age difference of the units in a cluster. This command also displays information about a number of HA-related parameters for each cluster unit. You can enter the command from the primary unit CLI or you can enter the command from a subordinate unit after using `execute ha manage` to log into a subordinate unit CLI. The information displayed by the command is relative to the unit that you enter the command from.

For example, for a cluster of two FortiGate-5001SX units with no changes to the default HA configuration except to enable interface monitoring for port5, entering the `diagnose sys ha dump 1` command from the primary unit CLI displays information similar to the following:

```
diagnose sys ha dump 1
      HA information.
vcluster id=1, nentry=2, state=work,
  digest=fe.21.14.b3.e1.8d...
ventry idx=0,id=1,FG50012205400050,prio=128,0,claimed=0,
  override=0,flag=1,time=0,mon=0.
  mondev=port5,50
ventry idx=1,id=1,FG50012204400045,prio=128,0,claimed=0,
  override=0,flag=0,time=194,mon=0.
```

The command displays one `ventry` line for each cluster unit. The first `ventry` in the example contains information for the cluster unit that you are logged into. The other `ventry` lines contain information for the subordinate units (in the example there is only one subordinate unit). The `mondev` entry displays the interface monitoring configuration.

The `time` field is always 0 for the unit that you are logged into. The `time` field for the other cluster unit is the age difference between the unit that you are logged into and the other cluster unit. The age difference is in the form seconds/10.

In the example, the age of the primary unit is 19.4 seconds more than the age of the subordinate unit. The age difference is less than 5 minutes (less than 300 seconds or `time` is less than 3000) so age has no affect on primary unit selection. The cluster selected the unit with the highest serial number to be the primary unit.

If you use `execute ha manage 1` to log into the subordinate unit CLI and enter `diagnose sys ha dump 1` you get results similar to the following:

```
diagnose sys ha dump 1
      HA information.
vcluster id=1, nentry=2, state=standby,
  digest=fe.21.14.b3.e1.8d...
ventry idx=1,id=1,FG50012204400045,prio=128,0,claimed=0,
  override=0,flag=1,time=0,mon=0.
  mondev=port5,50
ventry idx=0,id=1,FG50012205400050,prio=128,0,claimed=0,
  override=0,flag=0,time=-194,mon=0.
```

The `time` for the primary unit is -194, indicating that age of the subordinate unit is 19.4 seconds less than the age of the primary unit.

If port5 (the monitored interface) of the primary unit is disconnected, the cluster renegotiates and the former subordinate unit becomes the primary unit. When you log into the new primary unit CLI and enter `diagnose sys ha dump 1` you could get results similar to the following:

```
diagnose sys ha dump 1
      HA information.
vcluster id=1, nentry=2, state=work,
  digest=9e.70.74.a2.5e.4a...
ventry idx=0,id=1,FG50012204400045,prio=128,0,claimed=0,
  override=0,flag=1,time=0,mon=0.
  mondev=port5,50
ventry idx=1,id=1,FG50012205400050,prio=128,-50,claimed=0,
  override=0,flag=0,time=58710,mon=0.
```

The command results show that the age of the new primary unit is 5871.0 seconds more than the age of the new subordinate unit.

If port5 of the former primary unit is reconnected the cluster will not select a new primary unit because the age of the primary unit will still be 5871.0 seconds more than the age of the subordinate unit. When you log into the primary unit CLI and enter `diagnose sys ha dump 1` you get results similar to the following:

```
diagnose sys ha dump 1
      HA information.
vcluster id=1, nentry=2, state=work,
  digest=9e.70.74.a2.5e.4a...
ventry idx=0,id=1,FG50012204400045,prio=128,0,claimed=0,
  override=0,flag=1,time=0,mon=0.
  mondev=port5,50
ventry idx=1,id=1,FG50012205400050,prio=128,0,claimed=0,
  override=0,flag=0,time=58710,mon=0.
```

Resetting the age of all cluster units

In some cases, age differences among cluster units can result in the wrong cluster unit or the wrong virtual cluster becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units when it rejoins the cluster. Since age takes precedence over priority the priority of this cluster unit will not be a factor in primary unit selection.

This problem also affects virtual cluster VDOM partitioning in a similar way. After a reboot of one of the units in a virtual cluster configuration, traffic for all VDOMs could continue to be processed by the cluster unit that did not reboot. This can happen because the age of both virtual clusters on the unit that did not reboot is greater than the age of both virtual clusters on the unit that rebooted.

One way to resolve this issue is to reboot all of the cluster units at the same time so that the age of all of the cluster units is reset. However, rebooting cluster units may interrupt or at least slow down traffic. If you would rather not reboot all of the cluster units you can instead use the following command to reset the ages of all of the cluster units.

```
diagnose sys ha reset-uptime
```

This command resets the age of all cluster units so age is no longer a factor in primary unit selection and device priority is used to select the primary unit.



The `diagnose sys ha reset-uptime` command should only be used as a temporary solution. The command resets the HA age internally and does not affect the up time displayed for cluster units using the `diagnose sys ha dump 1` command or the up time displayed on the Dashboard or cluster members list. To make sure the actual up time for cluster units is the same as the HA age you should reboot the cluster units during a maintenance window.

Primary unit selection and device priority

A cluster unit with the highest device priority becomes the primary unit when the cluster starts up or renegotiates. By default, the device priority for all cluster units is 128. You can change the device priority to control which FortiGate unit becomes the primary unit during cluster negotiation. All other factors that influence primary unit selection either cannot be configured (age and serial number) or are synchronized among all cluster units (interface monitoring). You can set a different device priority for each cluster unit. During negotiation, if all monitored interfaces are connected, and all cluster units enter the cluster at the same time (or have the same age), the cluster with the highest device priority becomes the primary unit.

A higher device priority does not affect primary unit selection for a cluster unit with the most failed monitored interfaces or with an age that is higher than all other cluster units because failed monitored interfaces and age are used to select a primary unit before device priority.

Increasing the device priority of a cluster unit does not always guarantee that this cluster unit will become the primary unit. During cluster operation, an event that may affect primary unit selection may not always result in the cluster renegotiating. For example, when a unit joins a functioning cluster, the cluster will not renegotiate. So if a unit with a higher device priority joins a cluster the new unit becomes a subordinate unit until the cluster renegotiates.



Enabling the `override` HA CLI keyword makes changes in device priority more effective by causing the cluster to negotiate more often to make sure that the primary unit is always the unit with the highest device priority. For more information about `override`, see “HA override” on page 2008.

Controlling primary unit selection by changing the device priority

You set a different device priority for each cluster unit to control the order in which cluster units become the primary unit when the primary unit fails.

To change the device priority from the web-based manager go to *Config > System > HA* and change the *Device Priority*.

Enter the following CLI command to change the device priority to 200:

```
config system ha
  set priority 200
end
```

The device priority is not synchronized among cluster units. In a functioning cluster you change device priority to change the priority of any unit in the cluster. Whenever you change the device priority of a cluster unit, when the cluster negotiates, the unit with the highest device priority becomes the primary unit.

The following example shows how to change the device priority of a subordinate unit to 255 so that this subordinate unit becomes the primary unit. This example involves connecting to the cluster CLI and using the `execute ha manage 0` command to connect to the highest priority subordinate unit. After you enter the following commands the cluster renegotiates and selects a new primary unit.

```
execute ha manage 1
config system ha
    set priority 255
end
```

If you have three units in a cluster you can set the device priorities as shown in [Table 113](#). When the cluster starts up, cluster unit A becomes the primary unit because it has the highest device priority. If unit A fails, unit B becomes the primary unit because unit B has a higher device priority than unit C.

Table 113: Example device priorities for a cluster of three FortiGate units

Cluster unit	Device priority
A	200
B	100
C	50

Normally, when configuring HA you do not have to change the device priority of any of the cluster units. If all cluster units have the same device priority, when the cluster first starts up the FGCP negotiates to select the cluster unit with the highest serial number to be the primary unit.

Clusters also function normally if all units have the same device priority. However, you can use the device priority if you want to control the roles that individual units play in the cluster. For example, if you want the same unit to always become the primary unit, set this unit device priority higher than the device priority of other cluster units. Also, if you want a cluster unit to always become a subordinate unit, set this cluster unit device priority lower than the device priority of other cluster units.

The device priority range is 0 to 255. The default device priority is 128.

If you are configuring a virtual cluster, if you have added virtual domains to both virtual clusters, you can set the device priority that the cluster unit has in virtual cluster 1 and virtual cluster 2. If a FortiGate unit has different device priorities in virtual cluster 1 and virtual cluster 2, the FortiGate unit may be the primary unit in one virtual cluster and the subordinate unit in the other. For more information, see [“Virtual clustering and load balancing or VDOM partitioning” on page 2091](#).

Primary unit selection and FortiGate unit serial number

The cluster unit with the highest serial number is more likely to become the primary unit. When first configuring FortiGate units to be added to a cluster, if you do not change the device priority of any cluster unit, then the cluster unit with the highest serial number always becomes the primary unit.

Age does take precedence over serial number, so if a cluster unit takes longer to join a cluster for some reason (for example if one cluster unit is powered on after the others), that cluster unit will not become the primary unit because the other units have been in the cluster longer.

Device priority and failed monitored interfaces also take precedence over serial number. A higher device priority means a higher priority. So if you set the device priority of one unit higher or if a monitored interface fails, the cluster will not use the FortiGate serial number to select the primary unit.

Points to remember about primary unit selection

Some points to remember about primary unit selection:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored interfaces > Age > Device Priority > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered if a cluster unit fails or if a monitored interface fails.
- If the HA age difference is more than 5 minutes, the cluster unit that is operating longer becomes the primary unit.
- If HA age difference is less than 5 minutes, the device priority and FortiGate serial number selects the cluster unit to become the primary unit.
- Every time a monitored interface fails the HA age of the cluster unit is reset to 0.
- Every time a cluster unit restarts the HA age of the cluster unit is reset to 0.

HA override

The HA `override` CLI keyword is disabled by default. When `override` is disabled a cluster may not renegotiate when an event occurs that affects primary unit selection. For example, when `override` is disabled a cluster will not renegotiate when you change a cluster unit device priority or when you add a new cluster unit to a cluster. This is true even if the unit added to the cluster has a higher device priority than any other unit in the cluster. Also, when `override` is disabled a cluster does not negotiate if the new unit added to the cluster has a failed or disconnected monitored interface.



For a virtual cluster configuration, `override` is enabled by default for both virtual clusters when you enable virtual cluster 2. For more information, see [“Virtual clustering and HA override” on page 2090](#).

In most cases you should keep `override` disabled to reduce how often the cluster negotiates. Frequent negotiations may cause frequent traffic interruptions.

However, if you want to make sure that the same cluster unit always operates as the primary unit and if you are less concerned about frequent cluster negotiation you can enable `override`.

To enable `override`, select a cluster unit to always be the primary unit. Connect to this cluster unit CLI and use the `config system ha` CLI command to enable `override`.

For `override` to be effective, you must also set the device priority highest on the cluster unit with `override` enabled. To increase the device priority, from the CLI use the `config system ha` command and increase the value of the `priority` keyword to a number higher than the default priority of 128.

You can also increase the device priority from the web-based manager by going to *System > Config > HA*. To increase the device priority of the primary unit select edit for the primary or subordinate unit and set the *Device Priority* to a number higher than 128.



The `override` setting and device priority value are not synchronized to all cluster units.

With `override` enabled, the primary unit with the highest device priority will always become the primary unit. Whenever an event occurs that may affect primary unit selection, the cluster negotiates. For example, when `override` is enabled a cluster renegotiates when you change the device priority of any cluster unit or when you add a new cluster unit to a cluster.

This section also describes:

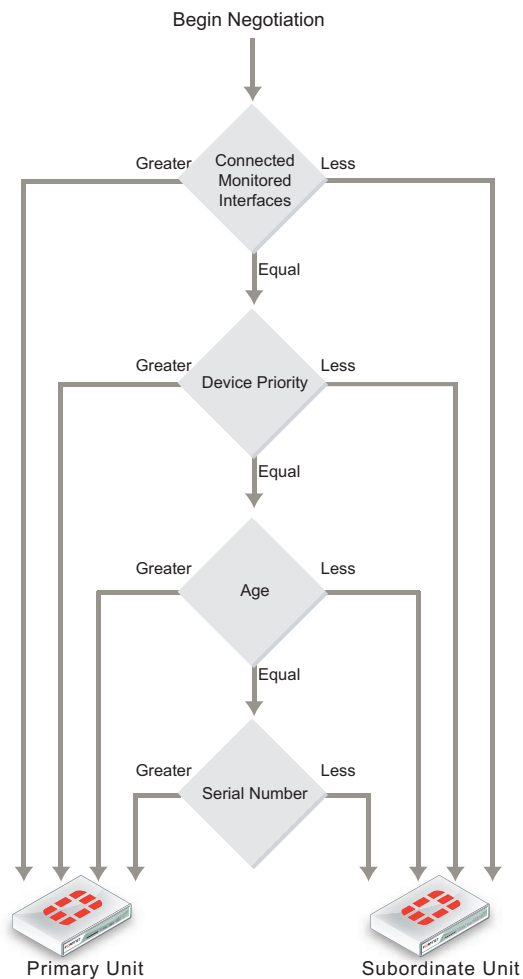
- [Override and primary unit selection](#)
- [Controlling primary unit selection using device priority and override](#)
- [Points to remember about primary unit selection when override is enabled](#)
- [Configuration changes can be lost if override is enabled](#)
- [Override and disconnecting a unit from a cluster](#)

Override and primary unit selection

Enabling `override` changes the order of primary unit selection. As shown in [Figure 196](#) if `override` is enabled, primary unit selection considers device priority before age and serial number. This means that if you set the device priority higher on one cluster unit, with `override` enabled this cluster unit becomes the primary unit even if its age and serial number are lower than other cluster units.

Similar to when `override` is disabled, when `override` is enabled primary unit selection checks for connected monitored interfaces first. So if interface monitoring is enabled, the cluster unit with the most disconnected monitored interfaces cannot become the primary unit, even if the unit has the highest device priority.

If all monitored interfaces are connected (or interface monitoring is not enabled) and the device priority of all cluster units is the same then age and serial number affect primary unit selection.

Figure 196: Selecting the primary unit with override enabled

Controlling primary unit selection using device priority and override

To configure one cluster unit to always become the primary unit you should set its device priority to be higher than the device priorities of the other cluster units and you should enable `override` for this cluster unit.

Using this configuration, when the cluster is operating normally the primary unit is always the unit with `override` enabled and with the highest device priority. If the primary unit fails the cluster renegotiates to select another cluster unit to be the primary unit. If the failed primary unit recovers, starts up again and rejoins the cluster, because `override` is enabled, the cluster renegotiates. Because the restarted primary unit has the highest device priority it once again becomes the primary unit.

In the same situation with `override` disabled, because the age of the failed primary unit is lower than the age of the other cluster units, when the failed primary unit rejoins the cluster it does not become the primary unit. Instead, even though the failed primary unit may have the highest device priority it becomes a subordinate unit because its age is lower than the age of all the other cluster units.

Points to remember about primary unit selection when override is enabled

Some points to remember about primary unit selection when `override` is enabled:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored Interfaces > Device Priority > Age > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered whenever an event occurs which may affect primary unit selection. For example negotiation occurs, when you change the device priority, when you add a new unit to a cluster, if a cluster unit fails, or if a monitored interface fails.
- Device priority is considered before age. Otherwise age is handled the same when `override` is enabled.

Configuration changes can be lost if override is enabled

In some cases, when `override` is enabled and you make configuration changes to an HA cluster these changes can be lost. For example, consider the following sequence:

- 1 A cluster of two FortiGate units is operating with `override` enabled.
 - FGT-A: Primary unit with device priority 200 and with `override` enabled
 - FGT-B: Subordinate unit with device priority 100 and with `override` disabled
 - If both units are operating, FGT-A always becomes the primary unit because FGT-A has the highest device priority.
- 2 FGT-A fails and FGT-B becomes the new primary unit.
- 3 The administrator makes configuration changes to the cluster.

The configuration changes are made to FGT-B because FGT-B is operating as the primary unit. These configuration changes are not synchronized to FGT-A because FGT-A is not operating.
- 4 FGT-A is restored and starts up again.
- 5 The cluster renegotiates and FGT-A becomes the new primary unit.
- 6 The cluster recognizes that the configurations of FGT-A and FGT-B are not the same.
- 7 The configuration of FGT-A is synchronized to FGT-B.

The configuration is always synchronized from the primary unit to the subordinate units.
- 8 The cluster is now operating with the same configuration as FGT-A. The configuration changes made to FGT-B have been lost.

The solution

When `override` is enabled, you can prevent configuration changes from being lost by doing the following:

- Verify that all cluster units are operating before making configuration changes (from the web-based manager go to *System > Config > HA* to view the cluster members list or from the FortiOS CLI enter `get system ha status`).
- Make sure the device priority of the primary unit is set higher than the device priorities of all other cluster units before making configuration changes.
- Disable `override` either permanently or until all configuration changes have been made and synchronized to all cluster units.

Override and disconnecting a unit from a cluster

A similar scenario to that described in “[Configuration changes can be lost if override is enabled](#)” may occur when `override` is enabled and you use the Disconnect from Cluster option from the web-based manager or the `execute ha disconnect` command from the CLI to disconnect a cluster unit from a cluster.

Configuration changes made to the cluster can be lost when you reconnect the disconnected unit to the cluster. You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. Otherwise, when the disconnected unit joins the cluster, if `override` is enabled, the cluster renegotiates and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units and any configuration changes made between when the unit was disconnected and reconnected are lost.

FortiGate HA compatibility with PPPoE and DHCP

FortiGate HA is not compatible with PPP protocols such as PPPoE. FortiGate HA is also not compatible with DHCP. If one or more FortiGate unit interfaces is dynamically configured using DHCP or PPPoE you cannot switch to operate in HA mode. Also, you cannot switch to operate in HA mode if one or more FortiGate unit interfaces is configured as a PPTP or L2TP client.



Configuring an interface for DHCP or PPPoE is only supported in NAT/Route mode. So, usually when configuring HA in Transparent mode an interface being configured for DHCP or PPPoE should not affect HA operation. However, in some cases you may not be able to enable HA if you had configured an interface for DHCP or PPPoE before switching to Transparent mode. So, if you are blocked from operating a Transparent mode FortiGate unit in HA and cannot find another reason for the problem, try switching the FortiGate unit back to NAT/Route mode and setting all interface modes to static before switching to Transparent mode and enabling HA. You could also enable HA before switching to Transparent mode.

You can configure a cluster to act as a DHCP server or a DHCP relay agent. In both active-passive and active-active clusters DHCP relay sessions are always handled by the primary unit. It is possible that a DHCP relay session could be interrupted by a failover. If this occurs the DHCP relay session is not resumed after the failover and the DHCP client may have to repeat the DHCP request.

When a cluster is operating as a DHCP server the primary unit responds to all DHCP requests and maintains the DHCP server address lease database. The cluster also dynamically synchronizes the DHCP server address lease database to the subordinate units. If a failover occurs, the new primary unit will have an up-to-date DHCP server address lease database. Synchronizing the DHCP address lease database prevents the new primary unit from responding incorrectly to new DHCP requests after a failover.

Also, it is possible that when FortiGate units first negotiate to form a cluster that a unit that ends up as a subordinate unit in the cluster will have information in its DHCP address lease database that the cluster unit operating as the primary unit does not have. This can happen if a FortiGate unit responds to DHCP requests while operating as a standalone unit and then when the cluster is formed this unit becomes a subordinate unit. Because of this possibility, after a cluster is formed the DHCP address lease databases of all of the cluster units are merged into one database which is then synchronized to all cluster units.

Hard disk configuration and HA

If your cluster units include hard disks, all cluster units must have identical hard disk configurations. This means each cluster unit must have same number of hard disks (including AMC and FortiGate Storage Module (FSM) hard disks) and also means that matching hard disks in each cluster unit must be the same size, have the same hard disk format, and have the same number of partitions.

In most cases the default hard disk configuration of the cluster units will be compatible. However, a hard disk formatted by an older FortiGate firmware version may not be compatible with a hard disk formatted by a more recent firmware version. Problems may also arise if you have used the `execute scsi-dev` command to add or change hard disk protections.

If a cluster unit CLI displays hard disk compatibility messages, you may need to use the `execute scsi-dev delete` command to delete partitions. You can also use the `execute formatlogdisk` command to reformat hard disks. In some cases after deleting all partitions and reformatting the hard disks, you may still see hard disk incompatibility messages. If this happens, contact Fortinet Customer Support for assistance.

HA Best practices

Fortinet suggests the following practices related to high availability:

- Use Active-Active HA to distribute TCP and UTM sessions among multiple cluster units. An active-active cluster may have higher throughput than a standalone FortiGate unit or than an active-passive cluster.
- Use a different host name on each FortiGate unit when configuring an HA cluster. Fewer steps are required to add host names to each cluster unit before configuring HA and forming a cluster.
- Enabling `load-balance-all` can increase device and network load since more traffic is load-balanced. This may be appropriate for use in a deployment using the firewall capabilities of the FortiGate unit and IPS but no other content inspection. See [“Load balancing UTM sessions and TCP sessions” on page 2221](#).
- An advantage of using session pickup is that non-UTM sessions will be picked up by the new primary unit after a failover. The disadvantage is that the cluster generates more heartbeat traffic to support session pickup as a larger portion of the session table must be synchronized. Session pickup should be configured only when required and is not recommended for use with SOHO FortiGate models. Session pickup should only be used if the primary heartbeat link is dedicated (otherwise the additional HA heartbeat traffic could affect network performance). See [“Session failover \(session pick-up\)” on page 2205](#).
- If you need to enable session pickup, consider enabling session pickup delay to improve performance by reducing the number of sessions that are synchronized. If possible, also consider enabling session synchronization or multiple FortiGate Interfaces. See [“Improving session synchronization performance” on page 2205](#) for more information.
- To avoid unpredictable results, when you connect a switch to multiple redundant or aggregate interfaces in an active-passive cluster you should configure separate redundant or aggregate interfaces on the switch; one for each cluster unit. See [“HA MAC addresses and 802.3ad aggregation” on page 2058](#).

- Use SNMP, syslog, or email alerts to monitor a cluster for failover messages. Alert messages about cluster failovers may help find and diagnose network problems quickly and efficiently. See [“Operating a cluster” on page 2127](#).

Heartbeat interfaces

Fortinet suggests the following practices related to heartbeat interfaces:

- Isolate heartbeat interfaces from user networks. Heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. If the cluster consists of two FortiGate units, connect the heartbeat interfaces directly using a crossover cable. For clusters with more than two units, connect heartbeat interfaces to a separate switch that is not connected to any network.
- If heartbeat traffic cannot be isolated from user networks, enable heartbeat message encryption and authentication to protect cluster information. See [“Enabling or disabling HA heartbeat encryption and authentication” on page 2177](#).
- Configure and connect multiple heartbeat interfaces so that if one heartbeat interface fails or becomes disconnected, HA heartbeat traffic can continue to be transmitted using the backup heartbeat interface. If heartbeat communication fails, all cluster members will think they are the primary unit resulting in multiple devices on the network with the same IP addresses and MAC addresses (condition referred to as *Split Brain*) and communication will be disrupted until heartbeat communication can be reestablished.
- Do not monitor dedicated heartbeat interfaces; monitor those interfaces whose failure should trigger a device failover.

Interface monitoring (port monitoring)

Fortinet suggests the following practices related to interface monitoring (also called port monitoring):

- Wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.
- Monitor interfaces connected to networks that process high priority traffic so that the cluster maintains connections to these networks if a failure occurs.
- Avoid configuring interface monitoring for all interfaces.
- Supplement interface monitoring with remote link failover. Configure remote link failover to maintain packet flow if a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails. See [“Remote link failover” on page 2200](#).

Troubleshooting

The following sections in this document contain troubleshooting information:

- [“Troubleshooting HA clusters” on page 2083](#)
- [“Troubleshooting virtual clustering” on page 2109](#)
- [“Troubleshooting full mesh HA” on page 2125](#)
- [“Troubleshooting layer-2 switches” on page 2233](#)

FGCP HA terminology

The following HA-specific terms are used in this document.

Cluster

A group of FortiGate units that act as a single virtual FortiGate unit to maintain connectivity even if one of the FortiGate units in the cluster fails.

Cluster unit

A FortiGate unit operating in a FortiGate HA cluster.

Device failover

Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device. See also [“Device failover, link failover, and session failover” on page 2000](#).

Failover

A FortiGate unit taking over processing network traffic in place of another unit in the cluster that suffered a device failure or a link failure.

Failure

A hardware or software problem that causes a FortiGate unit or a monitored interface to stop processing network traffic.

FGCP

The FortiGate clustering protocol (FGCP) that specifies how the FortiGate units in a cluster communicate to keep the cluster operating.

Full mesh HA

Full mesh HA is a method of removing single points of failure on a network that includes an HA cluster. FortiGate models that support redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA includes redundant connections between all network components. If any single component or any single connection fails, traffic switches to the redundant component or connection.

HA virtual MAC address

When operating in HA mode, all of the interfaces of the primary unit acquire the same HA virtual MAC address. All communications with the cluster must use this MAC address. The HA virtual MAC address is set according to the group ID.

Heartbeat

Also called FGCP heartbeat or HA heartbeat. The heartbeat constantly communicates HA status and synchronization information to make sure that the cluster is operating properly.

Heartbeat device

An ethernet network interface in a cluster that is used by the FGCP for heartbeat communications among cluster units.

Heartbeat failover

If an interface functioning as the heartbeat device fails, the heartbeat is transferred to another interface also configured as an HA heartbeat device.

Hello state

In the hello state a cluster unit has powered on in HA mode, is using HA heartbeat interfaces to send hello packets, and is listening on its heartbeat interfaces for hello packets from other FortiGate units. Hello state may appear in HA log messages.

High availability

The ability that a cluster has to maintain a connection when there is a device or link failure by having another unit in the cluster take over the connection, without any loss of connectivity. To achieve high availability, all FortiGate units in the cluster share session and configuration information.

Interface monitoring

You can configure interface monitoring (also called port monitoring) to monitor FortiGate interfaces to verify that the monitored interfaces are functioning properly and connected to their networks. If a monitored interface fails or is disconnected from its network the interface leaves the cluster and a link failover occurs. For more information about interface monitoring, see [“Link failover” on page 2194](#).

Link failover

Link failover means that if a monitored interface fails, the cluster reorganizes to re-establish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic. See also [“Device failover, link failover, and session failover” on page 2000](#).

Load balancing

Also known as active-active HA. All units in the cluster process network traffic. The FGCP employs a technique called unicast load balancing in which a given interface of all cluster units has the same virtual MAC address. The primary unit is associated with the cluster HA virtual MAC address and cluster IP address. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit can process packets itself, or propagate them to subordinate units according to a load balancing schedule.

Monitored interface

An interface that is monitored by a cluster to make sure that it is connected and operating correctly. The cluster monitors the connectivity of this interface for all cluster units. If a monitored interface fails or becomes disconnected from its network, the cluster will compensate.

Primary unit

Also called the primary cluster unit, this cluster unit controls how the cluster operates. The primary unit sends hello packets to all cluster units to synchronize session information, synchronize the cluster configuration, and to synchronize the cluster routing table. The hello packets also confirm for the subordinate units that the primary unit is still functioning.

The primary unit also tracks the status of all subordinate units. When you start a management connection to a cluster, you connect to the primary unit.

In an active-passive cluster, the primary unit processes all network traffic. If a subordinate unit fails, the primary unit updates the cluster configuration database.

In an active-active cluster, the primary unit receives all network traffic and re-directs this traffic to subordinate units. If a subordinate unit fails, the primary unit updates the cluster status and redistributes load balanced traffic to other subordinate units in the cluster.

The FortiGate firmware uses the term master to refer to the primary unit.

Session failover

Session failover means that a cluster maintains active network sessions after a device or link failover. FortiGate HA does not support session failover by default. To enable session failover you must change the HA configuration to select Enable Session Pick-up. See also [“Device failover, link failover, and session failover” on page 2000](#).

Session pickup

If you enable session pickup for a cluster, if the primary unit fails or a subordinate unit in an active-active cluster fails, all communication sessions with the cluster are maintained or picked up by the cluster after the cluster negotiates to select a new primary unit.

If session pickup is not a requirement of your HA installation, you can disable this option to save processing resources and reduce the network bandwidth used by HA session synchronization. In many cases interrupted sessions will resume on their own after a failover even if session pickup is not enabled. You can also enable session pickup delay to reduce the number of sessions that are synchronized by session pickup.

Standby state

A subordinate unit in an active-passive HA cluster operates in the standby state. In a virtual cluster, a subordinate virtual domain also operates in the standby state. The standby state is actually a hot-standby state because the subordinate unit or subordinate virtual domain is not processing traffic but is monitoring the primary unit session table to take the place of the primary unit or primary virtual domain if a failure occurs.

In an active-active cluster all cluster units operate in a work state.

When standby state appears in HA log messages this usually means that a cluster unit has become a subordinate unit in an active-passive cluster or that a virtual domain has become a subordinate virtual domain.

State synchronization

The part of the FGCP that maintains connections after failover.

Subordinate unit

Also called the subordinate cluster unit, each cluster contains one or more cluster units that are not functioning as the primary unit. Subordinate units are always waiting to become the primary unit. If a subordinate unit does not receive hello packets from the primary unit, it attempts to become the primary unit.

In an active-active cluster, subordinate units keep track of cluster connections, keep their configurations and routing tables synchronized with the primary unit, and process network traffic assigned to them by the primary unit. In an active-passive cluster, subordinate units do not process network traffic. However, active-passive subordinate units do keep track of cluster connections and do keep their configurations and routing tables synchronized with the primary unit.

The FortiGate firmware uses the terms slave and subsidiary unit to refer to a subordinate unit.

Virtual clustering

Virtual clustering is an extension of the FGCP for FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

Work state

The primary unit in an active-passive HA cluster, a primary virtual domain in a virtual cluster, and all cluster units in an active-active cluster operate in the work state. A cluster unit operating in the work state processes traffic, monitors the status of the other cluster units, and tracks the session table of the cluster.

When work state appears in HA log messages this usually means that a cluster unit has become the primary unit or that a virtual domain has become a primary virtual domain.

HA web-based manager options

Go to *System > Config > HA* to change HA options.

You can configure HA options for a FortiGate unit with virtual domains (VDOMs) enabled by logging into the web-based manager as the global admin administrator and going to *System > Config > HA*.

If already operating in HA mode, go to *System > Config > HA* to display the cluster members list (see [“Cluster members list” on page 2147](#)).

Go to *System > Config > HA* and select *View HA Statistics* to view statistics about cluster operation. See [“Viewing HA statistics” on page 2150](#).



If your cluster uses virtual domains, you are configuring HA virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below and see [“Configuring and connecting virtual clusters” on page 2089](#).



HA is not compatible with PPP protocols such as PPPoE. HA is also not compatible with DHCP. If one or more FortiGate interfaces is dynamically configured using DHCP or PPPoE, you cannot switch to operate in HA mode. You also cannot switch to operate in HA mode if one or more FortiGate interfaces is configured as a PPTP or L2TP client or if the FortiGate unit is configured for standalone session synchronization.

Mode	<p>Select an HA mode for the cluster or return the FortiGate unit in the cluster to standalone mode. When configuring a cluster, you must set all members of the HA cluster to the same HA mode. You can select <i>Standalone</i> (to disable HA), <i>Active-Passive</i>, or <i>Active-Active</i>.</p> <p>If virtual domains are enabled you can select <i>Active-Passive</i> or <i>Standalone</i>.</p>
Device Priority	<p>Optionally set the device priority of the cluster FortiGate unit. Each FortiGate unit in a cluster can have a different device priority. During HA negotiation, the FortiGate unit with the highest device priority usually becomes the primary unit. See “Primary unit selection” on page 2000.</p> <p>In a virtual cluster configuration, each cluster FortiGate unit can have two different device priorities, one for each virtual cluster. During HA negotiation, the FortiGate unit with the highest device priority in a virtual cluster becomes the primary FortiGate unit for that virtual cluster.</p> <p>Changes to the device priority are not synchronized. You can accept the default device priority when first configuring a cluster. When the cluster is operating you can change the device priority for different cluster units as required.</p>
Group Name	<p>Enter a name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.</p> <p>When the cluster is operating you can change the group name, if required.</p>
Password	<p>Enter a password to identify the cluster. The maximum password length is 15 characters. The password must be the same for all cluster FortiGate units before the cluster FortiGate units can form a cluster.</p> <p>The default is no password. When the cluster is operating, you can add a password, if required. Two clusters on the same network must have different passwords.</p>
Enable Session pickup	<p>Select to enable session pickup so that if the primary unit fails, sessions are picked up by the cluster unit that becomes the new primary unit.</p> <p>You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce HA CPU usage and reduce HA heartbeat network bandwidth usage.</p> <p>Session pickup is disabled by default. You can accept the default setting for session pickup and later choose to enable session pickup after the cluster is operating. See “Session failover (session pick-up)” on page 2205.</p>

Port Monitor	<p>Select to enable or disable monitoring FortiGate interfaces to verify the monitored interfaces are functioning properly and are connected to their networks. See “Link failover” on page 2194.</p> <p>If a monitored interface fails or is disconnected from its network, the interface leaves the cluster and a link failover occurs. The link failover causes the cluster to reroute the traffic being processed by that interface to the same interface of another cluster FortiGate unit that still has a connection to the network. This other cluster FortiGate unit becomes the new primary unit.</p> <p>Port monitoring (also called interface monitoring) is disabled by default. Leave port monitoring disabled until the cluster is operating and then only enable port monitoring for connected interfaces.</p> <p>You can monitor up to 16 interfaces. This limit only applies to FortiGate units with more than 16 physical interfaces.</p>
Heartbeat Interface	<p>Select to enable or disable HA heartbeat communication for each interface in the cluster and set the heartbeat interface priority. The heartbeat interface with the highest priority processes all heartbeat traffic. If two or more heartbeat interfaces have the same priority, the heartbeat interface with the lowest hash map order value processes all heartbeat traffic. The web-based manager lists interfaces in alphanumeric order:</p> <ul style="list-style-type: none"> • port1 • port2 through 9 • port10 <p>Hash map order sorts interfaces in the following order:</p> <ul style="list-style-type: none"> • port1 • port10 • port2 through port9 <p>The default heartbeat interface configuration is different for each unit. This default configuration usually sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration or change it as required.</p> <p>The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0.</p> <p>You must select at least one heartbeat interface. If heartbeat communication is interrupted, the cluster stops processing traffic. See “HA heartbeat and communication between cluster units” on page 2170.</p> <p>You can select up to 8 heartbeat interfaces. This limit only applies to units with more than 8 physical interfaces.</p>
VDOM partitioning	<p>If you are configuring virtual clustering, you can set the virtual domains to be in virtual cluster 1 and the virtual domains to be in virtual cluster 2. The root virtual domain must always be in virtual cluster 1. See “Configuring and connecting virtual clusters” on page 2089.</p>



Configuring and connecting HA clusters

This chapter contains general procedures and descriptions as well as detailed configuration examples that describe how to configure FortiGate HA clusters.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

This chapter contains the following sections:

- [About the procedures in this chapter](#)
- [Example: NAT/Route mode active-passive HA configuration](#)
- [Example: Transparent mode active-active HA configuration](#)
- [Example: advanced Transparent mode active-active HA configuration](#)
- [Example: converting a standalone FortiGate unit to a cluster](#)
- [Example: adding a new unit to an operating cluster](#)
- [Example: replacing a failed cluster unit](#)
- [Example: HA and 802.3ad aggregated interfaces](#)
- [Example: HA and redundant interfaces](#)
- [Troubleshooting HA clusters](#)

About the procedures in this chapter

The procedures in this chapter describe some of many possible sequences of steps for configuring HA clustering. As you become more experienced with FortiOS HA you may choose to use a different sequence of configuration steps.

For simplicity, many of these procedures assume that you are starting with new FortiGate units set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

Example: NAT/Route mode active-passive HA configuration

This section describes a simple HA network topology that includes an HA cluster of two FortiGate-620B units in NAT/Route mode installed between an internal network and the Internet.


- [Example NAT/Route mode HA network topology](#)
- [General configuration steps](#)
- [Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - web-based manager](#)

- [Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI](#)

Example NAT/Route mode HA network topology

Figure 197 shows a typical FortiGate-620B HA cluster consisting of two FortiGate-620B units (620_ha_1 and 620_ha_2) connected to the same internal (port2) and external (port1) networks.

Figure 197: Example NAT/Route mode HA network topology



Port3 and port4 are the default FortiGate-620B heartbeat interfaces. Because the cluster consists of two FortiGate units, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

- 1 Configure the FortiGate units for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
- 2 Connect the cluster to the network.
- 3 Confirm that the cluster units are operating as a cluster and add basic configuration settings to the cluster.
 - View cluster status from the web-based manager or CLI.
 - Add a password for the admin administrative account.
 - Change the IP addresses and netmasks of the internal and external interfaces.
 - Add a default route.

Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - web-based manager

Use the following procedures to configure two FortiGate-620B units for NAT/Route HA operation using the FortiGate web-based manager. These procedures assume you are starting with two FortiGate-620B units with factory default settings.



Give each cluster unit a unique host name to make the individual units easier to identify when they are part of a functioning cluster. The default FortiGate unit host name is the FortiGate serial number. You may want to change this host name to something more meaningful for your network.

To configure the first FortiGate-620B unit (host name 620_ha_1)

- 1 Power on the first FortiGate unit.
- 2 On your management computer with an Ethernet connection, set the static IP address to 192.168.1.2 and the netmask to 255.255.255.0.
- 3 On a management computer, start a web browser and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).
The FortiGate login is displayed.
- 4 Type *admin* in the *Name* field and select *Login*.
The FortiGate dashboard is displayed.
- 5 On the *System Information* dashboard widget beside *Host Name*, select *Change*.
- 6 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
-----------------	----------

- 7 Select OK.
- 8 Go to *System > Config > HA* and change the following settings:

Mode	Active-Passive
Group Name	example1.com
Password	HA_pass_1



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each unit in the cluster.

9 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

10 Power off the first FortiGate unit (620_ha_1).

To configure the second FortiGate-620B unit (host name 620_ha_2)

- 1 Power on the second FortiGate unit.
- 2 On a management computer, start a web browser and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).
The FortiGate login is displayed.
- 3 Type *admin* in the *Name* field and select *Login*.
The FortiGate dashboard is displayed.
- 4 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 5 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_2
-----------------	----------

- 6 Select OK.
- 7 Go to *System > Config > HA* and change the following settings:

Mode	Active-Passive
Group Name	example1.com
Password	HA_pass_1

- 8 Select OK.
The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.
To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.
- 9 Power off the second FortiGate unit.

To connect the cluster to the network

- 1 Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2 Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 3 Connect the port3 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate unit. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

- 1 Start Internet Explorer and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).

The FortiGate Login is displayed.

- 2 Type `admin` in the *Name* field and select Login.

The FortiGate dashboard is displayed.

The System Information dashboard widget shows the *Cluster Name* (example1.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

Figure 198: Sample FortiGate-620B System Information dashboard widget

System Information	
Cluster Name	example1.com
Cluster Members	620_ha_2/FG600B3908600825 (Master) 620_ha_1/FG600B3908600705 (Slave)
Serial Number	FG600B3908600825
Operation Mode	NAT [Change]
HA Status	Active-Passive [Configure]
System Time	Wed Feb 9 14:35:11 2011 [Change]
Firmware Version	v4.0,build0415,110126 (Interim) [Update]
System Configuration	Last Backup: N/A [Backup] [Restore]
Current Administrator	admin [Change Password] /4 in Total [Details]
Uptime	13 day(s) 7 hour(s) 34 min(s)
Virtual Domain	Disabled [Enable]

- 3 Go to *System > Config > HA* to view the cluster members list.

The list shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

Figure 199: Sample FortiGate-620B cluster members list

HA Cluster		Cluster Member	Hostname	Role	Priority	View HA Statistics
			620_ha_2	MASTER	128	
			620_ha_1	SLAVE	128	

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster to connect to its network. The following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.
- 3 For *admin*, select the *Change Password* icon
- 4 Enter and confirm a new password.
- 5 Select OK.
- 6 Go to *System > Network > Interface*.
- 7 Edit the *port2* interface and change *IP/Netmask* to 10.11.101.100/24.
- 8 Select OK.



After changing the IP address of the port1 interface you may have to change the IP address of your management computer and then reconnect to the port1 interface using the 172.20.120.141 IP address.

- 9 Edit the *port1* interface and change *IP/Netmask* to 172.20.120.141/24.
- 10 Select OK.
- 11 Go to *Router > Static*.
- 12 Change the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	port1
Distance	10

- 13 Select OK.

Configuring a NAT/Route mode active-passive cluster of two FortiGate-620B units - CLI

Use the following procedures to configure two FortiGate-620B units for NAT/Route HA operation using the FortiGate CLI. These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the first FortiGate-620B unit (host name 620_ha_1)

- 1 Power on the FortiGate unit.
- 2 Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
- 3 Start HyperTerminal (or any terminal emulation program), enter a name for the connection, and select OK.

- 4 Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
- 5 Select the following port settings and select OK.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 6 Press Enter to connect to the FortiGate CLI.

The FortiGate unit CLI login prompt appears.

If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

- 7 Type `admin` and press Enter twice.
- 8 Change the host name for this FortiGate unit.

```
config system global
  set hostname 620_ha_1
end
```

- 9 Configure HA settings.

```
config system ha
  set mode a-p
  set group-name example1.com
  set password HA_pass_1
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c

- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

10 Display the HA configuration (optional).

```
get system ha
  group-id           : 0
  group-name         : example1.com
  mode               : a-p
  password           : *
  hbdev              : "port3" 50 "port4" 50
  session-sync-dev   :
  route-ttl          : 10
  route-wait         : 0
  route-hold         : 10
  sync-config        : enable
  encryption         : disable
  authentication     : disable
  hb-interval        : 2
  hb-lost-threshold  : 6
  helo-holddown      : 20
  arps               : 5
  arps-interval      : 8
  session-pickup     : disable
  link-failed-signal : disable
  uninterruptable-upgrade: enable
  ha-mgmt-status     : disable
  ha-eth-type        : 8890
  hc-eth-type        : 8891
```

```

l2ep-eth-type      : 8893
subsecond          : disable
vcluster2          : disable
vcluster-id        : 1
override           : disable
priority           : 128
monitor            :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom                : "root"

```

11 Power off the FortiGate unit.

To configure the second FortiGate-620B unit (host name 620_ha_2)

- 1 Power on the FortiGate unit.
- 2 Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
- 3 Start HyperTerminal, enter a name for the connection, and select OK.
- 4 Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
- 5 Select the following port settings and select OK.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 6 Press Enter to connect to the FortiGate CLI.
The FortiGate unit CLI login prompt appears.
- 7 Type `admin` and press Enter twice.
- 8 Change the host name for this FortiGate unit.

```

config system global
  set hostname 620_ha_2
end

```

- 9 Configure HA settings.

```

config system ha
  set mode a-p
  set group-name example1.com
  set password HA_pass_1
end

```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

10 Display the HA configuration (optional).

```

get system ha
  group-id           : 0
  group-name         : example1.com
  mode               : a-p
  password           : *
  hbdev              : "port3" 50 "port4" 50
  session-sync-dev   :
  route-ttl          : 10
  route-wait         : 0
  route-hold         : 10
  sync-config        : enable
  encryption         : disable
  authentication     : disable
  hb-interval        : 2
  hb-lost-threshold  : 6
  helo-holddown      : 20
  arps               : 5
  arps-interval      : 8
  session-pickup      : disable
  link-failed-signal : disable
  uninterruptable-upgrade: enable
  ha-mgmt-status     : disable
  ha-eth-type        : 8890
  hc-eth-type        : 8891
  l2ep-eth-type      : 8893
  subsecond          : disable
  vcluster2          : disable
  vcluster-id        : 1
  override           : disable
  priority           : 128
  monitor            :
  pingserver-monitor-interface:
  pingserver-failover-threshold: 0
  pingserver-flip-timeout: 60
  vdom               : "root"

```

11 Power off the FortiGate unit.**To connect the cluster to the network**

- 1** Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2** Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 3** Connect the port3 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4** Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5** Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1 Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.

If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.

If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2          FG600B3908600825 0
Slave :128 620_ha_1          FG600B3908600705 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to add some basic settings to the cluster so that it can connect to the network.

1 Log into the primary unit CLI.**2** Add a password for the admin administrative account.

```
config system admin
edit admin
set password <password_str>
end
```

3 Configure the port1 and port2 interfaces.

```

config system interface
  edit port1
    set ip 172.20.120.141/24
  next
  edit port2
    set ip 10.11.101.100/24
end

```

4 Add a default route.

```

config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device port1
  end

```

Example: Transparent mode active-active HA configuration

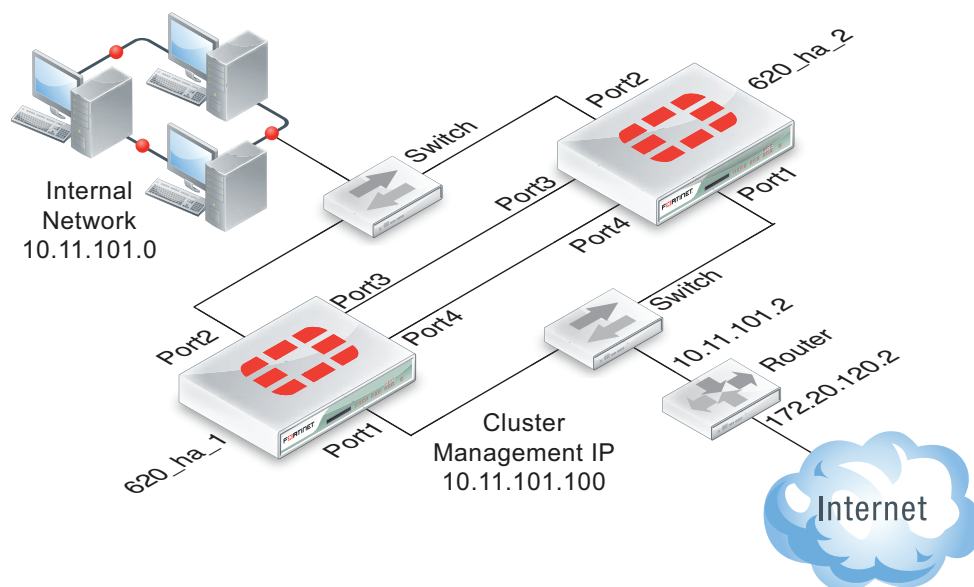
This section describes a simple HA network topology that includes an HA cluster of two FortiGate-620B units installed between an internal network and the Internet and running in Transparent mode.

- [Example Transparent mode HA network topology](#)
- [General configuration steps](#)

Example Transparent mode HA network topology

Figure 200 shows a Transparent mode FortiGate-620B HA cluster consisting of two FortiGate-620B units (620_ha_1 and 620_ha_2) installed between the Internet and internal network. The topology includes a router that performs NAT between the internal network and the Internet. The cluster management IP address is 10.11.101.100.

Figure 200: Transparent mode HA network topology



Port3 and port4 are the default FortiGate-620B heartbeat interfaces. Because the cluster consists of two FortiGate units, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

General configuration steps

This section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

In this example, the configuration steps are identical to the NAT/Route mode configuration steps until the cluster is operating. When the cluster is operating, you can switch to Transparent mode and add basic configuration settings to cluster.

General configuration steps

- 1 Configure the FortiGate units for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
- 2 Connect the cluster to the network.
- 3 Confirm that the cluster units are operating as a cluster.
- 4 Switch the cluster to Transparent mode and add basic configuration settings to the cluster.
 - Switch to Transparent mode, add the management IP address and a default route.
 - Add a password for the admin administrative account.
 - View cluster status from the web-based manager or CLI.

Configuring a Transparent mode active-active cluster of two FortiGate-620B units - web-based manager

Use the following procedures to configure the FortiGate-620B units for HA operation using the FortiGate web-based manager. These procedures assume you are starting with two FortiGate-620B units with factory default settings.



Waiting until you have established the cluster to switch to Transparent mode means fewer configuration steps because you can switch the mode of the cluster in one step.

To configure the first FortiGate-620B unit (host name 620_ha_1)

- 1 Power on the first FortiGate unit.
- 2 Set the IP address of a management computer with an Ethernet connection to the static IP address 192.168.1.2 and a netmask of 255.255.255.0.
- 3 On a management computer, start a web browser and browse to the address <https://192.168.1.99> (remember to include the "s" in https://).
The FortiGate login is displayed.
- 4 Type *admin* in the *Name* field and select *Login*.
The FortiGate dashboard is displayed.
- 5 On the *System Information* dashboard widget, beside *Host Name* select *Change*.

- 6 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
-----------------	----------

- 7 Select OK.

- 8 Go to *System > Config > HA* and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

9 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

10 Power off the first FortiGate unit.

To configure the second FortiGate-620B unit (host name 620_ha_2)

- 1 Power on second FortiGate unit.
- 2 On a management computer, start Internet Explorer and browse to the address <https://192.168.1.99> (remember to include the “s” in https://).
The FortiGate login is displayed.
- 3 Type *admin* in the *Name* field and select *Login*.
The FortiGate dashboard is displayed.
- 4 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 5 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_2
-----------------	----------

- 6 Select OK.
- 7 Go to *System > Config > HA* and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2

- 8 Select OK.
The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.
To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.
- 9 Power off the second FortiGate unit.

To connect the cluster to the network

- 1 Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2 Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 3 Connect the port3 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To switch the cluster to Transparent mode

Switching from NAT/Route to Transparent mode involves adding the Transparent mode management IP address and default route.



Since configuration changes are synchronized to all cluster units, switching the cluster to operate in Transparent mode once the cluster is operating is the same as switching an individual FortiGate unit to Transparent mode. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

- 1 Start a web browser and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).
- The FortiGate Login is displayed.
- 2 Type admin in the Name field and select Login.
- 3 Under System Information, beside *Operation Mode* select *Change*.
- 4 Set Operation Mode to Transparent.
- 5 Configure basic Transparent mode settings.

Operation Mode	Transparent
Management IP/Mask	10.11.101.100/24
Default Gateway	10.11.101.2

- 6 Select Apply.
- The cluster switches to operating in Transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate unit. You could have performed the following configuration steps separately on each FortiGate unit before you connected them to form a cluster.

- 1 Start Internet Explorer and browse to the address `https://10.11.101.100` (remember to include the “s” in `https://`).
- The FortiGate Login is displayed.
- 2 Type admin in the Name field and select Login.
- The FortiGate dashboard is displayed.
- The System Information dashboard widget shows the *Cluster Name* (example2.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.
- 3 Go to *System > Config > HA* to view the cluster members list.
- The list shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster. Note that the following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.
- 3 For *admin*, select the *Change Password* icon
- 4 Enter and confirm a new password.
- 5 Select OK.



You added a default gateway when you switched to Transparent mode so you don't need to add a default route as part of the basic configuration of the cluster at this point.

Configuring a Transparent mode active-active cluster of two FortiGate-620B units - CLI

Use the following procedures to configure the FortiGate-620B units for Transparent mode HA operation using the FortiGate CLI.

To configure each FortiGate unit for HA operation

- 1 Power on the FortiGate unit.
- 2 Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
- 3 Start HyperTerminal, enter a name for the connection, and select OK.
- 4 Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
- 5 Select the following port settings and select OK.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 6 Press Enter to connect to the FortiGate CLI.
The FortiGate unit CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.
- 7 Type `admin` and press Enter twice.
- 8 Change the host name for this FortiGate unit. For example:
`config system global`

```
    set hostname 620_ha_1
end
```

9 Configure HA settings.

```
config system ha
    set mode a-a
    set group-name example2.com
    set password HA_pass_2
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

10 Display the HA configuration (optional).

```
get system ha
group-id          : 0
```

```

group-name          : example2.com
mode                : a-a
password            : *
hbdev               : "port3" 50 "port4" 50
session-sync-dev    :
route-ttl           : 10
route-wait          : 0
route-hold          : 10
sync-config         : enable
encryption          : disable
authentication      : disable
hb-interval         : 2
hb-lost-threshold   : 6
helo-holddown       : 20
arps                : 5
arps-interval       : 8
session-pickup      : disable
link-failed-signal  : disable
uninterruptable-upgrade: enable
ha-mgmt-status      : disable
ha-eth-type         : 8890
hc-eth-type         : 8891
l2ep-eth-type       : 8893
subsecond           : disable
vcluster2           : disable
vcluster-id         : 1
override            : disable
priority            : 128
monitor             :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom                : "root"

```

11 Power off the FortiGate unit.

To configure the second FortiGate-620B unit (host name 620_ha_2)

- 1** Power on the FortiGate unit.
- 2** Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
- 3** Start HyperTerminal, enter a name for the connection, and select OK.
- 4** Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select OK.
- 5** Select the following port settings and select OK.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6 Press Enter to connect to the FortiGate CLI.

The FortiGate unit CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate unit and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

7 Type `admin` and press Enter twice.**8** Change the host name for this FortiGate unit.

```
config system global
  set hostname 620_ha_2
end
```

9 Configure HA settings.

```
config system ha
  set mode a-a
  set group-name example2.com
  set password HA_pass_2
end
```

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

10 Display the HA configuration (optional).

```
get system ha
  group-id           : 0
  group-name         : example2.com
  mode               : a-a
  password           : *
  hbdev              : "port3" 50 "port4" 50
  session-sync-dev   :
  route-ttl          : 10
  route-wait         : 0
  route-hold         : 10
  sync-config        : enable
  encryption         : disable
  authentication     : disable
  hb-interval        : 2
  hb-lost-threshold  : 6
  helo-holddown      : 20
  arps               : 5
  arps-interval      : 8
  session-pickup     : disable
  link-failed-signal : disable
  uninterruptable-upgrade: enable
  ha-mgmt-status     : disable
  ha-eth-type        : 8890
  hc-eth-type        : 8891
  l2ep-eth-type      : 8893
  subsecond          : disable
  vcluster2          : disable
```

```

vcluster-id      : 1
override        : disable
priority        : 128
monitor         :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom            : "root"

```

11 Power off the FortiGate unit.

To connect the cluster to the network

- 1 Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2 Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 3 Connect the port3 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To connect to the cluster CLI and switch the cluster to Transparent mode

- 1 Determine which cluster unit is the primary unit.
 - Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
 - Enter the command `get system status`.
 If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.
 If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode. See [“Troubleshooting the initial cluster configuration” on page 2084](#).

- 2 Change to transparent mode.

```

config system settings
  set opmode transparent
  set manageip 192.168.20.3/24
  set gateway 192.168.20.1

```

end

The cluster switches to Transparent Mode, and your administration session is disconnected.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (192.168.20.3).

To view cluster status

Use the following steps to view cluster status from the CLI.

1 Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.

If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.

If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2          FG600B3908600825 0
Slave :128 620_ha_1          FG600B3908600705 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add a password for the admin administrative account

1 Add a password for the admin administrative account.

```
config system admin
edit admin
```

```

set password <psswrd>
end

```

Example: advanced Transparent mode active-active HA configuration

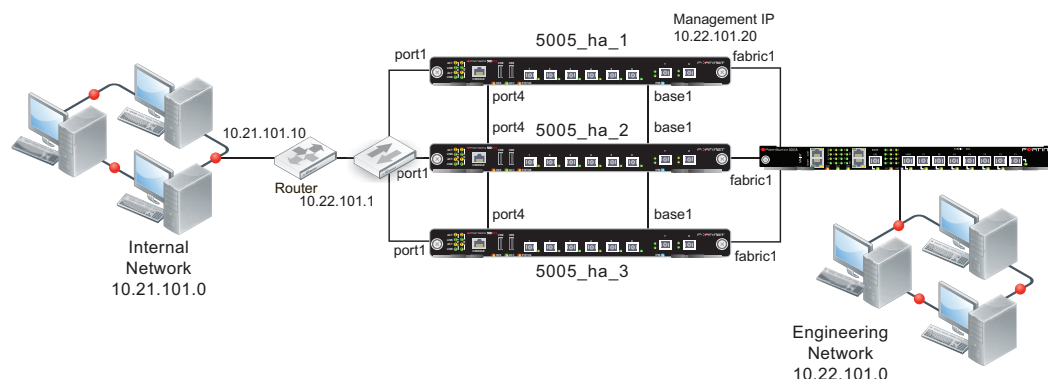
This section describes a more complex HA network topology that includes an HA cluster of three FortiGate-5002FA2 units running in Transparent mode and installed between an internal network and an engineering network.

- [Example Transparent mode HA network topology](#)
- [General configuration steps](#)

Example Transparent mode HA network topology

Figure 201 shows a Transparent mode FortiGate-5005FA2 HA cluster consisting of three FortiGate-5005FA2 units (5005_ha_1, 5005_ha_2, and 5005_ha_3) installed in a FortiGate-5000 series chassis with one FortiSwitch-5003A board. The cluster applies virus scanning to traffic passing between an engineering network and an internal network. The topology includes a router that performs NAT between the internal network and the engineering network. The cluster is connected to the engineering network with an management IP address of 10.22.101.20. This IP address is on the engineering network subnet.

Figure 201: Transparent mode HA network topology



By default fabric1 and fabric2 are the FortiGate-5005FA2 heartbeat interfaces. This example changes the heartbeat configuration to use the base1 and port4 interfaces for the heartbeat. The base1 connection is handled using the base backplane channel switched by the FortiSwitch-5003A board. The port4 connection is handled by connecting the port4 interfaces together using a switch.

The cluster connects to the engineering network using fabric1. The FortiSwitch-5003A board provides switching for the fabric1 interfaces and the fabric1 connection to the engineering network.

Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - web-based manager

These procedures assume you are starting with three FortiGate-5005FA2 units with factory default settings but not installed in chassis slots and a FortiSwitch-5003A board installed in chassis slot 1. The chassis is powered on. This configuration works for a FortiGate-5050 chassis or for a FortiGate-5140 chassis. No configuration changes to the FortiSwitch-5003A board are required.

To configure the FortiGate-5005FA2 units

- 1 Power on the first FortiGate unit by inserting it into chassis slot 5.
- 2 Connect port1 to the network and log into the web-based manager.
- 3 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 4 Enter a new Host Name for this FortiGate unit.

New Name	5005_ha_1
-----------------	-----------

- 5 Select OK.
- 6 Go to *System > Network > Interface* and select *Show backplane interfaces*.
- 7 Make sure the administrative status and link status is for base1 and fabric1.

You can edit the interface to set the administrative status to up. The link status will be up if the administrative status is up and the FortiGate-5005FA2 board can connect to the FortiSwitch-5003A board.

- 8 Go to *System > Config > HA* and change the following settings:

Mode	Active-Active	
Group Name	example3.com	
Password	HA_pass_3	
Heartbeat Interface		
	Enable	Priority
base1	Select	50
fabric1	Clear check box	0
fabric2	Clear check box	0
port4	Select	50

9 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-5005FA2 interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- port1 interface virtual MAC: 00-09-0f-09-00-04
- port2 interface virtual MAC: 00-09-0f-09-00-05
- port3 interface virtual MAC: 00-09-0f-09-00-06
- port4 interface virtual MAC: 00-09-0f-09-00-07
- port5 interface virtual MAC: 00-09-0f-09-00-08
- port6 interface virtual MAC: 00-09-0f-09-00-09
- port7 interface virtual MAC: 00-09-0f-09-00-0a
- port8 interface virtual MAC: 00-09-0f-09-00-0b

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HWaddr      00:09:0f:09:00:04
Permanent_HWaddr    00:09:0f:71:0a:dc
.
.
.
```

10 Power off the first FortiGate unit.**11** Repeat these steps for the second and third FortiGate units, with the following difference.

Set the second FortiGate unit host name to:

New Name	5005_ha_2
-----------------	-----------

Set the third FortiGate unit host name to:

New Name	5005_ha_3
-----------------	-----------

As you insert and configure each FortiGate unit, they will negotiate and join the cluster using the base1 interface for HA heartbeat communication.

To connect the cluster to the network

- 1 Connect the port1 interfaces of the cluster to a switch that can connect to the router and the internal network.
- 2 Connect the port4 interfaces of the cluster units together using a switch.
These interfaces become the backup heartbeat interface.
- 3 Connect one of the FortiSwitch-5003A front panel fabric interfaces (for example, F3) to the engineering network.

To switch the cluster to operate in Transparent mode

Switching from NAT/Route to Transparent mode also involves adding the Transparent mode management IP address and default route.

- 1 Log into the web-based manager.
- 2 Under System Information, beside *Operation Mode* select *Change*.
- 3 Set *Operation Mode* to *Transparent*.
- 4 Configure basic Transparent mode settings.

Operation Mode	Transparent
Management IP/Mask	10.22.101.20/24
Default Gateway	10.22.101.1

- 5 Select Apply.
The cluster switches to operating in Transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change. You must login again using the new TP address.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

- 1 View the system dashboard.
The System Information dashboard widget shows the *Cluster Name* (example3.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.
- 2 Go to *System > Config > HA* to view the cluster members list.
The list shows three cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster. The following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.

- 3 For *admin*, select the *Change Password* icon
- 4 Enter and confirm a new password.
- 5 Select OK.

The default route was changed when you switched to Transparent mode.

Configuring a Transparent mode active-active cluster of three FortiGate-5005FA2 units - CLI

Use the following procedures to configure the three FortiGate-5005FA2 units for Transparent mode HA operation using the FortiGate CLI.

To configure the FortiGate-5005FA2 units

- 1 Power on the first FortiGate unit by inserting it into chassis slot 5.
- 2 Connect port1 to the network and log into the CLI.

You can also use a console connection.

- 3 Change the host name for this FortiGate unit. For example:

```
config system global
  set hostname 5005_ha_1
end
```

- 4 Enable showing backplane interfaces.

```
config system global
  set show-backplane-intf enable
end
```

- 5 Make sure the administrative status and link status is up for base1 and fabric1.

Enter `get system interface` to view the status of these interfaces.

You can use the following commands to set the administrative status to up for these interfaces.

```
config system interface
  edit base1
    set status up
  next
  edit fabric1
    set status up
end
```

- 6 Configure HA settings.

```
config system ha
  set mode a-a
  set group-name example3.com
  set password HA_pass_3
  set hbdev base1 50 port4 50
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- port1 interface virtual MAC: 00-09-0f-09-00-04
- port2 interface virtual MAC: 00-09-0f-09-00-05
- port3 interface virtual MAC: 00-09-0f-09-00-06
- port4 interface virtual MAC: 00-09-0f-09-00-07
- port5 interface virtual MAC: 00-09-0f-09-00-08
- port6 interface virtual MAC: 00-09-0f-09-00-09
- port7 interface virtual MAC: 00-09-0f-09-00-0a
- port8 interface virtual MAC: 00-09-0f-09-00-0b

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HWaddr      00:09:0f:09:00:04
Permanent_HWaddr    00:09:0f:71:0a:dc
.
.
.
```

7 Display the HA configuration (optional).

```
get system ha
group-id             : 0
group-name           : example3.com
mode                 : a-a
password             : *
hbdev                : "base1" 50 "port4" 50
session-sync-dev     :
route-ttl            : 10
route-wait           : 0
route-hold           : 10
sync-config          : enable
encryption           : disable
authentication       : disable
```

```

hb-interval          : 2
hb-lost-threshold    : 20
helo-holddown        : 20
arps                 : 5
arps-interval        : 8
session-pickup       : disable
link-failed-signal   : disable
uninterruptable-upgrade: enable
ha-mgmt-status       : disable
ha-eth-type          : 8890
hc-eth-type          : 8891
l2ep-eth-type        : 8893
subsecond            : disable
vcluster2            : disable
vcluster-id          : 1
override             : disable
priority             : 128
schedule             : round-robin
monitor              :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom                 : "root"
load-balance-all    : disable

```

8 Repeat these steps for the second and third FortiGate units.

Set the second FortiGate unit host name to:

```

config system global
  set hostname 5005_ha_2
end

```

Set the third FortiGate unit host name to:

```

config system global
  set hostname 5005_ha_3
end

```

As you insert and configure each FortiGate unit they will negotiate and join the cluster using the base1 interface for HA heartbeat communication.

To connect the cluster to the network

- 1 Connect the port1 interfaces of the cluster to a switch that can connect to the router and the internal network.
- 2 Connect the port4 interfaces of the cluster units together using a switch.
These interfaces become the backup heartbeat interface.
- 3 Connect one of the FortiSwitch-5003A front panel fabric interfaces (for example, F3) to the engineering network.

To switch the cluster to Transparent mode

- 1 Log into the cluster CLI.
- 2 Change to Transparent mode.


```

config system settings
  set opmode transparent
  set manageip 10.22.101.20/24

```

```
set gateway 10.22.101.1
end
```

The cluster switches to Transparent Mode.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (10.22.101.20).

To view cluster status

Use the following steps to view cluster status from the CLI.

- 1 Log into the CLI.
- 2 To verify the HA status of the cluster unit that you logged into, enter the CLI command `get system status`. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

- 3 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 5005
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
load_balance: disable
schedule: round robin
Master:128 5005_ha_1          FG5A253E07600124 0
Slave :128 5005_ha_2          FG5A253E06500088 1
Slave :128 5005_ha_3          FG5A253E06500099 2
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG5A253E07600124
Slave :1 FG5A253E06500088
Slave :2 FG5A253E06500099
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add a password for the admin administrative account

- 1 Add a password for the admin administrative account.

```
config system admin
edit admin
set password <psswr>
```

end

Example: converting a standalone FortiGate unit to a cluster

You can convert an already configured and installed FortiGate unit into a cluster by configuring this FortiGate unit to be a primary unit and then adding subordinate units.

General configuration steps:

- Configure the original FortiGate unit for HA operation.
- Set the HA Device Priority of the original FortiGate unit to 255 to make sure that this FortiGate unit becomes the primary unit after cluster negotiation and synchronization.
- Back up the configuration of the original FortiGate unit.
- Configure one or more new FortiGate units with the same HA configuration as the original FortiGate unit with one exception. Keep the Unit Priority at the default setting, which is 128.
- Connect the FortiGate units to form a cluster and connect the cluster to your network.

When you power on all of the FortiGate units in the cluster, the original FortiGate unit becomes the primary unit. Its configuration is synchronized to all of the subordinate units. The entire cluster now operates with the original FortiGate unit configuration. No further configuration changes are required.

The new FortiGate units must:

- Have the same hardware configuration as the original FortiGate unit. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the original FortiGate unit.
- Be set to the same operating mode (NAT or Transparent) as the original FortiGate unit.
- Be operating in single VDOM mode.

In addition to one or more new FortiGate units, you need sufficient switches to connect all of the FortiGate interfaces in the cluster. Generally you will need one switch per interface, as it will have to connect that same interface on all cluster units. That is, all port1 interfaces use the port1 switch, port2 interfaces use the port2 switch, and so on. Intelligent switches that can be partitioned can reduce your switch requirements.

Converting a FortiGate unit to a primary unit and adding in the subordinate unit or units results in a brief service interruption as you disconnect and reconnect FortiGate interfaces and as the cluster negotiates. Therefore, conversion should only be done during off peak hours.

To configure the original FortiGate unit for HA operation

- 1 Connect to the FortiGate unit web-based manager.
- 2 Go to *System > Config > HA*.
- 3 Configure the FortiGate unit for HA operation.

Mode	Active-Active
Device Priority	255
Group Name	example4.com
Password	HA_pass_4

You can make other HA configuration changes after the cluster is operating.

4 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)).

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

5 Configure the new FortiGate units with the same HA configuration as the original FortiGate unit. The one exception is to keep the device priorities of the new FortiGate units at 128 to ensure the original FortiGate unit will become the primary unit in the new cluster.

Mode	Active-Active
Device Priority	128
Group Name	example4.com
Password	HA_pass_4

6 Configure the other FortiGate units to the same operation mode as the original FortiGate unit.

There is no need to make any other configuration changes (including network configuration changes) to the other FortiGate units.

7 Optionally power off all of the cluster units.

If you don't power off all of the units they may not negotiate to form a cluster when they are connected together.

8 Connect the cluster to your network.

For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 2025](#).

9 Power on all of the cluster units.

As the units start they change their MAC addresses and then negotiate to choose the primary unit and the subordinate units. This negotiation occurs with no user intervention and normally takes less than a minute.

The original the FortiGate unit becomes the primary unit because the device priority of the original FortiGate unit is higher than the device priority of the other FortiGate units. The configuration of the original FortiGate unit is synchronized to all the cluster units. As a result, the cluster is quickly up and running and configured for your network. No further configuration changes are required.

Example: adding a new unit to an operating cluster

This procedure describes how to add a new FortiGate unit to a functioning cluster. Adding a new unit to a cluster does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the new cluster unit.

You can use this procedure to add as many units as required to the cluster.

To add a new unit to a functioning cluster

- 1 Install the same firmware build on the new cluster unit as is running on the cluster.
- 2 Configure the new cluster unit for HA operation with the same HA configuration as the other units in the cluster.
- 3 If the cluster is running in Transparent mode, change the operating mode of the new cluster unit to Transparent mode.
- 4 Connect the new cluster unit to the cluster.
For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 2025](#).
- 5 Power on the new cluster unit.

When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the new unit configuration with the configuration of the primary unit.

You can add a new unit to a functioning cluster at any time. The new cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.

Example: replacing a failed cluster unit

This procedure describes how to remove a failed cluster unit from a cluster and add a new one to replace it. You can also use this procedure to remove a failed unit from a cluster, repair it and add it back to the cluster. Replacing a failed does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the replacement unit.

You can use this procedure to replace more than one cluster unit.

To replace a failed cluster unit

- 1 Disconnect the failed unit from the cluster and the network.
If you maintain other connections between the network and the still functioning cluster unit or units and between remaining cluster units network traffic will continue to be processed.
- 2 Repair the failed cluster unit, or obtain a replacement unit with the exact same hardware configuration as the failed cluster unit.
- 3 Install the same firmware build on the repaired or replacement unit as is running on the cluster.
- 4 Configure the repaired or replacement unit for HA operation with the same HA configuration as the cluster.
- 5 If the cluster is running in Transparent mode, change the operating mode of the repaired or replacement cluster unit to Transparent mode.
- 6 Connect the repaired or replacement cluster unit to the cluster.
For example, for a configuration similar to the FortiGate-620B cluster configuration described in this chapter, see [“To connect the cluster to the network” on page 2025](#).

7 Power on the repaired or replacement cluster unit.

When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the repaired or replacement unit configuration with the configuration of the primary unit.

You can add a repaired or replacement unit to a functioning cluster at any time. The repaired or replacement cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.
- Be set to the same operating mode (NAT or Transparent) as the cluster.
- Be operating in single VDOM mode.


Example: HA and 802.3ad aggregated interfaces

On FortiGate models that support it you can use 802.3ad link aggregation to combine two or more interfaces into a single aggregated interface. 802.3ad Link Aggregation and its management protocol, Link Aggregation Control Protocol (LACP) are a method for combining multiple physical links into a single logical link. This increases both potential throughput and network resiliency. Using LACP, traffic is distributed among the physical interfaces in the link, potentially resulting in increased performance.

This example describes how to configure an HA cluster consisting of two FortiGate-620B units with two aggregated 1000 Mb connections to the Internet using port1 and port2 and two aggregated 1000 Mb connections to the internal network using port3 and port4. The aggregated interfaces are also configured as HA monitored interfaces.

Each of the aggregate links connects to a different switch. Each switch is configured for link aggregation (2x1000Mb).

Figure 202: Example cluster with aggregate interfaces



HA interface monitoring, link failover, and 802.3ad aggregation

When monitoring the aggregated interface, HA interface monitoring treats the aggregated link as a single interface and does not monitor the individual physical interfaces in the link. HA interface monitoring registers the link to have failed only if all the physical interfaces in the link have failed. If only some of the physical interfaces in the link fail or become disconnected, HA considers the link to be operating normally.

HA MAC addresses and 802.3ad aggregation

If a configuration uses the Link Aggregate Control Protocol (LACP) (either passive or active), LACP is negotiated over all of the interfaces in any link. For a standalone FortiGate unit, the FortiGate LACP implementation uses the MAC address of the first interface in the link to uniquely identify that link. For example, a link consisting of port1 and port2 interfaces would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. An aggregate interface in a cluster acquires the virtual MAC address that would have been acquired by the first interface in the aggregate.

Link aggregation, HA failover performance, and HA mode

To operate an active-active or active-passive cluster with aggregated interfaces and for best performance of a cluster with aggregated interfaces, the switches used to connect the cluster unit aggregated interfaces together should support configuring multiple Link Aggregation (LAG) groups.

For example, the cluster shown in [Figure 202](#) should be configured into two LAG groups on the external switch: one for the port1 and port2 aggregated interface of 620_ha_1 and a second one for the port1 and port2 aggregate interface of 620_ha_2. You should also be able to do the same on the internal switch for the port3 and port4 aggregated interfaces of each cluster unit.

As a result, the subordinate unit aggregated interfaces would participate in LACP negotiation while the cluster is operating. In an active-active mode cluster, packets could be redirected to the subordinate unit interfaces. As well, in active-active or active-passive mode, after a failover the subordinate unit can become a primary unit without having to perform LACP negotiation before it can process traffic. Performing LACP negotiation causes a minor failover delay.

However if you cannot configure multiple LAG groups on the switches, due to the primary and subordinate unit interfaces having the same MAC address, the switch will put all of the interfaces into the same LAG group which would disrupt the functioning of the cluster. To prevent this from happening, you must change the FortiGate aggregated interface configuration to prevent subordinate units from participating in LACP negotiation.

For example, use the following command to prevent subordinate units from participating in LACP negotiation with an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```

As a result of this setting, subordinate unit aggregated interfaces cannot accept packets. This means that you cannot operate the cluster in active-active mode because in active-active mode the subordinate units must be able to receive and process packets. Also, failover may take longer because after a failover the subordinate unit has to perform LACP negotiation before being able to process network traffic.

Also, it may also be necessary to configure the switch to use Passive or even Static mode for LACP to prevent the switch from sending packets to the subordinate unit interfaces, which won't be able to process them.

Finally, in some cases depending on the LACP configuration of the switches, you may experience delayed failover if the FortiGate LACP configuration is not compatible with the switch LACP configuration. For example, in some cases setting the FortiGate LACP mode to static reduces the failover delay because the FortiGate unit does not perform LACP negotiation. However there is a potential problem with this configuration because static LACP does not send periodic LAC Protocol Data Unit (LACPDU) packets to test the connections. So a non-physical failure (for example, if a device is not responding because its too busy) may not be detected and packets could be lost or delayed.

General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

- 1 Configure the FortiGate units for HA operation.
 - Change each unit's host name.
 - Configure HA.
- 2 Connect the cluster to the network.
- 3 View cluster status.
- 4 Add basic configuration settings and configure the aggregated interfaces.
 - Add a password for the admin administrative account.
 - Add the aggregated interfaces.
 - Disable `lacp-ha-slave` so that the subordinate unit does not send LACP packets.
 - Add a default route.

You could also configure aggregated interfaces in each FortiGate unit before the units form a cluster.

- 5 Configure HA port monitoring for the aggregated interfaces.

Configuring active-passive HA cluster that includes aggregated interfaces - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the web-based manager.
- 2 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 3 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
----------	----------

- 4 Select OK.

- 5 Go to *System > Config > HA* and change the following settings.

Mode	Active-Passive	
Group Name	example5.com	
Password	HA_pass_5	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for a aggregated interface, you must change the HA heartbeat configuration to not use those interfaces.

6 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7 Power off the first FortiGate unit.

- 8 Repeat these steps for the second FortiGate unit.

Set the second FortiGate unit host name to:

New Name	620_ha_2
-----------------	----------

To connect the cluster to the network

- 1 Connect the port1 and port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
Configure the switch so that the port1 and port2 of 620_ha_1 make up an aggregated interface and port1 and port2 of 620_ha_2 make up a second aggregated interface.
- 2 Connect the port3 and port4 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
Configure the switch so that the port3 and port4 of 620_ha_1 make up an aggregated interface and port3 and port4 of 620_ha_2 make up another aggregated interface.
- 3 Connect the port5 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.
The units negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.
When negotiation is complete, the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

- 1 View the system dashboard.
The System Information dashboard widget shows the *Cluster Name* (example5.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.
- 2 Go to *System > Config > HA* to view the cluster members list.
The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.
- 3 For *admin*, select the *Change Password* icon.
- 4 Enter and confirm a new password.
- 5 Select OK.

- 6 Go to *Router > Static* and temporarily delete the default route.
You cannot add an interface to a aggregated interface if any settings (such as the default route) are configured for it.
- 7 Go to *System > Network > Interface* and select *Create New* to add the aggregate interface to connect to the Internet.
- 8 Set *Type* to *802.3ad Aggregate* and configure the aggregate interface to be connected to the Internet:

Name	Port1_Port2
Physical Interface Members	
Selected Interfaces	port1, port2
IP/Netmask	172.20.120.141/24

- 9 Select OK.
- 10 Select *Create New* to add the aggregate interface to connect to the internal network.
- 11 Set *Type* to *802.3ad Aggregate* and configure the aggregate interface to be connected to the Internet:

Name	Port3_Port4
Physical Interface Members	
Selected Interfaces	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12 Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13 Connect to the CLI and enter the following command to disable sending LACP packets from the subordinate unit:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  next
  edit Port3_Port4
    set lacp-ha-slave disable
end
```

14 Go to *Router > Static*.**15** Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

16 Select OK.

To configure HA port monitoring for the aggregate interfaces

- 1 Go to *System > Config > HA*.
- 2 In the cluster members list, edit the primary unit.
- 3 Configure the following port monitoring for the aggregate interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

- 4 Select OK.

Configuring active-passive HA cluster that includes aggregate interfaces - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the CLI.
- 2 Change the host name for this FortiGate unit:

```
config system global
    set hostname 620_ha_1
end
```

- 3 Configure HA settings.

```
config system ha
    set mode a-p
    set group-name example5.com
    set password HA_pass_5
    set hbdev port5 50 port6 50
```

end

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```


4 Display the HA configuration (optional).

```

get system ha
group-id          : 0
group-name        : example5.com
mode              : a-p
password          : *
hbdev             : "port5" 50 "port6" 50
session-sync-dev  :
route-ttl         : 10
route-wait        : 0
route-hold        : 10
sync-config       : enable
encryption        : disable
authentication    : disable
hb-interval       : 2
hb-lost-threshold : 20
helo-holddown     : 20
arps              : 5
arps-interval     : 8
session-pickup    : disable
link-failed-signal : disable
uninterruptable-upgrade: enable
ha-mgmt-status    : disable
ha-eth-type       : 8890
hc-eth-type       : 8891
l2ep-eth-type     : 8893
subsecond         : disable
vcluster2         : disable
vcluster-id       : 1
override          : disable
priority          : 128
monitor           :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom              : "root"

```

5 Repeat these steps for the other FortiGate unit.

Set the other FortiGate unit host name to:

```

config system global
  set hostname 620_ha_2
end

```

To connect the cluster to the network

- 1 Connect the port1 and port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.

Configure the switch so that the port1 and port2 of 620_ha_1 make up an aggregated interface and port1 and port2 of 620_ha_2 make up another aggregated interface.

- 2 Connect the port3 and port4 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.

Configure the switch so that the port3 and port4 of 620_ha_1 make up an interfaced and port3 and port4 of 620_ha_2 make up another aggregated interface.

- 3 Connect the port5 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.
When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

- 1 Log into the CLI.
- 2 Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

- 3 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2          FG600B3908600825 0
Slave :128 620_ha_1          FG600B3908600705 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings and the aggregate interfaces.

- 1 Add a password for the admin administrative account.
`config system admin`

```
edit admin
    set password <psswr>
end
```

2 Temporarily delete the default route.

You cannot add an interface to an aggregate interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
    delete 1
end
```

3 Add the aggregate interfaces:

```
config system interface
    edit Port1_Port2
        set type aggregate
        set lacp-ha-slave disable
        set member port1 port2
        set ip 172.20.120.141/24
        set vdom root
    next
    edit Port3_Port4
        set type aggregate
        set lacp-ha-slave disable
        set member port3 port4
        set ip 10.11.101.100/24
        set vdom root
```

```
end
```

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4 Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the aggregate interfaces


1 Configure HA port monitoring for the aggregate interfaces.

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Example: HA and redundant interfaces

On FortiGate models that support it you can combine two or more interfaces into a single redundant interface. A redundant interface consists of two or more physical interfaces. Traffic is processed by the first physical interface in the redundant interface. If that physical interface fails, traffic fails over to the next physical interface. Redundant interfaces don't have the benefit of improved performance that aggregate interfaces can have, but they do provide failover if a physical interface fails or is disconnected.

Figure 203: Example cluster with a redundant interfaces



This example describes how to configure an HA cluster consisting of two FortiGate-620B units with a redundant interface connection to the Internet and to an internal network. The connection to the Internet uses port1 and port2. The connection to the internal network uses port3 and port4. The HA heartbeat uses port5 and port6.

The redundant interfaces are also configured as HA monitored interfaces.

HA interface monitoring, link failover, and redundant interfaces

HA interface monitoring monitors the redundant interface as a single interface and does not monitor the individual physical interfaces in the redundant interface. HA interface monitoring registers the redundant interface to have failed only if all the physical interfaces in the redundant interface have failed. If only some of the physical interfaces in the redundant interface fail or become disconnected, HA considers the redundant interface to be operating normally.

HA MAC addresses and redundant interfaces

For a standalone FortiGate unit a redundant interface has the MAC address of the first physical interface added to the redundant interface configuration. A redundant interface consisting of port1 and port2 would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. A redundant interface in a cluster acquires the virtual MAC address that would have been acquired by the first physical interface added to the redundant interface configuration.

Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode

HA assigns the same virtual MAC addresses to the subordinate unit interfaces as are assigned to the corresponding primary unit interfaces. Consider a cluster of two FortiGate units operating in active-passive mode with a redundant interface consisting of port1 and port2. You can connect multiple redundant interfaces to the same switch if you configure the switch so that it defines multiple separate redundant interfaces and puts the redundant interfaces of each cluster unit into separate redundant interfaces. In this configuration, each cluster unit forms a separate redundant interface with the switch.

However, if the switch is configured with a single four-port redundant interface configuration, because the same MAC addresses are being used by both cluster units, the switch adds all four interfaces (port1 and port2 from the primary unit and port1 and port2 from the subordinate unit) to the same redundant interface.

To avoid unpredictable results, when you connect a switch to multiple redundant interfaces in an active-passive cluster you should configure separate redundant interfaces on the switch; one for each cluster unit.

Connecting multiple redundant interfaces to one switch while operating in active-active HA mode

In an active-active cluster, all cluster units send and receive packets. To operate a cluster with redundant interfaces in active-active mode, with multiple redundant interfaces connected to the same switch, you must separate the redundant interfaces of each cluster unit into different redundant interfaces on the connecting switch.

General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

- 1 Configure the FortiGate units for HA operation.
 - Change each unit's host name.
 - Configure HA.
- 2 Connect the cluster to the network.
- 3 View cluster status.
- 4 Add basic configuration settings and configure the redundant interfaces.
 - Add a password for the admin administrative account.
 - Add the redundant interfaces.
 - Add a default route.

You could also configure redundant interfaces in each FortiGate unit before they form a cluster.

- 5 Configure HA port monitoring for the redundant interfaces.

Configuring active-passive HA cluster that includes redundant interfaces - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the web-based manager.
- 2 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 3 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
-----------------	----------

- 4 Select OK.
- 5 Go to *System > Config > HA* and change the following settings.

Mode	Active-Passive	
Group Name	example6.com	
Password	HA_pass_6	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

6 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7 Power off the first FortiGate unit.

- 8 Repeat these steps for the second FortiGate unit.

Set the second FortiGate unit host name to:

New Name	620_ha_2
-----------------	----------

To connect the cluster to the network

- 1 Connect the port1 and port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
Configure the switch so that the port1 and port2 of 620_ha_1 make up a redundant interface and port1 and port2 of 620_ha_2 make up another redundant interface.
- 2 Connect the port3 and port4 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
Configure the switch so that the port3 and port4 of 620_ha_1 make up a redundant interface and port3 and port4 of 620_ha_2 make up another redundant interface.
- 3 Connect the port5 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.
When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

- 1 View the system dashboard.
The System Information dashboard widget shows the *Cluster Name* (example5.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.
- 2 Go to *System > Config > HA* to view the cluster members list.
The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.
- 3 For *admin*, select the *Change Password* icon
- 4 Enter and confirm a new password.
- 5 Select OK.

- 6 Go to *Router > Static* and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
- 7 Go to *System > Network > Interface* and select *Create New* to add the redundant interface to connect to the Internet.
- 8 Set *Type* to *Redundant Interface* and configure the redundant interface to be connected to the Internet:

Name	Port1_Port2
Physical Interface Members	
Selected Interfaces	port1, port2
IP/Netmask	172.20.120.141/24

- 9 Select OK.
- 10 Select *Create New* to add the redundant interface to connect to the internal network.
- 11 Set *Type* to *Redundant Interface* and configure the redundant interface to be connected to the Internet:

Name	Port3_Port4
Physical Interface Members	
Selected Interfaces	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12 Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13 Go to *Router > Static*.**14** Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

15 Select OK.**To configure HA port monitoring for the redundant interfaces**

- 1** Go to *System > Config > HA*.
- 2** In the cluster members list, edit the primary unit.
- 3** Configure the following port monitoring for the redundant interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4 Select OK.

Configuring active-passive HA cluster that includes redundant interfaces - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the CLI.
- 2 Change the host name for this FortiGate unit:

```
config system global
  set hostname 620_ha_1
end
```

- 3 Configure HA settings.

```
config system ha
  set mode a-p
  set group-name example6.com
  set password HA_pass_6
  set hbdev port5 50 port6 50
```

end

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4 Display the HA configuration (optional).

```

get system ha
group-id          : 0
group-name        : example6.com
mode              : a-p
password          : *
hbdev             : "port5" 50 "port6" 50
session-sync-dev  :
route-ttl         : 10
route-wait        : 0
route-hold        : 10
sync-config       : enable
encryption        : disable
authentication    : disable
hb-interval       : 2
hb-lost-threshold : 20
helo-holddown     : 20
arps              : 5
arps-interval     : 8
session-pickup    : disable
link-failed-signal : disable
uninterruptable-upgrade: enable
ha-mgmt-status    : disable
ha-eth-type       : 8890
hc-eth-type       : 8891
l2ep-eth-type     : 8893
subsecond         : disable
vcluster2         : disable
vcluster-id       : 1
override          : disable
priority          : 128
monitor           :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom              : "root"

```

5 Repeat these steps for the other FortiGate unit.

Set the other FortiGate unit host name to:

```

config system global
  set hostname 620_ha_2
end

```

To connect the cluster to the network

- 1** Connect the port1 and port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.

Configure the switch so that the port1 and port2 of 620_ha_1 make up a redundant interface and port1 and port2 of 620_ha_2 make up another redundant interface.

- 2** Connect the port3 and port4 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.

Configure the switch so that the port3 and port4 of 620_ha_1 make up a redundant interface and port3 and port4 of 620_ha_2 make up another redundant interface.

- 3 Connect the port5 interfaces of 620_ha_1 and 620_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 4 Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 5 Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.
When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

- 1 Log into the CLI.
- 2 Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

- 3 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2          FG600B3908600825 0
Slave :128 620_ha_1          FG600B3908600705 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings and the redundant interfaces.

- 1 Add a password for the admin administrative account.
`config system admin`

```
edit admin
  set password <psswrd>
end
```

2 Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
  delete 1
end
```

3 Add the redundant interfaces:

```
config system interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
    set ip 172.20.120.141/24
    set vdom root
  next
  edit Port3_Port4
    set type redundant
    set member port3 port4
    set ip 10.11.101.100/24
    set vdom root
```


end

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4 Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the redundant interfaces

1 Configure HA port monitoring for the redundant interfaces.

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Troubleshooting HA clusters

This section describes some HA clustering troubleshooting techniques.

Before you set up a cluster

Before you set up a cluster ask yourself the following questions about the FortiGate units that you are planning to use to create a cluster.

- 1 Do all the FortiGate units have the same hardware configuration? Including the same hard disk configuration and the same AMC cards installed in the same slots?
- 2 Do all FortiGate units have the same firmware build?
- 3 Are all FortiGate units set to the same operating mode (NAT or Transparent)?
- 4 Are all the FortiGate units operating in single VDOM mode?
- 5 If the FortiGate units are operating in multiple VDOM mode do they all have the same VDOM configuration?



In some cases you may be able to form a cluster if different FortiGate units have different firmware builds, different VDOM configurations, and are in different operating modes. However, if you encounter problems they may be resolved by installing the same firmware build on each unit, and give them the same VDOM configuration and operating mode.

Troubleshooting the initial cluster configuration

This section describes how to check a cluster when it first starts up to make sure that it is configured and operating correctly. This section assumes you have already configured your HA cluster.

To verify that a cluster can process traffic and react to a failure

- 1 Add a basic security policy configuration and send network traffic through the cluster to confirm connectivity.

For example, if the cluster is installed between the Internet and an internal network, set up a basic internal to external security policy that accepts all traffic. Then from a PC on the internal network, browse to a website on the Internet or ping a server on the Internet to confirm connectivity.

- 2 From your management PC, set ping to continuously ping the cluster, and then start a large download, or in some other way establish ongoing traffic through the cluster.
- 3 While traffic is going through the cluster, disconnect the power from one of the cluster units.

You could also shut down or restart a cluster unit.

Traffic should continue with minimal interruption.

- 4 Start up the cluster unit that you disconnected.

The unit should re-join the cluster with little or no affect on traffic.

- 5 Disconnect a cable for one of the HA heartbeat interfaces.

The cluster should keep functioning, using the other HA heartbeat interface.

- 6 If you have port monitoring enabled, disconnect a network cable from a monitored interface.

Traffic should continue with minimal interruption.

To verify the cluster configuration - web-based manager

- 1 Log into the cluster web-based manager.
- 2 Check the system dashboard to verify that the System Information widget displays all of the cluster units.

- 3 Check the cluster member graphic to verify that the correct cluster unit interfaces are connected.
- 4 Go to *System > Config > HA* and verify that all of the cluster units are displayed on the cluster members list.
- 5 From the cluster members list, edit the primary unit (master) and verify the cluster configuration is as expected.

To troubleshoot the cluster configuration - web-based manager

- 1 Connect to each cluster unit web-based manager and verify that the HA configurations are the same.

To connect to each web-based manager, you may need to disconnect some units from the network to connect to the other if the units have the same IP address.

- 2 If the configurations are the same, try re-entering the cluster *Password* on each cluster unit in case you made an error typing the password when configuring one of the cluster units.
- 3 Check that the correct interfaces of each cluster unit are connected.

Check the cables and interface LEDs.

Use the Unit Operation dashboard widget, system network interface list, or cluster members list to verify that each interface that should be connected actually is connected.

If Link is down re-verify the physical connection. Try replacing network cables or switches as required.

To verify the cluster configuration - CLI

- 1 Log into each cluster unit CLI.

You can use the console connection if you need to avoid the problem of units having the same IP address.

- 2 Enter the command `get system status`.

Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

- 3 Verify that the `get system ha status` command displays all of the cluster units.
- 4 Enter the `get system ha` command to verify that the HA configuration is correct and the same for each cluster unit.

To troubleshoot the cluster configuration - CLI

- 1 Try using the following command to re-enter the cluster password on each cluster unit in case you made an error typing the password when configuring one of the cluster units.

```
config system ha
  set password <password>
end
```

- 2 Check that the correct interfaces of each cluster unit are connected.

Check the cables and interface LEDs.

Use `get hardware nic <interface_name>` command to confirm that each interface is connected. If the interface is connected the command output should contain a `Link: up` entry similar to the following:

```
get hardware nic port1
.
.
.
Link: up
.
.
.
```

If Link is down, re-verify the physical connection. Try replacing network cables or switches as required.

More troubleshooting information

Much of the information in this HA guide can be useful for troubleshooting HA clusters. Here are some links to sections with more information.

- If sessions are lost after a failover you may need to change route-ttl to keep synchronized routes active longer. See [“Change how long routes stay in a cluster unit routing table” on page 2192](#).
- To control which cluster unit becomes the primary unit, you can change the device priority and enable override. See [“Controlling primary unit selection using device priority and override” on page 2010](#).
- Changes made to a cluster can be lost if override is enabled. See [“Configuration changes can be lost if override is enabled” on page 2011](#).
- In some cases, age differences among cluster units result in the wrong cluster unit becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units. You can resolve this problem by resetting the age of one or more cluster units. See [“Resetting the age of all cluster units” on page 2005](#). You can also adjust how sensitive the cluster is to age differences. This can be useful if large age differences cause problems. See [“Cluster age difference margin \(grace period\)” on page 2003](#) and [“Changing the cluster age difference margin” on page 2003](#).
- If one of the cluster units needs to be serviced or removed from the cluster for other reasons, you can do so without affecting the operation of the cluster. See [“Disconnecting a cluster unit from a cluster” on page 2164](#).
- The web-based manager and CLI will not allow you to configure HA if:
 - You have configured a FortiGate interface to get its IP address using DHCP or PPPoE. See [“FortiGate HA compatibility with PPPoE and DHCP” on page 2012](#).
 - You have enabled VRRP. See [“VRRP” on page 2237](#).
 - You have enabled TCP session synchronization. See [“TCP session synchronization” on page 2243](#).

- Some third-party network equipment may prevent HA heartbeat communication, resulting in a failure of the cluster or the creation of a split brain scenario. For example, some switches use packets with the same Ethertype as HA heartbeat packets use for internal functions and when used for HA heartbeat communication the switch generates CRC errors and the packets are not forwarded. See [“Heartbeat packet Ethertypes” on page 2174](#).
- Very busy clusters may not be able to send HA heartbeat packets quickly enough, also resulting in a split brain scenario. You may be able to resolve this problem by modifying HA heartbeat timing. See [“Modifying heartbeat timing” on page 2175](#).
- Very busy clusters may suffer performance reductions if session pickup is enabled. If possible you can disable this feature to improve performance. If you require session pickup for your cluster, several options are available for improving session pickup performance. See [“Improving session synchronization performance” on page 2205](#).
- If it takes longer than expected for a cluster to failover you can try changing how the primary unit sends gratuitous ARP packets. See [“Changing how the primary unit sends gratuitous ARP packets after a failover” on page 2178](#).
- You can also improve failover times by configuring the cluster for subsecond failover. See [“Subsecond failover” on page 2200](#) and [“Failover performance” on page 2214](#).
- When you first put a FortiGate unit in HA mode you may lose connectivity to the unit. This occurs because HA changes the MAC addresses of all FortiGate unit interfaces, including the one that you are connecting to. The cluster MAC addresses also change if you change some HA settings such as the cluster group ID. The connection will be restored in a short time as your network and PC updates to the new MAC address. To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.
- Since HA changes all cluster unit MAC addresses, if your network uses MAC address filtering you may have to make configuration changes to account for the HA MAC addresses.
- A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID. See [“Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain” on page 2182](#).
- The cluster CLI displays `slave is not in sync` messages if there is a synchronization problem between the primary unit and one or more subordinate units. See [“How to diagnose HA out of sync messages” on page 2189](#).
- If you have configured dynamic routing and the new primary unit takes too long to update its routing table after a failover you can configure graceful restart and also optimize how routing updates are synchronized. See [“Configuring graceful restart for dynamic routing failover” on page 2191](#) and [“Controlling how the FGCP synchronizes routing updates” on page 2192](#).

- Some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur after a link failover if the switch does not detect the failure and does not clear its MAC forwarding table. See [“Updating MAC forwarding tables when a link failover occurs” on page 2198](#).
- If a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails you can enable remote link failover to maintain communication. See [“Remote link failover” on page 2200](#).
- If you find that some cluster units are not running the same firmware build you can reinstall the correct firmware build on the cluster to upgrade all cluster units to the same firmware build. See [“Synchronizing the firmware build running on a new cluster unit” on page 2153](#).



Configuring and connecting virtual clusters

This chapter provides an introduction to virtual clustering and also contains general procedures and configuration examples that describe how to configure FortiGate HA virtual clustering.

This chapter contains the following sections:

- [Virtual clustering overview](#)
- [Configuring HA for virtual clustering](#)
- [Example: virtual clustering with two VDOMs and VDOM partitioning](#)
- [Example: inter-VDOM links in a virtual clustering configuration](#)
- [Troubleshooting virtual clustering](#)

Virtual clustering overview

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGate units operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

[Figure](#) shows an example virtual cluster configuration consisting of two FortiGate-620B units. The virtual cluster has two virtual domains, root and Eng_vdm.

The root virtual domain includes the port1 and port2 interfaces. The Eng_vdm virtual domain includes the port5 and port6 interfaces. The port3 and port4 interfaces (not shown in the diagram) are the HA heartbeat interfaces.



FortiGate virtual clustering is limited to a cluster of 2 FortiGate units with multiple VDOMs enabled. If you want to create a cluster of more than 2 FortiGate units operating with multiple VDOMS you could consider other solutions that either do not include multiple VDOMs in one cluster or employ a feature such as standalone session synchronization. See [“TCP session synchronization” on page 2243](#).

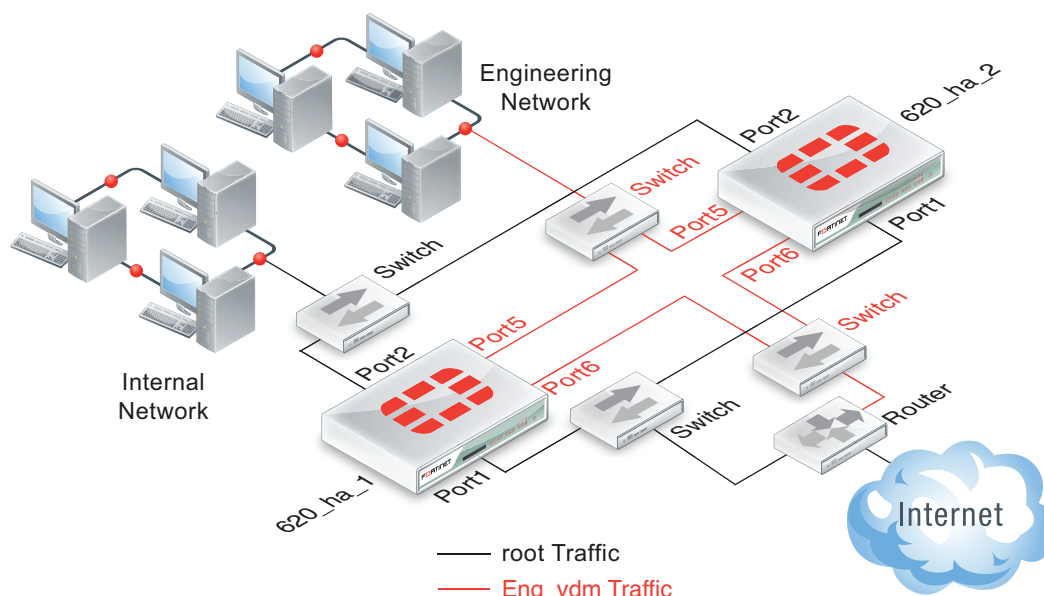
Virtual clustering and failover protection

Virtual clustering operates on a cluster of two (and only two) FortiGate units with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGate units in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

Figure 204: Example virtual cluster of two FortiGate-620B units



Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM partitioning from the web-based manager by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI `config system ha` command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.



If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.

For more information about HA override see [“HA override” on page 2008](#).

Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved.

From the web-based manager you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.

From the CLI you configure VDOM partitioning by setting the HA mode to `a-p`. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1. You do the same for virtual cluster 2 by entering the `config secondary-vcluster` command.

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit.

If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

Configuring HA for virtual clustering

If your cluster uses VDOMs, you are configuring virtual clustering. Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below.

To configure HA options for a cluster with VDOMs enabled:

- Log into the global web-based manager and go to *System > Config > HA*.
- From the CLI, log into the Global Configuration:

The following example shows how to configure active-active virtual clustering:

```

config global
  config system ha
    set mode a-a
    set group-name vexample1.com
    set password vHA_pass_1
  end
end

```

The following example shows how to configure active-passive virtual clustering:

```

config global
  config system ha
    set mode a-p
    set group-name vexample1.com
    set password vHA_pass_1
  end
end

```

The following example shows how to configure VDOM partitioning for virtual clustering. In the example, the FortiGate unit is configured with three VDOMs (domain_1, domain_2, and domain_3) in addition to the root VDOM. The example shows how to set up a basic HA configuration that sets the device priority of virtual cluster 1 to 200. The example also shows how to enable `vcluster2`, how to set the device priority of virtual cluster 2 to 100 and how to add the virtual domains `domain_2` and `domain_3` to virtual cluster 2.

When you enable multiple VDOMs, `vcluster2` is enabled by default. Even so the command to enable `vcluster2` is included in this example in case for some reason it has been disabled. When `vcluster2` is enabled, `override` is also enabled.

The result of this configuration would be that the cluster unit that you are logged into becomes the primary unit for virtual cluster 1. This cluster unit processes all traffic for the root and `domain_1` virtual domains.

```

config global
  config system ha
    set mode a-p
    set group-name vexample1.com
    set password vHA_pass_1
    set priority 200
    set vcluster2 enable
  config secondary-vcluster
    set vdom domain_2 domain_3
    set priority 100
  end
end
end

```

The following example shows how to use the `execute ha manage` command to change the device priorities for virtual cluster 1 and virtual cluster 2 for the other unit in the cluster. The commands set the device priority of virtual cluster 1 to 100 and virtual cluster 2 to 200.

The result of this configuration would be that the other cluster unit becomes the primary unit for virtual cluster 2. This other cluster unit would process all traffic for the `domain_2` and `domain_3` virtual domains.

```

config global
  execute ha manage 1
  config system ha
    set priority 100

```

```
        set vcluster2 enable
        config secondary-vcluster
            set priority 200
        end
    end
end
end
```

Example: virtual clustering with two VDOMs and VDOM partitioning

This section describes how to configure the example virtual clustering configuration shown in [Figure 205](#). This configuration includes two virtual domains, root and Eng_vdm and includes VDOM partitioning that sends all root VDOM traffic to 620_ha_1 and all Eng_vdm VDOM traffic to 620_ha_2. The traffic from the internal network and the engineering network is distributed between the two FortiGate units in the virtual cluster. If one of the cluster units fails, the remaining unit will process traffic for both VDOMs.

The procedures in this example describe some of many possible sequences of steps for configuring virtual clustering. For simplicity many of these procedures assume that you are starting with new FortiGate units set to the factory default configuration. However, this is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

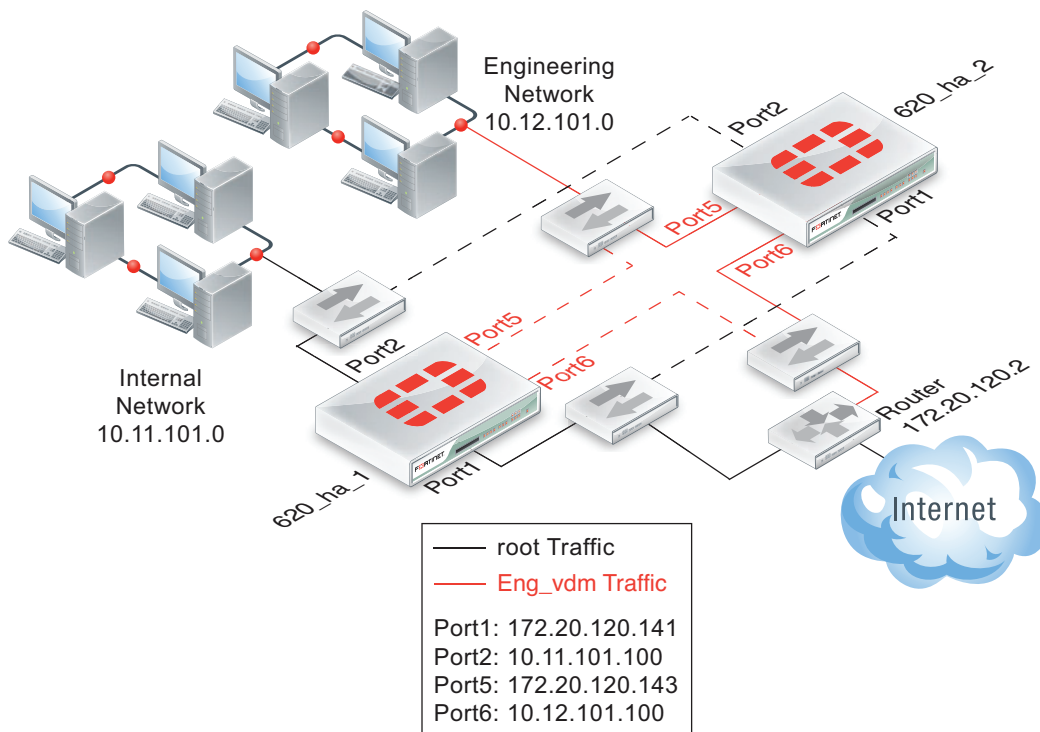
Example virtual clustering network topology

[Figure 205](#) shows a typical FortiGate-620B HA virtual cluster consisting of two FortiGate-620B units (620_ha_1 and 620_ha_2) connected to an internal network, an engineering network and the Internet. To simplify the diagram the heartbeat connections are not shown.

The traffic from the internal network is processed by the root VDOM, which includes the port1 and port2 interfaces. The traffic from the engineering network is processed by the Eng_vdm VDOM, which includes the port5 and port6 interfaces. VDOM partitioning is configured so that all traffic from the internal network is processed by 620_ha_1 and all traffic from the engineering network is processed by 620_ha_2.

This virtual cluster uses the default FortiGate-620B heartbeat interfaces (port3 and port4).

Figure 205: Example virtual cluster of two FortiGate-620B units showing VDOM partitioning



General configuration steps

The section includes web-based manager and CLI procedures. These procedures assume that the FortiGate-620B units are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

- 1 Configure the FortiGate units for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
- 2 Connect the cluster to the network.
- 3 Configure VDOM settings for the cluster:
 - Enable multiple VDOMs.
 - Add the Eng_vdm VDOM.
 - Add port5 and port6 to the Eng_vdm.
- 4 Configure VDOM partitioning.
- 5 Confirm that the cluster units are operating as a virtual cluster and add basic configuration settings to the cluster.
 - View cluster status from the web-based manager or CLI.
 - Add a password for the admin administrative account.
 - Change the IP addresses and netmasks of the port1, port2, port5, and port6 interfaces.
 - Add a default routes to each VDOM.

Configuring virtual clustering with two VDOMs and VDOM partitioning - web-based manager

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the web-based manager.
- 2 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 3 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
-----------------	----------

- 4 Select OK.
- 5 Go to *System > Config > HA* and change the following settings.

Mode	Active-Passive
Group Name	vexample2.com
Password	vHA_pass_2

6 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7 Power off the first FortiGate unit.

- 8 Repeat these steps for the second FortiGate unit.

Set the second FortiGate unit host name to:

New Name	620_ha_2
-----------------	----------

To connect the cluster to the network

- 1 Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2 Connect the port5 interfaces of 620_ha_1 and 620_ha_2 to switch connected to the Internet.
You could use the same switch for the port1 and port5 interfaces.
- 3 Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 4 Connect the port6 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the engineering network.
- 5 Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 6 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 7 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete you can continue.

To configure VDOM settings for the cluster

- 1 Log into the web-based manager.
- 2 On the *System Information* dashboard widget, beside *Virtual Domain* select *Enable*.
- 3 Select OK and then log back into the web-based manager.
- 4 Go to *System > VDOM* and select *Create New* to add a new VDOM.

Name	Eng_vdm
-------------	---------

- 5 Go to *System > Network > Interface*.
- 6 Edit the *port5* interface, add it to the Eng_vdm VDOM and configure other interface settings:

Alias	Engineering_external
Virtual Domain	Eng_vdm
IP/Netmask	172.20.120.143/24

- 7 Select OK.
- 8 Edit the *port6* interface, add it to the Eng_vdm VDOM and configure other interface settings:

Alias	Engineering_internal
Virtual Domain	Eng_vdm

IP/Netmask	10.120.101.100/24
Administrative Access	HTTPS, PING, SSH

- 9 Select OK.

To add a default route to each VDOM

- 1 Go to *System > VDOM* and Enter the root VDOM.
- 2 Go to *Router > Static*.
- 3 Change the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	port1
Distance	10

- 4 Select *Global*.
- 5 Go to *System > VDOM* and Enter the Eng_vdm VDOM.
- 6 Go to *Router > Static*.
- 7 Change the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	port5
Distance	10

To configure VDOM partitioning

- 1 Go to *System > Config > HA*.
The cluster members shows two cluster units in Virtual Cluster 1.
- 2 Edit the cluster unit with the *Role* of *MASTER*.
- 3 Change *VDOM partitioning* to move the Eng_vdm to the *Virtual Cluster 2* list.
- 4 Select OK.

- 5 Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit to the following:

	Device Priority	
Host Name	Virtual Cluster 1	Virtual Cluster 2
620_ha_1	200	100
620_ha_2	100	200

You can do this by editing the HA configurations of each cluster unit in the cluster members list and changing device priorities.

Since the device priority of Virtual Cluster 1 is highest for 620_ha_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by 620_ha_1.

Since the device priority of Virtual Cluster 2 is highest for 620_ha_2 and since the Eng_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng_vdm VDOM is processed by 620_ha_2.



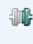
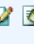












To view cluster status and verify the VDOM partitioning configuration

- 1 Log into the web-based manager.
- 2 Go to *System > Config > HA*.

The cluster members list should show the following:

- Virtual Cluster 1 contains the root VDOM.
- 620_ha_1 is the primary unit (master) for Virtual Cluster 1.
- Virtual Cluster 2 contains the Eng_vdm VDOM.
- 620_ha_2 is the primary unit (master) for Virtual Cluster 2.

Figure 206: Example virtual clustering cluster members list

Virtual Cluster 1						View HA Statistics
Virtual Domains: root						
	Cluster Member	Hostname	Role	Priority		
		620_ha_2	MASTER	128		
		620_ha_1	SLAVE	128		
Virtual Cluster 2						
Virtual Domains: Eng_vdm						
	Cluster Member	Hostname	Role	Priority		
		620_ha_2	MASTER	128		
		620_ha_1	SLAVE	128		

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for the root VDOM is processed by 620_ha_1 and traffic for the Eng_vdm is processed by 620_ha_2.

- 1 Log into the web-based manager by connecting to port2 using IP address 10.11.101.100.

You will log into 610_ha_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by 610_ha_1. You can confirm that you have logged into 610_ha_1 by checking the HTML title displayed by your web browser. The title will include the 610_ha_1 host name. Also on the System Information dashboard widget displays the serial number of the 610_ha_1 FortiGate unit.

- 2 Log into the web-based manager by connecting to port6 using IP address 10.12.101.100.

You will log into 610_ha_2 because port6 is in the Eng_vdm VDOM and all traffic for this VDOM is processed by 610_ha_2.

- 3 Add security policies to the root virtual domain that allows communication from the internal network to the Internet and connect to the Internet from the internal network.
- 4 Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.

The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620_ha_1 unit.

- 5 Add security policies to the Eng_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.

- 6 Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.

The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620_ha_2 unit.

Configuring virtual clustering with two VDOMs and VDOM partitioning - CLI

These procedures assume you are starting with two FortiGate-620B units with factory default settings.

To configure the FortiGate-620B units for HA operation

- 1 Power on the first FortiGate-620B unit and log into the CLI.
- 2 Change the host name for this FortiGate unit:

```
config system global
  set hostname 620_ha_1
end
```

- 3 Configure HA settings.

```
config system ha
  set mode a-p
  set group-name vexample2.com
  set password vHA_pass_2
```

end

The FortiGate unit negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4 Display the HA configuration (optional).

```
get system ha
```

```

group-id           : 0
group-name         : vexample2.com
mode               : a-p
password           : *
hbdev              : "port3" 50 "port4" 50
session-sync-dev   :
route-ttl          : 10
route-wait         : 0
route-hold         : 10
sync-config        : enable
encryption         : disable
authentication     : disable
hb-interval        : 2
hb-lost-threshold  : 20
helo-holddown      : 20
arps               : 5
arps-interval      : 8
session-pickup     : disable
link-failed-signal : disable
uninterruptable-upgrade: enable
ha-mgmt-status     : disable
ha-eth-type        : 8890
hc-eth-type        : 8891
l2ep-eth-type      : 8893
subsecond          : disable
vcluster2          : disable
vcluster-id        : 1
override           : disable
priority           : 128
monitor            :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-flip-timeout: 60
vdom               : "root"

```

5 Power off the first FortiGate unit.

6 Repeat these steps for the second FortiGate unit.

Set the other FortiGate unit host name to:

```

config system global
    set hostname 620_ha_2
end

```

To connect the cluster to the network

- 1** Connect the port1 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the Internet.
- 2** Connect the port5 interfaces of 620_ha_1 and 620_ha_2 to switch connected to the Internet.
You could use the same switch for port1 and port5.
- 3** Connect the port2 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the internal network.
- 4** Connect the port6 interfaces of 620_ha_1 and 620_ha_2 to a switch connected to the engineering network.

- 5 Connect the port3 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 6 Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
- 7 Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.
When negotiation is complete you can continue.

To configure VDOM settings for the cluster

- 1 Log into the CLI.
- 2 Enter the following command to enable multiple VDOMs for the cluster.

```
config system global
    set vdom-admin enable
end
```
- 3 Log back into the CLI.
- 4 Enter the following command to add the Eng_vdm VDOM:

```
config vdom
    edit Eng_vdm
end
```
- 5 Edit the port5 interface, add it to the Eng_vdm VDOM and configure other interface settings:

```
config global
    config system interface
        edit port5
            set vdom Eng_vdm
            set alias Engineering_external
            set ip 172.20.12.143/24
        next
        edit port6
            set vdom Eng_vdm
            set alias Engineering_internal
            set ip 10.120.101.100/24
        end
    end
end
```

To add a default route to each VDOM

- 1 Enter the following command to add default routes to the root and Eng_vdm VDOMs.

```
config vdom
    edit root
        config router static
            edit 1
                set dst 0.0.0.0/0.0.0.0
                set gateway 172.20.120.2
                set device port1
            end
        next
    edit Eng_vdm
        config router static
            edit 1
```

```

        set dst 0.0.0.0/0.0.0.0
        set gateway 172.20.120.2
        set device port5
    end
end

```

To configure VDOM partitioning

- 1 Enter the `get system ha status` command to view cluster unit status:

For example, from the 620_ha_2 cluster unit CLI:

```

config global
  get system ha status
  Model: 620
  Mode: a-p
  Group: 0
  Debug: 0
  ses_pickup: disable
  Master:128 620_ha_2          FG600B3908600825 0
  Slave :128 620_ha_1          FG600B3908600705 1
  number of vcluster: 1
  vcluster 1: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705

```

This command output shows that VDOM partitioning has not been configured because only virtual cluster 1 is shown. The command output also shows that the 620_ha_2 is the primary unit for the cluster and for virtual cluster 1 because this cluster unit has the highest serial number

- 2 Enter the following commands to configure VDOM partitioning:

```

config global
  config system ha
    set vcluster2 enable
    config secondary-vcluster
      set vdom Eng_vdm
    end
  end
end

```

- 3 Enter the `get system ha status` command to view cluster unit status:

For example, from the 620_ha_2 cluster unit CLI:

```

config global
  get system ha status
  Model: 620
  Mode: a-p
  Group: 0
  Debug: 0
  ses_pickup: disable
  Master:128 620_ha_2          FG600B3908600825 0
  Slave :128 620_ha_1          FG600B3908600705 1
  number of vcluster: 2
  vcluster 1: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705
  vcluster 2: work 169.254.0.1

```

```
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

This command output shows VDOM partitioning has been configured because both virtual cluster 1 and virtual cluster 2 are visible. However the configuration is not complete because 620_ha_2 is the primary unit for both virtual clusters. The command output shows this because under both vcluster entries the `Master` entry shows FG600B3908600825, which is the serial number of 620_ha_2. As a result of this configuration, 620_ha_2 processes traffic for both VDOMs and 620_ha_1 does not process any traffic.

- 4 Change the Virtual Cluster 1 and Virtual Cluster 2 device priorities for each cluster unit so that 620_ha_1 processes virtual cluster 1 traffic and 620_ha_2 processes virtual cluster 2 traffic.

Since the root VDOM is in virtual cluster 1 and the Eng_vdm VDOM is in virtual cluster 2 the result of this configuration will be that 620_ha_1 will process all root VDOM traffic and 620_ha_2 will process all Eng_vdm traffic. You make this happen by changing the cluster unit device priorities for each virtual cluster. You could use the following settings:

Host Name	Device Priority	
	Virtual Cluster 1	Virtual Cluster 2
620_ha_1	200	100
620_ha_2	100	200

Since the device priority is not synchronized you can edit the device priorities of each virtual cluster on each FortiGate unit separately. To do this:

- Log into the CLI and note the FortiGate unit you have actually logged into (for example, by checking the host name displayed in the CLI prompt).
- Change the virtual cluster 1 and 2 device priorities for this cluster unit.
- Then use the `execute ha manage` command to log into the other cluster unit CLI and set its virtual cluster 1 and 2 device priorities.

Enter the following commands from the 620_ha_1 cluster unit CLI:

```
config global
  config system ha
    set priority 200
    config secondary-vcluster
      set priority 100
    end
  end
end
```

Enter the following commands from the 620_ha_2 cluster unit CLI:

```
config global
  config system ha
    set priority 100
    config secondary-vcluster
      set priority 200
    end
  end
end
```

end



The cluster may renegotiate during this step resulting in a temporary loss of connection to the CLI and a temporary service interruption.

Since the device priority of Virtual Cluster 1 is highest for 620_ha_1 and since the root VDOM is in Virtual Cluster 1, all traffic for the root VDOM is processed by 620_ha_1.

Since the device priority of Virtual Cluster 2 is highest for 620_ha_2 and since the Eng_vdm VDOM is in Virtual Cluster 2, all traffic for the Eng_vdm VDOM is processed by 620_ha_2.

To verify the VDOM partitioning configuration

- 1 Log into the 620_ha_2 cluster unit CLI and enter the following command:

```
config global
get system ha status
Model: 620
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Slave :100 620_ha_2          FG600B3908600825 0
Master:200 620_ha_1          FG600B3908600705 1
number of vcluster: 2
vcluster 1: standby 169.254.0.2
Slave :1 FG600B3908600825
Master:0 FG600B3908600705
vcluster 2: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows that 620_ha_1 is the primary unit for virtual cluster 1 (because the command output show the `Master` of virtual cluster 1 is the serial number of 620_ha_1) and that 620_ha_2 is the primary unit for virtual cluster 2.

If you enter the same command from the 620_ha_1 CLI the same information is displayed but in a different order. The command always displays the status of the cluster unit that you are logged into first.

```
config global
get system ha status
Model: 620
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:200 620_ha_1          FG600B3908600705 1
Slave :100 620_ha_2          FG600B3908600825 0
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:0 FG600B3908600705
Slave :1 FG600B3908600825
vcluster 2: standby 169.254.0.1
Slave :1 FG600B3908600705
Master:0 FG600B3908600825
```


To test the VDOM partitioning configuration

You can do the following to confirm that traffic for the root VDOM is processed by 620_ha_1 and traffic for the Eng_vdm is processed by 620_ha_2. These steps assume the cluster is operating correctly.

- 1 Log into the CLI by connecting to port2 using IP address 10.11.101.100.
You will log into 610_ha_1 because port2 is in the root VDOM and all traffic for this VDOM is processed by 610_ha_1. You can confirm that you have logged into 610_ha_1 by checking the host name in the CLI prompt. Also the `get system status` command displays the status of the 610_ha_1 cluster unit.
- 2 Log into the web-based manager or CLI by connecting to port6 using IP address 10.12.101.100.
You will log into 610_ha_2 because port6 is in the Eng_vdm VDOM and all traffic for this VDOM is processed by 610_ha_2.
- 3 Add security policies to the root virtual domain that allow communication from the internal network to the Internet and connect to the Internet from the internal network.
- 4 Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620_ha_1 unit.
- 5 Add security policies to the Eng_vdm virtual domain that allow communication from the engineering network to the Internet and connect to the Internet from the engineering network.
- 6 Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*.
The statistics display shows more active sessions, total packets, network utilization, and total bytes for the 620_ha_2 unit.

Example: inter-VDOM links in a virtual clustering configuration

In a virtual domain configuration you can use inter-VDOM links to route traffic between two virtual domains operating in a single FortiGate unit without using physical interfaces. Adding an inter-VDOM link has the affect of adding two interfaces to the FortiGate unit and routing traffic between the virtual domains using the inter-VDOM link interfaces.

In a virtual clustering configuration inter-VDOM links can only be made between virtual domains that are in the same virtual cluster. So, if you are planning on configuring inter-VDOM links in a virtual clustering configuration, you should make sure the virtual domains that you want to link are in the same virtual cluster.

For example, [Table 114](#) and [Table 115](#) show an example virtual clustering configuration where each virtual cluster contains three virtual domains. In this configuration you can configure inter-VDOM links between root and vdom_1 and between vdom_2 and vdom_3. But, you cannot configure inter-VDOM links between root and vdom_2 or between vdom_1 and vdom_3 (and so on).

Table 114: Virtual Cluster 1 configuration

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
root vdom_1	Priority 200	Priority 100
	Role Primary	Role Subordinate

Table 115: Virtual Cluster 2 configuration

Virtual Domains	Hostname	
	FortiGate_A	FortiGate_B
vdom_2 vdom_3	Priority 100	Priority 200
	Role Subordinate	Role Primary

Configuring inter-VDOM links in a virtual clustering configuration

Configuring inter-VDOM links in a virtual clustering configuration is very similar to configuring inter-VDOM links for a standalone FortiGate unit. The main difference the `config system vdom-link` command includes the `vcluster` keyword. The default setting for `vcluster` is `vcluster1`. So you only have to use the `vcluster` keyword if you are added an inter-VDOM link to virtual cluster 2.

To add an inter-VDOM link to virtual cluster 1

This procedure describes how to create an inter-VDOM link to virtual cluster 1 that results in a link between the root and vdom_1 virtual domains.



Inter-VDOM links are also called internal point-to-point interfaces.

- 1 Add an inter-VDOM link called `vc1link`.

```
config global
  config system vdom-link
    edit vc1link
  end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc1link0` and `vc1link1`. These interfaces appear in all CLI and web-based manager interface lists. These interfaces can only be added to virtual domains in virtual cluster 1.

- 2 Bind the `vc1link0` interface to the root virtual domain and bind the `vc1link1` interface to the vdom_1 virtual domain.

```

config system interface
    edit vc1link0
        set vdom root
    next
    edit vc1link1
        set vdom vdom_1
    end

```

To add an inter-VDOM link to virtual cluster 2

This procedure describes how to create an inter-VDOM link to virtual cluster 2 that results in a link between the vdom_2 and vdom_3 virtual domains.

3 Add an inter-VDOM link called vc2link.

```

config global
    config system vdom-link
        edit vc2link
            set vcluster vcluster2
        end
    end

```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called vc2link0 and vc2link1. These interfaces appear in all CLI and web-based manager interface lists. These interfaces can only be added to virtual domains in virtual cluster 2.

4 Bind the vc2link0 interface to the vdom_2 virtual domain and bind the vc2link1 interface to the vdom_3 virtual domain.

```

config system interface
    edit vc2link0
        set vdom vdom_2
    next
    edit vc2link1
        set vdom vdom_3
    end

```

Troubleshooting virtual clustering

Troubleshooting virtual clusters is similar to troubleshooting any cluster (see [“Troubleshooting HA clusters” on page 2083](#)). This section describes a few testing and troubleshooting techniques for virtual clustering.

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for different VDOMs will be distributed among both FortiGate units in the virtual cluster. These steps assume the cluster is otherwise operating correctly.

1 Log into the web-based manager or CLI using the IP addresses of interfaces in each VDOM.

Confirm that you have logged into the FortiGate unit that should be processing traffic for that VDOM by checking the HTML title displayed by your web browser or the CLI prompt. Both of these should include the host name of the cluster unit that you have logged into. Also on the system Dashboard, the System Information widget displays the serial number of the FortiGate unit that you logged into. From the CLI the `get system status` command displays the status of the cluster unit that you logged into.

- 2 To verify that the correct cluster unit is processing traffic for a VDOM:
 - Add security policies to the VDOM that allow communication between the interfaces in the VDOM.
 - Optionally enable traffic logging and other monitoring for that VDOM and these security policies.
 - Start communication sessions that pass traffic through the VDOM.
 - Log into the web-based manager and go to *Config > System > HA* and select *View HA Statistics*. Verify that the statistics display shows more active sessions, total packets, network utilization, and total bytes for the unit that should be processing all traffic for the VDOM.
 - Optionally check traffic logging and the Top Sessions Widget for the FortiGate unit that should be processing traffic for that VDOM to verify that the traffic is being processed by this FortiGate unit.



Configuring and operating FortiGate full mesh HA

This chapter provides an introduction to full mesh HA and also contains general procedures and configuration examples that describe how to configure FortiGate full mesh HA.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

This chapter contains the following sections:

- [Full mesh HA overview](#)
- [Example: full mesh HA configuration](#)
- [Troubleshooting full mesh HA](#)

Full mesh HA overview

When two or more FortiGate units are connected to a network in an HA cluster the reliability of the network is improved because the HA cluster replaces a single FortiGate unit as a single point of failure. With a cluster, a single FortiGate unit is replaced by a cluster of two or more FortiGate units.


However, even with a cluster, potential single points of failure remain. The interfaces of each cluster unit connect to a single switch and that switch provides a single connection to the network. If the switch fails or if the connection between the switch and the network fails service is interrupted to that network.

The HA cluster does improve the reliability of the network because switches are not as complex components as FortiGate units, so are less likely to fail. However, for even greater reliability, a configuration is required that includes redundant connections between the cluster the networks that it is connected to.

FortiGate models that support 802.3ad Aggregate or Redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster.

This redundant configuration can be achieved using FortiGate 802.3ad Aggregate or Redundant interfaces and a full mesh HA configuration. In a full mesh HA configuration, you connect an HA cluster consisting of two or more FortiGate units to the network using 802.3ad Aggregate or Redundant interfaces and redundant switches. Each 802.3ad Aggregate or Redundant interface is connected to two switches and both of these switches are connected to the network.

The resulting full mesh configuration, an example is shown in [Figure 207](#), includes redundant connections between all network components. If any single component or any single connection fails, traffic automatically switches to the redundant component and connection and traffic flow resumes.

Figure 207: Single points of failure in a standalone and HA network configuration

Full mesh HA and redundant heartbeat interfaces

A full mesh HA configuration also includes redundant HA heartbeat interfaces. At least two heartbeat interfaces should be selected in the HA configuration and both sets of HA heartbeat interfaces should be connected. The HA heartbeat interfaces do not have to be configured as redundant interfaces because the FGCP handles failover between heartbeat interfaces.

Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces


Full mesh HA is supported for both redundant interfaces and 802.3ad aggregate interfaces. In most cases you would simply use redundant interfaces. However, if your switches support 802.3ad aggregate interfaces and split multi-trunking you can use aggregate interfaces in place of redundant interfaces for full mesh HA. One advantage of using aggregate interfaces is that all of the physical interfaces in the aggregate interface can send and receive packets. As a result, using aggregate interfaces may increase the bandwidth capacity of the cluster.

Usually redundant and aggregate interfaces consist of two physical interfaces. However, you can add more than two physical interfaces to a redundant or aggregate interface. Adding more interfaces can increase redundancy protection. Adding more interfaces can also increase bandwidth capacity if you are using 802.3ad aggregate interfaces.

Example: full mesh HA configuration

Figure 207 shows a full mesh HA configuration with a cluster of two FortiGate-620b units. This section describes the FortiGate configuration settings and network components required for a full mesh HA configuration. This section also contains example steps for setting up this full mesh HA configuration. The procedures in this section describe one of many possible sequences of steps for configuring full mesh HA. As you become more experienced with FortiOS, HA, and full mesh HA you may choose to use a different sequence of configuration steps.

Figure 208: Full Mesh HA configuration



For simplicity these procedures assume that you are starting with two new FortiGate units set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

These procedures describe how to configure a cluster operating in NAT/Route mode because NAT/Route is the default FortiGate operating mode. However, the steps are the same if the cluster operates in Transparent mode. You can either switch the cluster units to operate in Transparent mode before beginning these procedures, or you can switch the cluster to operate in Transparent mode after HA is configured and the cluster is connected and operating.

FortiGate-620B full mesh HA configuration

The two FortiGate-620B units (620_ha_1 and 620_ha_2) can be operating in NAT/Route or Transparent mode. Aside from the standard HA settings, the FortiGate-620B configuration includes the following:

- The port5 and port6 interfaces configured as heartbeat interfaces. A full mesh HA configuration also includes redundant HA heartbeat interfaces.
- The port1 and port2 interfaces added to a redundant interface. Port1 is the active physical interface in this redundant interface. To make the port1 interface the active physical interface it should appear above the port2 interface in the redundant interface configuration.
- The port3 and port4 interfaces added to a redundant interface. Port3 is the active physical interface in this redundant interface. To make the port3 interface the active physical interface it should appear above the port4 interface in the redundant interface configuration.

Full mesh switch configuration

The following redundant switch configuration is required:

- Two redundant switches (Sw3 and Sw4) connected to the internal network. Establish an interswitch-link (ISL) between them.
- Two redundant switches (Sw1 and Sw2) connected to the Internet. Establish an interswitch-link (ISL) between them.

Full mesh network connections

Make the following physical network connections for 620_ha_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

Make the following physical network connections for 620_ha_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

How packets travel from the internal network through the full mesh cluster and to the Internet

If the cluster is operating in active-passive mode and 620_ha_2 is the primary unit, all packets take the following path from the internal network to the internet:

- 1 From the internal network to Sw4. Sw4 is the active connection to 620_ha_2; which is the primary unit. The primary unit receives all packets.

- 2 From Sw4 to the 620_ha_2 port3 interface. Active connection between Sw4 and 620_ha_2. Port3 is the active member of the redundant interface.
- 3 From 620_ha_2 port3 to 620_ha_2 port1. Active connection between 620_ha_2 and Sw2. Port1 is the active member of the redundant interface.
- 4 From Sw2 to the external router and the Internet.

Configuring FortiGate-620B units for HA operation - web-based manager

Each FortiGate-620B unit in the cluster must have the same HA configuration.

To configure the FortiGate-620B units for HA operation

- 1 Connect to the web-based manager of one of the FortiGate-620B units.
- 2 On the *System Information* dashboard widget, beside *Host Name* select *Change*.
- 3 Enter a new Host Name for this FortiGate unit.

New Name	620_ha_1
-----------------	----------

- 4 Go to *System > Config > HA* and change the following settings.

Mode	Active-Active	
Group Name	Rexample1.com	
Password	RHA_pass_1	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

5 Select OK.

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see “[Cluster virtual MAC addresses](#)” on page 2177). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

6 Power off the first FortiGate unit.

- 7 Repeat these steps for the second FortiGate unit.

Set the second FortiGate unit host name to:

New Name	620_ha_2
-----------------	----------

To connect the cluster to your network

- 1 Make the following physical network connections for 620_ha_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

- 2 Make the following physical network connections for 620_ha_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

- 3 Connect Sw3 and Sw4 to the internal network.

- 4 Connect Sw1 and Sw2 to the external router.

- 5 Enable ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.

- 6 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

- 1 View the system dashboard.

The System Information dashboard widget shows the *Cluster Name* (Rexample1.com) and the host names and serial numbers of the *Cluster Members*. The Unit Operation widget shows multiple cluster units.

- 2 Go to *System > Config > HA* to view the cluster members list.

The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

- 1 Log into the cluster web-based manager.
- 2 Go to *System > Admin > Administrators*.
- 3 For *admin*, select the *Change Password* icon

- 4 Enter and confirm a new password.
- 5 Select OK.
- 6 Go to *Router > Static* and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
- 7 Go to *System > Network > Interface* and select *Create New* and configure the redundant interface to connect to the Internet.

Name	Port1_Port2
Type	Redundant
Physical Interface Members	
Selected Interfaces	port1, port2
IP/Netmask	172.20.120.141/24

- 8 Select OK.
- 9 Select *Create New* and configure the redundant interface to connect to the internal network.

Name	Port3_Port4
Type	Redundant
Physical Interface Members	
Selected Interfaces	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

10 Select OK.

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Notice that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

11 Go to *Router > Static*.**12** Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

13 Select OK.**To configure HA port monitoring for the redundant interfaces**

- 1** Go to *System > Config > HA*.
- 2** In the cluster members list, edit the primary unit.
- 3** Configure the following port monitoring for the redundant interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4 Select OK.

Configuring FortiGate-620B units for HA operation - CLI

Each FortiGate-620B unit in the cluster must have the same HA configuration. Use the following procedure to configure the FortiGate-620B units for HA operation.

To configure the FortiGate-620B units for HA operation

- 1 Connect to the CLI of one of the FortiGate-620B units.
- 2 Enter a new Host Name for this FortiGate unit.

```
config system global
  set hostname 620_ha_1
end
```

- 3 Configure HA settings.

```
config system ha
  set mode a-a
  set group-name Rexample1.com
  set password RHA_pass_1
  set hbdev port5 50 port6 50
```

end

The FortiGate unit negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate unit as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate unit interfaces (see [“Cluster virtual MAC addresses” on page 2177](#)). The MAC addresses of the FortiGate-620B interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate unit (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate unit interface. For example, use the following command to view the port1 interface virtual MAC address (Current_HWaddr) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4 Power off the first FortiGate unit.

- 5 Repeat these steps for the second FortiGate unit.

Set the other FortiGate unit host name to:

```
config system global
    set hostname 620_ha_2
end
```

To connect the cluster to your network

- 1 Make the following physical network connections for 620_ha_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

- 2 Make the following physical network connections for 620_ha_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

- 3 Connect Sw3 and Sw4 to the internal network.

- 4 Connect Sw1 and Sw2 to the external router.

- 5 Enable ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.

- 6 Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

- 1 Log into the CLI.

- 2 Enter `get system status` to verify the HA status of the cluster unit that you logged into.

If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit.

If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit.

If the command output includes `Current HA mode: standalone` the cluster unit is not operating in HA mode.

- 3 Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
Model: 620
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
Master:128 620_ha_2          FG600B3908600825 0
Slave :128 620_ha_1          FG600B3908600705 1
number of vcluster: 1
```



```
vcluster 1: work 169.254.0.1
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

The command output shows both cluster units, their host names, their roles in the cluster, and their priorities. You can use this command to confirm that the cluster is operating normally. For example, if the command shows only one cluster unit then the other unit has left the cluster for some reason.

- 4 Use the `execute ha manage` command to connect to the other cluster unit's CLI and use these commands to verify cluster status.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGate units are not functioning as a cluster. See [“Troubleshooting HA clusters” on page 2083](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings. Some steps use the CLI and some the web-based manager.

- 1 Log into the cluster CLI.
- 2 Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <password_str>
  end
```

- 3 Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.

```
config router static
  delete 1
end
```

- 4 Go to *System > Network > Interface* and select *Create New* to add the redundant interface to connect to the Internet.

- 5 Add the redundant interface to connect to the Internet.

```
config sysetem interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
  end
```

- 6 Add the redundant interface to connect to the internal network.

```
config sysetem interface
  edit Port3_Port4
    set type redundant
    set member port3 port4
```

```
end
```

The virtual MAC addresses of the FortiGate-620B interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

7 Go to *Router* > *Static*.

8 Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the redundant interfaces

1 Enter the following command to configure port monitoring for the redundant interfaces:

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Troubleshooting full mesh HA

Troubleshooting full mesh HA clusters is similar to troubleshooting any cluster (see [“Troubleshooting HA clusters” on page 2083](#) or [“Troubleshooting virtual clustering” on page 2109](#)). The configuration and operation of a full mesh HA cluster is very similar to the configuration and operation of a standard cluster. The only differences relate to the configuration, connection, and operation of the redundant interfaces and redundant switches.

- Make sure the redundant interfaces and switches are connected correctly. With so many connections it is possible to make mistakes or for cables to become disconnected.
- Confirm that the configuration of the cluster unit 802.3ad Aggregate or Redundant interfaces is correct according to the configuration procedures in this chapter.
- In some configurations with some switch hardware, MAC-learning delays on the inter-switch links on the surrounding topologies may occur. The delays occur if the gratuitous ARP packets sent by the cluster after a failover are delayed by the switches before being sent across the inter-switch link. If this happens the surrounding topologies may be delayed in recognizing the failover and will keep sending packets to the MAC address of the failed primary unit resulting in lost traffic. Resolving this problem may require changing the configuration of the switch or replacing them with switch hardware that does not delay the gratuitous ARP packets.



Operating a cluster

With some exceptions, you can operate a cluster in much the same way as you operate a standalone FortiGate unit. This chapter describes those exceptions and also the similarities involved in operating a cluster instead of a standalone FortiGate unit.

This chapter contains the following sections:

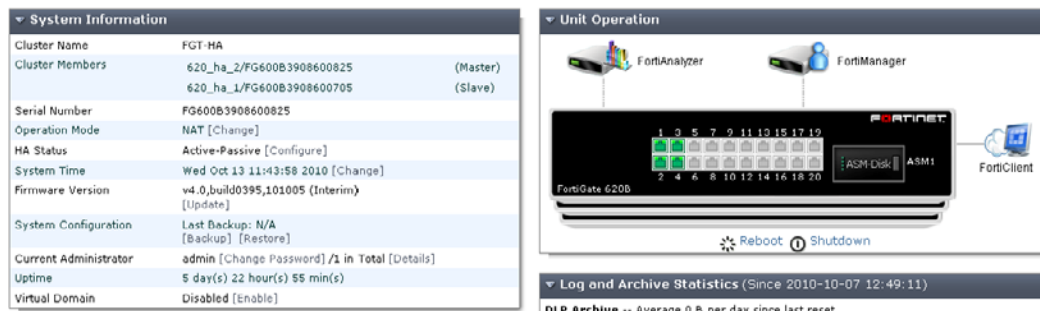
- [Operating a cluster](#)
- [Operating a virtual cluster](#)
- [Managing individual cluster units using a reserved management interface](#)
- [The primary unit acts as a router for subordinate unit management traffic](#)
- [Clusters and FortiGuard services](#)
- [Clusters and logging](#)
- [Clusters and SNMP](#)
- [Clusters and file quarantine](#)
- [Cluster members list](#)
- [Virtual cluster members list](#)
- [Viewing HA statistics](#)
- [Changing the HA configuration of an operating cluster](#)
- [Changing the HA configuration of an operating virtual cluster](#)
- [Changing the subordinate unit host name and device priority](#)
- [Upgrading cluster firmware](#)
- [Downgrading cluster firmware](#)
- [Backing up and restoring the cluster configuration](#)
- [Monitoring cluster units for failover](#)
- [Viewing cluster status from the CLI](#)
- [Disconnecting a cluster unit from a cluster](#)
- [Adding a disconnected FortiGate unit back to its cluster](#)

Operating a cluster

The configurations of all of the FortiGate units in a cluster are synchronized so that the cluster units can simulate a single FortiGate unit. Because of this synchronization, you manage the HA cluster instead of managing the individual cluster units. You manage the cluster by connecting to the web-based manager using any cluster interface configured for HTTPS or HTTP administrative access. You can also manage the cluster by connecting to the CLI using any cluster interface configured for SSH or telnet administrative access.

The cluster web-based manager dashboard displays the cluster name, the host name and serial number of each cluster member, and also shows the role of each unit in the cluster. The roles can be master (primary unit) and slave (subordinate units). The dashboard also displays a cluster unit front panel illustration.

Figure 209: Example cluster web-based manager dashboard



You can also go to *System > Config > HA* to view the cluster members list. This includes status information for each cluster unit. You can also use the cluster members list for a number of cluster management functions including changing the HA configuration of an operating cluster, changing the host name and device priority of a subordinate unit, and disconnecting a cluster unit from a cluster. See [“Cluster members list” on page 2147](#).

You can use log messages to view information about the status of the cluster. See [“Viewing and managing log messages for individual cluster units” on page 2136](#). You can use SNMP to manage the cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration information and receive traps.

You can configure a reserved management interface to manage individual cluster units. You can use this interface to access the web-based manager or CLI and to configure SNMP management for individual cluster units. See [“Managing individual cluster units using a reserved management interface” on page 2129](#).


You can manage individual cluster units by using SSH, telnet, or the CLI console on the web-based manager dashboard to connect to the CLI of the cluster. From the CLI you can use the `execute ha manage` command to connect to the CLI of any unit in the cluster.

You can also manage individual cluster units by using a null-modem cable to connect to any cluster unit CLI. From there you can use the `execute ha manage` command to connect to the CLI of each unit in the cluster.

Operating a virtual cluster

Managing a virtual cluster is very similar to managing a cluster that does not contain multiple virtual domains. Most of the information in this chapter applies to managing both kinds of clusters. This section describes what is different when managing a virtual cluster.

If virtual domains are enabled, the cluster web-based manager dashboard displays the cluster name and the role of each cluster unit in virtual cluster 1 and virtual cluster 2.

Figure 210: Example virtual clustering web-based manager dashboard

The configuration and maintenance options that you have when you connect to a virtual cluster web-based manager or CLI depend on the virtual domain that you connect to and the administrator account that you use to connect.

If you connect to a cluster as the administrator of a virtual domain, you connect directly to the virtual domain. Since HA virtual clustering is a global configuration, virtual domain administrators cannot see HA configuration options. However, virtual domain administrators see the host name of the cluster unit that they are connecting to on the web browser title bar or CLI prompt. This host name is the host name of the primary unit for the virtual domain. Also, when viewing log messages by going to *Log & Report > Log Access* virtual domain administrator can select to view log messages for either of the cluster units.

If you connect to a virtual cluster as the admin administrator you connect to the global web-based manager or CLI. Even so, you are connecting to an interface and to the virtual domain that the interface has been added to. The virtual domain that you connect to does not make a difference for most configuration and maintenance operations. However, there are a few exceptions. You connect to the FortiGate unit that functions as the primary unit for the virtual domain. So the host name displayed on the web browser title bar and on the CLI is the host name of this primary unit.

Managing individual cluster units using a reserved management interface

You can provide direct management access to all cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. Configuration changes to the reserved management interface are not synchronized to other cluster units.

The reserved management interface provides direct management access to each cluster unit and gives each cluster unit a different identity on your network. This simplifies using external services, such as SNMP, to separately monitor and manage each cluster unit.



The reserved management interface is not assigned an HA virtual MAC address like other cluster interfaces. Instead the reserved management interface retains the permanent hardware address of the physical interface unless you change it using the `config system interface` command.

The reserved management interface and IP address should not be used for managing a cluster using FortiManager. To correctly manage a FortiGate HA cluster with FortiManager use the IP address of one of the cluster unit interfaces. If you use a reserved management interface, FortiManager will assume it is connected to a single FortiGate unit instead of a cluster.

If you enable SNMP administrative access for the reserved management interface you can use SNMP to monitor each cluster unit using the reserved management interface IP address. To monitor each cluster unit using SNMP, just add the IP address of each cluster unit's reserved management interface to the SNMP server configuration. You must also enable direct management of cluster members in the cluster SNMP configuration.

If you enable HTTPS or HTTP administrative access for the reserved management interfaces you can connect to the web-based manager of each cluster unit. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. From the subordinate units the web-based manager has the same features as the primary unit except that unit-specific information is displayed for the subordinate unit, for example:


- The *Dashboard System Information* widget displays the subordinate unit serial number but also displays the same information about the cluster as the primary unit
- On the Cluster members list (go to *System > Config > HA*) you can change the HA configuration of the subordinate unit that you are logged into. For the primary unit and other subordinate units you can change only the host name and device priority.
- Log Access displays the logs of the subordinate that you are logged into first. You use the HA Cluster list to view the log messages of other cluster units including the primary unit.

If you enable SSH or TELNET administrative access for the reserved management interfaces you can connect to the CLI of each cluster unit. The CLI prompt contains the host name of the cluster unit that you have connected to. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. You can also use the `execute ha manage` command to connect to other cluster unit CLIs.

The reserved management interface is available in NAT/Route and in Transparent mode. It is also available if the cluster is operating with multiple VDOMs. In Transparent mode you cannot normally add an IP address to an interface. However, you can add an IP address to the reserved management interface.

Configuring the reserved management interface and SNMP remote management of individual cluster units

This example describes how to configure SNMP remote management of individual cluster units using the HA reserved management interface. The configuration consists of two FortiGate-620B units already operating as a cluster. In the example, the port8 interface of each cluster unit is connected to the internal network using the switch and configured as the reserved management interface.

Figure 211: SNMP remote management of individual cluster units**To configure the reserved management interface - web-based manager**

- 1 Go to *System > Config > HA*.
- 2 Edit the primary unit.
- 3 Select *Reserve Management Port for Cluster Member* and select port8.
- 4 Select OK.

To configure the reserved management interface - CLI

From the CLI you can also configure a default route that is only used by the reserved management interface.

- 1 Log into the CLI of any cluster unit.
- 2 Enter the following command to enable the reserved management interface, set port8 as the reserved interface, and add a default route of 10.11.101.100 for the reserved management interface.

```
config system ha
    set ha-mgmt-status enable
    set ha-mgmt-interface port8
    set ha-mgmt-interface-gateway 10.11.101.100
end
```

The reserved management interface default route is not synchronized to other cluster units.

To change the primary unit reserved management interface configuration - web-based manager

You can change the IP address of the primary unit reserved management interface from the primary unit web-based manager. Configuration changes to the reserved management interface are not synchronized to other cluster units.

- 1 From a PC on the internal network, browse to <http://10.11.101.100> and log into the cluster web-based manager.

This logs you into the primary unit web-based manager.

You can identify the primary unit from its serial number or host name that appears on the System Information dashboard widget.

- 2 Go to *System > Network > Interface* and edit the port8 interface as follows:

Alias	primary_reserved
IP/Netmask	10.11.101.101/24
Administrative Access	Ping, SSH, HTTPS, SNMP

- 3 Select OK.

You can now log into the primary unit web-based manager by browsing to <https://10.11.101.101>. You can also log into this primary unit CLI by using an SSH client to connect to 10.11.101.101.

To change subordinate unit reserved management interface configuration - CLI

At this point you cannot connect to the subordinate unit reserved management interface because it does not have an IP address. Instead, this procedure describes connecting to the primary unit CLI and using the `execute ha manage` command to connect to subordinate unit CLI to change the port8 interface. You can also use a serial connection to the cluster unit CLI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

- 1 Connect to the primary unit CLI and use the `execute ha manage` command to connect to a subordinate unit CLI.

You can identify the subordinate unit from its serial number or host name. The host name appears in the CLI prompt.

- 2 Enter the following command to change the port8 IP address to 10.11.101.102 and set management access to HTTPS, ping, SSH, and SNMP.

```
config system interface
  edit port8
    set ip 10.11.101.102/24
    set allowaccess https ping ssh snmp
end
```

You can now log into the subordinate unit web-based manager by browsing to <https://10.11.101.102>. You can also log into this subordinate unit CLI by using an SSH client to connect to 10.11.101.102.

To configure the cluster for SNMP management using the reserved management interfaces - CLI

This procedure describes how to configure the cluster to allow the SNMP server to get status information from the primary unit and the subordinate unit. The SNMP configuration is synchronized to all cluster units. To support using the reserved management interfaces, you must add at least one HA direct management host to an SNMP community. If your SNMP configuration includes SNMP users with user names and passwords you must also enable HA direct management for SNMP users.

- 1 Enter the following command to add an SNMP community called `Community` and add a host to the community for the reserved management interface of each cluster unit. The host includes the IP address of the SNMP server (10.11.101.20).

```
config system snmp community
  edit 1
    set name Community
    config hosts
      edit 1
        set ha-direct enable
        set ip 10.11.101.20
      end
    end
```

- 2 Enter the following command to add an SNMP user for the reserved management interface.

```
config system snmp user
  edit 1
    set ha-direct enable
    set notify-hosts 10.11.101.20
  end
```

Configure other settings as required.

To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses

From the command line of an SNMP manager, you can use the following SNMP commands to get CPU, memory and network usage information for each cluster unit. In the examples, the community name is `Community`. The commands use the MIB field names and OIDs listed in [Table 117 on page 2144](#).

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.101 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.101
1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.102 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.102
1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

The primary unit acts as a router for subordinate unit management traffic

HA uses routing and inter-VDOM links to route subordinate unit management traffic through the primary unit to the network. Similar to a standalone FortiGate unit, subordinate units may generate their own management traffic, including:


- DNS queries.
- FortiGuard Web Filtering rating requests.
- Log messages to be sent to a FortiAnalyzer unit, to a syslog server, or to the FortiGuard Analysis and Management Service.
- Log file uploads to a FortiAnalyzer unit.
- Quarantine file uploads to a FortiAnalyzer unit.
- SNMP traps.
- Communication with remote authentication servers (RADIUS, LDAP, TACACS+ and so on)

Subordinate units send this management traffic over the HA heartbeat link to the primary unit. The primary unit forwards the management traffic to its destination. The primary unit also routes replies back to the subordinate unit in the same way.

HA uses a hidden VDOM called `vsys_ha` for HA operations. The `vsys_ha` VDOM includes the HA heartbeat interfaces, and all communication over the HA heartbeat link goes through the `vsys_ha` VDOM. To provide communication from a subordinate unit to the network, HA adds hidden inter-VDOM links between the primary unit management VDOM and the primary unit `vsys_ha` VDOM. By default, root is the management VDOM.

Management traffic from the subordinate unit originates in the subordinate unit `vsys_ha` VDOM. The `vsys_ha` VDOM routes the management traffic over the HA heartbeat link to the primary unit `vsys_ha` VDOM. This management traffic is then routed to the primary unit management VDOM and from there out onto the network.

DNS queries and FortiGuard Web Filtering and Email Filter requests are still handled by the HA proxy so the primary unit and subordinate units share the same DNS query cache and the same FortiGuard Web Filtering and Email Filter cache. In a virtual clustering configuration, the cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering, Email Filtering, and DNS query cache.

Figure 212: Subordinate unit management traffic path

Cluster communication with RADIUS and LDAP servers

In an active-passive cluster, only the primary unit processes traffic, so the primary unit communicates with RADIUS or LDAP servers. In a cluster that is operating in active-active mode, subordinate units send RADIUS and LDAP requests to the primary unit over the HA heartbeat link and the primary unit routes them to their destination. The primary unit relays the responses back to the subordinate unit.

Clusters and FortiGuard services

This section describes how various FortiGate HA clustering configurations communicate with the FDN.

In an operating cluster, the primary unit communicates directly with the FortiGuard Distribution Network (FDN). Subordinate units also communicate directly with the FDN but as described in [“The primary unit acts as a router for subordinate unit management traffic” on page 2134](#), all communication between subordinate units and the FDN is routed through the primary unit.

You must register and licence all of the units in a cluster for all required FortiGuard services, both because all cluster units communicate with the FDN and because any cluster unit could potentially become the primary unit.

FortiGuard and active-passive clusters

For an active-passive cluster, only the primary unit processes traffic. Even so, all cluster units communicate with the FDN. Only the primary unit sends FortiGuard Web Filtering and Antispam requests to the FDN. All cluster units receive FortiGuard Antivirus, IPS, and application control updates from the FDN.

In an active-passive cluster the FortiGuard Web Filter and Email Filter caches are located on the primary unit in the same way as for a standalone FortiGate unit. The caches are not shared among cluster units so after a failover the new primary unit must build up new caches.

In an active-passive cluster all cluster units also communicate with the FortiGuard Analysis and Management Service (FAMS).

FortiGuard and active-active clusters

For an active-active cluster, both the primary unit and the subordinate units process traffic. Communication between the cluster units and the FDN is the same as for active-passive clusters with the following exception.

Because the subordinate units process traffic, they may also be making FortiGuard Web Filtering and Email Filter requests. The primary unit receives all such requests from the subordinate units and relays them to the FDN and then relays the FDN responses back to the subordinate units. The FortiGuard Web Filtering and Email Filtering URL caches are maintained on the primary unit. The primary unit caches are used for primary and subordinate unit requests.

FortiGuard and virtual clustering

For a virtual clustering configuration the management virtual domain of each cluster unit communicates with the FDN. The cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering and Email Filtering caches. All FortiGuard Web Filtering and Email Filtering requests are proxied by the management VDOM of the cluster unit that is the primary unit for the management virtual domain.

Clusters and logging

This section describes the log messages that provide information about how HA is functioning, how to view and manage logs for each unit in a cluster, and provides some example log messages that are recorded during specific cluster events.

You configure logging for a cluster in the same way as you configuring logging for a standalone FortiGate unit. Log configuration changes made to the cluster are synchronized to all cluster units.

All cluster units record log messages separately to the individual cluster unit's log disk, to the cluster unit's system memory, or both. You can view and manage log messages for each cluster unit from the cluster web-based manager Log Access page.

When remote logging is configured, all cluster units send log messages to remote FortiAnalyzer units or other remote servers as configured. HA uses routing and inter-VDOM links to route subordinate unit log traffic through the primary unit to the network. See [“The primary unit acts as a router for subordinate unit management traffic” on page 2134](#).

When you configure a FortiAnalyzer unit to receive log messages from a FortiGate cluster, you should add a cluster to the FortiAnalyzer unit configuration so that the FortiAnalyzer unit can receive log messages from all cluster units.

Viewing and managing log messages for individual cluster units

This section describes how to view and manage log messages for an individual cluster unit.

To view HA cluster log messages

- 1 Log into the cluster web-based manager.
- 2 Go to *Log&Report > Log Access* and select Memory or Disk.
For each log display, the *HA Cluster* list displays the serial number of the cluster unit for which log messages are displayed. The serial numbers are displayed in order in the list.
- 3 Set *HA Cluster* to the serial number of one of the cluster units to display log messages for that unit.

You can view logs saved to memory or logs saved to the hard disk for the cluster unit.

About HA event log messages

HA event log messages always include the host name and serial number of the cluster unit that recorded the message. HA event log messages also include the HA state of the unit and also indicate when a cluster unit switches (or moves) from one HA state to another. Cluster units can operate in the HA states listed in [Table 116](#):

Table 116: HA states

Hello	A FortiGate unit configured for HA operation has started up and is looking for other FortiGate units with which to form a cluster.
Work	In an active-passive cluster a cluster unit is operating as the primary unit. In an active-active cluster unit is operating as the primary unit or a subordinate unit.
Standby	In an active-passive cluster the cluster unit is operating as a subordinate unit.

HA log Event log messages also indicate the virtual cluster that the cluster unit is operating in as well as the member number of the unit in the cluster. If virtual domains are not enabled, all clusters unit are always operating in virtual cluster 1. If virtual domains are enabled, a cluster unit may be operating in virtual cluster 1 or virtual cluster 2. The member number indicates the position of the cluster unit in the cluster members list. Member 0 is the primary unit. Member 1 is the first subordinate unit, member 2 is the second subordinate unit, and so on.

The following log message indicates that the cluster unit with host name 5005_ha_2 and serial number FG5A253E06500088 has become the primary unit because it is operating in the work state as member 0.

```
2010-01-13 13:45:32 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=work vcluster_member=0 hostname=5005_ha_2
sn=FG5A253E06500088
```

The following log message indicates that the cluster unit with host name 5005_ha_1 and serial number FG5A253E06500088 has become the first subordinate unit in an active-passive cluster because it is operating in the standby state as member 1.

```
2010-01-13 14:28:39 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=standby vcluster_member=1 hostname=5005_ha_2
sn=FG5A253E06500088
```

The following log message indicates that the cluster unit with host name 5005_ha_1 and serial number FG5A253E07600124 has become the first subordinate unit in an active-active cluster because it is operating in the work state as member 1.

```
2010-01-13 14:23:58 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=work vcluster_member=1 hostname=5005_ha_1
sn=FG5A253E07600124
```

The following log message indicates that the FortiGate unit was disconnected from a cluster. The message shows that the cluster unit was disconnected over the telnet link between cluster units and the HA mode was changed from active-active to standalone.

```
2010-01-13 13:45:09 log_id=0104032140 type=event subtype=admin vd=root
pri=notice user="FGT_ha_admin" ui=telnet(169.254.0.2) old=A-A new=standalone
msg="User FGT_ha_admin changed HA mode from A-A to standalone"
```


HA log messages

See the [FortiGate Log Message Reference](#) for a listing of and descriptions of the HA log messages.

Example log messages

This section displays some log message sequences when specific cluster events occur.

Unit changing to HA mode and becoming the primary unit

- 1 2010-01-14 13:24:56 log_id=0104032140 type=event subtype=admin vd=root pri=notice user="admin" ui=GUI(10.21.101.100) old=standalone new=A-P msg="User admin changed HA mode from standalone to A-P"
The administrator changed the HA mode to active-passive.
- 2 2010-01-14 13:25:09 log_id=0105037899 type=event subtype=ha pri=notice vd="root" msg="HA device(interface) peerinfo" ha_role=slave devintfname=port4
The the cluster unit received heartbeat packets at the port4 interface.
- 3 2010-01-14 13:25:09 log_id=0105037894 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster detected member join" vcluster=1 ha_group=0
The the cluster unit detected another unit that it could form a cluster with.
- 4 2010-01-14 13:25:11 log_id=0105037892 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster's member state moved" vcluster=1 vcluster_state=work vcluster_member=0 hostname=5005_ha_1 sn=FG5A253E07600124
The the cluster unit's state changed to work, meaning that the cluster unit became the primary unit.

Unit changing to HA mode and becoming a subordinate unit

- 1 2010-01-13 14:57:04 log_id=0104032140 type=event subtype=admin vd=root pri=notice user="admin" ui=GUI(10.21.101.100) old=standalone new=A-P msg="User admin changed HA mode from standalone to A-P"
The administrator changed the HA mode to active-passive.
- 2 2010-01-13 14:57:07 log_id=0105037899 type=event subtype=ha pri=notice vd="root" msg="HA device(interface) peerinfo" ha_role=slave devintfname=port4
The the cluster unit received heartbeat packets at the port4 interface.
- 3 2010-01-13 14:57:07 log_id=0105037894 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster detected member join" vcluster=1 ha_group=0
The the cluster unit detected another unit that it could form a cluster with.
- 4 2010-01-13 14:57:09 log_id=0105037892 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster's member state moved" vcluster=1 vcluster_state=standby vcluster_member=1 hostname=5005_ha_1 sn=FG5A253E07600124
The the cluster unit's state changed to standby, meaning that the cluster unit became a subordinate unit.

Primary unit fails (and is removed from cluster)

These messages are recorded by a subordinate unit (which becomes the primary unit).

- 1 2010-01-13 14:49:44 log_id=0105037901 type=event subtype=ha pri=critical
vd="root" msg="Heartbeat device(interface) down" ha_role=slave
hbdn_reason=neighbor info lost devintfname=port4

The subordinate unit loses communication with the primary unit.

- 2 2010-01-13 14:50:00 log_id=0105037893 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster detected member dead" vcluster=1 ha_group=0

The subordinate unit determines that the primary unit is no longer operating.

- 3 2010-01-13 14:50:02 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=work vcluster_member=0 hostname=5005_ha_2
sn=FG5A253E06500088

The subordinate unit negotiates to form a cluster and then begins to operate as the primary unit.

Subordinate unit fails (and is removed from cluster)

These messages are recorded by the primary unit.

- 1 2010-01-14 13:01:03 log_id=0105037901 type=event subtype=ha pri=critical
vd="root" msg="Heartbeat device(interface) down" ha_role=master
hbdn_reason=neighbor info lost devintfname=port4

The primary unit loses contact with the subordinate unit.

- 2 2010-01-14 13:01:19 log_id=0105037893 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster detected member dead" vcluster=1 ha_group=0

The primary unit determines that the subordinate unit is no longer operating.

- 3 2010-01-14 13:01:21 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=work vcluster_member=0 hostname=5005_ha_1
sn=FG5A253E07600124

The primary unit negotiates to form a cluster and then continues to operate as the primary unit.

New unit added to cluster

These messages are recorded by the primary unit.

- 1 2010-01-13 14:57:07 log_id=0105037899 type=event subtype=ha pri=notice
vd="root" msg="HA device(interface) peerinfo" ha_role=master devintfname=port4

The cluster unit received heartbeat packets at the port4 interface.

- 2 2010-01-13 14:57:07 log_id=0105037894 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster detected member join" vcluster=1 ha_group=0

The cluster unit detected another unit that it could form a cluster with.

- 3 2010-01-13 14:57:09 log_id=0105037892 type=event subtype=ha pri=notice
vd="root" msg="Virtual cluster's member state moved" vcluster=1
vcluster_state=work vcluster_member=0 hostname=5005_ha_2
sn=FG5A253E06500088

The primary unit negotiates to form a cluster with the new cluster unit and then continues to operate as the primary unit.

Unit removed from cluster

These log messages appear on the unit that was removed from the cluster.:

- 1 2010-01-13 14:49:37 log_id=0104032121 type=event subtype=admin vd=root pri=notice user="admin" ui=GUI(10.21.101.100) intf="port1" field=access old="https+ping+ssh" new="https+ping+ssh+snmp+http+telnet" msg="User admin changed the access setting of interface port1 from GUI(10.21.101.100)"
- 2 2010-01-13 14:49:37 log_id=0104032121 type=event subtype=admin vd=root pri=notice user="admin" ui=GUI(10.21.101.100) intf="port1" field=ip old=10.21.101.102:255.255.255.0 new=10.21.101.103:255.255.255.0 msg="User admin changed the ip setting of interface port1 from GUI(10.21.101.100)"
- 3 2010-01-13 14:49:37 log_id=0104032140 type=event subtype=admin vd=root pri=notice user="admin" ui=GUI(10.21.101.100) old=A-P new=standalone msg="User admin changed HA mode from A-P to standalone"

Link failure of a monitored interface

These log messages, recorded by the primary unit, show the monitored port1 interface failed or was disconnected and the primary unit becoming a subordinate unit:

- 1 2010-01-14 16:59:39 log_id=0100020099 type=event subtype=system vd=root pri=information action=interface-stat-change status=DOWN msg="Link monitor: Interface port1 was turned down"
- 2 2010-01-14 16:59:39 log_id=0105037898 type=event subtype=ha pri=warning vd="root" msg="HA device(interface) fail" ha_role=master devintfname=port1
- 3 2010-01-14 16:59:41 log_id=0105037892 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster's member state moved" vcluster=1 vcluster_state=standby vcluster_member=1 hostname=620_ha_1 sn=FG600B3908600705

Link failure of a monitored interface fixed

These log messages show the monitored port1 being reconnected and the cluster unit becoming the primary unit.

- 1 2010-01-14 16:59:58 log_id=0100020099 type=event subtype=system vd=root pri=information action=interface-stat-change status=UP msg="Link monitor: Interface port1 was turned up"
- 2 2010-01-14 16:59:58 log_id=0105037897 type=event subtype=ha pri=notice vd="root" msg="HA device(interface) ready" ha_role=slave devintfname=port1
- 3 2010-01-14 17:00:00 log_id=0105037892 type=event subtype=ha pri=notice vd="root" msg="Virtual cluster's member state moved" vcluster=1 vcluster_state=work vcluster_member=0 hostname=620_ha_1 sn=FG600B3908600705

Configuration changes synchronized from primary unit to subordinate unit

The following event log message is written by the primary unit when the admin administrator adds security policy with ID=3 by connecting to the web-based manager from a management PC with IP address 172.20.120.14 using HTTPS or HTTP:

```
2009-11-13 09:11:45 log_id=0104032126 type=event subtype=admin vd=root
pri=notice user="admin" ui=GUI(172.20.120.14) seq=1 sintf="external"
dintf="internal" saddr="all" daddr="all" act=accept nat=no iptype=ipv4
schd="always" svr="ANY" log=no idbased=no msg="User admin added IPv4 firewall
policy 3 from GUI(172.20.120.11)"
```

When incremental synchronization makes the same change to a subordinate unit the subordinate unit writes the following log message:

```
2009-11-13 09:11:45 log_id=0104032126 type=event subtype=admin vd=root
pri=notice user="admin" ui=ha_daemon seq=1 sintf="external" dintf="internal"
saddr="all" daddr="all" act=accept nat=no iptype=ipv4 schd="always" svr="ANY"
log=no idbased=no msg="User admin added IPv4 firewall policy 3 from
GUI(172.20.120.11)"
```

Notice that the two messages are identical (including the log IDs) except that on the subordinate unit the ui (user interface) is ha_daemon. ha_daemon is the name of the user interface used by the HA synchronization process to make incremental synchronization configuration changes.

Configuration changes synchronized from subordinate unit to primary unit

The following event log message is written by a subordinate unit after the admin administrator logs into the subordinate unit CLI using the `execute ha manage` command and adds security policy 6.

```
2009-11-13 09:14:45 log_id=0104032126 type=event subtype=admin vd=root
pri=notice user="admin" ui=telnet(169.254.0.1) seq=6 sintf="external"
dintf="internal" saddr="all" daddr="all" act=accept nat=no iptype=ipv4
schd="always" svr="ANY" log=no idbased=no msg="User admin added IPv4 firewall
policy 6 from telnet(169.254.0.1)"
```

Notice the user interface is telnet(169.254.0.1). 169.254.0.1 is the IP address of the HA heartbeat interface of the primary unit. The log message shows that the `execute ha manage` command sets up a telnet session from the primary unit to the subordinate unit over the HA heartbeat link. Note that the IP address could be 169.254.0.2 if the cluster renegotiated.

When incremental synchronization makes the same change to the primary unit, the primary unit writes the following log message:

```
2009-11-13 09:14:45 log_id=0104032126 type=event subtype=admin vd=root
pri=notice user="admin" ui=ha_daemon seq=6 sintf="external" dintf="internal"
saddr="all" daddr="all" act=accept nat=no iptype=ipv4 schd="always" svr="ANY"
log=no idbased=no msg="User admin added IPv4 firewall policy 6 from ha_daemon"
```

Notice again that the messages are identical (including the log ID) except for the user interface.

Fortigate HA message "HA master heartbeat interface <intf_name> lost neighbor information"

The following HA log messages may be recorded by an operating cluster:

```
2009-02-16 11:06:34 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=critical vd=root msg="HA slave heartbeat interface internal lost
neighbor information"

2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=notice vd=root msg="Virtual cluster 1 of group 0 detected new joined
HA member"

2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event
subtype=ha pri=notice vd=root msg="HA master heartbeat interface internal get peer
information"
```

These log messages indicate that the cluster units could not connect to each other over the HA heartbeat link for the period of time that is given by `hb-interval` x `hb-lost-threshold`, which is 1.2 seconds with the default values.

To diagnose this problem

- 1 Check all heartbeat interface connections including cables and switches to make sure they are connected and operating normally.

- 2 Use the following commands to display the status of the heartbeat interfaces.

```
get hardware nic <heartbeat_interface_name>
diagnose hardware deviceinfo nic <heartbeat_interface_name>
```

The status information may indicate the interface status and link status and also indicate if a large number of errors have been detected.

- 3 If the log message only appear during peak traffic times, increase the tolerance for missed HA heartbeat packets by using the following commands to increase the lost heartbeat threshold and heartbeat interval:

```
config system ha
    set hb-lost-threshold 12
    set hb-interval 4
end
```

These settings multiply by 4 the loss detection interval. You can use higher values as well.

- 4 Optionally disable session-pickup to reduce the processing load on the heartbeat interfaces.
- 5 Instead of disabling session-pickup you can enable `session-pickup-delay` to reduce the number of sessions that are synchronized. With this option enabled only sessions that are active for more than 30 seconds are synchronized.

It may be useful to monitor CPU and memory usage to check for low memory and high CPU usage. You can configure event logging to monitor CPU and memory usage. You can also enable the CPU over usage and memory low SNMP events.

Once this monitoring is in place, try and determine if there have been any changes in the network or an increase of traffic recently that could be the cause. Check to see if the problem happens frequently and if so what the pattern is.

To monitor the CPU of the cluster units and troubleshoot further, use the following procedure and commands:

```
get system performance status
get sys performance top 2
diagnose sys top 2
```

These commands repeated at frequent intervals will show the activity of the CPU and the number of sessions.

Search the [Fortinet Knowledge Base](https://docs.fortinet.com) for articles about monitoring CPU and Memory usage.

If the problem persists, gather the following information (a console connection might be necessary if connectivity is lost) and provide it to Technical Support when opening a ticket:

- Debug log from the web-based manager: *System > Maintenance > Advanced > debug log*
- CLI command output:
`diag sys top 2` (keep it running for 20 seconds)

```
get sys perf status (repeat this command multiple times to get good samples)
get sys ha status
diag sys ha status
diag sys ha dump all
diag sys ha dump 2
diag sys ha dump 3
diag netlink dev list
diag hardware dev nic <Heartbeat port Name>
execute log filter category event
execute log display
```

Clusters and SNMP

You can use SNMP to manage a cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration and status information and receive traps.

You configure SNMP for a cluster in the same way as configuring SNMP for a standalone FortiGate unit. SNMP configuration changes made to the cluster are shared by all cluster units.

Each cluster unit sends its own traps and SNMP manager systems can use SNMP get commands to query each cluster unit separately. To set SNMP get queries to each cluster unit you must create a special get command that includes the serial number of the cluster unit.

Alternatively you can use the HA reserved management interface feature to give each cluster unit a different management IP address. Then you can create an SNMP get command for each cluster unit that just includes the management IP address and does not have to include the serial number. See [“Managing individual cluster units using a reserved management interface” on page 2129](#).

For a list of HA MIB fields and OIDs, see [“Fortinet MIBs” on page 434](#).

SNMP get command syntax for the primary unit

Normally, to get configuration and status information for a standalone FortiGate unit or for a primary unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to the following:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> |
<MIB_field>}
```

where:

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is `public`.

<address_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself. The HA MIB fields and OIDs are listed in [Table 117](#).

Table 117: SNMP field names and OIDs

MIB field	OID	Description
fgHaSystemMode	.1.3.6.1.4.1.12356.101.13.1.1.0	HA mode (standalone, a-a, or a-p)
fgHaGroupId	.1.3.6.1.4.1.12356.101.13.1.2.0	The HA priority of the cluster unit. Default 128.
fgHaPriority	.1.3.6.1.4.1.12356.101.13.1.3.0	The HA priority of the cluster unit. Default 128.
fgHaOverride	.1.3.6.1.4.1.12356.101.13.1.4.0	Whether HA override is disabled or enabled for the cluster unit.
fgHaAutoSync	.1.3.6.1.4.1.12356.101.13.1.5.0	Whether automatic HA synchronization is disabled or enabled.
fgHaSchedule	.1.3.6.1.4.1.12356.101.13.1.6.0	The HA load balancing schedule. Set to none unless operating in a-p mode.
fgHaGroupName	.1.3.6.1.4.1.12356.101.13.1.7.0	The HA group name.
fgHaStatsIndex	.1.3.6.1.4.1.12356.101.13.2.1.1.1.1	The cluster index of the cluster unit. 1 for the primary unit, 2 to x for the subordinate units.
fgHaStatsSerial	.1.3.6.1.4.1.12356.101.13.2.1.1.2.1	The serial number of the cluster unit.
fgHaStatsCpuUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.3.1	The cluster unit's current CPU usage.
fgHaStatsMemUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.4.1	The cluster unit's current Memory usage.
fgHaStatsNetUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.5.1	The cluster unit's current Network bandwidth usage.
fgHaStatsSesCount	.1.3.6.1.4.1.12356.101.13.2.1.1.6.1	The cluster unit's current session count.
fgHaStatsPktCount	.1.3.6.1.4.1.12356.101.13.2.1.1.7.1	The cluster unit's current packet count.
fgHaStatsByteCount	.1.3.6.1.4.1.12356.101.13.2.1.1.8.1	The cluster unit's current byte count.
fgHaStatsIdsCount	.1.3.6.1.4.1.12356.101.13.2.1.1.9.1	The number of attacks reported by the IPS for the cluster unit.
fgHaStatsAvCount	.1.3.6.1.4.1.12356.101.13.2.1.1.10.1	The number of viruses reported by the antivirus system for the cluster unit.
fgHaStatsHostname	.1.3.6.1.4.1.12356.101.13.2.1.1.11.1	The hostname of the cluster unit.

To get the HA priority for the primary unit

The following SNMP get command gets the HA priority for the primary unit. The community name is `public`. The IP address of the cluster interface configured for SNMP management access is 10.10.10.1. The HA priority MIB field is `fgHaPriority` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.1.3.0. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgHaPriority
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.1.3.0
```

SNMP get command syntax for any cluster unit

To get configuration status information for a specific cluster unit (for the primary unit or for any subordinate unit), the SNMP manager must add the serial number of the cluster unit to the SNMP get command after the community name. The community name and the serial number are separated with a dash. The syntax for this SNMP get command would be:

```
snmpget -v2c -c <community_name>-<fgt_serial> <address_ipv4>
{<OID> | <MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. All units in the cluster have the same community name. The most commonly used community name is `public`.

`<fgt_serial>` is the serial number of any cluster unit. For example, FGT4002803033172. You can specify the serial number of any cluster unit, including the primary unit, to get information for that unit.

`<address_ipv4>` is the IP address of the FortiGate interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see [“Fortinet MIBs” on page 434](#).

If the serial number matches the serial number of a subordinate unit, the SNMP get request is sent over the HA heartbeat link to the subordinate unit. After processing the request, the subordinate unit sends the reply back over the HA heartbeat link back to the primary unit. The primary unit then forwards the response back to the SNMP manager.

If the serial number matches the serial number of the primary unit, the SNMP get request is processed by the primary unit. You can actually add a serial number to the community name of any SNMP get request. But normally you only need to do this for getting information from a subordinate unit.

To get the CPU usage for a subordinate unit

The following SNMP get command gets the CPU usage for a subordinate unit in a FortiGate-5001SX cluster. The subordinate unit has serial number FG50012205400050. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1. The HA status table MIB field is `fgHaStatsCpuUsage` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.3.1. The first command uses the MIB field name and the second uses the OID for this table:

```
snmpget -v2c -c public-FG50012205400050 10.10.10.1
fgHaStatsCpuUsage
snmpget -v2c -c public-FG50012205400050 10.10.10.1
1.3.6.1.4.1.12356.101.13.2.1.1.3.1
```


FortiGate SNMP recognizes the community name with syntax `<community_name>-<fgt_serial>`. When the primary unit receives an SNMP get request that includes the community name followed by serial number, the FGCP extracts the serial number from the request. Then the primary unit redirects the SNMP get request to the cluster unit with that serial number. If the serial number matches the serial number of the primary unit, the SNMP get is processed by the primary unit.

Getting serial numbers of cluster units

The following SNMP get commands use the MIB field name `fgHaStatsSerial.<index>` to get the serial number of each cluster unit. Where `<index>` is the cluster unit's cluster index and 1 is the cluster index of the primary unit, 2 is the cluster index of the first subordinate unit, and 3 is the cluster index of the second subordinate unit.

The OID for this MIB field is `1.3.6.1.4.1.12356.101.13.2.1.1.2.1`. The community name is `public`. The IP address of the FortiGate interface is `10.10.10.1`.

The first command uses the MIB field name and the second uses the OID for this table and gets the serial number of the primary unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.1
snmpget -v2c -c public 10.10.10.1
1.3.6.1.4.1.12356.101.13.2.1.1.2.1
```

The second command uses the MIB field name and the second uses the OID for this table and gets the serial number of the first subordinate unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.2
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.2.2
```

SNMP get command syntax - reserved management interface enabled

To get configuration and status information for any cluster unit where you have enabled the HA reserved management interface feature and assigned IP addresses to the management interface of each cluster unit, an SNMP manager would use the following get command syntax:

```
snmpget -v2c -c <community_name> <mgt_address_ipv4> {<OID> |
<MIB_field>}
```

where:

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community names to a FortiGate SNMP configuration. The most commonly used community name is `public`.

`<mgt_address_ipv4>` is the IP address of the FortiGate HA reserved management interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see your FortiGate unit's online help.

See [“To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses” on page 2133](#).

Clusters and file quarantine

You can configure file quarantine for a cluster in the same way as configuring file quarantine for a standalone FortiGate unit. Quarantine configuration changes made to the cluster are shared by all cluster units.

In an active-active cluster, both the primary unit and the subordinate units accept antivirus sessions and may quarantine files. In an active-passive cluster, only the primary unit quarantines files. Multiple cluster units in an active-passive cluster may have quarantined files if different cluster units have been the primary unit.

All cluster units quarantine files separately to their own hard disk. You can go to *Log&Report > Archive Access > Quarantine* to view and manage the quarantine file list for each cluster unit.

All cluster units can also quarantine files to a FortiAnalyzer unit. When you configure a FortiAnalyzer unit to receive quarantine files from a cluster, you should add each cluster unit to the FortiAnalyzer device configuration so that the FortiAnalyzer unit can receive quarantine files from all cluster units.

Cluster members list

Display the cluster members list to view the status of the FortiGate units in an operating cluster. To display the cluster members list, go to *System > Config > HA*.

From the cluster members list you can also:

- View HA statistics (see “Viewing HA statistics” on page 2150).
- View and optionally change the HA configuration of the operating cluster (see “Changing the HA configuration of an operating cluster” on page 2151).
- View and optionally change the host name and device priority of a subordinate unit (see “Changing the subordinate unit host name and device priority” on page 2152).
- Disconnect a cluster unit from a cluster (see “Disconnecting a cluster unit from a cluster” on page 2164).
- Download the Debug log for any cluster unit. You can send this debug log file to Fortinet Technical Support to help diagnose problems with the cluster or with individual cluster units.

Figure 213: Example cluster members list

Up and Down Arrows

Download Debug Log

Edit

Disconnect from Cluster

HA Cluster

Cluster Member

FortiGate 620B

FortiGate 620B

Cluster Member	Hostname	Role	Priority	
620_ha_2	MASTER	128		
620_ha_1	SLAVE	128		

View HA Statistics	Display the serial number, status, and monitor information for each cluster unit. See “Viewing HA statistics” on page 2150.
Up and down arrows	Change the order in which cluster members are listed. The operation of the cluster or of the units in the cluster are not affected. All that changes is the order in which cluster units are displayed on the cluster members list.

Cluster member	Illustrations of the front panels of the cluster units. If the network jack for an interface is shaded green, the interface is connected. Pause the mouse pointer over each illustration to view the cluster unit host name, serial number, and how long the unit has been operating (up time). ^a The list of monitored interfaces is also displayed.
Hostname	<p>The host name of the FortiGate unit. The default host name of the FortiGate unit is the FortiGate unit serial number.</p> <ul style="list-style-type: none"> To change the primary unit host name, go to the system dashboard and select Change beside the current host name in the System Information widget. To change a subordinate unit host name, from the cluster members list select the edit icon for a subordinate unit.
Role	<p>The status or role of the cluster unit in the cluster.</p> <ul style="list-style-type: none"> Role is MASTER for the primary (or master) unit Role is SLAVE for all subordinate (or backup) cluster units
Priority	<p>The device priority of the cluster unit. Each cluster unit can have a different device priority. During HA negotiation, the unit with the highest device priority becomes the primary unit.</p> <p>The device priority range is 0 to 255. The default device priority is 128.</p>
Disconnect from cluster	Disconnect the cluster unit from the cluster. See “Disconnecting a cluster unit from a cluster” on page 2164 .
Edit	<p>Select Edit to change a cluster unit HA configuration.</p> <ul style="list-style-type: none"> For a primary unit, select Edit to change the cluster HA configuration. You can also change the device priority of the primary unit. For a primary unit in a virtual cluster, select Edit to change the virtual cluster HA configuration. You can also change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit. For a subordinate unit, select Edit to change the subordinate unit host name and device priority. See “Changing the subordinate unit host name and device priority” on page 2152. For a subordinate unit in a virtual cluster, select Edit to change the subordinate unit host name. In addition you can change the device priority for the subordinate unit for the selected virtual cluster.
Download debug log	Download an encrypted debug log to a file. You can send this debug log file to Fortinet Technical Support to help diagnose problems with the cluster or with individual cluster units.

^aIn this case, with the mouse position showing a pop-up, it would be worth showing the mouse cursor itself in the screen capture.

Virtual cluster members list

If virtual domains are enabled, you can display the cluster members list to view the status of the operating virtual clusters. The virtual cluster members list shows the status of both virtual clusters including the virtual domains added to each virtual cluster.

To display the virtual cluster members list for an operating cluster log in as the admin administrator, select **Global Configuration** and go to **System > Config > HA**.

Figure 214: Example FortiGate-5001SX virtual cluster members list

The screenshot displays the FortiGate-5001SX virtual cluster members list. It shows two virtual clusters, Virtual Cluster 1 and Virtual Cluster 2. Each cluster has a table of members with columns for Hostname, Role, and Priority. Virtual Cluster 1 has members 620_ha_2 (MASTER, Priority 120) and 620_ha_1 (SLAVE, Priority 128). Virtual Cluster 2 has members 620_ha_2 (MASTER, Priority 128) and 620_ha_1 (SLAVE, Priority 128). The interface includes icons for 'Up and Down Arrows', 'Disconnect from Cluster', 'Edit', and 'Download Debug Log'. A tooltip for port14 is visible over the 620_ha_2 member of Virtual Cluster 1.

The fields and functions of the virtual cluster members list are the same as the fields and functions described in “[Cluster members list](#)” on page 2147 with the following exceptions.

- When you select the edit icon for a primary unit in a virtual cluster, you can change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit and you can edit the VDOM partitioning configuration of the cluster.
- When you select the edit icon for a subordinate unit in a virtual cluster, you can change the device priority for the subordinate unit for the selected virtual cluster.

Also, the HA cluster members list changes depending on the cluster unit. For the virtual cluster described in the “[Example: virtual clustering with two VDOMs and VDOM partitioning](#)” on page 2093 if you connect to port5 using you are connecting to 620b_ha_2 (620b_ha_2 is displayed on the web browser title bar or in the CLI prompt).

If you connect to port1 you are connecting to 620b_ha_1 (620b_ha_2 is displayed on the web browser title bar or in the CLI prompt).

Viewing HA statistics

From the cluster members list you can select View HA statistics to display the serial number, status, and monitor information for each cluster unit. To view HA statistics, go to *System > Config > HA* and select View HA Statistics.

Figure 215: Example HA statistics (active-passive cluster)

Refresh every none Back to HA monitor >>						
Unit	Status	Up Time	Monitor			
620_ha_2 FG600B3908600825	✓	5 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		22 hours	0%	42	74875	0
		57 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		17 seconds	10%	30 Kbps	26981277	0
620_ha_1 FG600B3908600705	✓	5 days	CPU Usage	Active Sessions	Total Packets	Virus Detected
		22 hours	0%	21	12115	0
		48 minutes	Memory Usage	Network Utilization	Total Bytes	Intrusion Detected
		58 seconds	10%	19 Kbps	930358	0

Refresh every	Select to control how often the web-based manager updates the HA statistics display.
Back to HA monitor	Close the HA statistics list and return to the cluster members list.
Serial No.	Use the serial number ID to identify each FortiGate unit in the cluster. The cluster ID matches the FortiGate unit serial number.
Status	Indicates the status of each cluster unit. A green check mark indicates that the cluster unit is operating normally. A red X indicates that the cluster unit cannot communicate with the primary unit.
Up Time	The time in days, hours, minutes, and seconds since the cluster unit was last started.
Monitor	Displays system status information for each cluster unit.
CPU Usage	The current CPU status of each cluster unit. The web-based manager displays CPU usage for core processes only. CPU usage for management processes (for example, for HTTPS connections to the web-based manager) is excluded.
Memory Usage	The current memory status of each cluster unit. The web-based manager displays memory usage for core processes only. Memory usage for management processes (for example, for HTTPS connections to the web-based manager) is excluded.
Active Sessions	The number of communications sessions being processed by the cluster unit.
Total Packets	The number of packets that have been processed by the cluster unit since it last started up.
Virus Detected	The number of viruses detected by the cluster unit.

Network Utilization	The total network bandwidth being used by all of the cluster unit interfaces.
Total Bytes	The number of bytes that have been processed by the cluster unit since it last started up.
Intrusion Detected	The number of intrusions or attacks detected by Intrusion Protection running on the cluster unit.

Changing the HA configuration of an operating cluster

To change the configuration settings of an operating cluster, go to *System > Config > HA* to display the cluster members list. Select Edit for the master (or primary) unit in the cluster members list to display the HA configuration page for the cluster.

You can use the HA configuration page to check and fine tune the configuration of the cluster after the cluster is up and running. For example, if you connect or disconnect cluster interfaces you may want to change the Port Monitor configuration.

Any changes you make on this page, with the exception of changes to the device priority, are first made to the primary unit configuration and then synchronized to the subordinate units. Changing the device priority only affects the primary unit.

Changing the HA configuration of an operating virtual cluster

To change the configuration settings of the primary unit in a functioning cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to *System > Config > HA* to display the cluster members list. Select Edit for the master (or primary) unit in virtual cluster 1 or virtual cluster 2 to display the HA configuration page for the virtual cluster.

You can use the virtual cluster HA configuration page to check and fine tune the configuration of both virtual clusters after the cluster is up and running. For example, you may want to change the Port Monitor configuration for virtual cluster 1 and virtual cluster 2 so that each virtual cluster monitors its own interfaces.

You can also use this configuration page to move virtual domains between virtual cluster 1 and virtual cluster 2. Usually you would distribute virtual domains between the two virtual clusters to balance the amount of traffic being processed by each virtual cluster.

Any changes you make on this page, with the exception of changes to the device priorities, are first made to the primary unit configuration and then synchronized to the subordinate unit.

You can also adjust device priorities to configure the role of this cluster unit in the virtual cluster. For example, to distribute traffic to both cluster units in the virtual cluster configuration, you would want one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. You can create this configuration by setting the device priorities. The cluster unit with the highest device priority in virtual cluster 1 becomes the primary unit for virtual cluster 1. The cluster unit with the highest device priority in virtual cluster 2 becomes the primary unit in virtual cluster 2.

Changing the subordinate unit host name and device priority

To change the host name and device priority of a subordinate unit in an operating cluster, go to *System > Config > HA* to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

To change the host name and device priority of a subordinate unit in an operating cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to *System > Config > HA* to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

You can change the host name (Peer) and device priority (Priority) of this subordinate unit. These changes only affect the configuration of the subordinate unit.

The device priority is not synchronized among cluster members. In a functioning cluster you can change device priority to change the priority of any unit in the cluster. The next time the cluster negotiates, the cluster unit with the highest device priority becomes the primary unit.

The device priority range is 0 to 255. The default device priority is 128.

Upgrading cluster firmware

You can upgrade the FortiOS firmware running on an HA cluster in the same manner as upgrading the firmware running on a standalone FortiGate unit. During a normal firmware upgrade, the cluster upgrades the primary unit and all subordinate units to run the new firmware image. The firmware upgrade takes place without interrupting communication through the cluster.



Upgrading cluster firmware to a new major release (for example upgrading from 3.0 MRx to 4.0 MRx) is supported for clusters. Make sure you are taking an upgrade path described in the release notes. Even so you should back up your configuration and only perform such a firmware upgrade during a maintenance window.

To upgrade the firmware without interrupting communication through the cluster, the cluster goes through a series of steps that involve first upgrading the firmware running on the subordinate units, then making one of the subordinate units the primary unit, and finally upgrading the firmware on the former primary unit. These steps are transparent to the user and the network, but depending upon your HA configuration may result in the cluster selecting a new primary unit.

The following sequence describes in detail the steps the cluster goes through during a firmware upgrade and how different HA configuration settings may affect the outcome.

- 1 The administrator uploads a new firmware image from the web-based manager or CLI.
- 2 If the cluster is operating in active-active mode load balancing is turned off.
- 3 The cluster upgrades the firmware running on all of the subordinate units.
- 4 Once the subordinate units have been upgraded, a new primary unit is selected.

This primary unit will be running the new upgraded firmware.

- 5 The cluster now upgrades the firmware of the former primary unit.

If the age of the new primary unit is more than 300 seconds (5 minutes) greater than the age of all other cluster units, the new primary unit continues to operate as the primary unit.

This is the intended behavior but does not usually occur because the age difference of the cluster units is usually less than the cluster age difference margin of 300 seconds. So instead, the cluster negotiates again to select a primary unit as described in [“Primary unit selection” on page 2000](#).

You can keep the cluster from negotiating again by reducing the cluster age difference margin using the `ha-uptime-diff-margin` option. However, you should be cautious when reducing the age or other problems may occur. For information about the cluster age difference margin, see [“Cluster age difference margin \(grace period\)” on page 2003](#)). For more information about changing the cluster age margin, see [“Changing the cluster age difference margin” on page 2003](#).

- 6 If the cluster is operating in active-active mode, load balancing is turned back on.

Changing how the cluster processes firmware upgrades

By default cluster firmware upgrades proceed as uninterruptable upgrades that do not interrupt traffic flow. If required, you can use the following CLI command to change how the cluster handles firmware upgrades. You might want to change this setting if you are finding uninterruptable upgrades take too much time.

```
config system ha
  set uninterruptable-upgrade disable
end
```

`uninterruptable-upgrade` is enabled by default. If you disable `uninterruptable-upgrade` the cluster still upgrades the firmware on all cluster units, but all cluster units are upgraded at once; which takes less time but interrupts communication through the cluster.

Synchronizing the firmware build running on a new cluster unit

If the firmware build running on a FortiGate unit that you add to a cluster is older than the cluster firmware build, you may be able to use the following steps to synchronize the firmware running on the new cluster unit.

This procedure describes re-installing the same firmware build on a cluster to force the cluster to upgrade all cluster units to the same firmware build.

Due to firmware upgrade and synchronization issues, in some cases this procedure may not work. In all cases it will work to install the same firmware build on the new unit as the one that the cluster is running before adding the new unit to the cluster.

To synchronize the firmware build running on a new cluster unit

- 1 Obtain a firmware image that is the same as build already running on the cluster.
- 2 Connect to the cluster using the web-based manager.
- 3 Go to *System > Dashboard > Status*.
- 4 Select *Update* beside *Firmware Version*.
You can also install a newer firmware build.
- 5 Select OK.

After the firmware image is uploaded to the cluster, the primary unit upgrades all cluster units to this firmware build.

Downgrading cluster firmware

For various reasons you may need to downgrade the firmware that a cluster is running. You can use the information in this section to downgrade the firmware version running on a cluster.

In most cases you can downgrade the firmware on an operating cluster using the same steps as for a firmware upgrade. A warning message appears during the downgrade but the downgrade usually works and after the downgrade the cluster continues operating normally with the older firmware image.

Downgrading between some firmware versions, especially if features have changed between the two versions, may not always work without the requirement to fix configuration issues after the downgrade.

Only perform firmware downgrades during maintenance windows and make sure you back up your cluster configuration before the downgrade.

If the firmware downgrade that you are planning may not work without configuration loss or other problems, you can use the following downgrade procedure to make sure your configuration is not lost after the downgrade.

To downgrade cluster firmware

This example shows how to downgrade the cluster shown in [Figure 197 on page 2022](#). The cluster consists of two cluster units (620_ha_1 and 620_ha_2). The port1 and port2 interfaces are connected networks and the port3 and port4 interfaces are connected together for the HA heartbeat.

This example, describes separating each unit from the cluster and downgrading the firmware for the standalone FortiGate units. There are several ways you could disconnect units from the cluster. This example describes using the disconnect from cluster function described in [“Disconnecting a cluster unit from a cluster” on page 2164](#).

- 1 Go to *System > Maintenance > Backup & Restore* and backup the cluster configuration.
From the CLI use `execute backup config`.
- 2 Go to *System > Config > HA* and for 620_ha_1 select *Disconnect from cluster* icon.
- 3 Select the port2 interface and enter an IP address and netmask of 10.11.101.101/24 and select OK.

From the CLI you can enter the following command (FG600B3908600705 is the serial number of the cluster unit) to be able to manage the standalone FortiGate unit by connecting to the port2 interface with IP address and netmask 10.11.101.101/24.

```
execute ha disconnect FG600B3908600705 port2 10.11.101.101/24
```

After 620_ha_1 is disconnected, 620_ha_2 continues processing traffic.

- 4 Connect to the 620_ha_1 web-based manager or CLI using IP address 10.11.101.101/24 and follow normal procedures to downgrade standalone FortiGate unit firmware.
- 5 When the downgrade is complete confirm that the configuration of 620_ha_1 is correct.
- 6 Set the HA mode of 620_ha_2 to Standalone and follow normal procedures to downgrade standalone FortiGate unit firmware.
Network communication will be interrupted for a short time during the downgrade.
- 7 When the downgrade is complete confirm that the configuration of 620_ha_2 is correct.

- 8 Set the HA mode of 620_ha_2 to Active-Passive or the required HA mode.
- 9 Set the HA mode of 620_ha_1 to the same mode as 620_ha_2.

If you have not otherwise changed the HA settings of the cluster units and if the firmware downgrades have not affected the configurations the units should negotiate and form cluster running the downgraded firmware.

Backing up and restoring the cluster configuration

You can backup the configuration of the primary unit by logging into the web-based manager or CLI and following normal configuration backup procedures.

The following configuration settings are not synchronized to all cluster units:

- HA override and priority
- The interface configuration of the HA reserved management interface (`config system interface`)
- The HA reserved management interface default route (`ha-mgmt-interface-gateway`)
- The FortiGate unit host name.

To backup these configuration settings for each cluster unit you must log into each cluster unit and backup its configuration.

If you need to restore the configuration of the cluster including the configuration settings that are not synchronized you should first restore the configuration of the primary unit and then restore the configuration of each cluster unit. Alternatively you could log into each cluster unit and manually add the configuration settings that were not restored.

Monitoring cluster units for failover

If the primary unit in the cluster fails, the units in the cluster renegotiate to select a new primary unit. Failure of the primary unit results in the following:

- If SNMP is enabled, the new primary unit sends HA trap messages. The messages indicate a cluster status change, HA heartbeat failure, and HA member down. For more info about HA and SNMP, see [“Clusters and SNMP” on page 2143](#).
- If event logging is enabled and HA activity event is selected, the new primary unit records log messages that show that the unit has become the primary unit. See [“Example log messages” on page 2138](#) for some example message sequences when a failover occurs.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGate units. The failed primary unit no longer appears on the Cluster Members list.
- The host name and serial number of the primary unit changes. You can see these changes when you log into the web-based manager or CLI.
- The cluster info displayed on the dashboard, cluster members list or from the `get system ha status` command changes.

If a subordinate unit fails, the cluster continues to function normally. Failure of a subordinate unit results in the following:

- If event logging is enabled and HA activity event is selected, the primary unit records log messages that show that a subordinate has been removed from the cluster. See [“Example log messages” on page 2138](#) for some example message sequences.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGate units. The failed unit no longer appears on the Cluster Members list.

Viewing cluster status from the CLI

Use the `get system ha status` command to display information about an HA cluster. The command displays general HA configuration settings. The command also displays information about how the cluster unit that you have logged into is operating in the cluster.

Usually you would log into the primary unit CLI using SSH or telnet. In this case the `get system ha status` command displays information about the primary unit first, and also displays the HA state of the primary unit (the primary unit operates in the work state). However, if you log into the primary unit and then use the `execute ha manage` command to log into a subordinate unit, (or if you use a console connection to log into a subordinate unit) the `get system status` command displays information about this subordinate unit first, and also displays the HA state of this subordinate unit. The state of a subordinate unit is work for an active-active cluster and standby for an active-passive cluster.

For a virtual cluster configuration, the `get system ha status` command displays information about how the cluster unit that you have logged into is operating in virtual cluster 1 and virtual cluster 2. For example, if you connect to the cluster unit that is the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2, the output of the `get system ha status` command shows virtual cluster 1 in the work state and virtual cluster 2 in the standby state. The `get system ha status` command also displays additional information about virtual cluster 1 and virtual cluster 2.

The command display includes the following fields.

Fields	Description
Model	The FortiGate model number.
Mode	The HA mode of the cluster: a-a or a-p.
Group	The group ID of the cluster.
Debug	The debug status of the cluster.
ses_pickup	The status of session pickup: enable or disable.
load balance	The status of the <code>load-balance-all</code> keyword: enable or disable. Relevant to active-active clusters only.
schedule	The active-active load balancing schedule. Relevant to active-active clusters only.

Fields	Description
Master Slave	<p>Master displays the device priority, host name, serial number, and cluster index of the primary (or master) unit.</p> <p>Slave displays the device priority, host name, serial number, and cluster index of the subordinate (or slave, or backup) unit or units.</p> <p>The list of cluster units changes depending on how you log into the CLI. Usually you would use SSH or telnet to log into the primary unit CLI. In this case the primary unit would be at the top the list followed by the other cluster units.</p> <p>If you use <code>execute ha manage</code> or a console connection to log into a subordinate unit CLI, and then enter <code>get system ha status</code> the subordinate unit that you have logged into appears at the top of the list of cluster units.</p>
number of vcluster	<p>The number of virtual clusters. If virtual domains are not enabled, the cluster has one virtual cluster. If virtual domains are enabled the cluster has two virtual clusters.</p>
vcluster 1 Master Slave	<p>The HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 1. If virtual domains are not enabled, <code>vcluster 1</code> displays information for the cluster. If virtual domains are enabled, <code>vcluster 1</code> displays information for virtual cluster 1.</p> <p>The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 1 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 1.</p> <p><code>vcluster 1</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 1. The list includes the cluster index and serial number of each cluster unit in virtual cluster 1. The cluster unit that you have logged into is at the top of the list.</p> <p>If virtual domains are not enabled and you connect to the primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the primary unit.</p> <p>If virtual domains are not enabled and you connect to a subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you have logged into.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the virtual cluster 1 primary unit.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>

Fields	Description
vcluster 2 Master Slave	<p><code>vcluster 2</code> only appears if virtual domains are enabled.</p> <p><code>vcluster 2</code> displays the HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 2. The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 2 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 2.</p> <p><code>vcluster 2</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 2. The list includes the cluster index and serial number of each cluster unit in virtual cluster 2. The cluster unit that you have logged into is at the top of the list.</p> <p>If you connect to the virtual cluster 2 primary unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>work</code>. The display lists the cluster units starting with the virtual cluster 2 primary unit.</p> <p>If you connect to the virtual cluster 2 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>standby</code>. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>

Examples

The following example shows `get system ha status` output for a cluster of two FortiGate-5001SX units operating in active-active mode. The cluster group ID, session pickup, load balance all, and the load balancing schedule are all set to the default values. The device priority of the primary unit is also set to the default value. The device priority of the subordinate unit has been reduced to 100. The host name of the primary unit is `5001_Slot_4`. The host name of the subordinate unit is `5001_Slot_3`.

The command output was produced by connecting to the primary unit CLI (host name `5001_Slot_4`).

```

Model: 5000
Mode: a-a
Group: 0
Debug: 0
ses_pickup: disable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4      FG50012204400045 1
Slave :100 5001_Slot_3     FG50012205400050 0
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050

```

The following command output was produced by using `execute HA manage 0` to log into the subordinate unit CLI of the cluster shown in the previous example. The host name of the subordinate unit is `5001_Slot_3`.

```

Model: 5000
Mode: a-a
Group: 0
Debug: 0

```

```

ses_pickup: disable
load_balance: disable
schedule: round robin
Slave :100 5001_Slot_3      FG50012205400050 0
Master:128 5001_Slot_4      FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.2
Slave :1 FG50012205400050
Master:0 FG50012204400045

```

The following example shows `get system ha status` output for a cluster of three FortiGate-5001 units operating in active-passive mode. The cluster group ID is set to 20 and session pickup is enabled. Load balance all and the load balancing schedule are set to the default value. The device priority of the primary unit is set to 200. The device priorities of the subordinate units are set to 128 and 100. The host name of the primary unit is 5001_Slot_5. The host names of the subordinate units are 5001_Slot_3 and 5001_Slot_4.

```

Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:200 5001_Slot_5      FG50012206400112 0
Slave :100 5001_Slot_3      FG50012205400050 1
Slave :128 5001_Slot_4      FG50012204400045 2
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG50012206400112
Slave :1 FG50012204400045
Slave :2 FG50012205400050

```

The following example shows `get system ha status` output for a cluster of two FortiGate-5001 units with virtual clustering enabled. This command output was produced by logging into the primary unit for virtual cluster 1 (hostname: 5001_Slot_4, serial number FG50012204400045).

The virtual clustering output shows that the cluster unit with host name 5001_Slot_4 and serial number FG50012204400045 is operating as the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2.

For virtual cluster 1 the cluster unit that you have logged into is operating in the work state and the serial number of the primary unit for virtual cluster 1 is FG50012204400045. For virtual cluster 2 the cluster unit that you have logged into is operating in the standby state and the serial number of the primary unit for virtual cluster 2 is FG50012205400050.

```

Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4      FG50012204400045 1
Slave :100 5001_Slot_3      FG50012205400050 0
number of vcluster: 2

```

```

vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
vcluster 2: standby 169.254.0.1
Slave :1 FG50012204400045
Master:0 FG50012205400050

```

The following example shows `get system ha status` output for the same cluster as shown in the previous example after using `execute ha manage 0` to log into the primary unit for virtual cluster 2 (hostname: 5001_Slot_3, serial number FG50012205400050).

```

Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Slave :100 5001_Slot_3      FG50012205400050 0
Master:128 5001_Slot_4      FG50012204400045 1
number of vcluster: 2
vcluster 1: standby 169.254.0.2
Slave :1 FG50012205400050
Master:0 FG50012204400045
vcluster 2: work 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045

```

The following example shows `get system ha status` output for a virtual cluster configuration where the cluster unit with hostname: 5001_Slot_4 and serial number FG50012204400045 is the primary unit for both virtual clusters. This command output is produced by logging into cluster unit with host name 5001_Slot_4 and serial number FG50012204400045.

```

Model: 5000
Mode: a-p
Group: 20
Debug: 0
ses_pickup: enable
load_balance: disable
schedule: round robin
Master:128 5001_Slot_4      FG50012204400045 1
Slave :100 5001_Slot_3      FG50012205400050 0
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050
vcluster 2: work 169.254.0.2
Master:0 FG50012204400045
Slave :1 FG50012205400050

```

About the HA cluster index and the execute ha manage command

When a cluster starts up, the FortiGate Cluster Protocol (FGCP) assigns a cluster index and a HA heartbeat IP address to each cluster unit based on the serial number of the cluster unit. The FGCP selects the cluster unit with the highest serial number to become the primary unit. The FGCP assigns a cluster index of 0 and an HA heartbeat IP address of 169.254.0.1 to this unit. The FGCP assigns a cluster index of 1 and an HA heartbeat IP address of 169.254.0.2 to the cluster unit with the second highest serial number. If the cluster contains more units, the cluster unit with the third highest serial number is assigned a cluster index of 2 and an HA heartbeat IP address of 169.254.0.3, and so on. You can display the cluster index assigned to each cluster unit using the `get system ha status` command. Also when you use the `execute ha manage` command you select a cluster unit to log into by entering its cluster index.

The cluster index and HA heartbeat IP address only change if a unit leaves the cluster or if a new unit joins the cluster. When one of these events happens, the FGCP resets the cluster index and HA heartbeat IP address of each cluster unit according to serial number in the same way as when the cluster first starts up.

Each cluster unit keeps its assigned cluster index and HA heartbeat IP address even as the units take on different roles in the cluster. After the initial cluster index and HA heartbeat IP addresses are set according to serial number, the FGCP checks other primary unit selection criteria such as device priority and monitored interfaces. Checking these criteria could result in selecting a cluster unit without the highest serial number to operate as the primary unit.

Even if the cluster unit without the highest serial number now becomes the primary unit, the cluster indexes and HA heartbeat IP addresses assigned to the individual cluster units do not change. Instead the FGCP assigns a second cluster index, which could be called the operating cluster index, to reflect this role change. The operating cluster index is 0 for the primary unit and 1 and higher for the other units in the cluster. By default both sets of cluster indexes are the same. But if primary unit selection selects the cluster unit that does not have the highest serial number to be the primary unit then this cluster unit is assigned an operating cluster index of 0. The operating cluster index is used by the FGCP only. You can display the operating cluster index assigned to each cluster unit using the `get system ha status` command. There are no CLI commands that reference the operating cluster index.



Even though there are two cluster indexes there is only one HA heartbeat IP address and the HA heartbeat address is not affected by a change in the operating cluster index.

Using the execute ha manage command

When you use the CLI command `execute ha manage <index_integer>` to connect to the CLI of another cluster unit, the `<index_integer>` that you enter is the cluster index of the unit that you want to connect to.

Using get system ha status to display cluster indexes

You can display the cluster index assigned to each cluster unit using the CLI command `get system ha status`. The following example shows the information displayed by the `get system ha status` command for a cluster consisting of two FortiGate-5001SX units operating in active-passive HA mode with virtual domains not enabled and without virtual clustering.

```
get system ha status
Model: 5000
```

```

Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0
Slave :128 5001_slot_11 FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045

```

In this example, the cluster unit with serial number FG50012205400050 has the highest serial number and so has a cluster index of 0 and the cluster unit with serial number FG50012204400045 has a cluster index of 1. From the CLI of the primary (or master) unit of this cluster you can connect to the CLI of the subordinate (or slave) unit using the following command:

```
execute ha manage 1
```

This works because the cluster unit with serial number FG50012204400045 has a cluster index of 1.

The `get system ha status` command output shows two similar lists of indexes and serial numbers. The listing on the sixth and seventh lines of the command output are the cluster indexes assigned according to cluster unit serial number. These are the cluster indexes that you enter when using the `execute ha manage` command. The cluster indexes shown in the last two lines of the command output are the operating cluster indexes that reflect how the cluster units are actually operating in the cluster. In this example both sets of cluster indexes are the same.

The last three lines of the command output display the status of vcluster 1. In a cluster consisting of two cluster units operating without virtual domains enabled all clustering actually takes place in virtual cluster 1. HA is designed to work this way to support virtual clustering. If this cluster was operating with virtual domains enabled, adding virtual cluster 2 is similar to adding a new copy of virtual cluster 1. Virtual cluster 2 is visible in the `get system ha status` command output when you add virtual domains to virtual cluster 2.

The HA heartbeat IP address displayed on line 8 is the HA heartbeat IP address of the cluster unit that is actually operating as the primary unit. For a default configuration this IP address will always be 169.254.0.1 because the cluster unit with the highest serial number will be the primary unit. This IP address changes if the operating primary unit is not the primary unit with the highest serial number.

Example: actual and operating cluster indexes do not match

This example shows `get system ha status` command output for same cluster of two FortiGate-5001SX units. However, in this example the device priority of the cluster unit with the serial number FG50012204400045 is increased to 200. As a result the cluster unit with the lowest serial number becomes the primary unit. This means the actual and operating cluster indexes of the cluster units do not match.

```

get system ha status
Model: 5000
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0

```



```
Slave :200 5001_slot_11 FG50012204400045 1
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master:1 FG50012205400050
Slave :0 FG50012204400045
```

The actual cluster indexes have not changed but the operating cluster indexes have. Also, the HA heartbeat IP address displayed for vcluster 1 has changed to 169.254.0.2.

Virtual clustering example output

The `get system ha status` command output is the same if a cluster is operating with virtual clustering turned on but with all virtual domains in virtual cluster 1. The following `get system ha status` command output example shows the same cluster operating as a virtual cluster with virtual domains in virtual cluster 1 and added to virtual cluster 2. In this example the cluster unit with serial number FG50012204400045 is the primary unit for virtual cluster 1 and the cluster unit with serial number FG50012205400050 is the primary unit for virtual cluster 2.

```
get system ha status
Model: 5000
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:128 5001_slot_7 FG50012205400050 0
Slave :200 5001_slot_11 FG50012204400045 1
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master:1 FG50012205400050
Slave :0 FG50012204400045
vcluster 2: standby 169.254.0.1
Master:0 FG50012205400050
Slave :1 FG50012204400045
```

This example shows three sets of indexes. The indexes in lines six and seven are still used by the `execute ha manage` command. The indexes on lines ten and eleven are for the primary and subordinate units in virtual cluster 1 and the indexes on the last two lines are for virtual cluster 2.

Managing individual cluster units

The following procedure describes how to use SSH to log into the primary unit CLI and from there to use the `execute ha manage` command to connect to the CLI of any other unit in the cluster. The procedure is very similar if you use telnet, or the web-based manager dashboard CLI console.

You can use the `execute ha manage` command from the CLI of any cluster unit to log into the CLI of another the cluster unit. Usually you would use this command from the CLI of the primary unit to log into the CLI of a subordinate unit. However, if you have logged into a subordinate unit CLI, you can use this command to log into the primary unit CLI, or the CLI of another subordinate unit.

Using SSH or telnet or the web-based manager dashboard CLI console you can only log into the primary unit CLI. Using a direct console connection you can log into any cluster unit. In both cases you can use `execute ha manage` to connect to the CLI of other cluster units.



You log into the subordinate unit using the `FGT_ha_admin` administrator account. This built-in administrator account gives you read and write permission on the subordinate unit. Normally this built-in administrative account is not visible, however `FGT_ha_admin` does appear in event log messages.

- 1 Use SSH to connect to the cluster and log into the primary unit CLI.
Connect to any cluster interface configured for SSH administrative access to log into the cluster.
- 2 Enter the following command followed by a space and type a question mark (?):
`execute ha manage`
The CLI displays a list of all the subordinate units in the cluster. Each cluster unit is numbered, starting at 1. The information displayed for each cluster unit includes the unit serial number and the host name of the unit.
- 3 Complete the command with the number of the subordinate unit to log into. For example, to log into subordinate unit 1, enter the following command:
`execute ha manage 1`
Press Enter to connect to and log into the CLI of the selected subordinate unit. If this subordinate unit has a different host name, the CLI prompt changes to this host name.
You can use CLI commands to manage this subordinate unit. If you make changes to the configuration of any cluster unit (primary or subordinate unit) these changes are synchronized to all cluster units.
- 4 You can now use the `execute ha manage` command to connect to any other cluster unit (including the primary unit). You can also use the `exit` command to return to the primary unit CLI.

Disconnecting a cluster unit from a cluster

Use the following procedures to disconnect a cluster unit from a functioning cluster without disrupting the operation of the cluster. You can disconnect a cluster unit if you need to use the disconnected FortiGate unit for another purpose, such as to act as a standalone firewall.

You can use the following procedures for a standard cluster and for a virtual clustering configuration. To use the following procedures from a virtual cluster you must be logged in as the admin administrator and you must have selected Global Configuration.

When you disconnect a cluster unit you must assign an IP address and netmask to one of the interfaces of the disconnected unit. You can disconnect any unit from the cluster even the primary unit. After the unit is disconnected, the cluster responds as if the disconnected unit has failed. The cluster may renegotiate and may select a new primary unit.

When the cluster unit is disconnected the HA mode is changed to standalone. In addition, all interface IP addresses of the disconnected unit are set to 0.0.0.0 except for the interface that you configure.

Otherwise the configuration of the disconnected unit is not changed. The HA configuration of the disconnected unit is not changed either (except to change the HA mode to Standalone).

To disconnect a cluster unit from a cluster - web-based manager

- 1 Go to *System > Config > HA* to view the cluster members list.
- 2 Select the Disconnect from cluster icon for the cluster unit to disconnect from the cluster.
- 3 Select the interface that you want to configure. You also specify the IP address and netmask for this interface. When the FortiGate unit is disconnected, all management access options are enabled for this interface.
- 4 Specify an IP address and netmask for the interface. You can use this IP address to connect to the interface to configure the disconnected FortiGate unit.
- 5 Select OK.

The FortiGate unit is disconnected from the cluster and the cluster may renegotiate and select a new primary unit. The selected interface of the disconnected unit is configured with the specified IP address and netmask.

To disconnect a cluster unit from a cluster - CLI

- 1 Enter the following command to disconnect a cluster unit with serial number FGT5002803033050. The internal interface of the disconnected unit is set to IP address 1.1.1.1 and netmask 255.255.255.0.

```
execute ha disconnect FGT5002803033050 internal 1.1.1.1  
255.255.255.0
```

Adding a disconnected FortiGate unit back to its cluster

If you disconnect a FortiGate unit from a cluster, you can re-connect the disconnected FortiGate unit to the cluster by setting the HA mode of the disconnected unit to match the HA mode of the cluster. Usually the disconnected unit rejoins the cluster as a subordinate unit and the cluster automatically synchronizes its configuration.



You do not have to change the HA password on the disconnected unit unless the HA password has been changed after the unit was disconnected. Disconnecting a unit from a cluster does not change the HA password.



You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. You should also make sure that the HA `override` CLI option is not enabled on the disconnected unit. Otherwise, when the disconnected unit joins the cluster, the cluster will renegotiate and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units. This configuration change might disrupt the operation of the cluster.

The following procedure assumes that the disconnected FortiGate unit is correctly physically connected to your network and to the cluster but is not running in HA mode and not part of the cluster.

Before you start this procedure you should note the device priority of the primary unit.

To add a disconnected FortiGate unit back to its cluster - web-based manager

- 1 Log into the disconnected FortiGate unit.
If virtual domains are enabled, log in as the admin administrator and select Global Configuration.

- 2 Go to *System > Config > HA*.
- 3 Change Mode to match the mode of the cluster.
- 4 If required, change the group name and password to match the cluster.
- 5 Set the Device Priority lower than the device priority of the primary unit.
- 6 Select OK.

The disconnected FortiGate unit joins the cluster.

To add a disconnected FortiGate unit back to its cluster - CLI

- 1 Log into the CLI of the FortiGate unit to be added back to the cluster.
- 2 Enter the following command to access the global configuration and add the FortiGate unit back to a cluster operating in active-passive mode and set the device priority to 50 (a low number) so that this unit will not become the primary unit:

```
config global
  config system ha
    set mode a-p
    set priority 50
  end
end
```

You may have to also change the group name, group id and password. However if you have not changed these for the cluster or the FortiGate unit after it was disconnected from the cluster you should not have to adjust them now.



HA and failover protection

In FortiGate active-passive HA, the FortiGate Clustering Protocol (FGCP) provides failover protection. This means that an active-passive cluster can provide FortiGate services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiGate unit. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in a mission-critical environment.

The FGCP supports three kinds of failover protection. Device failover automatically replaces a failed device and restarts traffic flow with minimal impact on the network. Link failover maintains traffic flow if a link fails. Session failover resumes communication sessions with minimal loss of data if a device or link failover occurs.

This chapter describes how FGCP failover protection works and provides detailed NAT/Route and Transparent mode packet flow descriptions.

This chapter contains the following sections:

- [About active-passive failover](#)
- [About active-active failover](#)
- [Device failover](#)
- [HA heartbeat and communication between cluster units](#)
- [Cluster virtual MAC addresses](#)
- [Synchronizing the configuration](#)
- [Synchronizing routing table updates](#)
- [Synchronizing IPsec VPN SAs](#)
- [Link failover](#)
- [Subsecond failover](#)
- [Remote link failover](#)
- [Session failover \(session pick-up\)](#)
- [WAN optimization and HA](#)
- [Failover and attached network equipment](#)
- [Monitoring cluster units for failover](#)
- [NAT/Route mode active-passive cluster packet flow](#)
- [Transparent mode active-passive cluster packet flow](#)
- [Failover performance](#)

About active-passive failover

To achieve failover protection in an active-passive cluster, one of the cluster units functions as the primary unit, while the rest of the cluster units are subordinate units, operating in an active stand-by mode. The cluster IP addresses and HA virtual MAC addresses are associated with the cluster interfaces of the primary unit. All traffic directed at the cluster is actually sent to and processed by the primary unit.

While the cluster is functioning, the primary unit functions as the FortiGate network security device for the networks that it is connected to. In addition, the primary unit and subordinate units use the HA heartbeat to keep in constant communication. The subordinate units report their status to the cluster unit and receive and store connection and state table updates.

Device failure

If the primary unit encounters a problem that is severe enough to cause it to fail, the remaining cluster units negotiate to select a new primary unit. This occurs because all of the subordinate units are constantly waiting to negotiate to become primary units. Only the heartbeat packets sent by the primary unit keep the subordinate units from becoming primary units. Each received heartbeat packet resets negotiation timers in the subordinate units. If this timer is allowed to run out because the subordinate units do not receive heartbeat packets from the primary unit, the subordinate units assume that the primary unit has failed, and negotiate to become primary units themselves.

Using the same FCGP negotiation process that occurs when the cluster starts up, after they determine that the primary unit has failed, the subordinate units negotiate amongst themselves to select a new primary unit. The subordinate unit that wins the negotiation becomes the new primary unit with the same MAC and IP addresses as the former primary unit. The new primary unit then sends gratuitous ARP packets out all of its interfaces to inform attached switches to send traffic to the new primary unit. Sessions then resume with the new primary unit.

Link failure

If a primary unit interface fails or is disconnected while a cluster is operation, a link failure occurs. When a link failure occurs the cluster units negotiate to select a new primary unit. Since the primary unit has not stopped operating, it participates in the negotiation. The link failure means that a new primary unit must be selected and the cluster unit with the link failure joins the cluster as a subordinate unit.

Just as for a device failover, the new primary unit sends gratuitous arp packets out all of its interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary unit.

If a subordinate unit experiences a device failure its status in the cluster does not change. However, in future negotiations a cluster unit with a link failure is unlikely to become the primary unit.

Session failover

If you enable session failover (also called session pickup) for the cluster, during cluster operation the primary unit informs the subordinate units of changes to the primary unit connection and state tables, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster.

After a failover the new primary unit recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary unit and are handled according to their last known state.

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed.

Primary unit recovery

If a primary unit recovers after a device or link failure, it will operate as a subordinate unit, unless the `override` CLI keyword is enabled and its device priority is set higher than the unit priority of other cluster units (see [“HA override” on page 2008](#)).

About active-active failover

HA failover in a cluster running in active-active mode is similar to active-passive failover described above. Active-active subordinate units are constantly waiting to negotiate to become primary units and, if session failover is enabled, continuously receive connection state information from the primary unit. If the primary unit fails, or one of the primary unit interfaces fails, the cluster units use the same mechanisms to detect the failure and to negotiate to select a new primary unit. If session failover is enabled, the new primary unit also maintains communication sessions through the cluster using the shared connection state table.

Active-active HA load balances sessions among all cluster units. For session failover, the cluster must maintain all of these sessions. To load balance sessions, the functioning cluster uses a load balancing schedule to distribute sessions to all cluster units. The shared connection state table tracks the communication sessions being processed by all cluster units (not just the primary unit). After a failover, the new primary unit uses the load balancing schedule to re-distribute all of the communication sessions recorded in the shared connection state table among all of the remaining cluster units. The connections continue to be processed by the cluster, but possibly by a different cluster unit, and are handled according to their last known state.

Device failover

The FGCP provides transparent device failover. Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

In the case of FortiOS HA, the device is the primary unit. If the primary unit fails, device failover ensures that one of the subordinate units in the cluster automatically takes the place of the primary unit and can continue processing network traffic in the same way as the failed primary unit.



Device failover does not maintain communication sessions. After a device failover, communication sessions have to be restarted. To maintain communication sessions, you must enable session failover. See [“Session failover \(session pick-up\)” on page 2205](#).

FortiGate HA device failover is supported by the HA heartbeat, virtual MAC addresses, configuration synchronization, route synchronization and IPsec VPN SA synchronization.

The HA heartbeat makes sure that the subordinate units detect a primary unit failure. If the primary unit fails to respond on time to HA heartbeat packets the subordinate units assume that the primary unit has failed and negotiate to select a new primary unit.

The new primary unit takes the place of the failed primary unit and continues functioning in the same way as the failed primary unit. For the new primary unit to continue functioning like the failed primary unit, the new primary unit must be able to reconnect to network devices and the new primary unit must have the same configuration as the failed primary unit.

FortiGate HA uses virtual MAC addresses to reconnect the new primary unit to network devices. The FGCP causes the new primary unit interfaces to acquire the same virtual MAC addresses as the failed primary unit. As a result, the new primary unit has the same network identity as the failed primary unit.

The new primary unit interfaces have different physical connections than the failed primary unit. Both the failed and the new primary unit interfaces are connected to the same switches, but the new primary unit interfaces are connected to different ports on these switches. To make sure that the switches send packets to the new primary unit, the new primary unit interfaces send gratuitous ARP packets to the connected switches. These gratuitous ARP packets notify the switches that the primary unit MAC and IP addresses are on different switch ports and cause the switches to send packets to the ports connected to the new primary unit. In this way, the new primary unit continues to receive packets that would otherwise have been sent to the failed primary unit.

Configuration synchronization means that the new primary unit always has the same configuration as the failed primary unit. As a result the new primary unit operates in exactly the same way as the failed primary unit. If configuration synchronization were not available the new primary unit may not process network traffic in the same way as the failed primary unit.

Route synchronization synchronizes the primary unit routing table to all subordinate units so that after a failover the new primary unit does not have to form a completely new routing table. IPsec VPN SA synchronization synchronizes IPsec VPN security associations (SAs) and other IPsec session data so that after a failover the new primary unit can resume IPsec tunnels without having to establish new SAs.

HA heartbeat and communication between cluster units

The HA heartbeat keeps cluster units communicating with each other. The heartbeat consists of hello packets that are sent at regular intervals by the heartbeat interface of all cluster units. These hello packets describe the state of the cluster unit and are used by other cluster units to keep all cluster units synchronized.

HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The default time interval between HA heartbeats is 200 ms. The FGCP uses link-local IP4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

For best results, isolate the heartbeat devices from your user networks by connecting the heartbeat devices to a separate switch that is not connected to any network. If the cluster consists of two FortiGate units you can connect the heartbeat device interfaces directly using a crossover cable. Heartbeat packets contain sensitive information about the cluster configuration. Heartbeat packets may also use a considerable amount of network bandwidth. For these reasons, it is preferable to isolate heartbeat packets from your user networks.

On startup, a FortiGate unit configured for HA operation broadcasts HA heartbeat hello packets from its HA heartbeat interface to find other FortiGate units configured to operate in HA mode. If two or more FortiGate units operating in HA mode connect with each other, they compare HA configurations (HA mode, HA password, and HA group ID). If the HA configurations match, the units negotiate to form a cluster.

While the cluster is operating, the HA heartbeat confirms that all cluster units are functioning normally. The heartbeat also reports the state of all cluster units, including the communication sessions that they are processing.

Heartbeat interfaces

A heartbeat interface is an Ethernet network interface in a cluster that is used by the FGCP for HA heartbeat communications between cluster units.

To change the HA heartbeat configuration go to *System > Config > HA* and select the *FortiGate interfaces to use as HA heartbeat interfaces*.

From the CLI enter the following command to make port4 and port5 HA heartbeat interfaces and give both interfaces a heartbeat priority of 150:

```
config system ha
  set hbdev port4 150 port5 150
end
```

The following example shows how to change the default heartbeat interface configuration so that the port4 and port1 interfaces can be used for HA heartbeat communication and to give the port4 interface the highest heartbeat priority so that port4 is the preferred HA heartbeat interface.

```
config system ha
  set hbdev port4 100 port1 50
end
```

By default, for most FortiGate models two interfaces are configured to be heartbeat interfaces. You can change the heartbeat interface configuration as required. For example you can select additional or different heartbeat interfaces. You can also select only one heartbeat interface.

In addition to selecting the heartbeat interfaces, you also set the *Priority* for each heartbeat interface. In all cases, the heartbeat interface with the highest priority is used for all HA heartbeat communication. If the interface fails or becomes disconnected, the selected heartbeat interface that has the next highest priority handles all heartbeat communication.

If more than one heartbeat interface has the same priority, the heartbeat interface with the highest priority that is also highest in the heartbeat interface list is used for all HA heartbeat communication. If this interface fails or becomes disconnected, the selected heartbeat interface with the highest priority that is next highest in the list handles all heartbeat communication.

The default heartbeat interface configuration sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration if one or both of the default heartbeat interfaces are connected. You can select different heartbeat interfaces, select more heartbeat interfaces and change heartbeat priorities according to your requirements.

For the HA cluster to function correctly, you must select at least one heartbeat interface and this interface of all of the cluster units must be connected together. If heartbeat communication is interrupted and cannot failover to a second heartbeat interface, the cluster units will not be able to communicate with each other and more than one cluster unit may become a primary unit. As a result the cluster stops functioning normally because multiple devices on the network may be operating as primary units with the same IP and MAC addresses creating a kind of split brain scenario.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0. The higher the number the higher the priority.

In most cases you can maintain the default heartbeat interface configuration as long as you can connect the heartbeat interfaces together. Configuring HA heartbeat interfaces is the same for virtual clustering and for standard HA clustering.

You can enable heartbeat communications for physical interfaces, but not for VLAN subinterfaces, IPsec VPN interfaces, redundant interfaces, or for 802.3ad aggregate interfaces. You cannot select these types of interfaces in the heartbeat interface list.

Selecting more heartbeat interfaces increases reliability. If a heartbeat interface fails or is disconnected, the HA heartbeat fails over to the next heartbeat interface.

You can select up to 8 heartbeat interfaces. This limit only applies to FortiGate units with more than 8 physical interfaces.

HA heartbeat traffic can use a considerable amount of network bandwidth. If possible, enable HA heartbeat traffic on interfaces used only for HA heartbeat traffic or on interfaces connected to less busy networks.

Connecting HA heartbeat interfaces

For most FortiGate models if you do not change the heartbeat interface configuration, you can isolate the default heartbeat interfaces of all of the cluster units by connecting them all to the same switch. Use one switch per heartbeat interface. If the cluster consists of two units you can connect the heartbeat interfaces together using crossover cables. For an example of how to connect heartbeat interfaces, see [“Connecting a FortiGate HA cluster” on page 1996](#).

HA heartbeat and data traffic are supported on the same cluster interface. In NAT/Route mode, if you decide to use heartbeat interfaces for processing network traffic or for a management connection, you can assign the interface any IP address. This IP address does not affect HA heartbeat traffic.

In Transparent mode, you can connect the heartbeat interface to your network and enable management access. You would then establish a management connection to the interface using the Transparent mode management IP address. This configuration does not affect HA heartbeat traffic.

Heartbeat interfaces and FortiGate switch interfaces

You can configure a FortiGate interface that contains an internal switch as an HA heartbeat interface. However this configuration is not recommended for two reasons:

- For security reasons and to save network bandwidth you should keep HA heartbeat traffic off of your internal network, and internal switch interfaces are usually intended to be connected to your internal network.
- Heartbeat packets may be lost if the switch interface is processing high volumes of traffic. Losing heartbeat packets may lead to unnecessary and repeated failovers.

Heartbeat packets and heartbeat interface selection

HA heartbeat hello packets are constantly sent by all of the enabled heartbeat interfaces. Using these hello packets, each cluster unit confirms that the other cluster units are still operating. The FGCP selects one of the heartbeat interfaces to be used for communication between the cluster units. The FGCP selects the heartbeat interface for heartbeat communication based on the linkfail states of the heartbeat interfaces, on the priority of the heartbeat interfaces, and on the interface index.

The FGCP checks the linkfail state of all heartbeat interfaces to determine which ones are connected. The FGCP selects one of these connected heartbeat interfaces to be the one used for heartbeat communication. The FGCP selects the connected heartbeat interface with the highest priority for heartbeat communication.

If more than one connected heartbeat interface has the highest priority the FGCP selects the heartbeat interface with the lowest interface index. The web-based manager lists the FortiGate unit interfaces in alphabetical order. This order corresponds to the interface index order with lowest index at the top and highest at the bottom. If more than one heartbeat interface has the highest priority, the FGCP selects the interface that is highest in the heartbeat interface list (or first in alphabetical order) for heartbeat communication.

If the interface that is processing heartbeat traffic fails or becomes disconnected, the FGCP uses the same criteria to select another heartbeat interface for heartbeat communication. If the original heartbeat interface is fixed or reconnected, the FGCP again selects this interface for heartbeat communication.

The HA heartbeat communicates cluster session information, synchronizes the cluster configuration, synchronizes the cluster routing table, and reports individual cluster member status. The HA heartbeat constantly communicates HA status information to make sure that the cluster is operating properly.

Interface index and display order

The web-based manager and CLI display interface names in alphanumeric order. For example, the sort order for a FortiGate unit with 10 interfaces (named port1 through port10) places port10 at the bottom of the list:

- port1
- port2 through 9
- port10

However, interfaces are indexed in hash map order, rather than purely by alphabetic order or purely by interface number value comparisons. As a result, the list is sorted primarily alphabetical by interface name (for example, base1 is before port1), then secondarily by index numbers:

- port1
- port10
- port2 through port9

HA heartbeat interface IP addresses

The FGCP uses link-local IP4 addresses ([RFC 3927](#)) in the 169.254.0.x range for HA heartbeat interface IP addresses and for inter-VDOM link interface IP addresses. When a cluster initially starts up, the primary unit heartbeat interface IP address is 169.254.0.1. Subordinate units are assigned heartbeat interface IP addresses in the range 169.254.0.2 to 169.254.0.63. HA inter-VDOM link interfaces on the primary unit are assigned IP addresses 169.254.0.65 and 169.254.0.66.

The ninth line of the following CLI command output shows the HA heartbeat interface IP address of the primary unit.

```
get system ha status
Model: 620
Mode: a-p
Group: 0
Debug: 0
ses_pickup: disable
Master:150 head_office_upper FG600B3908600825 1
Slave :150 head_office_lower FG600B3908600705 0
number of vcluster: 1
vcluster 1: work 169.254.0.1
```

```
Master:0 FG600B3908600825
Slave :1 FG600B3908600705
```

You can also use the `execute traceroute` command from the subordinate unit CLI to display HA heartbeat IP addresses and the HA inter-VDOM link IP addresses. For example, use `execute ha manage 1` to connect to the subordinate unit CLI and then enter the following command to trace the route to an IP address on your network:

```
execute traceroute 172.20.20.10
traceroute to 172.20.20.10 (172.20.20.10), 32 hops max, 72 byte packets
 1  169.254.0.1  0 ms  0 ms  0 ms
 2  169.254.0.66 0 ms  0 ms  0 ms
 3  172.20.20.10 0 ms  0 ms  0 ms
```

Both HA heartbeat and data traffic are supported on the same FortiGate interface. All heartbeat communication takes place on a separate VDOM called `vsys_ha`. Heartbeat traffic uses a virtual interface called `port_ha` in the `vsys_ha` VDOM. Data and heartbeat traffic use the same physical interface, but they're logically separated into separate VDOMs.

Heartbeat packet Ethertypes

Normal IP packets are 802.3 packets that have an Ethernet type (Ethertype) field value of 0x0800. Ethertype values other than 0x0800 are understood as level2 frames rather than IP packets.

By default, HA heartbeat packets use the following Ethertypes:

- HA heartbeat packets for NAT/Route mode clusters use Ethertype 0x8890. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ethertype of these packets using the `ha-eth-type` option of the `config system ha` command.
- HA heartbeat packets for Transparent mode clusters use Ethertype 0x8891. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ethertype of these packets using the `hc-eth-type` option of the `config system ha` command.
- HA telnet sessions between cluster units over HA heartbeat links use Ethertype 0x8893. The telnet sessions are used to synchronize the cluster configurations. Telnet sessions are also used when an administrator uses the `execute ha manage` command to connect from one cluster unit CLI to another. You can change the Ethertype of these packets using the `l2ep-eth-type` option of the `config system ha` command.

Because heartbeat packets are recognized as level2 frames, the switches and routers on your heartbeat network that connect to heartbeat interfaces must be configured to allow them. If level2 frames are dropped by these network devices, heartbeat traffic will not be allowed between the cluster units.

Some third-party network equipment may use packets with these Ethertypes for other purposes. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGate units connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

Alternatively, you can use the following CLI options to change the Ethertypes of the HA heartbeat packets:

```
config system ha
    set ha-eth-type <ha_ethertype_4-digit_hex>
    set hc-eth-type <hc_ethertype_4-digit_hex>
    set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For example, use the following command to change the Ethertype of the HA heartbeat packets from 0x8890 to 0x8895 and to change the Ethertype of HA Telnet session packets from 0x8891 to 0x889f:

```
config system ha
    set ha-eth-type 8895
    set l2ep-eth-type 889f
end
```

Modifying heartbeat timing

In an HA cluster, if a cluster unit CPU becomes very busy, the cluster unit may not be able to send heartbeat packets on time. If heartbeat packets are not sent on time other units in the cluster may think that the cluster unit has failed and the cluster will experience a failover.

A cluster unit CPU may become very busy if the cluster is subject to a syn flood attack, if network traffic is very heavy, or for other similar reasons. You can use the following CLI commands to configure how the cluster times HA heartbeat packets:

```
config system ha
    set hb-interval <interval_integer>
    set hb-lost-threshold <threshold_integer>
    set helo-holddown <holddown_integer>
end
```

Changing the lost heartbeat threshold

The lost heartbeat threshold is the number of consecutive heartbeat packets that are not received from another cluster unit before assuming that the cluster unit has failed. The default value is 6, meaning that if the 6 heartbeat packets are not received from a cluster unit then that cluster unit is considered to have failed. The range is 1 to 60 packets.

If the primary unit does not receive a heartbeat packet from a subordinate unit before the heartbeat threshold expires, the primary unit assumes that the subordinate unit has failed.

If a subordinate unit does not receive a heartbeat packet from the primary unit before the heartbeat threshold expires, the subordinate unit assumes that the primary unit has failed. The subordinate unit then begins negotiating to become the new primary unit.

The lower the `hb-lost-threshold` the faster a cluster responds when a unit fails. However, sometimes heartbeat packets may not be sent because a cluster unit is very busy. This can lead to a false positive failure detection. To reduce these false positives you can increase the `hb-lost-threshold`.

Use the following CLI command to increase the lost heartbeat threshold to 12:

```
config system ha
```

```
set hb-lost-threshold 12
end
```

Changing the heartbeat interval

The heartbeat interval is the time between sending HA heartbeat packets. The heartbeat interval range is 1 to 20 (100*ms). The heartbeat interval default is 2 (200 ms).

A heartbeat interval of 2 means the time between heartbeat packets is 200 ms. Changing the heartbeat interval to 5 changes the time between heartbeat packets to 500 ms (5 * 100ms = 500ms).

The HA heartbeat packets consume more bandwidth if the heartbeat interval is short. But if the heartbeat interval is very long, the cluster is not as sensitive to topology and other network changes.

Use the following CLI command to increase the heartbeat interval to 10:

```
config system ha
set hb-interval 10
end
```

The heartbeat interval combines with the lost heartbeat threshold to set how long a cluster unit waits before assuming that another cluster unit has failed and is no longer sending heartbeat packets. By default, if a cluster unit does not receive a heartbeat packet from a cluster unit for 6 * 200 = 1200 milliseconds or 1.2 seconds the cluster unit assumes that the other cluster unit has failed.

You can increase both the heartbeat interval and the lost heartbeat threshold to reduce false positives. For example, increasing the heartbeat interval to 20 and the lost heartbeat threshold to 30 means a failure will be assumed if no heartbeat packets are received after 30 * 2000 milliseconds = 60,000 milliseconds, or 60 seconds.

Use the following CLI command to increase the heartbeat interval to 20 and the lost heartbeat threshold to 30:

```
config system ha
set hb-lost-threshold 20
set hb-interval 30
end
```

Changing the time to wait in the hello state

The hello state hold-down time is the number of seconds that a cluster unit waits before changing from hello state to work state. After a failure or when starting up, cluster units operate in the hello state to send and receive heartbeat packets so that all the cluster units can find each other and form a cluster. A cluster unit should change from the hello state to work state after it finds all of the other FortiGate units to form a cluster with. If for some reason all cluster units cannot find each other during the hello state then some cluster units may be joining the cluster after it has formed. This can cause disruptions to the cluster and affect how it operates.

One reason for a delay in all of the cluster units joining the cluster could be the cluster units are located at different sites or if for some other reason communication is delayed between the heartbeat interfaces.

If cluster units are joining your cluster after it has started up or if it takes a while for units to join the cluster you can increase the time that the cluster units wait in the hello state. The hello state hold-down time range is 5 to 300 seconds. The hello state hold-down time default is 20 seconds.

Use the following CLI command to increase the time to wait in the hello state to 1 minute (60 seconds):

```
config system ha
    set hello-holddown 60
end
```

Enabling or disabling HA heartbeat encryption and authentication

You can enable HA heartbeat encryption and authentication to encrypt and authenticate HA heartbeat packets. HA heartbeat packets should be encrypted and authenticated if the cluster interfaces that send HA heartbeat packets are also connected to your networks.

If HA heartbeat packets are not encrypted the cluster password and changes to the cluster configuration could be exposed and an attacker may be able to sniff HA packets to get cluster information. Enabling HA heartbeat message authentication prevents an attacker from creating false HA heartbeat messages. False HA heartbeat messages could affect the stability of the cluster.

HA heartbeat encryption and authentication are disabled by default. Enabling HA encryption and authentication could reduce cluster performance. Use the following CLI command to enable HA heartbeat encryption and authentication.

```
config system ha
    set authentication enable
    set encryption enable
end
```

HA authentication and encryption uses AES-128 for encryption and SHA1 for authentication.

Cluster virtual MAC addresses

When a cluster is operating, the FGCP assigns virtual MAC addresses to each primary unit interface. HA uses virtual MAC addresses so that if a failover occurs, the new primary unit interfaces will have the same virtual MAC addresses and IP addresses as the failed primary unit. As a result, most network equipment would identify the new primary unit as the exact same device as the failed primary unit.

If the MAC addresses changed after a failover, the network would take longer to recover because all attached network devices would have to learn the new MAC addresses before they could communicate with the cluster.

If a cluster is operating in NAT/Route mode, the FGCP assigns a different virtual MAC address to each primary unit interface. VLAN subinterfaces are assigned the same virtual MAC address as the physical interface that the VLAN subinterface is added to. Redundant interfaces or 802.3ad aggregate interfaces are assigned the virtual MAC address of the first interface in the redundant or aggregate list.

If a cluster is operating in Transparent mode, the FGCP assigns a virtual MAC address for the primary unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.



A MAC address conflict can occur if two clusters are operating on the same network. See [“Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain” on page 2182](#) for more information.



Subordinate unit MAC addresses do not change. You can verify this by connecting to the subordinate unit CLI and using the `get hardware interface nic` command to display the MAC addresses of each FortiGate interface.

When the new primary unit is selected after a failover, the primary unit sends gratuitous ARP packets to update the devices connected to the cluster interfaces (usually layer-2 switches) with the virtual MAC address. Gratuitous ARP packets configure connected network devices to associate the cluster virtual MAC addresses and cluster IP address with primary unit physical interfaces and with the layer-2 switch physical interfaces. This is sometimes called using gratuitous ARP packets (sometimes called GARP packets) to train the network. The gratuitous ARP packets sent from the primary unit are intended to make sure that the layer-2 switch forwarding databases (FDBs) are updated as quickly as possible.

Sending gratuitous ARP packets is not required for routers and hosts on the network because the new primary unit will have the same MAC and IP addresses as the failed primary unit. However, since the new primary unit interfaces are connected to different switch interfaces than the failed primary unit, many network switches will update their FDBs more quickly after a failover if the new primary unit sends gratuitous ARP packets.

Changing how the primary unit sends gratuitous ARP packets after a failover

When a failover occurs it is important that the devices connected to the primary unit update their FDBs as quickly as possible to reestablish traffic forwarding.

Depending on your network configuration, you may be able to change the number of gratuitous ARP packets and the time interval between ARP packets to reduce the cluster failover time.

You cannot disable sending gratuitous ARP packets, but you can use the following command to change the number of packets that are sent. For example, enter the following command to send 20 gratuitous ARP packets:

```
config system ha
    set arps 20
end
```

You can use this command to configure the primary unit to send from 1 to 60 ARP packets. Usually you would not change the default setting of 5. In some cases, however, you might want to reduce the number of gratuitous ARP packets. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending a higher number gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully, you could reduce the number of gratuitous ARP packets that are sent to reduce the amount of traffic produced after a failover.

If failover is taking longer than expected, you may be able to reduce the failover time by increasing the number of gratuitous ARP packets sent.

You can also use the following command to change the time interval in seconds between gratuitous ARP packets. For example, enter the following command to change the time between ARP packets to 3 seconds:

```
config system ha
    set arps-interval 3
end
```


The time interval can be in the range of 1 to 20 seconds. The default is 8 seconds between gratuitous ARP packets. Normally you would not need to change the time interval. However, you could decrease the time to be able send more packets in less time if your cluster takes a long time to failover.

There may also be a number of reasons to set the interval higher. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully you could increase the interval to reduce the amount of traffic produced after a failover.

For more information about gratuitous ARP packets see [RFC 826](#) and [RFC 3927](#).

How the virtual MAC address is determined

The virtual MAC address is determined based on following formula:

00-09-0f-09-`<group-id_hex>`-`<vcluster_integer>``<idx>`

where

`<group-id_hex>` is the HA Group ID for the cluster converted to hexadecimal.

[Table 118](#) lists the virtual MAC address set for each group ID.

Table 118: HA group ID in integer and hexadecimal format

Integer Group ID	Hexadecimal Group ID
0	00
1	01
2	02
3	03
4	04
...	...
10	0a
11	0b
...	...
63	3f

`<vcluster_integer>` is 0 for virtual cluster 1 and 2 for virtual cluster 2. If virtual domains are not enabled, HA sets the virtual cluster to 1 and by default all interfaces are in the root virtual domain. Including virtual cluster and virtual domain factors in the virtual MAC address formula means that the same formula can be used whether or not virtual domains and virtual clustering is enabled.

`<idx>` is the index number of the interface. Interfaces are numbered from 0 to x (where x is the number of interfaces). Interfaces are numbered according to their has map order. See [“Interface index and display order” on page 2173](#). The first interface has an index of 0. The second interface in the list has an index of 1 and so on.



Only the `<idx>` part of the virtual MAC address is different for each interface. The `<vcluster_integer>` would be different for different interfaces if multiple VDOMs have been added.

Example virtual MAC addresses

An HA cluster with HA group ID unchanged (default=0) and virtual domains not enabled would have the following virtual MAC addresses for interfaces port1 to port12:

- port1 virtual MAC: 00-09-0f-09-00-00
- port10 virtual MAC: 00-09-0f-09-00-01
- port2 virtual MAC: 00-09-0f-09-00-02
- port3 virtual MAC: 00-09-0f-09-00-03
- port4 virtual MAC: 00-09-0f-09-00-04
- port5 virtual MAC: 00-09-0f-09-00-05
- port6 virtual MAC: 00-09-0f-09-00-06
- port7 virtual MAC: 00-09-0f-09-00-07
- port8 virtual MAC: 00-09-0f-09-00-08
- port9 virtual MAC: 00-09-0f-09-00-09
- port11 virtual MAC: 00-09-0f-09-00-0a
- port12 virtual MAC: 00-09-0f-09-00-0b

If the group ID is changed to 34 these virtual MAC addresses change to:

- port1 virtual MAC: 00-09-0f-09-22-00
- port10 virtual MAC: 00-09-0f-09-22-01
- port2 virtual MAC: 00-09-0f-09-22-02
- port3 virtual MAC: 00-09-0f-09-22-03
- port4 virtual MAC: 00-09-0f-09-22-04
- port5 virtual MAC: 00-09-0f-09-22-05
- port6 virtual MAC: 00-09-0f-09-22-06
- port7 virtual MAC: 00-09-0f-09-22-07
- port8 virtual MAC: 00-09-0f-09-22-08
- port9 virtual MAC: 00-09-0f-09-22-09
- port11 virtual MAC: 00-09-0f-09-22-0a
- port12 virtual MAC: 00-09-0f-09-22-0b

A cluster with virtual domains enabled where the HA group ID has been changed to 23, port5 and port 6 are in the root virtual domain (which is in virtual cluster1), and port7 and port8 are in the vdom_1 virtual domain (which is in virtual cluster 2) would have the following virtual MAC addresses:

port5 interface virtual MAC: 00-09-0f-09-23-05
port6 interface virtual MAC: 00-09-0f-09-23-06
port7 interface virtual MAC: 00-09-0f-09-23-27
port8 interface virtual MAC: 00-09-0f-09-23-28

Displaying the virtual MAC address

Every FortiGate unit physical interface has two MAC addresses: the current hardware address and the permanent hardware address. The permanent hardware address cannot be changed, it is the actual MAC address of the interface hardware. The current hardware address can be changed. The current hardware address is the address seen by the network. For a FortiGate unit not operating in HA, you can use the following command to change the current hardware address of the port1 interface:

```
config system interface
  edit port1
    set macaddr <mac_address>
  end
end
```

For an operating cluster, the current hardware address of each cluster unit interface is changed to the HA virtual MAC address by the FGCP. The `macaddr` option is not available for a functioning cluster. You cannot change an interface MAC address and you cannot view MAC addresses from the `system interface` CLI command.

You can use the `get hardware nic <interface_name_str>` command to display both MAC addresses for any FortiGate interface. This command displays hardware information for the specified interface. Depending on their hardware configuration, this command may display different information for different interfaces. You can use this command to display the current hardware address as `Current_HWaddr` and the permanent hardware address as `Permanent_HWaddr`. For some interfaces the current hardware address is displayed as `MAC`. The command displays a great deal of information about the interface so you may have to scroll the output to find the hardware addresses.



You can also use the `diagnose hardware deviceinfo nic <interface_str>` command to display both MAC addresses for any FortiGate interface.

Before HA configuration the current and permanent hardware addresses are the same. For example for one of the units in `Cluster_1`:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 02:09:0f:78:18:c9
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

During HA operation the current hardware address becomes the HA virtual MAC address, for example for the units in `Cluster_1`:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

The following command output for Cluster_2 shows the same current hardware address for port1 as for the internal interface of Cluster_2, indicating a MAC address conflict.

```
FG300A2904500238 # get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 00:09:0F:85:40:FD
.
.
.
```

Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain

A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID.

This section describes a topology that can result in packet loss, how to determine if packets are being lost, and how to correct the problem by changing the HA Group ID.



Packet loss on a network can also be caused by IP address conflicts. Finding and fixing IP address conflicts can be difficult. However, if you are experiencing packet loss and your network contains two FortiGate HA clusters you can use the information in this article to eliminate one possible source of packet loss.

Changing the HA group ID to avoid MAC address conflicts


Change the Group ID to change the virtual MAC address of all cluster interfaces. You can change the Group ID from the FortiGate CLI using the following command:

```
config system ha
    set group-id <id_integer>
end
```

Example topology

The topology below shows two clusters. The Cluster_1 internal interfaces and the Cluster_2 port 1 interfaces are both connected to the same broadcast domain. In this topology the broadcast domain could be an internal network. Both clusters could also be connected to the Internet or to different networks.

Figure 216: Example HA topology with possible MAC address conflicts



Ping testing for packet loss

If the network is experiencing packet loss, it is possible that you will not notice a problem unless you are constantly pinging both HA clusters. During normal operation of the network you also might not notice packet loss because the loss rate may not be severe enough to timeout TCP sessions. Also many common types of TCP traffic, such as web browsing, may not be greatly affected by packet loss. However, packet loss can have a significant effect on real time protocols that deliver audio and video data.

To test for packet loss you can set up two constant ping sessions, one to each cluster. If packet loss is occurring the two ping sessions should show alternating replies and timeouts from each cluster.

Cluster_1	Cluster_2
reply	timeout
reply	timeout
reply	timeout
timeout	reply
timeout	reply
reply	timeout
reply	timeout
timeout	reply
timeout	reply
timeout	reply
timeout	reply

Viewing MAC address conflicts on attached switches

If two HA clusters with the same virtual MAC address are connected to the same broadcast domain (L2 switch or hub), the MAC address will conflict and bounce between the two clusters. This example Cisco switch MAC address table shows the MAC address flapping between different interfaces (1/0/1 and 1/0/4).

```

1      0009.0f09.0002    DYNAMIC    Gi1/0/1
1      0009.0f09.0002    DYNAMIC    Gi1/0/4

```

Synchronizing the configuration

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit.

The following settings are not synchronized between cluster units:

- HA override.
- HA device priority.
- The virtual cluster priority.
- The FortiGate unit host name.
- The HA priority setting for a ping server (or dead gateway detection) configuration.
- The system interface settings of the HA reserved management interface.
- The HA default route for the reserved management interface, set using the `ha-mgt-interface-gateway` option of the `config system ha` command.

The primary unit synchronizes all other configuration settings, including the other HA configuration settings.

Disabling automatic configuration synchronization

In some cases you may want to use the following command to disable automatic synchronization of the primary unit configuration to all cluster units.

```

config system ha
    set sync-config disable
end

```

When this option is disabled the cluster no longer synchronizes configuration changes. If a device failure occurs, the new primary unit may not have the same configuration as the failed primary unit. As a result, the new primary unit may process sessions differently or may not function on the network in the same way.

In most cases you should not disable automatic configuration synchronization. However, if you have disabled this feature you can use the `execute ha synchronize` command to manually synchronize a subordinate unit's configuration to that of the primary unit.

You must enter `execute ha synchronize` commands from the subordinate unit that you want to synchronize with the primary unit. Use the `execute ha manage` command to access a subordinate unit CLI. See [“Viewing cluster status from the CLI” on page 2156](#).

For example, to access the first subordinate unit and force a synchronization at any time, even if automatic synchronization is disabled enter:

```

execute ha manage 0
execute ha synchronize start

```

You can use the following command to stop a synchronization that is in progress.

```

execute ha synchronize stop

```

You can use the following command to a synchronization all parts of the configuration:

```

execute ha synchronize all

```

Individual options are also available to synchronize parts of the configuration. For example, enter the following command to synchronize CA certificates:

```

execute ha synchronize ca

```

Incremental synchronization

When you log into the cluster web-based manager or CLI to make configuration changes, you are actually logging into the primary unit. All of your configuration changes are first made to the primary unit. Incremental synchronization then immediately synchronizes these changes to all of the subordinate units.

When you log into a subordinate unit CLI (for example using `execute ha manage`) all of the configuration changes that you make to the subordinate unit are also immediately synchronized to all cluster units, including the primary unit, using the same process.

Incremental synchronization also synchronizes other dynamic configuration information such as the DHCP server address lease database, routing table updates, IPsec SAs, MAC address tables, and so on. See [“FortiGate HA compatibility with PPPoE and DHCP” on page 2012](#) for more information about DHCP server address lease synchronization and [“Synchronizing routing table updates” on page 2191](#) for information about routing table updates.

Whenever a change is made to a cluster unit configuration, incremental synchronization sends the same configuration change to all other cluster units over the HA heartbeat link. An HA synchronization process running on the each cluster unit receives the configuration change and applies it to the cluster unit. The HA synchronization process makes the configuration change by entering a CLI command that appears to be entered by the administrator who made the configuration change in the first place.

Synchronization takes place silently, and no log messages are recorded about the synchronization activity. However, log messages can be recorded by the cluster units when the synchronization process enters CLI commands. You can see these log messages on the subordinate units if you enable event logging and set the minimum severity level to *Information* and then check the event log messages written by the cluster units when you make a configuration change. See [“Configuration change synchronized from primary unit to subordinate unit” on page 2140](#).

You can also see these log messages on the primary unit if you make configuration changes from a subordinate unit. See [“Configuration change synchronized from subordinate unit to primary unit” on page 2141](#).

Periodic synchronization

Incremental synchronization makes sure that as an administrator makes configuration changes, the configurations of all cluster units remain the same. However, a number of factors could cause one or more cluster units to go out of sync with the primary unit. For example, if you add a new unit to a functioning cluster, the configuration of this new unit will not match the configuration of the other cluster units. Its not practical to use incremental synchronization to change the configuration of the new unit.

Periodic synchronization is a mechanism that looks for synchronization problems and fixes them. Every minute the cluster compares the configuration file checksum of the primary unit with the configuration file checksums of each of the subordinate units. If all subordinate unit checksums are the same as the primary unit checksum, all cluster units are considered synchronized.

If one or more of the subordinate unit checksums is not the same as the primary unit checksum, the subordinate unit configuration is considered out of sync with the primary unit. The checksum of the out of sync subordinate unit is checked again every 15 seconds. This re-checking occurs in case the configurations are out of sync because an incremental configuration sequence has not completed. If the checksums do not match after 5 checks the subordinate unit that is out of sync retrieves the configuration from the primary unit. The subordinate unit then reloads its configuration and resumes operating as a subordinate unit with the same configuration as the primary unit.

The configuration of the subordinate unit is reset in this way because when a subordinate unit configuration gets out of sync with the primary unit configuration there is no efficient way to determine what the configuration differences are and to correct them. Resetting the subordinate unit configuration becomes the most efficient way to resynchronize the subordinate unit.

Synchronization requires that all cluster units run the same FortiOS firmware build. If some cluster units are running different firmware builds, then unstable cluster operation may occur and the cluster units may not be able to synchronize correctly.



Re-installing the firmware build running on the primary unit forces the primary unit to upgrade all cluster units to the same firmware build.

Console messages when configuration synchronization succeeds

When a cluster first forms, or when a new unit is added to a cluster as a subordinate unit, the following messages appear on the CLI console to indicate that the unit joined the cluster and had its configuring synchronized with the primary unit.

```
slave's configuration is not in sync with master's, sequence:0
slave's configuration is not in sync with master's, sequence:1
slave's configuration is not in sync with master's, sequence:2
slave's configuration is not in sync with master's, sequence:3
slave's configuration is not in sync with master's, sequence:4
slave starts to sync with master
logout all admin users
slave succeeded to sync with master
```

Console messages when configuration synchronization fails

If you connect to the console of a subordinate unit that is out of synchronization with the primary unit, messages similar to the following are displayed.

```
slave is not in sync with master, sequence:0. (type 0x3)
slave is not in sync with master, sequence:1. (type 0x3)
slave is not in sync with master, sequence:2. (type 0x3)
slave is not in sync with master, sequence:3. (type 0x3)
slave is not in sync with master, sequence:4. (type 0x3)
global compared not matched
```


If synchronization problems occur the console message sequence may be repeated over and over again. The messages all include a type value (in the example `type 0x3`). The type value can help Fortinet Support diagnose the synchronization problem.

Table 119: HA out of sync object messages and the configuration objects that they reference

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_CONFIGURATION = 0x03	/data/config
HA_SYNC_SETTING_AV = 0x10	
HA_SYNC_SETTING_VIR_DB = 0x11	/etc/vir
HA_SYNC_SETTING_SHARED_LIB = 0x12	/data/lib/libav.so
HA_SYNC_SETTING_SCAN_UNIT = 0x13	/bin/scanunitd
HA_SYNC_SETTING_IMAP_PRXY = 0x14	/bin/imapd
HA_SYNC_SETTING_SMTP_PRXY = 0x15	/bin/smtp
HA_SYNC_SETTING_POP3_PRXY = 0x16	/bin/pop3
HA_SYNC_SETTING_HTTP_PRXY = 0x17	/bin/thttp
HA_SYNC_SETTING_FTP_PRXY = 0x18	/bin/ftpd
HA_SYNC_SETTING_FCNI = 0x19	/etc/fcni.dat
HA_SYNC_SETTING_FDNI = 0x1a	/etc/fdnserver.dat
HA_SYNC_SETTING_FSCI = 0x1b	/etc/sci.dat
HA_SYNC_SETTING_FSAE = 0x1c	/etc/fsae_adgrp.cache
HA_SYNC_SETTING_IDS = 0x20	/etc/ids.rules
HA_SYNC_SETTING_IDSUSER_RULES = 0x21	/etc/idsuser.rules
HA_SYNC_SETTING_IDSCUSTOM = 0x22	
HA_SYNC_SETTING_IDS_MONITOR = 0x23	/bin/ipsmonitor
HA_SYNC_SETTING_IDS_SENSOR = 0x24	/bin/ipsengine
HA_SYNC_SETTING_NIDS_LIB = 0x25	/data/lib/libips.so
HA_SYNC_SETTING_WEBLISTS = 0x30	
HA_SYNC_SETTING_CONTENTFILTER = 0x31	/data/cmdb/webfilter.bword
HA_SYNC_SETTING_URLFILTER = 0x32	/data/cmdb/webfilter.urlfilter
HA_SYNC_SETTING_FSGD_OVRD = 0x33	/data/cmdb/webfilter.fgtd-ovrd
HA_SYNC_SETTING_FSGD_LRATING = 0x34	/data/cmdb/webfilter.fgtd-ovrd
HA_SYNC_SETTING_EMAILLISTS = 0x40	
HA_SYNC_SETTING_EMAILCONTENT = 0x41	/data/cmdb/spamfilter.bword

Table 119: HA out of sync object messages and the configuration objects that they reference (Continued)

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_EMAILBWLST = 0x42	/data/cmdb/spamfilter.emailbwl
HA_SYNC_SETTING_IPBWL = 0x43	/data/cmdb/spamfilter.ipbwl
HA_SYNC_SETTING_MHEADER = 0x44	/data/cmdb/spamfilter.mheader
HA_SYNC_SETTING_RBL = 0x45	/data/cmdb/spamfilter.rbl
HA_SYNC_SETTING_CERT_CONF = 0x50	/etc/cert/cert.conf
HA_SYNC_SETTING_CERT_CA = 0x51	/etc/cert/ca
HA_SYNC_SETTING_CERT_LOCAL = 0x52	/etc/cert/local
HA_SYNC_SETTING_CERT_CRL = 0x53	/etc/cert/crl
HA_SYNC_SETTING_DB_VER = 0x55	
HA_GET_DETAIL_CSUM = 0x71	
HA_SYNC_CC_SIG = 0x75	/etc/cc_sig.dat
HA_SYNC_CC_OP = 0x76	/etc/cc_op
HA_SYNC_CC_MAIN = 0x77	/etc/cc_main
HA_SYNC_FTGD_CAT_LIST = 0x7a	/migadmin/webfilter/ublock/ftgd/data/

Comparing checksums of cluster units

You can use the `diagnose sys ha showcsum` command to compare the configuration checksums of all cluster units. The output of this command shows checksums labelled `global` and `all` as well as checksums for each of the VDOMs including the `root` VDOM. The `get system ha-nonsync-csum` command can be used to display similar information; however, this command is intended to be used by FortiManager.

The primary unit and subordinate unit checksums should be the same. If they are not you can use the `execute ha synchronize` command to force a synchronization.

The following command output is for the primary unit of a cluster that does not have multiple VDOMs enabled:

```
diagnose sys ha showcsum
is_manage_master()=1, is_root_master()=1
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following command output is for a subordinate unit of the same cluster:

```
diagnose sys ha showcsum
is_manage_master()=0, is_root_master()=0
debugzone
```

```
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

```
checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following example shows using this command for the primary unit of a cluster with multiple VDOMs. Two VDOMs have been added named `test` and `Eng_vdm`.

From the primary unit:

```
config global
diagnose sys ha showcsum
is_manage_master()=1, is_root_master()=1
debugzone
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53
```

From the subordinate unit:

```
config global
diagnose sys ha showcsum
is_manage_master()=0, is_root_master()=0
debugzone
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

checksum
global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53
```

How to diagnose HA out of sync messages

This section describes how to use the commands `diagnose sys ha showcsum` and `diagnose debug` to diagnose the cause of HA out of sync messages.

If HA synchronization is not successful, use the following procedures on each cluster unit to find the cause.

To determine why HA synchronization does not occur

- 1 Connect to each cluster unit CLI by connected to the console port.
- 2 Enter the following commands to enable debugging and display HA out of sync messages.

```
diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application hataalk -1
diagnose debug application hasync -1
```

Collect the console output and compare the out of sync messages with the information in [Table 119 on page 2187](#).

- 3 Enter the following commands to turn off debugging.

```
diagnose debug disable
diagnose debug reset
```

To determine what part of the configuration is causing the problem

If the previous procedure displays messages that include sync object 0x30 (for example, `HA_SYNC_SETTING_CONFIGURATION = 0x03`) there is a synchronization problem with the configuration. Use the following steps to determine the part of the configuration that is causing the problem.

If your cluster consists of two cluster units, use this procedure to capture the configuration checksums for each unit. If your cluster consists of more than two cluster units, repeat this procedure for all cluster units that returned messages that include 0x30 sync object messages.

- 1 Connect to each cluster unit CLI by connected to the console port.
 - 2 Enter the following command to turn on terminal capture
- ```
diagnose debug enable
```
- 3 Enter the following command to stop HA synchronization.
- ```
execute ha sync stop
```
- 4 Enter the following command to display configuration checksums.
- ```
diagnose sys ha showcsum 1
```
- 5 Copy the output to a text file.
  - 6 Repeat for all affected units.
  - 7 Compare the text file from the primary unit with the text file from each cluster unit to find the checksums that do not match.

You can use a diff function to compare text files.

- 8 Repeat steps 4 to 7 for each checksum level:

```
diagnose sys ha showcsum 2
diagnose sys ha showcsum 3
diagnose sys ha showcsum 4
diagnose sys ha showcsum 5
diagnose sys ha showcsum 6
diagnose sys ha showcsum 7
diagnose sys ha showcsum 8
```

- 9 When the non-matching checksum is found, attempt to drill down further. This is possible for objects that have sub-components.

For example you can enter the following commands:

```
diagnose sys ha showcsum system.global
```

```
diagnose sys ha showcsum system.interface
```

Generally it is the first non-matching checksum in one of the levels that is the cause of the synchronization problem.

- 10** Attempt to can remove/change the part of the configuration that is causing the problem. You can do this by making configuration changes from the primary unit or subordinate unit CLI.

- 11** Enter the following commands to start HA configuration and stop debugging:

```
execute ha sync start
diagnose debug dis
diagnose debug reset
```

## Synchronizing routing table updates

In a functioning cluster, the primary unit keeps all subordinate unit routing tables up to date and synchronized with the primary unit. After a failover, because of these routing table updates the new primary unit does not have to populate its routing table before being able to route traffic. After a failover the new primary unit rebuilds its routing table, but having the synchronized routes already available means the table is rebuilt much faster than if no route information was available.

This section describes how clusters handle dynamic routing failover and also describes how to use CLI commands to control the timing of routing table updates of the subordinate unit routing tables from the primary unit.

### Configuring graceful restart for dynamic routing failover

When an HA failover occurs, neighbor routers will detect that the cluster has failed and remove it from the network until the routing topology stabilizes. During the time the routers may stop sending IP packets to the cluster and communications sessions that would normally be processed by the cluster may time out or be dropped. Also the new primary unit will not receive routing updates and so will not be able to build and maintain its routing database.

You can configure graceful restart (also called nonstop forwarding (NSF)) as described in [RFC3623](#) (Graceful OSPF Restart) to solve the problem of dynamic routing failover. If graceful restart is enabled on neighbor routers, they will keep sending packets to the cluster following the HA failover instead of removing it from the network. The neighboring routers assume that the cluster is experiencing a graceful restart.

After the failover, the new primary unit can continue to process communication sessions using the synchronized routing data received from the failed primary unit before the failover. This gives the new primary unit time to update its routing table after the failover.

You can use the following commands to enable graceful restart or NSF on Cisco routers:

```
router ospf 1
 log-adjacency-changes
 nsf ietf helper strict-lsa-checking
```

If the cluster is running BGP, use the following command to enable graceful restart for BGP:

```
config router bgp
 set graceful-restart enable
end
```

You can also add BGP neighbors and configure the cluster unit to notify these neighbors that it supports graceful restart.

```
config router bgp
 config neighbor
 edit <neighbor_address_Ipv4>
 set capability-graceful-restart enable
 end
 end
```

If the cluster is running OSPF, use the following command to enable graceful restart for OSPF:

```
config router ospf
 set restart-mode graceful-restart
end
```

To make sure the new primary unit keeps its synchronized routing data long enough to acquire new routing data, you should also increase the HA route time to live, route wait, and route hold values to 60 using the following CLI command:

```
config system ha
 set route-ttl 60
 set route-wait 60
 set route-hold 60
end
```

## Controlling how the FGCP synchronizes routing updates

You can use the following commands to control some of the timing settings that the FGCP uses when synchronizing routing updates from the primary unit to subordinate units and maintaining routes on the primary unit after a failover.

```
config system ha
 set route-hold <hold_integer>
 set route-ttl <ttl_integer>
 set route-wait <wait_integer>
end
```

## Change how long routes stay in a cluster unit routing table

Change the `route-ttl` time to control how long routes remain in a cluster unit routing table. The time to live range is 0 to 3600 seconds. The default time to live is 10 seconds.

The time to live controls how long routes remain active in a cluster unit routing table after the cluster unit becomes a primary unit. To maintain communication sessions after a cluster unit becomes a primary unit, routes remain active in the routing table for the route time to live while the new primary unit acquires new routes.

If `route-ttl` is set to 0 the primary unit must acquire all new routes before it can continue processing traffic. By default, `route-ttl` is set to 10 which may mean that only a few routes will remain in the routing table after a failover. Normally keeping `route-ttl` to 10 or reducing the value to 0 is acceptable because acquiring new routes usually occurs very quickly, especially if graceful restart is enabled, so only a minor delay is caused by acquiring new routes.

If the primary unit needs to acquire a very large number of routes, or if for other reasons, there is a delay in acquiring all routes, the primary unit may not be able to maintain all communication sessions.

You can increase the route time to live if you find that communication sessions are lost after a failover so that the primary unit can use synchronized routes that are already in the routing table, instead of waiting to acquire new routes.

## Change the time between routing updates

Change the `route-hold` time to change the time that the primary unit waits between sending routing table updates to subordinate units. The route hold range is 0 to 3600 seconds. The default route hold time is 10 seconds.

To avoid flooding routing table updates to subordinate units, set `route-hold` to a relatively long time to prevent subsequent updates from occurring too quickly. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Increasing the time between updates means that this data exchange will not have to happen so often.

The `route-hold` time should be coordinated with the `route-wait` time.

## Change the time the primary unit waits after receiving a routing update

Change the `route-wait` time to change how long the primary unit waits after receiving routing updates before sending the updates to the subordinate units. For quick routing table updates to occur, set `route-wait` to a relatively short time so that the primary unit does not hold routing table changes for too long before updating the subordinate units.

The `route-wait` range is 0 to 3600 seconds. The default `route-wait` is 0 seconds.

Normally, because the `route-wait` time is 0 seconds the primary unit sends routing table updates to the subordinate units every time its routing table changes.

Once a routing table update is sent, the primary unit waits the `route-hold` time before sending the next update.

Usually routing table updates are periodic and sporadic. Subordinate units should receive these changes as soon as possible so `route-wait` is set to 0 seconds. `route-hold` can be set to a relatively long time because normally the next route update would not occur for a while.

In some cases, routing table updates can occur in bursts. A large burst of routing table updates can occur if a router or a link on a network fails or changes. When a burst of routing table updates occurs, there is a potential that the primary unit could flood the subordinate units with routing table updates. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Setting `route-wait` to a longer time reduces the frequency of additional updates and prevents flooding of routing table updates from occurring.

# Synchronizing IPsec VPN SAs

The FGCP synchronizes IPsec security associations (SAs) between cluster members so that if a failover occurs, the cluster can resume IPsec sessions without having to establish new SAs. The result is improved failover performance because IPsec sessions are not interrupted to establish new SAs. Also, establishing a large number of SAs can reduce cluster performance.

The FGCP implements slightly different synchronization mechanisms for IKEv1 and IKEv2.

## Synchronizing SAs for IKEv1

When an SA is synchronized to the subordinate units, the sequence number is set to the maximum sequence number. After a failover, all inbound traffic that connects with the new primary unit and uses the SA will be accepted without needing to re-key. However, first outbound packet to use the SA causes the sequence number to overflow and so causes the new primary unit to re-key the SA.

Please note the following:

- The cluster synchronizes all IPsec SAs.
- IPsec SAs are not synchronized until the IKE process has finished synchronizing the ISAKMP SAs. This is required in for dialup tunnels since it is the synchronizing of the ISAKMP SA that creates the dialup tunnel.
- A dialup interface is created as soon as the phase1 is complete. This ensures that the when HA synchronizes phase1 information the dialup name is included.
- If the IKE process re-starts for any reason it deletes any dialup tunnels that exist. This forces the peer to re-key them.
- IPsec SA deletion happens immediately. Routes associated with a dialup tunnel that is being deleted are cleaned up synchronously as part of the delete, rather than waiting for the SA hard-expiry.
- The FGCP does not sync the IPsec tunnel MTU from the primary unit to the subordinate units. This means that after HA failover if the first packet received by the FortiGate unit arrives after the HA route has been deleted and before the new route is added and the packet is larger than the default MTU of 1024 then the FortiGate unit sends back an ICMP fragmentation required. However, as soon as routing is re-established then the MTU will be corrected and traffic will flow.

## Synchronizing SAs for IKEv2

Due to the way the IKEv2 protocol is designed the FGCP cannot use exactly the same solution that is used for synchronizing IKEv1 SAs, though it is similar.

For IKEv2, like IKEv1, the FGCP synchronizes IKE and ISAKMP SAs from the primary unit to the subordinate units. However, for IKEv2 the FGCP cannot actually use this IKE SA to send/receive IKE traffic because IKEv2 includes a sequence number in every IKE message and thus it would require synchronizing every message to the subordinate units to keep the sequence numbers on the subordinate units up to date.

After a failover when the new primary unit accepts incoming IKEv2 sessions, as in IKEv1, the primary unit uses the synchronized SA to decrypt the traffic before passing it through to its destination. For outgoing sessions, because the synchronized SA has an old sequence number, the primary unit negotiates a new SA. This is different from IKEv1 where the existing SA is re-keyed.

Normally for IKEv2 the new primary unit could just negotiate a CHILD\_SA using the synchronized SA. However, because the sequence numbers are not up-to-date, as noted above, the synchronized SA cannot be used and the primary unit must instead negotiate a whole new SA.

## Link failover

Link failover means that if a monitored interface fails, the cluster reorganizes to reestablish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

You configure monitored interfaces (also called interface monitoring or port monitoring) by selecting the interfaces to monitor as part of the cluster HA configuration.

You can monitor up to 16 interfaces. This limit only applies to FortiGate units with more than 16 physical interfaces. In a multiple VDOM configuration you can monitor up to 16 interfaces per virtual cluster.

The interfaces that you can monitor appear on the port monitor list. You can monitor all FortiGate interfaces including redundant interfaces and 802.3ad aggregate interfaces.



You cannot monitor the following types of interfaces (you cannot select the interfaces on the port monitor list):

- FortiGate interfaces that contain an internal switch.
- VLAN subinterfaces.
- IPsec VPN interfaces.
- Individual physical interfaces that have been added to a redundant or 802.3ad aggregate interface.
- FortiGate-5000 series backplane interfaces that have not been configured as network interfaces.

If you are configuring a virtual cluster you can create a different port monitor configuration for each virtual cluster. Usually for each virtual cluster you would monitor the interfaces that have been added to the virtual domains in each virtual cluster.



Wait until after the cluster is up and running to enable interface monitoring. You do not need to configure interface monitoring to get a cluster up and running and interface monitoring will cause failovers if for some reason during initial setup a monitored interface has become disconnected. You can always enable interface monitoring once you have verified that the cluster is connected and operating properly.



You should only monitor interfaces that are connected to networks, because a failover may occur if you monitor an unconnected interface.

#### To enable interface monitoring - web-based manager

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

- 1 Connect to the cluster web-based manager.
  - 2 Go to *System > Config > HA* and edit the primary unit (*Role* is *MASTER*).
  - 3 Select the *Port Monitor* check boxes for the *port1* and *port2* interfaces and select *OK*.
- The configuration change is synchronized to all cluster units.

#### To enable interface monitoring - CLI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

- 1 Connect to the cluster CLI.
  - 2 Enter the following command to enable interface monitoring for port1 and port2.
- ```
configure system ha
  set monitor port1 port2
end
```

The following example shows how to enable monitoring for the external, internal, and DMZ interfaces.

```
config system ha
  set monitor external internal dmz
end
```

With interface monitoring enabled, during cluster operation, the cluster monitors each cluster unit to determine if the monitored interfaces are operating and connected. Each cluster unit can detect a failure of its network interface hardware. Cluster units can also detect if its network interfaces are disconnected from the switch they should be connected to.



Cluster units cannot determine if the switch that its interfaces are connected to is still connected to the network. However, you can use remote IP monitoring to make sure that the cluster unit can connect to downstream network devices. See [“Remote link failover” on page 2200](#).

Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link between a network and the primary unit fails, to maintain communication with this network, the cluster must select a different primary unit; one that is still connected to the network. Unless another link failure has occurred, the new primary unit will have an active link to the network and will be able to maintain communication with it.

To support link failover, each cluster unit stores link state information for all monitored cluster units in a link state database. All cluster units keep this link state database up to date by sharing link state information with the other cluster units. If one of the monitored interfaces on one of the cluster units becomes disconnected or fails, this information is immediately shared with all cluster units.

If a monitored interface on the primary unit fails

If a monitored interface on the primary unit fails, the cluster renegotiates to select a new primary unit using the process described in [“Primary unit selection” on page 2000](#). Because the cluster unit with the failed monitored interface has the lowest monitor priority, a different cluster unit becomes the primary unit. The new primary unit should have fewer link failures.

After the failover, the cluster resumes and maintains communication sessions in the same way as for a device failure. See [“Device failover” on page 2169](#).

If a monitored interface on a subordinate unit fails

If a monitored interface on a subordinate unit fails, this information is shared with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster.

In an active-passive cluster after a subordinate unit link failover, the subordinate unit continues to function normally as a subordinate unit in the cluster.


In an active-active cluster after a subordinate unit link failure:

- The subordinate unit with the failed monitored interface can continue processing connections between functioning interfaces. However, the primary unit stops sending sessions to a subordinate unit that use any failed monitored interfaces on the subordinate unit.
- If session pickup is enabled, all sessions being processed by the subordinate unit failed interface that can be are failed over to other cluster units. Sessions that cannot be failed over are lost and have to be restarted.
- If session pickup is not enabled all sessions being processed by the subordinate unit failed interface are lost.

How link failover maintains traffic flow

Monitoring an interface means that the interface is connected to a high priority network. As a high priority network, the cluster should maintain traffic flow to and from the network, even if a link failure occurs. Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link that the primary unit has to a high priority network fails, to maintain traffic flow to and from this network, the cluster must select a different primary unit. This new primary unit should have an active link to the high priority network.

Figure 217: A link failure causes a cluster to select a new primary unit



If a monitored interface on the primary unit fails, the cluster renegotiates and selects the cluster unit with the highest monitor priority to become the new primary unit. The cluster unit with the highest monitor priority is the cluster unit with the most monitored interfaces connected to networks.

After a link failover, the primary unit processes all traffic and all subordinate units, even the cluster unit with the link failure, share session and link status. In addition all configuration changes, routes, and IPsec SAs are synchronized to the cluster unit with the link failure.

In an active-active cluster, the primary unit load balances traffic to all the units in the cluster. The cluster unit with the link failure can process connections between its functioning interfaces (for, example if the cluster has connections to an internal, external, and DMZ network, the cluster unit with the link failure can still process connections between the external and DMZ networks).

If a monitored interface on a subordinate unit fails, the subordinate unit shares this information with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster. In an active-active cluster, the subordinate unit can continue processing connections between functioning interfaces. The primary unit re-distributes traffic that was being processed by the failed interface of the subordinate unit to other cluster units. If session pickup is enabled, similar to a failover, some of these sessions continue while others must restart. See [“Session failover \(session pick-up\)” on page 2205](#).

Recovery after a link failover

If you find and correct the problem that caused a link failure (for example, re-connect a disconnected network cable) the cluster updates its link state database and the cluster unit continues to operate as a subordinate unit. In an active-active cluster the primary unit will begin load balancing sessions to the now reconnected interface.

If the `override` CLI keyword is enabled on this cluster unit and its device priority is set higher than the unit priority of other cluster units the cluster will renegotiate when the link failure is repaired and the cluster unit with the highest device priority becomes the primary unit.

Testing link failover

You can test link failure by disconnecting the network cable from a monitored interface of a cluster unit. If you disconnect a cable from a primary unit monitored interface the cluster should renegotiate and select one of the other cluster units as the primary unit. You can also verify that traffic received by the disconnected interface continues to be processed by the cluster after the failover.

If you disconnect a cable from a subordinate unit interface the cluster will not renegotiate.

Updating MAC forwarding tables when a link failover occurs

When a FortiGate HA cluster is operating and a monitored interface fails on the primary unit, the primary unit usually becomes a subordinate unit and another cluster unit becomes the primary unit. After a link failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables (also called arp tables) of the switches connected to the cluster. This is normal link failover operation (for more information, see [“Link failover” on page 2194](#)).

Some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur if the switch does not detect the failure and does not clear its MAC forwarding table.

To make sure the switch detects the failover and clears its MAC forwarding tables, you can use the following command to cause a cluster unit with a monitored interface link failure to shut down all of its interfaces (except the heartbeat interfaces) for one second after the failover occurs. Usually this means the interfaces of the former primary unit are shut down. When this happens the switch should be able to detect this failure and clear its MAC forwarding tables of the MAC addresses of the former primary unit. Since the new primary unit has sent or will send gratuitous ARP packets the switch can then update its MAC forwarding tables to for the new primary unit.

```
config system ha
    set link-failed-signal enable
end
```

Multiple link failures


Every time a monitored interface fails, the cluster repeats the processes described above. If multiple monitored interfaces fail on more than one cluster unit, the cluster continues to negotiate to select a primary unit that can provide the most network connections.

Example link failover scenarios

For the following examples, assume a cluster configuration consisting of two FortiGate units (FGT_1 and FGT_2) connected to three networks: internal using port2, external using port1, and DMZ using port3. In the HA configuration, the device priority of FGT_1 is set higher than the unit priority of FGT_2.

The cluster processes traffic flowing between the internal and external networks, between the internal and DMZ networks, and between the external and DMZ networks. If there are no link failures, FGT1 becomes the primary unit because it has the highest device priority.

Figure 218: Sample link failover scenario topology



Example: the port1 link on FGT_1 fails

If the port1 link on FGT_1 fails, FGT_2 becomes primary unit because it has fewer interfaces with a link failure. If the cluster is operating in active-active mode, the cluster load balances traffic between the internal network (port2) and the DMZ network (port3). Traffic between the Internet (port1) and the internal network (port2) and between the Internet (port1) and the DMZ network (port3) is processed by the primary unit only.

Example: port2 on FGT_1 and port1 on FGT_2 fail

If port2 on FGT_1 and port1 on FGT_2 fail, then FGT_1 becomes the primary unit. After both of these link failures, both cluster units have the same monitor priority. So the cluster unit with the highest device priority (FGT_1) becomes the primary unit.

Only traffic between the Internet (port1) and DMZ (port3) networks can pass through the cluster and the traffic is handled by the primary unit only. No load balancing will occur if the cluster is operating in active-active mode.

Subsecond failover

HA link failover supports subsecond failover (that is a failover time of less than one second). Subsecond failover is available for interfaces that can issue a link failure system call when the interface goes down. When an interface experiences a link failure and sends the link failure system call, the FGCP receives the system call and initiates a link failover.

For interfaces that do not support subsecond failover, port monitoring regularly polls the connection status of monitored interfaces. When a check finds that an interface has gone down, port monitoring causes a link failover. Subsecond failover results in a link failure being detected sooner because the system doesn't have to wait for the next poll to find out about the failure.

Subsecond failover requires interfaces that support sending the link failure system call. This functionality is available for:

- Interfaces with network processors (NP2, NP4, etc.)
- Interfaces with content processors (CP4, CP5, CP6, etc.)
- Interfaces in Fortinet Mezzanine Cards that include network and content processors (FMC-XD2, FMC-XG2, etc.)
- Accelerated interface modules (FortiGate-ASM-FB4, ADM-FB8, ADM-XB2, ADM-XD4, RTM-XD2 etc).
- Interfaces in security processor modules (FortiGate-ASM-CE4, ASM-XE2, etc)

Subsecond failover can accelerate HA failover to reduce the link failover time to less than one second under ideal conditions. Actual failover performance may vary depending on traffic patterns and network configuration. For example, some network devices may respond slowly to an HA failover.

No configuration changes are required to support subsecond failover. However, for best subsecond failover results, the recommended heartbeat interval is 100ms and the recommended lost heartbeat threshold is 5. (See [“Changing the heartbeat interval” on page 2176](#))

```
config system ha
    set hb-lost-threshold 5
    set hb-interval 1
end
```


For information about how to reduce failover times, see [“Failover performance” on page 2214](#).

Remote link failover

Remote link failover (also called remote IP monitoring) is similar to HA port monitoring and interface dead gateway detection. Port monitoring causes a cluster to failover if a monitored primary unit interface fails or is disconnected. Remote IP monitoring uses ping servers configured for FortiGate interfaces on the primary unit to test connectivity with IP addresses of network devices. Usually these would be IP addresses of network devices not directly connected to the cluster. For example, a downstream router. Remote IP monitoring causes a failover if one or more of these remote IP addresses does not respond to a ping server.

By being able to detect failures in network equipment not directly connected to the cluster, remote IP monitoring can be useful in a number of ways depending on your network configuration. For example, in a full mesh HA configuration, with remote IP monitoring, the cluster can detect failures in network equipment that is not directly connected to the cluster but that would interrupt traffic processed by the cluster if the equipment failed.

Figure 219: Example HA remote IP monitoring topology



In the simplified example topology shown in [Figure 219](#), the switch connected directly to the primary unit is operating normally but the link on the other side of the switches fails. As a result traffic can no longer flow between the primary unit and the Internet.

To detect this failure you can create a remote IP monitoring configuration consisting of a ping server dead gateway detection configuration for port2 of the cluster. The primary unit tests connectivity to 192.168.20.20. If the ping server cannot connect to 192.268.20.20 the cluster fails over and the subordinate unit becomes the new primary unit. The remote HA monitoring ping server on the new primary unit can connect to 192.168.20.20 so the failover maintains connectivity between the internal network and the Internet through the cluster.

To configure remote IP monitoring

- 1 Enter the following commands to configure HA remote monitoring for the example topology.
 - Enter the `pingserver-monitor-interface` keyword to enable HA remote IP monitoring on port2.
 - Leave the `pingserver-failover-threshold` set to the default value of 0. You can change this value if you do not want a failover to occur if only one ping server fails.
 - Enter the `pingserver-flip-timeout` keyword to set the flip timeout to 120 minutes. After a failover, if HA remote IP monitoring on the new primary unit also causes a failover, the flip timeout prevents the failover from occurring until the timer runs out. Setting the `pingserver-flip-timeout` to 120 means that remote IP monitoring can only cause a failover every 120 minutes. This flip timeout is required to prevent repeating failovers if remote IP monitoring causes a failover from all cluster units because none of the cluster units can connect to the monitored IP addresses.

```
config system ha
    set pingserver-monitor-interface port2
    set pingserver-failover-threshold 0
    set pingserver-flip-timeout 120
end
```

- 2 Enter the following commands to add the ping server for the port2 interface and to set the HA remote IP monitoring priority for this ping server.
 - Enter the `detectserver` keyword to add the ping server and set the ping server IP address to 192.168.20.20.
 - Leave the `ha-priority` keyword set to the default value of 1. You only need to change this priority if you change the HA ping server failover threshold. The `ha-priority` setting is not synchronized among cluster units.
 - Use the `interval` keyword to set the time between ping server pings and use the `failtime` keyword to set the number of times that the ping can fail before a failure is detected (the failover threshold). The following example reduces the failover threshold to 2 but keeps the ping interval at the default value of 5.

```
config router gwdetect
    edit port2
        set server 192.168.20.20
        set ha-priority 1
        set interval 5
        set failtime 2
    end
```



You can also do this from the web-based manager by going to *Router > Static > Settings*, selecting *Create New* to add a new dead gateway detection configuration, setting *Ping Server* to 192.168.20.20, *HA Priority* to 1, *Ping Interval* to 5, and *Failover Threshold* to 2.

Adding HA remote IP monitoring to multiple interfaces

You can enable HA remote IP monitoring on multiple interfaces by adding more interface names to the `pingserver-monitor-interface` keyword. If your FortiGate configuration includes VLAN interfaces, aggregate interfaces and other interface types, you can add the names of these interfaces to the `pingserver-monitor-interface` keyword to configure HA remote IP monitoring for these interfaces.

For example, enable remote IP monitoring for interfaces named port2, port20, and vlan_234:

```
config system ha
    set pingserver-monitor-interface port2 port20 vlan_234
    set pingserver-failover-threshold 10
    set pingserver-flip-timeout 120
end
```

Then configure ping servers for each of these interfaces. In the following example, default values are accepted for all settings other than the server IP address.

```
config router gwdetect
    edit port2
        set server 192.168.20.20
    next
    edit port20
        set server 192.168.20.30
    next
    edit vlan_234
        set server 172.20.12.10
    end
```

Changing the ping server failover threshold

By default the ping server failover threshold is 0 and the HA priority is 1 so any HA remote IP monitoring ping server failure causes a failover. If you have multiple ping servers you may want a failover to occur only if more than one of them has failed.

For example, you may have 3 ping servers configured on three interfaces but only want a failover to occur if two of the ping servers fail. To do this you must set the HA priorities of the ping servers and the HA ping server failover threshold so that the priority of one ping server is less than the failover threshold but the added priorities of two ping servers is equal to or greater than the failover threshold. Failover occurs when the HA priority of all failed ping servers reaches or exceeds the threshold.

For example, set the failover threshold to 10 and monitor three interfaces:

```
config system ha
    set pingserver-monitor-interface port2 port20 vlan_234
    set pingserver-failover-threshold 10
    set pingserver-flip-timeout 120
end
```

Then set the HA priority of each ping server to 5.

```
config router gwdetect
    edit port2
        set server 192.168.20.20
        set ha-priority 5
    next
    edit port20
        set server 192.168.20.30
        set ha-priority 5
    next
    edit vlan_234
        set server 172.20.12.10
        set ha-priority 5
    end
```

If only one of the ping servers fails, the total ping server HA priority will be 5, which is lower than the failover threshold so a failover will not occur. If a second ping server fails, the total ping server HA priority of 10 will equal the failover threshold, causing a failover.

By adding multiple ping servers to the remote HA monitoring configuration and setting the HA priorities for each, you can fine tune remote IP monitoring. For example, if it is more important to maintain connections to some networks you can set the HA priorities higher for these ping servers. And if it is less important to maintain connections to other networks you can set the HA priorities lower for these ping servers. You can also adjust the failover threshold so that if the cluster cannot connect to one or two high priority IP addresses a failover occurs. But a failover will not occur if the cluster cannot connect to one or two low priority IP addresses.

Monitoring multiple IP addresses from one interface

You can add multiple IP addresses to a single ping server to use HA remote IP monitoring to monitor more than one IP address from a single interface. If you add multiple IP addresses, the ping will be sent to all of the addresses at the same time. The ping server only fails when no responses are received from any of the addresses.

```
config router gwdetect
  edit port2
    set server 192.168.20.20 192.168.20.30 172.20.12.10
  end
```

Flip timeout

The HA remote IP monitoring configuration also involves setting a flip timeout. The flip timeout is required to reduce the frequency of failovers if, after a failover, HA remote IP monitoring on the new primary unit also causes a failover. This can happen if the new primary unit cannot connect to one or more of the monitored remote IP addresses. The result could be that until you fix the network problem that blocks connections to the remote IP addresses, the cluster will experience repeated failovers. You can control how often the failovers occur by setting the flip timeout. The flip timeout stops HA remote IP monitoring from causing a failover until the primary unit has been operating for the duration of the flip timeout.

If you set the flip timeout to a relatively high number of minutes you can find and repair the network problem that prevented the cluster from connecting to the remote IP address without the cluster experiencing very many failovers. Even if it takes a while to detect the problem, repeated failovers at relatively long time intervals do not usually disrupt network traffic.

Detecting HA remote IP monitoring failovers

Just as with any HA failover, you can detect HA remote IP monitoring failovers by using SNMP to monitor for HA traps. You can also use alert email to receive notifications of HA status changes and monitor log messages for HA failover log messages. In addition, FortiGate units send the critical log message `Ping Server is down` when a ping server fails. The log message includes the name of the interface that the ping server has been added to.

Session failover (session pick-up)

Session failover means that a cluster maintains active network TCP and IPsec VPN sessions after a device or link failover. Session failover does not failover UDP, multicast, ICMP, or SSL VPN sessions. In some cases UDP sessions may be maintained after a failover.

FortiGate HA does not support session failover by default. To enable session failover go to *System > Config > HA* and select *Enable Session Pick-up*.

From the CLI enter:

```
config system ha
    set session-pickup enable
end
```

To support session failover, when *Enable Session Pick-up* is selected, the FGCP maintains an HA session table for most TCP communication sessions being processed by the cluster and synchronizes this session table with all cluster units. If a cluster unit fails, the HA session table information is available to the remaining cluster units and these cluster units use this session table to resume most of the TCP sessions that were being processed by the failed cluster unit without interruption.

You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce HA CPU usage and reduce HA heartbeat network bandwidth usage.

If *Enable Session Pick-up* is not selected, the FGCP does not maintain an HA session table and most TCP sessions do not resume after a failover. After a device or link failover all sessions are briefly interrupted and must be restarted after the cluster renegotiates. Many protocols can successfully restart sessions without loss of data. Other protocols may experience data loss and some protocols may require sessions to be manually restarted.

Some sessions may resume after a failover whether or not enable session pick-up is selected:

- [“Session failover and UDP, ICMP, multicast and broadcast packets” on page 2208](#)
- [“FortiOS Carrier GTP session failover” on page 2209](#)
- [“Active-active HA subordinate units sessions can resume after a failover” on page 2209.](#)

Improving session synchronization performance

Two HA configuration options are available to reduce the performance impact of enabling session pickup. They include reducing the number of sessions that are synchronized by adding a session pickup delay and using more FortiGate interfaces for session synchronization.

Reducing the number of sessions that are synchronized

Enable the `session-pickup-delay` CLI option to reduce the number of sessions that are synchronized by synchronizing sessions only if they remain active for more than 30 seconds. Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
    set session-pickup-delay enable
```

```
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

Using multiple FortiGate interfaces for session synchronization

Using the `session-sync-dev` option you can select one or more FortiGate interfaces to use for synchronizing sessions as required for session pickup. Normally session synchronization occurs over the HA heartbeat link. Using this HA option means only the selected interfaces are used for session synchronization and not the HA heartbeat link. If you select more than one interface, session synchronization traffic is load balanced among the selected interfaces.

Moving sessions synchronization from the HA heartbeat interface reduces the bandwidth requirements of the HA heartbeat interface and may improve the efficiency and performance of the cluster, especially if the cluster is synchronizing a large number of sessions. Load balancing session synchronization among multiple interfaces can further improve performance and efficiency if the cluster is synchronizing a large number of sessions.

Use the following command to perform cluster session synchronization using the port10 and port12 interfaces.

```
config system ha
    set session-sync-dev port10 port12
end
```

Session synchronization packets use Ethertype 0x8892. The interfaces to use for session synchronization must be connected together either directly using the appropriate cable (possible if there are only two units in the cluster) or using switches. If one of the interfaces becomes disconnected the cluster uses the remaining interfaces for session synchronization. If all of the session synchronization interfaces become disconnected, session synchronization reverts back to using the HA heartbeat link. All session synchronization traffic is between the primary unit and each subordinate unit.

Since large amounts of session synchronization traffic can increase network congestion, it is recommended that you keep this traffic off of your network by using dedicated connections for it.

Session failover not supported for all sessions

Most of the features applied to sessions by FortiGate UTM functionality require the FortiGate unit to maintain very large amounts of internal state information for each session. The FGCP does not synchronize internal state information for the following UTM features, so the following types of sessions will not resume after a failover:

- Virus scanning of HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, and NNTP sessions,
- Web filtering and FortiGuard Web Filtering of HTTP and HTTPS sessions,
- Spam filtering of IMAP, IMAPS, POP3, POP3S, SMTP, and SMTPS sessions,
- DLP scanning of IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, SIP, SIMPLE, and SCCP sessions,

- DLP archiving of HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, SMTP, SMTPS, IM, NNTP, AIM, ICQ, MSN, Yahoo! IM, SIP, SIMPLE, and SCCP signal control sessions,



Active-active clusters can resume some of these sessions after a failover. See [“Active-active HA subordinate units sessions can resume after a failover”](#) on page 2209 for details.

If you use these features to protect most of the sessions that your cluster processes, enabling session failover may not actually provide significant session failover protection.

TCP sessions that are not being processed by these UTM features resume after a failover even if these sessions are accepted by security policies with UTM options configured.

Only TCP sessions that are actually being processed by these UTM features do not resume after a failover. For example:

- TCP sessions that are not virus scanned, web filtered, spam filtered, content archived, or are not SIP, SIMPLE, or SCCP signal traffic resume after a failover, even if they are accepted by a security policy with UTM options enabled. For example, SNMP TCP sessions resume after a failover because FortiOS does not apply any UTM options to SNMP sessions.
- TCP sessions for a protocol for which UTM features have not been enabled resume after a failover even if they are accepted by a security policy with UTM features enabled. For example, if you have not enabled any antivirus or content archiving settings for FTP, FTP sessions resume after a failover.

The following UTM features do not affect TCP session failover:

- IPS does not affect session failover. Sessions being scanned by IPS resume after a failover. After a failover; however, IPS can only perform packet-based inspection of resumed sessions; reducing the number of vulnerabilities that IPS can detect. This limitation only applies to in-progress resumed sessions.
- Application control does not affect session failover. Sessions that are being monitored by application control resume after a failover.
- Logging enabled from UTM features does not affect session failover. UTM logging writes event log messages for UTM events; such as when a virus is found by antivirus scanning, when Web Filtering blocks a URL, and so on. Logging does not enable features that would prevent sessions from being failed over, logging just reports on the activities of enabled features.

If more than one UTM feature is applied to a TCP session, that session will not resume after a failover as long as one of the UTM features prevents session failover. For example:

- Sessions being scanned by IPS and also being virus scanned do not resume after a failover.
- Sessions that are being monitored by application control and that are being DLP archived or virus scanned will not resume after a failover.

SIP session failover

The FGCP supports SIP session failover (also called stateful failover) for active-passive HA. To support SIP session failover, create a standard HA configuration and select *Enable Session Pick-up* option.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in-progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

Session failover and explicit web proxy, WCCP, and WAN optimization sessions

Similar to UTM sessions, the explicit web proxy, WCCP and WAN optimization features all require the FortiGate unit to maintain very large amounts of internal state information for each session. This information is not maintained and these sessions do not resume after a failover.

Session failover and SSL offloading and HTTP multiplexing

SSL offloading and HTTP multiplexing are both enabled from firewall virtual IPs and firewall load balancing. Similar to the features applied by UTM, SSL offloading and HTTP multiplexing requires the FortiGate unit to maintain very large amounts of internal state information for each session. Sessions accepted by security policies containing virtual IPs or virtual servers with SSL offloading or HTTP multiplexing enabled do not resume after a failover.

IPsec VPN and SSL VPN sessions

Session failover is supported for all IPsec VPN tunnels. To support IPsec VPN tunnel failover, when an IPsec VPN tunnel starts, the FGCP distributes the SA and related IPsec VPN tunnel data to all cluster units.

Session failover is not supported for SSL VPN tunnels.

PPTP and L2TP VPN sessions

PPTP and L2TP VPNs are supported in HA mode. For a cluster you can configure PPTP and L2TP settings and you can also add security policies to allow PPTP and L2TP pass through. However, the FGCP does not provide session failover for PPTP or L2TP. After a failover, all active PPTP and L2TP sessions are lost and must be restarted.

Session failover and UDP, ICMP, multicast and broadcast packets

The FGCP does not maintain a session table for UDP, ICMP, multicast, or broadcast packets. So the cluster does not specifically support failover of these packets.

Some UDP traffic can continue to flow through the cluster after a failover. This can happen if, after the failover, a UDP packet that is part of an already established communication stream matches a security policy. Then a new session will be created and traffic will flow. So after a short interruption, UDP sessions can appear to have failed over. However, this may not be reliable for the following reasons:

- UDP packets in the direction of the security policy must be received before reply packets can be accepted. For example, if a port1 -> port2 policy accepts UDP packets, UDP packets received at port2 destined for the network connected to port1 will not be accepted until the policy accepts UDP packets at port1 that are destined for the network connected to port2. So, if a user connects from an internal network to the Internet and starts receiving UDP packets from the Internet (for example streaming media), after a failover the user will not receive any more UDP packets until the user re-connects to the Internet site.

- UDP sessions accepted by NAT policies will not resume after a failover because NAT will usually give the new session a different source port. So only traffic for UDP protocols that can handle the source port changing during a session will continue to flow.

FortiOS Carrier GTP session failover

FortiOS Carrier HA supports GTP session failover. The primary unit synchronizes the GTP tunnel state to all cluster units after the GTP tunnel setup is completed. After the tunnel setup is completed, GTP sessions use UDP and HA does not synchronize UDP sessions to all cluster units. However, similar to other UDP sessions, after a failover, since the new primary unit will have the GTP tunnel state information, GTP UDP sessions using the same tunnel can continue to flow with some limitations.

The limitation on packets continuing to flow is that there has to be a security policy to accept the packets. For example, if the FortiOS Carrier unit has an internal to external security policy, GTP UDP sessions using an established tunnel that are received by the internal interface are accepted by the security policy and can continue to flow. However, GTP UDP packets for an established tunnel that are received at the external interface cannot flow until packets from the same tunnel are received at the internal interface.

If you have bi-directional policies that accept GTP UDP sessions then traffic in either direction that uses an established tunnel can continue to flow after a failover without interruption.

Active-active HA subordinate units sessions can resume after a failover

In an active-active cluster, subordinate units process sessions. After a failover, all cluster units that are still operating may be able to continue processing the sessions that they were processing before the failover. These sessions are maintained because after the failover the new primary unit uses the HA session table to continue to send session packets to the cluster units that were processing the sessions before the failover. Cluster units maintain their own information about the sessions that they are processing and this information is not affected by the failover. In this way, the cluster units that are still operating can continue processing their own sessions without loss of data.

The cluster keeps processing as many sessions as it can. But some sessions can be lost. Depending on what caused the failover, sessions can be lost in the following ways:

- A cluster unit fails (the primary unit or a subordinate unit). All sessions that were being processed by that cluster unit are lost.
- A link failure occurs. All sessions that were being processed through the network interface that failed are lost.

This mechanism for continuing sessions is not the same as session failover because:

- Only the sessions that can be are maintained.
- The sessions are maintained on the same cluster units and not re-distributed.
- Sessions that cannot be maintained are lost.

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended HA configuration for WAN optimization is active-passive mode. Also, when the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions. HA also does not support WAN optimization session failover.

In a cluster, the primary unit only stores web cache and byte cache databases. These databases are not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate units that it is participating with in WAN optimization tunnels.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time experienced by your network users may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Monitoring cluster units for failover

You can use logging and SNMP to monitor cluster units for failover. Both the primary and subordinate units can be configured to write log messages and send SNMP traps if a failover occurs. You can also log into the cluster web-based manager and CLI to determine if a failover has occurred. See [“Monitoring cluster units for failover” on page 2155](#).

NAT/Route mode active-passive cluster packet flow

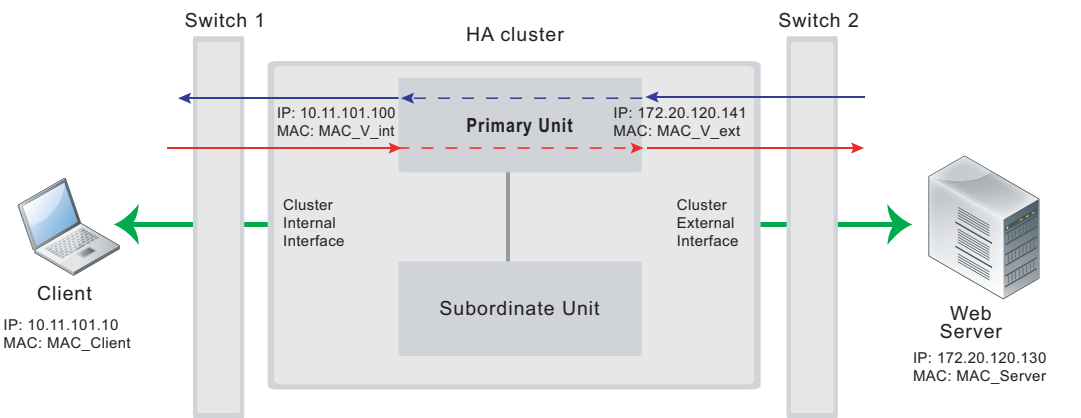
This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In an active-passive cluster operating in NAT/Route mode, four MAC addresses are involved in communication between the client and the web server when the primary unit processes the connection:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only know the cluster external virtual MAC address (MAC_V_ext). Cluster virtual MAC addresses are described in “Cluster virtual MAC addresses” on page 2177.

Figure 220: NAT/Route mode active-passive packet flow



Packet flow from client to web server

- 1 The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
- 2 The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
- 3 The client computer issues an ARP request to 10.11.101.100.
- 4 The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
- 5 The client’s request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

- 6 The primary unit processes the packet.
- 7 The primary unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_V_ext
Destination	172.20.120.130	MAC_Server

- 8 The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from web server to client

- 1 When the web server responds to the client’s packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
- 2 The web server issues an ARP request to 172.20.120.141.
- 3 The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.

- 4 The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

- 5 The primary unit processes the packet.
- 6 The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_V_int
Destination	10.11.101.10	MAC_Client

- 7 The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- 1 If the primary unit fails the subordinate unit becomes the primary unit.
- 2 The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
- 3 The new primary unit sends gratuitous ARP packets from the internal interface to the 10.11.101.0 network to associate its internal IP address with the internal virtual MAC address.
- 4 The new primary unit sends gratuitous ARP packets to the 172.20.120.0 to associate its external IP address with the external virtual MAC address.
- 5 Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Transparent mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.

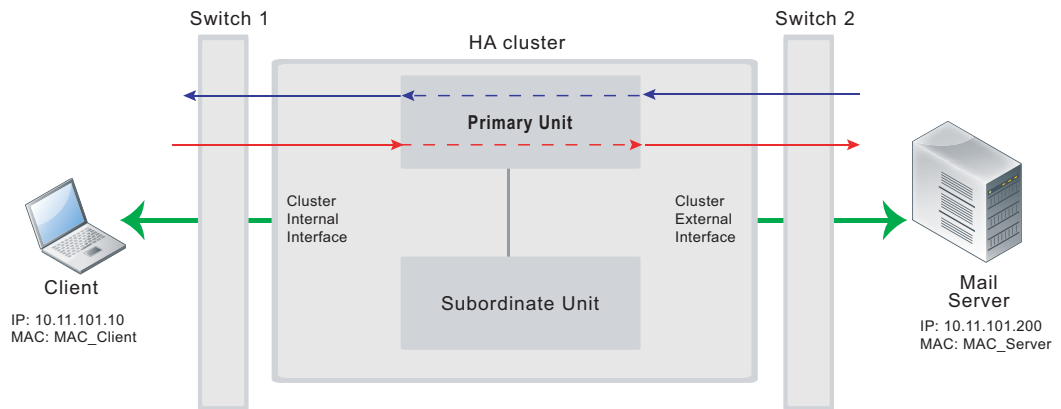
In an active-passive cluster operating in Transparent mode, two MAC addresses are involved in the communication between a client and a server when the primary unit processes a connection:

- Client MAC address (MAC_Client)
- Server MAC address (MAC_Server)

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and processed by the cluster.

The cluster's presence on the network is transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the HA virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the HA virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Figure 221: Transparent mode active-passive packet flow



Packet flow from client to mail server

- 1 The client computer requests a connection from 10.11.101.10 to 110.11.101.200.
- 2 The client computer issues an ARP request to 10.11.101.200.
- 3 The primary unit forwards the ARP request to the mail server.
- 4 The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
- 5 The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- 6 The primary unit processes the packet.
- 7 The primary unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- 8 The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from mail server to client

- 1 To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
- 2 The primary unit forwards the ARP request to the client computer.

- 3 The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 192.168.20.10. The primary unit returns the ARP response to the mail server.
- 4 The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- 5 The primary unit processes the packet.
- 6 The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- 7 The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- 1 If the primary unit fails, the subordinate unit negotiates to become the primary unit.
- 2 The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
- 3 The new primary unit sends gratuitous ARP packets to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
- 4 The new primary unit sends gratuitous ARP packets to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
- 5 Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Failover performance

This section describes the designed device and link failover times for a FortiGate cluster and also shows results of a failover performance test.

Device failover performance

By design FGCP device failover time is 2 seconds for a two-member cluster with ideal network and traffic conditions. If subsecond failover is enabled the failover time can drop below 1 second.

All cluster units regularly receive HA heartbeat packets from all other cluster units over the HA heartbeat link. If any cluster unit does not receive a heartbeat packet from any other cluster unit for 2 seconds, the cluster unit that has not sent heartbeat packets is considered to have failed.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions. Typically if subsecond failover is not enabled you can expect a failover time of 9 to 15 seconds depending on the cluster and network configuration. The failover time can also be increased by more complex configurations and or configurations with network equipment that is slow to respond.

You can change the `hb-lost-threshold` to increase or decrease the device failover time. See “[Modifying heartbeat timing](#)” on page 2175 for information about using `hb-lost-threshold`, and other heartbeat timing settings.

Link failover performance

Link failover time is controlled by how long it takes for a cluster to synchronize the cluster link database. When a link failure occurs, the cluster unit that experienced the link failure uses HA heartbeat packets to broadcast the updated link database to all cluster units. When all cluster units have received the updated database the failover is complete.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions.

Reducing failover times

You can do the following to help reduce failover times:

- Keep the network configuration as simple as possible with as few as possible network connections to the cluster.
- If possible operate the cluster in Transparent mode.
- Use high-performance switches to that the switches failover to interfaces connected to the new primary unit as quickly as possible.
- Use accelerated FortiGate interfaces. In some cases accelerated interfaces will reduce failover times.
- Make sure the FortiGate unit sends multiple gratuitous arp packets after a failover. In some cases, sending more gratuitous arp packets will cause connected network equipment to recognize the failover sooner. To send 10 gratuitous arp packets:

```
config system ha
    set arps 10
end
```

- Reduce the time between gratuitous arp packets. This may also caused connected network equipment to recognize the failover sooner. To send 50 gratuitous arp packets with 1 second between each packet:

```
config system ha
    set arp 50
    set arps-interval 1
end
```

- Reduce the number of lost heartbeat packets and reduce the heartbeat interval timers to be able to more quickly detect a device failure. To set the lost heartbeat threshold to 3 packets and the heartbeat interval to 100 milliseconds:

```
config system ha
    set hb-interval 3
    set hb-lost-threshold 1
end
```

- Reduce the hello state hold down time to reduce the amount of the time the cluster waits before transitioning from the hello to the work state. To set the hello state hold down time to 5 seconds:

```
config system ha
  set hello-holddown 5
end
```

- Enable sending a link failed signal after a link failover to make sure that attached network equipment responds as quickly as possible to a link failure. To enable the link failed signal:

```
config system ha
  set link-failed-signal enable
end
```



HA and load balancing

FGCP active-active load balancing distributes network traffic among all of the units in a cluster. Load balancing can improve cluster performance because the processing load is shared among multiple cluster units.

This chapter describes how active-active load balancing works and provides detailed NAT/Route and Transparent mode packet flow descriptions.

This chapter contains the following sections:

- [Load balancing overview](#)
- [Configuring load balancing settings](#)
- [NAT/Route mode active-active cluster packet flow](#)
- [Transparent mode active-active cluster packet flow](#)

Load balancing overview

In active-active HA, the FGCP uses unicast load balancing in which the primary unit is associated with the cluster HA virtual MAC address and cluster IP address. The primary unit is the only cluster unit to receive packets sent to the cluster.

An active-active HA cluster consists of a primary unit that processes communication sessions and one or more subordinate units that also process communication sessions. The primary unit receives all sessions and load balances sessions for security policies with UTM enabled to all cluster units. Because processing UTM sessions can be CPU and memory-intensive, load balancing UTM traffic may result in an active-active cluster having higher throughput than an active-passive cluster or a standalone FortiGate unit because resource-intensive UTM processing is distributed among all cluster units.

You can also enable the `load-balance-all` CLI keyword to have the primary unit load balance all TCP sessions. Load balancing TCP sessions is less likely to improve throughput because of extra overhead required for load balancing. So `load-balance-all` is disabled by default.

During active-active HA load balancing operation, when the primary unit receives the first packet of a UTM session (or a TCP session if `load-balance-all` is enabled) the primary unit uses the configured load balancing schedule to determine the cluster unit that will process the session. The primary unit stores the load balancing information for each active load balanced session in the cluster load balancing session table. Using the information in this table, the primary unit can then forward all of the remaining packets in each session to the appropriate cluster unit. The load balancing session table is synchronized among all cluster units.

UDP, ICMP, multicast, and broadcast sessions are never load balanced and are always processed by the primary unit. VoIP, IM, P2P, IPsec VPN, HTTPS, SSL VPN, HTTP multiplexing, SSL offloading, WAN optimization, explicit web proxy, and WCCP sessions are also always processed only by the primary unit.

In addition to load balancing, active-active HA also provides device and link failover protection similar to active-passive HA. If the primary unit fails, a subordinate unit becomes the primary unit and resumes operating the cluster. See [“Device failover” on page 2169](#) and [“Link failover” on page 2194](#) for more information.

Active-active HA provides session failover protection for all TCP sessions except UTM sessions. Active-active HA does not provide session failover for UTM sessions. Active-active HA also does not provide session failover for UDP, ICMP, multicast, and broadcast sessions. Protection profile sessions and all UDP, ICMP, multicast, and broadcast sessions are not failed over and must be restarted.

If a subordinate unit fails, the primary unit redistributes all TCP communications sessions among the remaining cluster units. Protection profile sessions that are in progress on the subordinate unit that failed are not failed over and must be restarted. All sessions being processed by the primary unit, including UDP, ICMP, multicast, and broadcast sessions, are not affected.

Because of the limitation of not supporting failover of UDP, ICMP, multicast, and broadcast sessions, active-active HA can be a less robust session failover solution than active-passive HA. See [“Session failover \(session pick-up\)” on page 2205](#) for more information about FortiGate session failover and its limitations.

Active-active HA does maintain as many UTM sessions as possible after a failover by continuing to process the UTM sessions that were being processed by the cluster units that are still operating. See [“Active-active HA subordinate units sessions can resume after a failover” on page 2209](#) for more information. Active-passive HA does not support maintaining UTM sessions after a failover.

Load balancing schedules

The load balancing schedule controls how the primary unit distributes packets to all cluster units. You can select from the following load balancing schedules.

None	No load balancing. Select <i>None</i> when the cluster interfaces are connected to load balancing switches. If you select <i>None</i> , the Primary unit does not load balance traffic and the subordinate units process incoming traffic that does not come from the Primary unit. For all other load balancing schedules, all traffic is received first by the Primary unit, and then forwarded to the subordinate units. The subordinate units only receive and process packets sent from the primary unit.
Hub	Load balancing if the cluster interfaces are connected to a hub. Traffic is distributed to cluster units based on the source IP and destination IP of the packet.
Least-Connection	If the cluster units are connected using switches, select <i>Least Connection</i> to distribute network traffic to the cluster unit currently processing the fewest connections.
Round-Robin	If the cluster units are connected using switches, select Round-Robin to distribute network traffic to the next available cluster unit.
Weighted Round-Robin	Similar to round robin, but weighted values are assigned to each of the units in a cluster based on their capacity and on how many connections they are currently processing. For example, the primary unit should have a lower weighted value because it handles scheduling and forwards traffic. Weighted round robin distributes traffic more evenly because units that are not processing traffic will be more likely to receive new connections than units that are very busy.

Random	If the cluster units are connected using switches, select Random to randomly distribute traffic to cluster units.
IP	Load balancing according to IP address. If the cluster units are connected using switches, select <i>IP</i> to distribute traffic to units in a cluster based on the source IP and destination IP of the packet.
IP Port	Load balancing according to IP address and port. If the cluster units are connected using switches, select IP Port to distribute traffic to units in a cluster based on the source IP, source port, destination IP, and destination port of the packet.

Once a packet has been propagated to a subordinate unit, all packets are part of that same communication session are also propagated to that same subordinate unit. Traffic is distributed according to communication session, not just according to individual packet.

Any subordinate unit that receives a forwarded packet processes it, without applying load balancing. Note that subordinate units are still considered to be active, because they perform routing, virus scanning, and other FortiGate unit tasks on their share of the traffic. Active subordinate units also share their session and link status information with all cluster units. The only things that active members do not do is make load balancing decisions.

Even though the primary unit is responsible for the load balancing process, the primary unit still acts like a FortiGate unit in that it processes packets, performing, routing, firewall, virus scanning, and other FortiGate unit tasks on its share of the traffic. Depending on the load balancing schedule used, the primary unit may assign itself a smaller share of the total load.

Selecting which packets are load balanced

The primary unit processes all UDP and ICMP traffic. By default, the primary unit also processes all TCP traffic and load balances virus scanning traffic among all cluster units. You can change the default configuration so that the cluster load balances both TCP traffic and virus scanning traffic among all cluster units.

Load balancing increases network bandwidth usage and also increases the load on the primary unit CPU. Because of this, in some network environments, load balancing TCP traffic may not result in an overall cluster performance increase. However, in other network environments, TCP load balancing may improve cluster performance.

If the cluster is configured to load balance virus scanning sessions, the primary unit uses the load balancing schedule to distribute HTTP, FTP, SMTP, POP3, and IMAP packets to be virus scanned, among the primary unit and the subordinate units. Load balancing virus scanning traffic is much more likely to increase cluster performance. Virus scanning is processor intensive for the cluster unit that is performing the virus scanning. Distributing virus scanning over the cluster units significantly reduces the processing load on the primary unit. As a result overall cluster performance should improve. See [“Load balancing UTM sessions and TCP sessions” on page 2221](#).

More about active-active failover

If a subordinate unit fails, the primary unit re-distributes the connections that the subordinate unit was processing among the remaining active cluster members. If the primary unit fails, the subordinate units negotiate to select a new primary unit. The new primary unit continues to distribute packets among the remaining active cluster units.

Failover works in a similar way if the cluster consists of only two units. If the primary unit fails the subordinate unit negotiates and becomes the new primary unit. If the subordinate unit fails, the primary unit processes all traffic. In both cases, the single remaining unit continues to function as a primary unit, maintaining the HA virtual MAC address for all of its interfaces.

HTTPS sessions, active-active load balancing, and proxy servers

To prevent HTTPS web filtering problems active-active HA does not load balance HTTPS sessions. The FortiGate unit identifies HTTPS sessions as all sessions received on the HTTPS TCP port. The default HTTPS port is 443. You can use the CLI command `config antivirus service` to configure the FortiGate unit to use a custom port for HTTPS sessions. If you change the HTTPS port using this CLI command, the FGCP stops load balancing all sessions that use the custom HTTPS port.

Normally you would not change the HTTPS port. However, if your network uses a proxy server for HTTPS traffic you may have to use the `config antivirus service` command to configure your cluster to use a custom HTTPS port. If your network uses a proxy server you might also use the same port for both HTTP and HTTPS traffic. In this case you would use `config antivirus service` to configure the FortiGate unit to use custom ports for both HTTP and HTTPS traffic.

Using the same port for HTTP and HTTPS traffic can cause problems with active-active clusters because active-active clusters always load balance HTTP traffic. If both HTTP and HTTPS use the same port, the active-active cluster cannot tell the difference between HTTP and HTTPS traffic and will load balance both HTTP and HTTPS traffic.

As mentioned above, load balancing HTTPS traffic may cause problems with HTTPS web filtering. To avoid this problem, you should configure your proxy server to use different ports for HTTP and HTTPS traffic. Then use the `config antivirus service` command to configure your cluster to also use different ports for HTTP and HTTPS.

Using FortiGate network processor interfaces to accelerate active-active HA performance

Many FortiGate models and FortiGate AMC modules include network processors that can provide hardware acceleration for active-active HA load balancing by offloading load balancing from the primary unit CPU. HA load balancing can be accelerated by NP1 network processors (called FA2 interfaces), NP2 network processors and NP4 network processors.

In some cases, performance of the primary unit can be reduced by active-active HA load balancing. Primary unit CPU cycles and bus bandwidth are required to receive, calculate load balancing schedules, and send balanced packets to the subordinate units. In very busy active-active clusters the primary unit may not be able to keep up with the processing load. This can result in lost traffic and can also cause the primary unit to delay sending heartbeat packets possibly reducing the stability and reliability of the active-active HA cluster.

Adding network processors to busy cluster unit interfaces increases load balancing performance by offloading load balancing to the network processors. The first packet of every new session is received by the primary unit and the primary unit uses its load balancing schedule to select the cluster unit that will process the new session. This information is passed back to the network processor and all subsequent packets of the same sessions are received by the primary unit interface network processor which sends the packet directly to a subordinate unit without using the primary unit CPU. Load balancing is effectively offloaded from the primary unit to the network processor resulting in a faster and more stable active-active cluster.

Using network processors to accelerate load balancing is especially useful if the `load-balance-all` option is enabled and the cluster is load balancing all TCP sessions because this could mean that the cluster is load balancing an excessive number of sessions.

To take advantage of network processor load balancing acceleration, connect the cluster unit interfaces with network processors to the busiest networks. Connect non-accelerated interfaces to less busy networks. No special FortiOS or HA configuration is required. Network processor acceleration of active-active HA load balancing is supported for any active-active HA configuration or active-active HA load balancing schedule.

Configuring load balancing settings

This section describes how to configure the following load balancing settings:

- [Selecting a load balancing schedule](#)
- [Load balancing UTM sessions and TCP sessions](#)
- [Configuring weighted-round-robin weights](#)

Selecting a load balancing schedule

You can select the load balancing schedule when initially configuring the cluster and you can change the load balancing schedule at any time while the cluster is operating without affecting cluster operation.

You can select a load balancing schedule from the CLI. Use the following command to select a load balancing schedule:

```
config system ha
    set schedule {hub | ip | ipport | leastconnection | none |
        random | round-robin | weight-round-robin}
end
```

Load balancing UTM sessions and TCP sessions

By default a FortiGate active-active cluster load balances UTM sessions among all cluster units. UTM processing applies protocol recognition, virus scanning, IPS, web filtering, email filtering, data leak prevention (DLP), application control, and VoIP content scanning and protection to HTTP, HTTPS, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, NNTP, SIP, SIMPLE, and SCCP sessions accepted by security policies. By load balancing this resource-intensive UTM processing among all cluster units, an active-active HA cluster may provide better UTM performance than a standalone FortiGate unit. Other features enabled in security policies such as Endpoint security, traffic shaping and authentication (identity-based policies) have no effect active-active load balancing.

All other sessions are processed by the primary unit. Using the CLI, you can configure the cluster to load balance TCP sessions among all cluster units in addition to UTM sessions. All UDP, ICMP, multicast, and broadcast sessions are never load balanced, but are always processed by the primary unit.

Use the following command to enable load balancing UTM and TCP sessions.

```
config system ha
    set load-balance-all enable
end
```

Enabling `load-balance-all` to load balance TCP sessions may not improve throughput because the cluster requires additional overhead to load balance sessions. The primary unit receives all sessions and load balances some TCP sessions to the subordinate units. Load balancing UTM sessions can improve performance because UTM session performance is limited by CPU performance. However, load balancing a non-UTM session usually requires about as much overhead as just processing it.

If your active-active cluster is processing TCP sessions and not performing UTM, you can enable `load-balance-all` and monitor network performance to see if it improves. If performance is not improved, you should change the HA mode to active-passive since active-active HA is not providing any benefit.

Configuring weighted-round-robin weights

You can configure weighted round-robin load balancing for a cluster and configure the static weights for each of the cluster units according to their priority in the cluster. When you set `schedule` to `weight-round-robin` you can use the `weight` option to set the static weight of each cluster unit. The static weight is set according to the priority of each unit in the cluster. A FortiGate HA cluster can contain up to four FortiGate units so you can set up to 4 static weights.

The priority of a cluster unit is determined by its device priority, the number of monitored interfaces that are functioning, its age in the cluster and its serial number. Priorities are used to select a primary unit and to set an order of all of the subordinate units. Thus the priority order of a cluster unit can change depending on configuration settings, link failures and so on. Since weights are also set using this priority order the weights are independent of specific cluster units but do depend on the role of the each unit in the cluster.

You can use the following command to display the priority order of units in a cluster. The following example displays the priority order for a cluster of 5 FortiGate-620B units:

```
get system ha status
  Model: 620
  Mode: a-p
  Group: 0
  Debug: 0
  ses_pickup: disable
  Master:150 head_office_cla FG600B3908600825 0
  Slave :150 head_office_clb FG600B3908600705 1
  Slave :150 head_office_clc FG600B3908600702 2
  Slave :150 head_office_cld FG600B3908600605 3
  Slave :150 head_office_cle FG600B3908600309 4
  number of vcluster: 1
  vcluster 1: work 169.254.0.1
  Master:0 FG600B3908600825
  Slave :1 FG600B3908600705
  Slave :2 FG600B3908600702
  Slave :3 FG600B3908600605
  Slave :4 FG600B3908600309
```

The cluster units are listed in priority order starting at the 6th output line. The primary unit always has the highest priority and is listed first followed by the subordinate units in priority order. The last 5 output lines list the cluster units in vcluster 1 and are not always in priority order. For more information about the `get system ha status` command, see [“Viewing cluster status from the CLI” on page 2156](#).

The default static weight for each cluster unit is 40. This means that sessions are distributed evenly among all cluster units. You can use the `set weight` command to change the static weights of cluster units to distribute sessions to cluster units depending on their priority in the cluster. The weight can be between 0 and 255. Increase the weight to increase the number of connections processed by the cluster unit with that priority.

You set the weight for each unit separately. For the example cluster of 5 FortiGate-620B units you can set the weight for each unit as follows:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set weight 0 5
  set weight 1 10
  set weight 2 15
  set weight 3 20
  set weight 4 30
end
```

If you enter the `get` command to view the HA configuration the output for `weight` would be:

```
weight 5 10 15 20 30 40 40 40 40 40 40 40 40 40 40
```

This configuration has the following results if the output of the `get system ha status` command is that shown above:

- The first five connections are processed by the primary unit (host name head_office_cla, priority 0, weight 5). From the output of the
- The next 10 connections are processed by the first subordinate unit (host name head_office_clb, priority 1, weight 10)
- The next 15 connections are processed by the second subordinate unit (host name head_office_clc, priority 2, weight 15)
- The next 20 connections are processed by the third subordinate unit (host name head_office_cld, priority 3, weight 20)
- The next 30 connections are processed by the fourth subordinate unit (host name head_office_cle, priority 4, weight 30)

Dynamically optimizing weighted load balancing according to how busy cluster units are

In conjunction with using static weights to load balance sessions among cluster units you can configure a cluster to load balance sessions according to individual cluster unit CPU usage, memory usage, and number of UTM proxy sessions. If any of these system loading indicators increases above configured thresholds, weighted load balancing sends fewer new sessions to the busy unit until it recovers.

High CPU or memory usage indicates that a unit is under increased load and may not be able to process more sessions. UTM proxy use is also a good indicator of how busy a cluster unit is since processing high numbers of UTM proxy sessions can quickly reduce overall cluster unit performance.

For example, if you set a CPU usage high watermark, when a cluster unit's CPU usage reaches the high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to a low watermark threshold. When this happens the cluster resumes load balancing sessions to the cluster unit as normal.

You can set different high and low watermark thresholds for CPU usage and memory usage, and for the number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP UTM proxy sessions. For each loading indicator you set a high watermark threshold a low watermark threshold and a weight. When you first enable this feature the weighted load balancing configuration is synchronized to all cluster units. Subsequent changes to the weighted load balancing configuration are not synchronized so you can configure different weights on each cluster unit.

If a cluster unit's CPU usage reaches the high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to the low watermark threshold. When this happens the cluster resumes sending sessions to at the previous rate.

The CPU usage, memory usage, and UTM proxy weights determine how the cluster load balances sessions when a high watermark threshold is reached and also affect how the cluster load balances sessions when multiple cluster units reach different high watermark thresholds at the same time. For example, you might be less concerned about a cluster unit reaching the memory usage high watermark threshold than reaching the CPU usage high watermark threshold. If this is the case you can set the weight lower for memory usage. Then, if one cluster unit reaches the CPU usage high watermark threshold and a second cluster unit reaches the memory usage high watermark threshold the cluster will load balance more sessions to the unit with high memory usage and fewer sessions to the cluster unit with high CPU usage. As a result, reaching the CPU usage high watermark will have a greater affect on how sessions are redistributed than reaching the memory usage high watermark.

When a high watermark threshold is reached, the corresponding weight is subtracted from the static weight of the cluster unit. The lower the weight the fewer the number of sessions that are load balanced to that unit. Subsequently when the low watermark threshold is reached, the static weight of the cluster returns to its configured value. For the weights to all be effective the weights assigned to the load indicators should usually be lower than or equal to the static weights assigned to the cluster units.

Use the following command to set thresholds and weights for CPU and memory usage and UTM proxy sessions:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold <weight> <low> <high>
  set memory-threshold <weight> <low> <high>
  set http-proxy-threshold <weight> <low> <high>
  set ftp-proxy-threshold <weight> <low> <high>
  set imap-proxy-threshold <weight> <low> <high>
  set nntp-proxy-threshold <weight> <low> <high>
  set pop3-proxy-threshold <weight> <low> <high>
  set smtp-proxy-threshold <weight> <low> <high>
end
```

For each option, the weight range is 0 to 255 and the default weight is 5. The low and high watermarks are a percent (0 to 100). The default low and high watermarks are 0 which means they are disabled. The high watermark must be greater than the low watermark.

For CPU and memory usage the low and high watermarks are compared with the percentage CPU and memory use of the cluster unit. For each of the UTM proxies the high and low watermarks are compared to a number that represents percent of the max number of proxy sessions being used by a proxy. This number is calculated using the formula:

```
proxy usage = (current sessions * 100) / max sessions
```

where:

`current sessions` is the number of active sessions for the proxy type.

`max sessions` is the session limit for the proxy type. The session limit depends on the FortiGate unit and its configuration.

You can use the following command to display the maximum and current number of sessions for a UTM proxy:

```
get test { ftpd | http | imap | nntp | pop3 | smtp } 4
```

You can use the following command to display the maximum number of session and the and current number of sessions for all of the proxies:

```
get test proxyworker 4
```

The command output includes lines similar to the following:

```
get test http 4
HTTP Common
Current Connections          5000/8032
```


In the example, 5000 is the current number of proxy connections being used by HTTP and 8032 is the maximum number of proxy sessions allowed. For this example the proxy usage would be:

```
proxy usage = (5000 * 100) / 8032
proxy usage = 62%
```

Example weighted load balancing configuration

Consider a cluster of three FortiGate-620B units with host names 620_ha_1, 620_ha_2, and 620_ha_3 as shown in Figure 222. This example describes how to configure weighted load balancing settings for CPU and memory usage for the cluster and then to configure UTM proxy weights to send most HTTP and POP3 proxy sessions to different cluster units.

Figure 222: Example HA weighted load balancing configuration



Connect to the cluster CLI and use the following command to set the CPU usage threshold weight to 30, low watermark to 60, and high watermark to 80. This command also sets the memory usage threshold weight to 10, low watermark to 60, and high watermark to 90.


```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold 30 60 80
  set memory-threshold 10 60 90
end
```

The static weights for the cluster units remain at the default values of 40. Since this command changes the mode to `a-a` and the schedule to `weight-round-robin` for the first time, the weight settings are synchronized to all cluster units.

As a result of this configuration, if the CPU usage of any cluster unit (for example, 620_ha_1) reaches 80% the static weight for that cluster unit is reduced from 40 to 10 and only 10 of every 120 new sessions are load balanced to this cluster unit. If the memory usage of 620_ha_1 also reaches 90% the static weight further reduces to 0 and no new sessions are load balanced to 620_ha_1. Also, if the memory usage of 620_ha_2 reaches 90% the static weight of 620_ha_2 reduces to 30 and 30 of every 120 new sessions are load balanced to 620_ha_2.

Now that you have established the weight load balancing configuration for the entire cluster you can monitor the cluster to verify that processing gets distributed evenly to all cluster units. From the web-based manager you can go to *System > Config > HA > View HA Statistics* and see the CPU usage, active sessions, memory usage and other statistics for all of the cluster units. If you notice that one cluster unit is more or less busy than others you can adjust the dynamic weights separately for each cluster unit.

For example, in some active-active clusters the primary unit may tend to be busier than other cluster units because in addition to processing sessions the primary unit also receives all packets sent to the cluster and performs load balancing to distribute the sessions to other cluster units. To reduce the load on the primary unit you could reduce the CPU and memory usage high watermark thresholds for the primary unit so that fewer sessions are distributed to the primary unit. You could also reduce the primary unit's high watermark setting for the proxies to distribute more proxy sessions to other cluster units.



Note that this would only be useful if you are using device priorities and override settings to make sure the same unit always becomes the primary unit. See [“Controlling primary unit selection using device priority and override” on page 2010](#).

If the example cluster is configured for 620_ha_2 to be the primary unit, connect to the 620_ha_2's CLI and enter the following command to set CPU usage, memory usage, and proxy usage high watermark thresholds lower.

```
config system ha
  set cpu-threshold 30 60 70
  set memory-threshold 30 60 70
  set http-proxy-threshold 30 60 70
  set ftp-proxy-threshold 30 60 70
  set imap-proxy-threshold 30 60 70
  set nntp-proxy-threshold 30 60 70
  set pop3-proxy-threshold 30 60 70
  set smtp-proxy-threshold 30 60 70
end
```

As a result, when any of these factors reaches 70% on the primary unit, fewer sessions will be processed by the primary unit, preventing the number of sessions being processed from rising.

NAT/Route mode active-active cluster packet flow


This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the Internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In NAT/Route mode, eight MAC addresses are involved in active-active communication between the client and the web server when the primary unit load balances packets to the subordinate unit:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only know the cluster external virtual MAC address (MAC_V_ext). The cluster virtual MAC address is described in [“Cluster virtual MAC addresses” on page 2177](#).

Figure 223: NAT/Route mode active-active packet flow



Packet flow from client to web server

- 1 The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
- 2 The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
- 3 The client computer issues an ARP request to 10.11.101.100.

- 4 The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
- 5 The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

- 6 The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_P_int
Destination	172.20.120.130	MAC_S_int

- 7 The subordinate unit recognizes that the packet has been forwarded from the primary unit and processes it.
- 8 The subordinate unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_S_ext
Destination	172.20.120.130	MAC_Server

- 9 The primary unit forwards further packets in the same session to the subordinate unit.
- 10 Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from web server to client

- 1 When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
- 2 The web server issues an ARP request to 172.20.120.141.
- 3 The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.
- 4 The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

- 5 The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_P_ext
Destination	172.20.120.141	MAC_S_ext

- 6 The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
- 7 The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_S_int
Destination	10.11.101.10	MAC_Client

- 8 The primary unit forwards further packets in the same session to the subordinate unit.
- 9 Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- 1 If the primary unit fails, the subordinate unit negotiates to become the primary unit.
- 2 The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
- 3 The new primary unit sends gratuitous ARP packets to the 10.10.101.0 network to associate its internal IP address with the internal virtual MAC address.
- 4 The new primary unit sends gratuitous ARP packets to the 172.20.120.0 network to associate its external IP address with the external virtual MAC address.
- 5 Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

Transparent mode active-active cluster packet flow

This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in Transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the Transparent mode cluster.


In Transparent mode, six MAC addresses are involved in active-active communication between a client and a server when the primary unit load balances packets to the subordinate unit:

- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and load balanced among cluster members.

The cluster's presence on the network and its load balancing are transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the external virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the internal virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Figure 224: Transparent mode active-active packet flow



Packet flow from client to mail server

- 1 The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
- 2 The client computer issues an ARP request to 10.11.101.200.
- 3 The primary unit forwards the ARP request to the mail server.
- 4 The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
- 5 The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- 6 The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_P_int
Destination	10.11.101.200	MAC_S_int

- 7 The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.

- 8 The subordinate unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_S_ext
Destination	10.11.101.200	MAC_Server

- 9 The primary unit forwards further packets in the same session to the subordinate unit.
 10 Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from mail server to client

- 1 To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
 2 The primary unit forwards the ARP request to the client computer.
 3 The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
 4 The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- 5 The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_P_ext
Destination	10.11.101.10	MAC_S_ext

- 6 The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
 7 The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_S_int
Destination	10.11.101.10	MAC_Client


- 8 The primary unit forwards further packets in the same session to the subordinate unit.
 9 Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- 1 If the primary unit fails the subordinate unit negotiates to become the primary unit.

- 2** The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
- 3** The new primary unit sends gratuitous ARP requests to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
- 4** The new primary unit sends gratuitous ARP requests to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
- 5** Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.



HA with third-party products


This chapter provides information about operating FortiGate clusters with third party products such as layer-2 and layer-3 switches. This chapter describes:

- Troubleshooting layer-2 switches
- Failover issues with layer-3 switches
- Changing spanning tree protocol settings for some switches
- Failover and attached network equipment
- Ethertype conflicts with third-party switches
- LACP, 802.3ad aggregation and third-party switches

Troubleshooting layer-2 switches

Issues may occur because of the way an HA cluster assigns MAC addresses to the primary unit. In a functioning HA cluster, all primary unit interfaces are assigned the same virtual MAC address. The last byte of the virtual MAC address is the hexadecimal equivalent of the group ID. See “[Cluster virtual MAC addresses](#)” on [page 2177](#) for more information about the HA group ID and the cluster virtual MAC address.

Figure 225: Typical HA configuration, each interface connected to a different switch



Assigning the virtual MAC addresses in this way results in two restrictions when installing HA clusters:

- Two clusters with the same group ID cannot connect to the same switch and cannot be installed on the same network unless they are separated by a router.
- Two or more interfaces on the same primary unit cannot be connected to the same switch unless the traffic is separated using VLANs and unless the switch is VLAN-aware.

Forwarding delay on layer 2 switches

You must ensure that if there is a switch between the FortiGate HA cluster and the network it is protecting and the switch has a forwarding delay (even if spanning tree is disabled) when one of its interfaces is activated then the forwarding delay should be set as low as possible. For example, some versions of Cisco IOS have a forwarding delay of 15 seconds even when spanning tree is disabled. If left at this default value then TCP session pickup can fail because traffic is not forwarded through the switch on HA failover.

Failover issues with layer-3 switches

After a failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables of the switches connected to the cluster. If the cluster is connected using layer-2 switches, the MAC forwarding tables (also called arp tables) are refreshed by the gratuitous ARP packets and the switches start directing packets to the new primary unit.

In some configurations that use layer-3 switches, after a failover, the layer-3 switches may not successfully re-direct traffic to the new primary unit. The possible reason for this is that the layer-3 switch might keep a table of IP addresses and interfaces and may not update this table for a relatively long time after the failover (the table is not updated by the gratuitous ARP packets). Until the table is updated, the layer-3 switch keeps forwarding packets to the now failed cluster unit. As a result, traffic stops and the cluster does not function.

As of the release date of this document, Fortinet has not developed a workaround for this problem. One possible solution would be to clear the forwarding table on the layer-3 switch.

The `config system ha link-failed-signal` command described in [“Updating MAC forwarding tables when a link failover occurs” on page 2198](#) can be used to resolve link failover issues similar to those described here.

Changing spanning tree protocol settings for some switches

Configuration changes may be required when you are running an active-active HA cluster that is connected to a switch that operates using the spanning tree protocol. For example, the following spanning tree parameters may need to be changed:

Maximum Age	The time that a bridge stores the spanning tree bridge control data unit (BPDU) before discarding it. A maximum age of 20 seconds means it may take 20 seconds before the switch changes a port to the listening state.
Forward Delay	The time that a connected port stays in listening and learning state. A forward delay of 15 seconds assumes a maximum network size of seven bridge hops, a maximum of three lost BPDUs and a hello-interval of 2 seconds.

For an active-active HA cluster to be compatible with the spanning tree algorithm, the FGCP requires that the sum of maximum age and forward delay should be less than 20 seconds. The maximum age and forward delay settings are designed to prevent layer 2 loops. If there is no possibility of layer 2 loops in the network, you could reduce the forward delay to the minimum value.

For some Dell 3348 switches the default maximum age is 20 seconds and the default forward delay is 15 seconds. In this configuration the switch cannot work with a FortiGate HA cluster. However, the switch and cluster are compatible if the maximum age is reduced to 10 seconds and the forward delay is reduced to 5 seconds.

Spanning Tree protocol (STP)

Spanning tree protocol is an IEEE 802.1 standard link management protocol that for media access control bridges. STP uses the spanning tree algorithm to provide path redundancy while preventing undesirable loops in a network that are created by multiple active paths between stations. Loops can be created if there are more than route between two hosts. To control path redundancy, STP creates a tree that spans all of the switches in an extended network. Using the information in the tree, the STP can force redundant paths into a standby, or blocked, state. The result is that only one active path is available at a time between any two network devices (preventing looping). Redundant links are used as backups if the initial link should fail. Without spanning tree in place, it is possible that two connections may be simultaneously live, which could result in an endless loop of traffic on the network.

Bridge Protocol Data Unit (BPDU)

BPDU are spanning tree data messages exchanged across switches within an extended network. BPDU packets contain information on ports, addresses, priorities and costs and ensure that the data ends up where it was intended to go. BPDU messages are exchanged across bridges to detect loops in a network topology. The loops are then removed by shutting down selected bridge interfaces and placing redundant switch ports in a backup, or blocked, state.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Ethertype conflicts with third-party switches

Some third-party network equipment may use packets with Ethertypes that are the same as the ethertypes used for HA heartbeat packets. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 heartbeat packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGate units connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8893, and 0x8891 to pass.

You can also use the following CLI commands to change the Ethertypes of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For more information, see [“Heartbeat packet Ethertypes” on page 2174](#).

LACP, 802.3ad aggregation and third-party switches

If a cluster contains 802.3ad aggregated interfaces you should connect the cluster to switches that support configuring multiple Link Aggregation (LAG) groups.


The primary and subordinate unit interfaces have the same MAC address, so if you cannot configure multiple LAG groups a switch may place all interfaces with the same MAC address into the same LAG group; disrupting the operation of the cluster.

You can change the FortiGate configuration to prevent subordinate units from participating in LACP negotiation. For example, use the following command to do this for an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```


This configuration prevents the subordinate unit interfaces from sending or receiving packets. Resulting in the cluster not being able to operate in active-active mode. As well, failover may be slower because after a failover the new primary unit has to perform LACP negotiation before being able to process network traffic.

For more information, see [“Example: HA and 802.3ad aggregated interfaces” on page 2057](#).



VRRP

A Virtual Router Redundancy Protocol (VRRP) configuration can be used as a high availability solution to make sure that a network maintains connectivity with the Internet (or with other networks) even if the default router for the network fails. Using VRRP, if a router or a FortiGate unit fails all traffic to this router transparently fails over to another router or FortiGate unit that takes over the role of the router or FortiGate unit that failed. If the failed router or FortiGate unit is restored, it will once again take over processing traffic for the network. VRRP is described by [RFC 3768](#).



To configure VRRP you create a VRRP group that contains two or more routers. Some or all of these routers can be FortiGate units. You can include different FortiGate models in the same VRRP group. The group members are configured to be the master router and one or more backup routers of the VRRP group. The network directs all traffic to the master's IP address and MAC address. If the master fails, VRRP dynamically shifts packet forwarding to a backup router. VRRP provides this redundancy without user intervention or additional configuration to any of the devices on the network.

The VRRP redundancy scheme means that devices on the network keep a single IP address for the default gateway and this IP address maps to a well-known virtual MAC address. If the VRRP master fails, one of the backup units becomes the new master and acquires virtual IP and MAC addresses that match the addresses of the master. The network then automatically directs all traffic to the backup unit. VRRP uses the broadcast capabilities of Ethernet networks. As long as one of the routers in a VRRP group is running, ARP requests for the default gateway IP address always receive replies. Additionally, hosts can send packets outside their subnet without interruption.

FortiGate units support VRRP and can be quickly and easily integrated into a network that has already deployed a group of routers using VRRP. You can also create a new VRRP configuration consisting of a FortiGate unit acting as a VRRP master with one or more VRRP-compatible routers acting as backup routers. Some or all of those backup routers can be FortiGate units.

During normal operation the VRRP master unit sends VRRP advertisement messages to the backup units. A backup unit will not attempt to become a master unit while it is receiving these messages. When a FortiGate unit operating as a VRRP master fails, a backup unit takes its place and continues processing network traffic. The backup unit assumes the master unit has failed if it stops receiving the advertisement messages from the master unit. The backup unit with the highest priority becomes the new master unit after a short delay. During this delay the new master unit sends gratuitous ARPs to the network to map the virtual router IP address to its MAC address. As a result, all packets sent to the default route IP address are sent to the new master unit. If the backup unit is a FortiGate unit, the network continues to benefit from FortiOS security features. If the backup unit is a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate unit is back on line.

During a VRRP failover, as the backup unit starts to forward traffic it will not have session information for all of the failed over in-progress sessions. If the backup unit is operating as a normal FortiGate unit it will not be able to forward this traffic because of the lack of session information. To resolve this problem, immediately after a failover and for a short time as its taking over traffic processing, the backup unit operates with asymmetric routing enabled. This allows the backup unit to re-create all of the in-progress sessions and add them to the session table. While operating with asymmetric routing enabled, the backup unit cannot apply security functions. When the start-time ends the backup unit disables asymmetric routing and returns to normal operation including applying security functions.

Adding a VRRP virtual router to a FortiGate interface

Use the following command to add a VRRP virtual router to the port10 interface of a FortiGate unit. This VRRP virtual router has a virtual router ID of 200, uses IP address 10.31.101.200 and has a priority of 255. Since this is the highest priority this interface is configured to be the master of the VRRP group with ID number 200.

```
config system interface
  edit port10
    config vrrp
      edit 200
        set vrip 10.31.101.200
        set priority 255
      end
    end
  end
```

VRRP virtual MAC address

The VRRP virtual MAC address (or virtual router MAC address) is a shared MAC address adopted by the VRRP master. If the master fails the same virtual MAC master fails over to the new master. As a result, all packets for VRRP routers can continue to use the same virtual MAC address. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

Each VRRP router is associated with its own virtual MAC address. The last part of the virtual MAC depends on the VRRP virtual router ID using the following format:

```
00-00-5E-00-01-<VRID_hex>
```

Where <VRID_hex> is the VRRP virtual router ID in hexadecimal format in internet standard bit-order. For more information about the format of the virtual MAC see RFC 3768.

Some examples:

- If the VRRP virtual router ID is 10 the virtual MAC would be 00-00-5E-00-01-05.
- If the VRRP virtual router ID is 200 the virtual MAC would be 00-00-5E-00-01-c8.

The VRRP virtual MAC address feature is disabled by default. When you enable the feature on a FortiGate interface, all of the VRRP routers added to that interface use their own VRRP virtual MAC address. Each virtual MAC address will be different because each virtual router has its own ID.

Use the following command to enable the VRRP virtual MAC address on the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  end
```

```
end
```

The port2 interface will now accept packets sent to the MAC addresses of the VRRP virtual routers added to this interface.

Using the VRRP virtual MAC address can improve network efficiency especially on large and complex LANs because when a failover occurs devices on the LAN do not have to learn a new MAC address for the new VRRP router.

If the VRRP virtual MAC address feature is disabled, the VRRP group uses the MAC address of the master. In the case of a FortiGate VRRP virtual router this is the MAC address of the FortiGate interface that the VRRP virtual routers are added to. If a master fails, when the new master takes over it sends gratuitous ARPs to associate the VRRP virtual router IP address with the MAC address of the new master (or the interface of the FortiGate unit that has become the new master). If the VRRP virtual MAC address is enabled the new master uses the same MAC address as the old master.

Configuring VRRP

To configure VRRP you must configure two or more FortiGate interfaces or routers with the same virtual router ID and IP address. Then these FortiGate units or routers can automatically join the same VRRP group. You must also assign priorities to each of the FortiGate units or routers in the VRRP group. One of the FortiGate units or routers must have the highest priority to become the master. The other FortiGate units or routers in the group are assigned lower priorities and become backup units. All of the units in the VRRP group should have different priorities. If the master unit fails, VRRP automatically fails over to the remaining unit in the group with the highest priority.


You configure VRRP from the FortiGate CLI by adding a VRRP virtual router to a FortiGate interface. You can add VRRP virtual routers to multiple FortiGate interfaces and you can add more than one virtual router to the same interface.

Example VRRP configuration: two FortiGate units in a VRRP group

This example includes a VRRP group consisting of two FortiGate units that connect an internal network to the Internet. As shown in [Figure 226](#), the internal network's default route is 10.31.101.120.

The FortiGate port2 interfaces connect to the internal network. A VRRP virtual router is added to each FortiGate unit's port2 interface. The virtual router IP address is 10.31.101.120 (the internal network's default route) and the virtual router's ID is 5. The VRRP priority of the master unit is set to 255 and the VRRP priority of the backup unit is 50. The port2 interface of each FortiGate unit should have an IP address that is different from the virtual router IP address and the port2 interface IP addresses should be different from each other.

This example also includes enabling the VRRP virtual MAC address on both FortiGate unit port2 interfaces so that the VRRP group uses the VRRP virtual MAC address.

Figure 226: Example VRRP configuration with two FortiGate units**To configure the FortiGate units for VRRP**

- 1 Select one of the FortiGate units to be the VRRP master and the other to be the backup unit.
- 2 From the master unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
    config vrrp
      edit 5
        set vrip 10.31.101.120
        set priority 255
      end
    end
  end
```

- 3 From the backup unit's CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
    config vrrp
      edit 5
        set vrip 10.31.101.120
        set priority 50
      end
    end
  end
```


Example VRRP configuration: VRRP load balancing two FortiGate units and two VRRP groups

In this configuration two VRRP groups are involved. Each FortiGate unit participates in both of them. One FortiGate unit is the master of one group and the other FortiGate unit is the master of the other group. The network distributes traffic between two different default routes (10.31.101.120 and 10.31.101.130). One VRRP group is configured with one of the default route IP addresses and the other VRRP group get the other default route IP address. So during normal operation both FortiGate units are processing traffic and the VRRP groups are used to load balance the traffic between the two FortiGate units.

If one of the FortiGate units fails, the remaining FortiGate unit becomes the master of both VRRP groups. The network sends all traffic for both default routes to this FortiGate unit. The result is a configuration that under normal operation load balances traffic between two FortiGate units, but if one of the FortiGate units fails, all traffic fails over to the unit that is still operating.

This example also includes enabling the VRRP virtual MAC address on both FortiGate unit port2 interfaces so that the VRRP groups use their VRRP virtual MAC addresses.

Figure 227: Example VRRP configuration with two FortiGate units and two VRRP groups



To configure the FortiGate units

- 1 Log into the CLI of FortiGate unit A.
- 2 Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate unit A:

```
config system interface
edit port2
set vrrp-virtual-mac enable
config vrrp
edit 50 (32)
set vrip 10.31.101.120
set priority 255
```



```
        next
        edit 100 (64)
            set vrip 10.31.101.130
            set priority 50
        end
    end
```

3 Log into the CLI of FortiGate unit B.

4 Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate unit B:

```
config system interface
    edit port2
        set vrrp-virtual-mac enable
    config vrrp
        edit 50
            set vrip 10.31.101.120
            set priority 50
        next
        edit 100
            set vrip 10.31.101.130
            set priority 255
        end
    end
end
```

Optional VRRP configuration settings

In addition to the basic configuration settings, you can change to the VRRP configuration to:

- Adjust the virtual router advertisement message interval between 1 and 255 seconds using the `adv-interval` option.
- Adjust the startup time using the `start-time` option. The default start time is 3 seconds and the range is 1 to 255 seconds. The start time is the maximum time that the backup unit waits between receiving advertisement messages from the master unit. If the backup unit does not receive an advertisement message during this time it assumes the master has failed and becomes the new master unit. In some cases the advertisement messages may be delayed. For example, some switches with spanning tree enabled may delay some of the advertisement message packets. If you find that backup units are attempting to become master units without the master unit failing, you can extend the start time to make sure the backup units wait long enough for the advertisement messages.
- Enable or disable individual virtual router configurations using the `status` option. Normally virtual router configurations are enabled but you can temporarily disable one if its not required.
- Enable or disable preempt mode using the `preempt` option. In preempt mode a higher priority backup unit can preempt a lower priority master unit. This can happen if a master has failed, a backup unit has become the master unit, and the failed master is restarted. Since the restarted unit will have a higher priority, if preempt mode is enabled the restarted unit will replace the current master unit. Preempt mode is enabled by default.
- Monitor the route to a destination IP address using the `vrdst` option.



TCP session synchronization

You can use the `config system session-sync` command to configure TCP session synchronization (also called standalone session synchronization) between two standalone FortiGate units. You can use this feature with external routers or load balancers configured to distribute or load balance TCP sessions between two peer FortiGate units. If one of the peers fails, session failover occurs and active TCP sessions fail over to the peer that is still operating. This failover occurs without any loss of data. As well, the external routers or load balancers will detect the failover and re-distribute all sessions to the peer that is still operating.



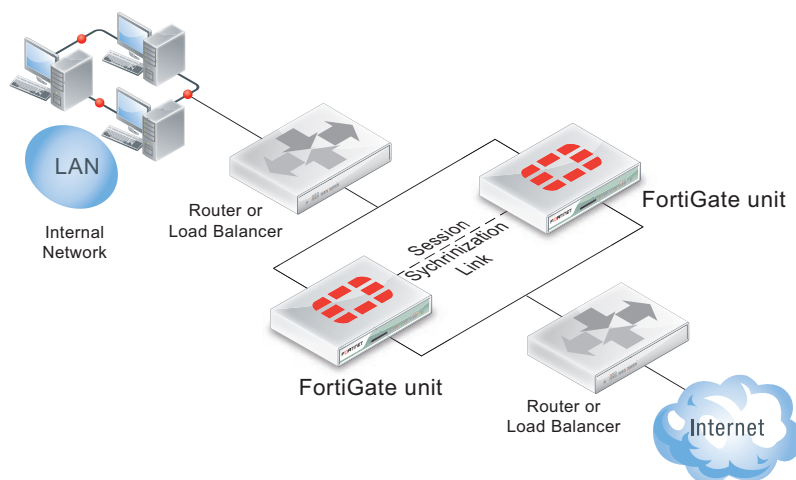
Standalone session synchronization between two standalone FortiGate units is also sometimes called TCP session synchronization or session synchronization between non-HA FortiGate units.



You cannot configure standalone session synchronization when HA is enabled.

Standalone session synchronization can be used instead of HA to provide TCP session synchronization between two peer FortiGate units. If the external load balancers direct all sessions to one peer the affect is similar to active-passive HA. If external load balancers or routers load balance traffic to both peers, the effect is similar to active-active HA. The load balancers should be configured so that all of the packets for any given session are processed by the same peer. This includes return packets.

Figure 228: Standalone session synchronization



By default, standalone session synchronization synchronizes all TCP sessions. You can optionally add filters to a configuration control which TCP sessions are synchronized. You can add filters to only synchronize packets from specified source and destination addresses, specified source and destination interfaces, and specified predefined firewall TCP services.

Unlike HA, standalone session synchronization does not include configuration synchronization. In fact, the configuration of the two peers is not identical because in most cases the peers would have different IP addresses. Also unlike HA, load balancing is done by external routers or load balancers. The FortiGate units only perform session synchronization and session failover.

Notes and limitations

Standalone session synchronization has the following limitations:

- Only TCP sessions accepted by security policies are synchronized. Due to their non-stateful nature, UDP and ICMP sessions don't need to be synchronized to naturally failover.
- Standalone session synchronization is a global configuration option. As a result you can only add one predefined firewall TCP service to a filter configuration. You cannot add custom services or service groups even if virtual domains are not enabled.
- You can only add one filter configuration to a given standalone session synchronization configuration. However, you can add multiple filters by adding multiple identical standalone session synchronization configurations, each one with a different filter configuration.
- Sessions accepted by security policies with UTM options configured are not synchronized.
- Sessions that include network address translation (NAT) applied by selecting NAT in security policies are not synchronized because the address translation binds to a FortiGate unit address and the peers have different IP addresses.
- Session synchronization is a CLI only configuration.
- Session synchronization is available for FortiGate units or virtual domains operating in NAT/Route or Transparent mode. NAT sessions are not synchronized in either mode. In NAT/Route mode, only sessions for route mode security policies are synchronized. In Transparent mode, only sessions for normal Transparent mode policies are synchronized.
- Session synchronization cannot be asymmetric. Session synchronization is stateful. So all of the packets of a given session must be processed on the same peer. This includes return packets. You must configure the load balancers so that they do not cause asymmetric routing.
- Session synchronization is supported for traffic on physical interfaces, VLAN interfaces, zones, and aggregate interfaces. Session synchronization has not been tested for inter-vdom links, accelerated interfaces (FA2 and NP2), between HA clusters, and for redundant interfaces.
- The names of the matching interfaces, including VLAN interfaces, aggregate interfaces and so on, must be the same on both peers.

Configuring session synchronization

You configure session synchronization for each virtual domain to be synchronized. If virtual domain configuration is not enabled, you configure session synchronization for the root virtual domain. When virtual domain configuration is enabled and you have added virtual domains you configure session synchronization for each virtual domain to be synchronized. You don't have to synchronize all of the virtual domains.

You must configure session synchronization on both peers. The session synchronization configurations of each peer should complement the other. In fact you can manage and configure both peers as separate FortiGate units. Using FortiManager, you can manage both peers as two separate FortiGate devices.

On each peer, configuring session synchronization consists of selecting the virtual domains to be synchronized using the `syncvd` field, selecting the virtual domain on the other peer that receives the synchronization packets using the `peervd` field, and setting IP address of the interface in the peer unit that receives the synchronization packets using the `peerip` field. The interface with the `peerip` must be in the `peervd` virtual domain.

The `syncvd` and `peervd` settings must be the same on both peers. However, the `peerip` settings will be different because the `peerip` setting on the first peer includes the IP address of an interface on the second peer. And the `peerip` setting on the second peer includes the IP address of an interface on the first peer.

Because session synchronization does not synchronize FortiGate configuration settings you must configure both peers separately. For session synchronization to work properly all session synchronized virtual domains must be added to both peers. The names of the matching interfaces in each virtual domain must also be the same; this includes the names of matching VLAN interfaces. Note that the index numbers of the matching interfaces and VLAN interfaces can be different. Also the VLAN IDs of the matching VLAN interfaces can be different.

As well, the session synchronized virtual domains should have the same security policies so that sessions can be resumed after a failover using the same security policies.

For a configuration example, see [“Basic example configuration” on page 2246](#).

Configuring the session synchronization link

When session synchronization is operating, the peers share session information over an Ethernet link between the peers similar to an HA heartbeat link. Usually you would use the same interface on each peer for session synchronization. You should connect the session synchronization interfaces directly without using a switch or other networking equipment. If possible use a crossover cable for the session synchronization link. For FortiGate-5000 systems you can use a backplane interface as the session synchronization link.

You can use different interfaces on each peer for session synchronization links. Also, if you have multiple sessions synchronization configurations, you can have multiple session synchronization links between the peers. In fact if you are synchronizing a lot of sessions, you may want to configure and connect multiple session synchronization links to distribute session synchronization traffic to these multiple links.

You cannot configure backup session synchronization links. Each configuration only includes one session synchronization link.

The session synchronization link should always be maintained. If session synchronization communication is interrupted and a failure occurs, sessions will not failover and data could be lost.

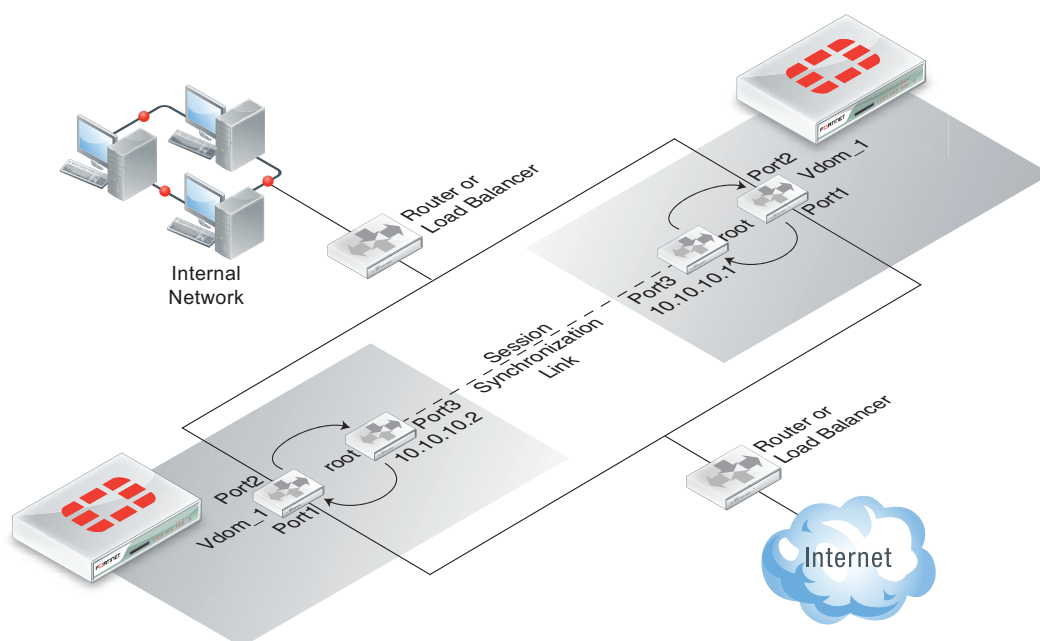
Session synchronization traffic can use a considerable amount of network bandwidth. If possible, session synchronization link interfaces should only be used for session synchronization traffic and not for data traffic.

Basic example configuration

The following configuration example shows how to configure a basic session synchronization configuration for two peer FortiGate units shown in [Figure 229 on page 2246](#). The host names of peers are peer_1 and peer_2. Both peers are configured with two virtual domains: root and vdom_1. All sessions processed by vdom_1 are synchronized. The synchronization link interface is port3 which is in the root virtual domain. The IP address of port3 on peer_1 is 10.10.10.1. The IP address of port3 on peer_2 is 10.10.10.2.

Also on both peers, port1 and port2 are added to vdom_1. On peer_1 the IP address of port1 is set to 192.168.20.1 and the IP address of port2 is set to 172.110.20.1. On peer_2 the IP address of port1 is set to 192.168.20.2 and the IP address of port2 is set to 172.110.20.2.

Figure 229: Example standalone session synchronization network configuration



To configure standalone session synchronization

- 1 Configure the load balancer or router to send all sessions to peer_1.
- 2 Configure the load balancer or router to send all traffic to peer_2 if peer_1 fails.
- 3 Use normal FortiGate configuration steps on peer_1:
 - Enable virtual domain configuration.
 - Add the vdom_1 virtual domain.
 - Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
 - Set the IP address of port1 to 192.168.20.1.
 - Set the IP address of port2 to 172.110.20.1.
 - Set the IP address of port3 to 10.10.10.1.
 - Add route mode security policies between port1 and port2 to vdom_1.
- 4 Enter the following commands to configure session synchronization for peer_1


```
config system session-sync
```

```
edit 1
  set peerip 10.10.10.2
  set peervd root
  set syncvd vdom_1
end
```

5 Use normal FortiGate configuration steps on peer_2:

- Enable virtual domain configuration.
- Add the vdom_1 virtual domain.
- Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
- Set the IP address of port1 to 192.168.20.2.
- Set the IP address of port2 to 172.110.20.2.
- Set the IP address of port3 to 10.10.10.1.
- Add route mode security policies between port1 and port2 to vdom_1.

6 Enter the following commands to configure session synchronization for peer_1

```
config system session-sync
  edit 1
    set peerip 10.10.10.1
    set peervd root
    set syncvd vdom_1
  end
end
```

To add a filter

You can add a filter to this basic configuration if you only want to synchronize some TCP sessions. For example you can enter the following commands on both FortiGate units to edit the standalone sessions configurations and add a filter so that only HTTP sessions are synchronized

```
config system session-sync
  edit 1
    config filter
      set service HTTP
    end
  end
end
```




Chapter 13 Traffic Shaping

With the ever-increasing demands on network systems for a number of protocols from email, to HTTP traffic both internally and externally to the internet, voice over IP, FTP and so on, traffic slow downs become a reality. As such, important traffic may be dropped or slowed to an unusable speed. This can cause the halting of business traffic and revenues.

Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. But there is a physical limitation to the amount of data which can be buffered and to the length of time.

FortiGate units can implement Quality of Service (QoS) by applying bandwidth limits and prioritization. Using traffic shaping, you can adjust how your FortiGate unit allocates resources to different traffic types to improve the performance and stability of latency sensitive or bandwidth intensive network applications.

This document discusses Quality of Service (QoS) and traffic shaping, describes FortiGate traffic shaping algorithms, and provides procedures and tips on how to configure traffic shaping on FortiGate units.

This FortiOS Handbook chapter contains the following chapters:

[The purpose of traffic shaping](#) describes traffic shaping theories and quality of service.

[Traffic shaping methods](#) describes the different methods of applying traffic shaping within FortiOS, and how to use TOS and Differentiated Services.

[Examples](#) provides some basic examples for the application of shapers.

[Troubleshooting](#) provides diagnose commands to use to troubleshoot traffic shapers to see if they are working correctly.



The purpose of traffic shaping

Traffic shaping, or traffic management, controls the bandwidth available and sets the priority of traffic processed by the policy to control the volume of traffic for a specific period (bandwidth throttling) or rate the traffic is sent (rate limiting).

Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. But there is a physical limitation to the amount of data which can be buffered and to the length of time. Once these thresholds have been surpassed, frames and packets will be dropped, and sessions will be affected in other ways.

A basic traffic shaping approach is to prioritize certain traffic flows over other traffic whose potential discarding is less advantageous. This would mean that you accept sacrificing certain performance and stability on low-priority traffic, to increase or guarantee performance and stability to high-priority traffic.

If, for example, you are applying bandwidth limitations to certain flows, you must accept the fact that these sessions can be limited and therefore negatively impacted.

Note that traffic shaping is effective for normal IP traffic at normal traffic rates. Traffic shaping is not effective during periods when traffic exceeds the capacity of the FortiGate unit. Because packets must be received by the FortiGate unit before they are subject to traffic shaping, if the FortiGate unit cannot process all of the traffic it receives, then dropped packets, delays, and latency are likely to occur.

To ensure that traffic shaping is working at its best, make sure that the interface Ethernet statistics show no errors, collisions or buffer overruns.

Accelerated interfaces (NP2, NP4, CE) affect traffic shaping. For more information, see the [FortiGate Hardware](#) Guide.

Quality of Service

Quality of service (QoS) is the capability to adjust some quality aspects of your overall network traffic. This can include such techniques as priority-based queuing and traffic policing. Because bandwidth is finite and because some types of traffic are slow, jitter or packet loss sensitive, bandwidth intensive, or operation critical, QoS can be a useful tool for optimizing the performance of the various applications on your network.

Before implementing QoS, organizations should first identify the types of traffic that are important to the organization, the types of traffic that use high amounts of bandwidth, and the types of traffic that are sensitive to latency or packet loss.

For example, a company might want to guarantee sufficient bandwidth for revenue producing e-commerce traffic. They need to ensure that transactions can be completed and that clients do not experience service delays and interruptions. At the same time, the company may need to ensure low latency for voice over IP (VoIP) traffic used by sales and customer support, while traffic latency and bursts may be less critical to the success of other network applications such as long term, resumable file transfers. Many organizations discover that QoS is especially important for managing their voice and streaming multi-media traffic. These types of traffic can rapidly consume bandwidth and are sensitive to latency.

Discovering the needs and relative importance of each traffic type on your network will help you to design an appropriate overall approach, including how you will configure each available QoS component technique. Some organizations discover that they only need to configure bandwidth limits for some services. Other organizations determine that they need to fully configure interface and security policy bandwidth limits for all services, and prioritize queuing of critical services relative to traffic rate.

You can implement QoS on FortiGate units using the following techniques:

Traffic policing	Drops packets that do not conform to bandwidth limitations.
Traffic shaping	Ensures that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows greater than the maximum rate are subject to traffic policing.
Queuing	Transmits packets in order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

When deciding how to configure QoS techniques, it can be helpful to know when FortiGate units employ each technique in the overall traffic processing flow, and the considerations that arise from those mechanisms.

Traffic policing

As traffic arrives (ingress) and departs (egress) on an interface, the FortiGate unit begins to process the traffic. In later phases of the network processing, such as enforcing maximum bandwidth use on sessions handled by a security policy, if the current rate for the destination interface or traffic regulated by that security policy is too high, the FortiGate unit may drop the packet. As a result, time spent on prior processing, such as web filtering, decryption or IPS, can be wasted on some packets that are not ultimately forwarded. This applies to VLAN interfaces as well as physical interfaces.

You can prevent this wasted effort on ingress by configuring the FortiGate unit to preemptively drop excess packets when they are received at the source interface, before most other traffic processing is performed:

```
config system interface
  edit <interface_name>
    set inbandwidth <rate_int>
  next
end
```

where <rate_int> is the bandwidth limit in Kb/s. Excess packets will be dropped. If inbandwidth is 0, the rate is not limited.

A similar command is available that can be performed on egress as well using the CLI commands:

```
config system interface
  edit <interface_name>
    set outbandwidth <rate_int>
  next
end
```

As with ingress, setting the rate to 0 (zero) sets the rate to unlimited.

Rate limiting traffic accepted by the interface enables you to restrict incoming traffic to rates that, while no longer the full capacity of the interface, at the traffic shaping point in the processing are more likely to result in acceptable rates of outgoing traffic per destination interface or all security policies. This conserves FortiGate processing resources for those packets that are more likely to be viable completely to the point of egress.

Excessive traffic policing can degrade network performance rather than improve it. For details on factors you may want to consider when configuring traffic policing, see [“Important considerations” on page 2258](#).

Bandwidth guarantee, limit, and priority interactions

After packet acceptance, the FortiGate unit classifies traffic and may apply traffic policing at additional points during processing. It may also apply additional QoS techniques, such as prioritization and traffic shaping. Traffic shaping consists of a mixture of traffic policing to enforce bandwidth limits, and priority queue adjustment to assist packets in achieving the guaranteed rate.

If you have configured prioritization, the FortiGate unit prioritizes egressing packets by distributing them among FIFO (first in, first out) queues associated with each possible priority number. Each physical interface has six priority queues. Virtual interfaces do not have their own queues, and instead use the priority queues of the physical interface to which they are bound.

Each physical interface's six queues are queue 0 to queue 5, where queue 0 is the highest priority queue. However, for the reasons described below, you may observe that your traffic uses only a subset of those six queues. Some traffic may always use a certain queue number. Some queuing may vary by the packet rate or mixture of services. Some queue numbers may be used only by through traffic for which you have configured traffic shaping in the security policy that applies to that traffic session. For example:

- Administrative access traffic will always use queue 0.
- Traffic matching security policies **without** traffic shaping may use queue 0, queue 1, or queue 2. Which queue will be used depends on the priority value you have configured for packets with that ToS (type of service) bit value, if you have configured ToS-based priorities.
- Traffic matching security policies **with** traffic shaping may use any queue. Which queue will be used depends on whether the packet rate is currently below the guaranteed bandwidth (queue 0), or above the guaranteed bandwidth. Packets at rates greater than the maximum bandwidth limit are dropped.
- If the global tos-based-priority is low (3), the priority in a traffic-shaper is medium (2) and a packet flows through a policy that refers to the shaper, the packet will be assigned the priority defined by the shaper, in this case medium (2).

Prioritization and traffic shaping behavior varies by your configuration, the service types and traffic volumes, and by whether the traffic is through traffic, or the traffic originates from or terminates at the FortiGate unit itself.

FortiGate traffic

Administrative access to the FortiGate through HTTPS or SSH, or IPsec tunnel negotiations, security policies do not apply, and therefore FortiGate units do not apply traffic shaping. Such traffic also uses the highest priority queue, queue 0. In other words:

packet priority = 0

Exceptions to this rule include traffic types that are connections related to a session governed by a security policy.

For example, if you have enabled scanning by FortiGuard antivirus, traffic from the sender technically terminates at the FortiGate proxy that scans that traffic type; the FortiGate unit initiates a second connection that transmits scanned content to its destination. Because the second connection's traffic is technically originating from the FortiGate proxy and therefore the FortiGate unit itself, it uses the highest priority queue, queue 0. However, this connection is logically associated with through traffic, and is therefore subject to possible bandwidth enforcement and guarantees in its governing security policy. In this way, it behaves partly like other through traffic.

Through traffic

For traffic passing through the FortiGate unit, the method a FortiGate unit uses to determine the priority queue varies by whether you have enabled Traffic Shaping. Packets may or may not use a priority queue directly or indirectly derived from the type of service (ToS) bit — sometimes used instead with differentiated services — in the packet's IP header.

If Traffic Shaping is not enabled in the security policy, the FortiGate unit neither limits nor guarantees bandwidth, and traffic for that session uses the priority queue determined directly by matching the ToS bit in its header with your configured values:

```
config system global
  set tos-based-priority {high | low | medium}
end
```

or, if you have configured a priority specifically for that TOS bit value:

```
config system tos-based-priority
  edit <id_int>
    set tos [0-15]
    set priority {high | low | medium}
  next
end
```

where `tos` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in the second location will override the global ToS-based priority. In other words:

packet priority = ToS-based priority

For example, you might specify that packets with a ToS bit value of 2 should use queue 0, the highest priority queue:

```
config system tos-based-priority
  edit 15
    set tos 2
    set priority high
  next
end
```

If Traffic Shaping is enabled in the security policy using shared traffic shapers, the FortiGate unit may instead or also subject packets to traffic policing, or priority queue increase in an effort to meet bandwidth guarantees configured in the shaper:


```

config firewall shaper traffic-shaper
edit <shaper_name>
...
set priority {high | medium | low}
set maximum-bandwidth <rate>
set guaranteed-bandwidth <rate>
end

```

where `high` has a priority value of 1 and `low` is 3, and `<rate>` is the bandwidth limit in kilobits per second.

Figure 230: Traffic queuing as packet rate increases



- If the current packet rate is less than Guaranteed Bandwidth, packets use priority queue 0. In other words:
packet priority = 0
- If the current packet rate is greater than Guaranteed Bandwidth but less than Maximum Bandwidth, the FortiGate unit assigns a priority queue by adding the numerical value of the security policy-based priority, where the value of High is 1, and Low is 3, with the numerical value of the ToS-based priority, where `high` has a priority value of 0 and `low` is 2. Because the two values are added, depending on the your configured ToS-based priorities, packets in this category could use queues from queue 1 to queue 5. In other words:
packet priority = ToS-based priority + security policy-based priority
For example, if you have enabled Traffic Shaping in the security policy, and the security policy's Traffic Priority is Low (value 3), and the priority normally applied to packets with that ToS bit is `medium` (value 1), then packets have a total packet priority of 4, and use priority queue 4.
- If the current packet rate exceeds Maximum Bandwidth, excess packets are dropped.

Calculation and regulation of packet rates

Packet rates specified for Maximum Bandwidth or Guaranteed Bandwidth are:

$$\text{rate} = \text{amount} / \text{time}$$

where rate is expressed in kilobits per second (Kb/s).

Burst size at any given instant cannot exceed the amount configured in Maximum Bandwidth. Packets in excess are dropped. Packets deduct from the amount of bandwidth available to subsequent packets and available bandwidth regenerates at a fixed rate. As a result, bandwidth available to a given packet may be less than the configured rate, down to a minimum of 0 Kb/s.


Rate calculation and behavior can alternatively be described using the token bucket metaphor, where:

- a traffic flow has an associated bucket, which represents burst size bounds, and is the size of your configured bandwidth limit
- the bucket receives tokens, which represent available bandwidth, at the fixed configured rate
- as time passes, tokens are added to the bucket, up to the capacity of the bucket; excess tokens are discarded
- when a packet arrives, the packet must deduct bandwidth tokens from the bucket equal to its packet size in order to egress
- packets cannot egress if there are insufficient tokens to pay for its egress; these nonconforming packets are dropped

Bursts are not redistributed over a longer interval, so bursts are propagated rather than smoothed, although their peak size is limited.

Maximum burst size is the capacity of the bucket (the configured bandwidth limit); actual size varies by the current number of tokens in the bucket, which may be less than bucket capacity, due to deductions from previous packets and the fixed rate at which tokens accumulate. A depleted bucket refills at the rate of your configured bandwidth limit. Bursts cannot borrow tokens from other time intervals. This behavior is illustrated in [Figure 231 on page 2256](#).

Figure 231: Bursts and bandwidth limits over time



By limiting traffic peaks and token regeneration in this way, the available bandwidth at a given moment may be less than bucket capacity, but your limit on the total amount per time interval is ensured. That is, total bandwidth use during each interval of 1 second is at most the integral of your configured rate.

You may observe that external clients, such as FTP or BitTorrent clients, initially report rates between Maximum Bandwidth and twice that of Maximum Bandwidth, depending on the size of their initial burst. This is notably so when a connection is initiated following a period of no network activity. The apparent discrepancy in rates is caused by a difference of perspective when delimiting time intervals. A burst from the client may initially consume all tokens in the bucket, and before the end of 1 second, as the bucket regenerates, be allowed to consume almost another bucket's worth of bandwidth. From the perspective of the client, this constitutes one time interval. From the perspective of the FortiGate unit, however, the bucket cannot accumulate tokens while full; therefore, the time interval for token regeneration begins **after** the initial burst, and does not contain the burst. These different points of reference result in an initial discrepancy equal to the size of the burst — the client's rate contains it, but the FortiGate unit's rate does not. If the connection is sustained to its limit and time progresses over an increasing number of intervals, however, this discrepancy decreases in importance relative to the bandwidth total, and the client's reported rate will eventually approach that of the FortiGate unit's configured rate limit.

For example, your Maximum Bandwidth might be 50 Kb/s and there has been no network activity for one or more seconds. The bucket is full. A burst from an FTP client immediately consumes 50 Kb. Because the bucket completely regenerates over 1 second, by the time almost another 1 second has elapsed from the initial burst, traffic can consume another 49.999 Kb, for a total of 99.999 Kb between the two points in time. From the vantage point of an external FTP client regulated by this bandwidth limit, it therefore initially appears that the bandwidth limit is 99.999 Kb/s, almost twice the configured limit of 50 Kb/s. However, bucket capacity only regenerates at your configured rate of 50 Kb/s, and so the connection can only consume a maximum of 50 Kb during each second thereafter. The result is that as bandwidth consumption is averaged over an increasing number of time intervals, each of which are limited to 50 Kb/s, the effects of the first interval's doubled bandwidth size diminishes proportionately, and the client's reported rate eventually approach your configured rate limit. This effect is illustrated in [Table 120 on page 2257](#).

Table 120: Effects of a 50 Kb/s limit on client reported rates

Total size transferred (Kb)	Time (s)	Rate reported by client (Kb/s)
99.999 (50 + 49.999)	1	99.999
149.999	2	74.999
199.999	3	66.666
249.999	4	62.499
299.999	5	59.998
349.999	6	58.333
...

Guaranteed Bandwidth can also be described using a token bucket metaphor. However, because this feature attempts to achieve or exceed a rate rather than limit it, the FortiGate unit does not discard non-conforming packets, as it does for Maximum Bandwidth; instead, when the flow does not achieve the rate, the FortiGate unit increases the packets' priority queue, in an effort to increase the rate.

Guaranteed and maximum bandwidth rates apply to the bidirectional total for all sessions controlled by the security policy. For example, an FTP connection may entail two separate connections for the data and control portion of the session; some packets may be reply traffic rather than initiating traffic. All packets for both connections are counted when calculating the packet rate for comparison with the guaranteed and maximum bandwidth rate.

Important considerations

In essence, by implementing QoS, you trade some performance and/or stability from traffic X by discarding packets or introducing latency in order to improve performance and stability of traffic Y. The best traffic shaping configuration for your network will appropriately balance the needs of each traffic flow by considering not only the needs of your particular organization, but also the resiliency and other characteristics of each particular service.

For example, you may find that web browsing traffic is both more resistant to interruptions or latency and less business critical than UDP or VoIP traffic, and so you might implement less restrictive QoS measures on UDP or VoIP traffic than on HTTP traffic.

An appropriate QoS configuration will also take into account the physical limits of your network devices, and the interactions of the aforementioned QoS mechanisms, described in [“Bandwidth guarantee, limit, and priority interactions” on page 2253](#).

You may choose to configure QoS differently based upon the hardware limits of your network and FortiGate unit. Traffic shaping may be less beneficial in extremely high-volume situations where traffic exceeds a network interface's or your FortiGate model's overall physical capacity. A FortiGate unit must have sufficient resources, such as memory and processing power, to process all traffic it receives, and to process it at the required rate; if it does not have this capacity, then dropped packets and increased latency are likely to occur. For example, if the total amount of memory available for queuing on a physical interface is frequently exceeded by your network's typical packet rates, frames and packets must be dropped. In such a situation, you might choose to implement QoS using a higher model FortiGate unit, or to configure an incoming bandwidth limit on each interface.

Incorrect traffic shaping configurations can actually further degrade certain network flows, because excessive discarding of packets or increased latency beyond points that can be gracefully handled by that protocol can create additional overhead at upper layers of the network, which may be attempting to recover from these errors. For example, a configuration might be too restrictive on the bandwidth accepted by an interface, and may therefore drop too many packets, resulting in the inability to complete or maintain a SIP call.

To optimize traffic shaping performance, first ensure that the network interface's Ethernet statistics are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```


If these are not clean, adjust FortiGate unit and settings of routers or other network devices that are connected to the FortiGate unit. For additional information, see [“Troubleshooting” on page 2285](#).

Once Ethernet statistics are clean, you may want to use only some of the available FortiGate QoS techniques, or configure them differently, based upon the nature of FortiGate QoS mechanisms described in [“Bandwidth guarantee, limit, and priority interactions” on page 2253](#). Configuration considerations include:

- For maximum bandwidth limits, ensure that bandwidth limits at the source interface and/or the security policy are not too low, which can cause the FortiGate unit to discard an excessive number of packets.
- For prioritization, consider the ratios of how packets are distributed between available queues, and which queue is used by which types of services. If you assign most packets to the same priority queue, it negates the effects of configuring prioritization. If you assign many high bandwidth services to high priority queues, lower priority queues may be starved for bandwidth and experience increased or indefinite latency. For example, you may want to prioritize a latency-sensitive service such as SIP over a bandwidth-intensive service such as FTP. Consider also that bandwidth guarantees can affect the queue distribution, assigning packets to queue 0 instead of their typical queue in high-volume situations.
- You may or may not want to guarantee bandwidth, because it causes the FortiGate unit to assign packets to queue 0 if the guaranteed packet rate is not currently being met. Comparing queuing behavior for lower-bandwidth and higher-bandwidth situations, this would mean that effects of prioritization only become visible as traffic volumes rise and exceed their guarantees. Because of this, you might want only some services to use bandwidth guarantees, to avoid the possibility that in high-volume situations all traffic uses the same queue, thereby negating the effects of configuring prioritization.
- For prioritization, configure prioritization for all through traffic. You may want to configure prioritization by either ToS-based priority or security policy priority, but not both. This simplifies analysis and troubleshooting.

Traffic subject to both security policy and ToS-based priorities will use a combined priority from both of those parts of the configuration, while traffic subject to only one of the prioritization methods will use only that priority. If you configure both methods, or if you configure either method for only a subset of your traffic, packets for which a combined priority applies will frequently receive a lower priority queue than packets for which you have only configured one priority method, or for which you have not configured prioritization.

For example, if both ToS-based priority and security policy priority both dictate that a packet should receive a “medium” priority, in the absence of bandwidth guarantees, a packet will use queue 3, while if only ToS-based priority had been configured, the packet would have used queue 1, and if only security policy-based priority had been configured, the packet would have used queue 2. If no prioritization had been configured at all, the packet would have used queue 0.

For example alternative QoS implementations that illustrate these considerations, see [“Examples” on page 2277](#)



Traffic shaping methods

In FortiOS, there are three types of traffic shaping configuration. Each has a specific function, and all can be used together in varying configurations. Policy shaping enables you to define the maximum bandwidth and guaranteed bandwidth set for a security policy. Per-IP shaping enables you to define traffic control on a more granular level. Application traffic shaping goes further, enabling traffic controls on specific applications or application groupings.

This chapter describes the types of traffic shapers and how to configure them in the web-based manager and the CLI.

Traffic shaping options

When configuring traffic shaping for your network, there are three different methods to control the flow of network traffic to ensure that the desired traffic gets through while also limiting the bandwidth that users use for other less important or bandwidth consuming traffic. The three shaping options are:

- shared policy shaping - bandwidth management by security policies
- per-IP shaping - bandwidth management by user IP addresses
- application control shaping - bandwidth management by application

Shared policy shaping and per IP shaping are enabled within the security policy, while the application control shaping is configured in *UTM Profiles > Application Control* and enabled in the security policy by selecting *UTM Profiles* and selecting the application control profile from the drop-down list.

The FortiGate unit offers three different traffic shaping options, all of which can be enabled at the same time within the same security policy. Generally speaking, the hierarchy for shapers in FortiOS is:

- Application Control shaper
- Security policy shaper
- Per-IP shaper

With this hierarchy, if an application control list has a traffic shaper defined, it will have precedence always over any other security policy shaper. For example, with the example above creating an application control for Facebook, the shaper defined for Facebook will supersede any security policy enabled traffic shapers. While the Facebook application may reach its maximum bandwidth, the user can still have the bandwidth room available from the shared shaper and, if enabled, the per-IP shaper.

Equally, any security policy shared shaper will have precedence over any per-IP shaper. However, traffic that exceeds any of these shapers will be dropped. For example, the policy shaper will take effect first, however, if the per-IP shaper limit is reached first, then traffic for that user will be dropped even if the shared shaper limit for the policy has not been exceeded.

Shared policy shaping

Traffic shaping by security policy enables you to control the maximum and/or guaranteed throughput for a selected security policy. When configuring a shaper, you can select to apply the bandwidth shaping per policy or for all policies. Depending on your selection, the FortiGate unit will apply the shaping rules differently.

Per policy

When selecting a shaper to be per policy, the FortiGate unit will apply the shaping rules defined to each security policy individually.

For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. Three of these have the shaper enabled. Each security policy has the same maximum bandwidth of 1000 Kb/s.

Per policy traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.

All policies

When selecting a shaper to be for all policies - *For All Policies Using This Shaper* - the FortiGate unit applies the shaping rules to all policies using the same shaper. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. All four have the shaper enabled. Each security policy must share the defined 1000 Kb/s, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kb/s, the remaining three must share 200 Kb/s. As policy 1 uses less bandwidth, it is opened up to the other policies to use as required. Once used, any other policies will encounter latency until free bandwidth opens from a policy currently in use.

Maximum and guaranteed bandwidth

The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper.

The guaranteed bandwidth ensures there is a consistent reserved bandwidth available for a given service or user. When setting the guaranteed bandwidth, ensure that the value is significantly less than the bandwidth capacity of the interface, otherwise no other traffic will pass through the interface or very little an potentially causing unwanted latency.



Setting Maximum Bandwidth to 0 (zero) provides unlimited bandwidth.

Traffic priority

Select a Traffic Priority of high, medium or low, so the FortiGate unit manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server needed to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority. The firewall provides bandwidth to low-priority connections only when bandwidth is not needed for high-priority connections.

Be sure to enable traffic shaping on all security policies. If you do not apply any traffic shaping rule to a policy, the policy is set to high priority by default.

Distribute security policies over all three priority queues.

VLAN, VDOM and virtual interfaces

Policy-based traffic shaping does not use queues directly. It shapes the traffic and if the packet is allowed by the security policy, then a priority is assigned. That priority controls what queue the packet will be put in upon egress. VLANs, VDOMs, aggregate ports and other virtual devices do not have queues and as such, traffic is sent directly to the underlying physical device where it is queued and affected by the physical ports.

This is also the case with IPsec connections.

Shared traffic shaper configuration settings

To configure a shared traffic shaper go to *Firewall Objects > Traffic Shaper > Shared* and select *Create New*.

Name	Enter a name for the traffic shaper.
Apply Shaper	<p>When selecting a shaper to be <i>Per Policy</i>, the FortiGate unit will apply the shaping rules defined to each security policy individually. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. Three of these have the shaper enabled. Each security policy has the same maximum bandwidth of 1000 Kb/s.</p> <p>Per policy traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.</p> <p>When selecting a shaper to be for all policies - <i>For All Policies Using This Shaper</i> - the FortiGate unit applies the shaping rules to all policies using the same shaper. For example, the shaper is set to be per policy with a maximum bandwidth of 1000 Kb/s. There are four security policies monitoring traffic through the FortiGate unit. All four have the shaper enabled. Each security policy must share the defined 1000 Kb/s, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kb/s, the remaining three must share 200 Kb/s. As policy 1 uses less bandwidth, it is opened up to the other policies to use as required. Once used, any other policies will encounter latency until free bandwidth opens from a policy currently in use</p>
Traffic Priority	<p>Select level of importance <i>Priority</i> so the FortiGate unit manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server needed to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority.</p> <p>If you do not apply any traffic shaping priority, the priority is set to high priority by default.</p>

Maximum Bandwidth	<p>The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper.</p> <p>Setting Maximum Bandwidth to 0 (zero) provides unlimited bandwidth</p>
Guaranteed Bandwidth	<p>The guaranteed bandwidth ensures there is a consistent reserved bandwidth available for a given service or user. When setting the guaranteed bandwidth, ensure that the value is significantly less than the bandwidth capacity of the interface, otherwise no other traffic will pass through the interface or very little an potentially causing unwanted latency.</p> <p>Setting Guaranteed Bandwidth to 0 (zero) provides unlimited bandwidth.</p>
DSCP	<p>Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see “Differentiated Services”.</p>

Example

The following steps creates a Per Policy traffic shaper called “Throughput” with a maximum traffic amount of 720,000 Kb/s, and a guaranteed traffic of 150,000 Kb/s with a high traffic priority.

To create the shared shaper - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaper > Shared* and select *Create New*.
- 2 Enter the *Name* *Throughput*.
- 3 Select *Per Policy*.
- 4 Select the *Maximum Bandwidth* check box and enter the value *120000*.
- 5 Select the *Guaranteed Bandwidth* check box and enter the value *150000*.
- 6 Set the *Traffic Priority* to *High*.
- 7 Select *OK*.

To create the shared shaper - CLI

```
config firewall shaper traffic-shaper
edit Throughput
set per-policy enable
set maximum-bandwidth 720000
set guaranteed-bandwidth 150000
set priority high
end
```

Per-IP shaping

Traffic shaping by IP enables you to apply traffic shaping to all source IP addresses in the security policy. As well as controlling the maximum bandwidth users of a selected policy, you can also define the maximum number of concurrent sessions.

Per-IP traffic shaping enables you limit the behavior of every member of a policy to avoid one user from using all the available bandwidth - it now is shared within a group equally. Using a per-IP shaper avoids having to create multiple policies for every user you want to apply a shaper



Per-IP traffic shaping is not supported over NP2 interfaces.

Per-IP traffic shaping configuration settings

To configure per-IP traffic shaping go to *Firewall Objects > Traffic Shaper > Per-IP*. and select *Create New*.

Name	Enter a name for the per-IP traffic shaper.
Maximum Bandwidth	The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the shaper. Setting Maximum Bandwidth to 0 (zero) provides unlimited bandwidth.
Maximum Concurrent Connections	Enter the maximum allowed concurrent connection.
Forward DSCP Reverse DSCP	Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see “Differentiated Services” .

Example

The following steps creates a Per-IP traffic shaper called “Accounting” with a maximum traffic amount of 720,000 Kb/s, and the number of concurrent sessions of 200.

To create the shared shaper - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaper > Per-IP*.
- 2 Select *Create New*.
- 3 Enter the *Name* *Accounting*.
- 4 Select the *Maximum Bandwidth* check box and enter the value 720000.

- 5 Select the *Maximum Concurrent Sessions* check box and enter the value 200.
- 6 Select *OK*.

To create the shared shaper - CLI

```
config firewall shaper per-ip-shaper
edit Accounting
    set max-bandwidth 720000
    set max-concurrent-sessions 200
end
```

Application control shaping

Traffic shaping is also possible for specific applications for both shared and per IP shaping. Through the *UTM Profiles > Application Control* feature, you can configure a specific application's maximum bandwidth. When configuring the application control features, if the application is set to pass, you can set the traffic shaping options. The shapers available are those set up in the *Firewall Objects > Traffic Shaping* menu.

For more information on configuring application control shapers, see the [UTM Guide](#).

Example

This example sets the traffic shaping definition for Facebook to a medium priority, a default traffic shaper.

To add traffic shaping for Facebook - web-based manager

- 1 Go to *UTM Objects > Application Control > Application Sensor*.
- 2 Select the *Create New* "Plus" icon in the upper right corner of the screen to create a new application group, and enter the name *Web*.
- 3 Select *OK*.
- 4 Select *Create New*.
- 5 Select *Web* from the *Category* drop-down list.
- 6 Select *Facebook* from the *Application* drop-down list.
- 7 Select *Monitor* for the *Action*.
- 8 Select *Traffic Shaping* and select medium-priority from the drop-down list.
- 9 Select *OK*.

To add traffic shaping for Facebook - CLI

```
config application list
edit web
    config entries
    edit 1
        set category 12
        set application 15832
        set action pass
        set shaper medium-priority
    end
end
end
```


Enabling in the security policy

All traffic shapers are enabled within a security policy, including the Application Control shapers. As such, the shapers are in effect after any DoS sensor policies, and before any routing or packet scanning occurs.

To enable traffic shaping - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New* or select an existing policy and select *Edit*.
- 3 Select *Traffic Shaping*.
- 4 Select the shaping option and select the shaper from the drop-down list.
- 5 Select *OK*.

To enable traffic shaping - CLI

```
config firewall policy
  edit <policy_number>
  ...
  set traffic-shaper <shaper_name>
  set per-ip-shaper <shaper_name>
end
```

Reverse direction traffic shaping

The shaper you select for the security policy (shared shaper) will affect the traffic in the direction defined in the policy. For example, if the source port is port 1 and the destination is port 3, the shaping affects the flow in this direction only. By selecting *Reverse Direction Traffic Shaping*, you can define the traffic shaper for the policy in the opposite direction. In this example, from port 3 to port 1.

To add a reverse shaper

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Traffic Shaping*.
- 3 Select *Shared Traffic Shaper Reverse Direction* and select the shaper from the list.
- 4 Select *OK*.

Setting the reverse direction only

There may be instances where you only need to have the traffic shaping for incoming connections. That is, the “reverse” direction to the typical traffic shaper. Using the CLI, you can enable a reverse-only shaper from the your configured list of shapers.

To add a reverse shaper - web-based manager

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Traffic Shaping*.
- 3 Select *Shared Traffic Shaper Reverse Direction* and select the shaper from the list.
- 4 Select *OK*.

To configure a reverse-only shaper - CLI

```
config firewall policy
  edit <policy_number>
```

```

...
set traffic-shaper-reverse <shaper_name>
end

```

Application control shaper

Application control shapers are in effect within the application control profile. Within the security policy options, select *UTM* then *Application Control* and select the application from the list.

Type of Service priority

Type of service (ToS) is an 8-bit field in the IP header that enables you to determine how the IP datagram should be delivered, using criteria of Delay, Throughput, Priority, Reliability, and Cost. Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority ToS is 0, the highest is 7 when bits 3, 4, and 5 are all set to 1. There are 4 other bits that are seldom used or reserved that are not included here.

Together these bits are the `tos` variable of the `tos-based-priority` command. The router tries to match the ToS of the datagram to the ToS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero ToS route. Using increased quality may increase the cost of delivery because better performance may consume limited network resources.

Each bit represents the priority as per RFC 1349:

- 1000 - minimize delay
- 0100 - maximize throughput
- 0010 - maximize reliability
- 0001 - minimize monetary cost

The TOS value is set in the CLI using the commands:

```

config system tos-based-priority
edit <sequence_number>
set tos [0-15]
set priority [high | medium | low]
end

```

Where `tos` is the value of the type of service bit in the IP datagram header with a value between 0 and 15, and `priority` is the priority of this type of service priority. These priority levels conform to the firewall traffic shaping priorities, as defined in [RFC 1349](#).

For example, if you want to configure the FortiGate unit so that reliability is the first priority, set the `tos` value to 4.

```

config system tos-based-priority
edit 1
set tos 4
set priority high
end

```

For a list of ToS values and their DSCP equivalents see [“Tos and DSCP mapping” on page 2275](#).

Example

```
config system tos-based-priority
edit 1
    set tos 1
    set priority low
next
edit 4
    set tos 4
    set priority medium
next
edit 6
    set tos 6
    set priority high
next
end
```

TOS in FortiOS

Traffic shaping and TOS follow the following sequence:

- 1 The CLI command `tos-based-priority` acts as a `tos-to-priority` mapping. FortiOS maps the TOS to a priority when it receives a packet.
- 2 Traffic shaping settings adjust the packet's priority according to the traffic.
- 3 Deliver the packet based on its priority.

Differentiated Services

Differentiated Services describes a set of end-to-end Quality of Service (QoS) capabilities. End-to-end QoS is the ability of a network to deliver service required by specific network traffic from one end of the network to another. By configuring differentiated services, you configure your network to deliver particular levels of service for different packets based on the QoS specified by each packet.

Differentiated Services (also called DiffServ) is defined by RFC 2474 and 2475 as enhancements to IP networking to enable scalable service discrimination in the IP network without the need for per-flow state and signaling at every hop. Routers that can understand differentiated services sort IP traffic into classes by inspecting the DS field in IPv4 header or the Traffic Class field in the IPv6 header.

You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet.

If the differentiated services feature is not enabled, the FortiGate unit treats traffic as if the DSCP value is set to the default (00), and will not change IP packets' DSCP field. DSCP values are also not applied to traffic if the traffic originates from a FortiGate unit itself.

The FortiGate unit applies the DSCP value and IPsec encryption to the differentiated services (formerly TOS) field in the first word of the IP header. The typical first word of an IP header, with the default DSCP value, is 4500:

- 4 for IPv4
- 5 for a length of five words
- 00 for the default DSCP value

You can change the packet's DSCP field for traffic initiating a session (forward) or for reply traffic (reverse) and enable each direction separately and configure it in the security policy.



Changes to DSCP values in a security policy effect new sessions. If traffic must use the new DSCP values immediately, clear all existing sessions.

DSCP is enabled using the CLI command:

```
config firewall policy
  edit <policy_number>
    ...
    set diffserv-forward enable
    set diffservcode-forward <binary_integer>
    set diffserv-reverse enable
    set diffservcode-rev <binary_integer>
  end
```

For more information on the different DSCP commands, see the examples below and the [CLI Reference](#).




If you only set `diffserv-forward` and `diffserv-reverse` without setting the corresponding `diffservcode` values, the FortiGate unit will reset the bits to zero.

For a list of DSCP values and their ToS equivalents see [“Tos and DSCP mapping” on page 2275](#).

DSCP values can also be defined within a shared shaper as a single value, and per-IP shaper for forward and reverse directions.

DSCP examples

For all the following DSCP examples, the FortiGate and client PC configuration is the following diagram and used firewall-based DSCP configurations.



Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through a FortiGate unit. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:


```
config firewall policy
  edit 2
    set srcintf port6
    set dstintf port3
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set diffserv-forward enable
    set diffservcode-forward 101110
  end
```

As a result, FortiGate A changes the DSCP field for outgoing traffic, but not to its reply traffic. The binary DSCP values used map to the following hexadecimal

TOS field values, which are observable by a sniffer (also known as a packet tracer):

- DSCP 000000 is TOS field 0x00
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)

If you performed an ICMP ping between User 1 and User 2, the following output illustrates the IP headers for the request and the reply by sniffers on each of FortiGate unit's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



Example


In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:

```
config firewall policy
  edit 2
    set srcintf port6
    set dstintf port3
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY"
    set diffserv-forward enable
    set diffserv-rev enable
    set diffservcode-forward 101110
    set diffservcode-rev 101111
  end
```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic. The binary DSCP values in map to the following hexadecimal TOS field values, which are observable by a sniffer (also known as a packet tracer):

- DSCP 000000 is TOS field 0x00
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field 0xbc

If you performed an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```
config firewall policy
  edit 2
```

```

set srcintf port6
set dstintf port3
set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY
set diffserv-forward enable
set diffserv-rev enable
set diffservcode-forward 101110
set diffservcode-rev 101111
end

```

FortiGate B contains the following configuration:

```


config firewall policy
edit 2
set srcintf wan2
set dstintf internal
set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY
set diffserv-rev enable
set diffservcode-rev 101101
end

```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, and FortiGate B changes the DSCP field only for reply traffic. The binary DSCP values in this configuration map to the following hexadecimal TOS field values:

- DSCP 000000 is TOS field 0x00
- DSCP 101101 is TOS field 0xb4
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field 0xbc

If you performed an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



Example

In this example, HTTPS and DNS traffic is sent from User 1 to FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```

config firewall policy
edit 2
set srcintf port6
set dstintf port3

```

```

set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY
set diffserv-forward enable
set diffserv-rev enable
set diffservcode-forward 101110
set diffservcode-rev 101111
end

```

FortiGate B contains the following configuration:

```


config firewall policy
edit 2
set srcintf wan2
set dstintf internal
set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY
set diffserv-rev enable
set diffservcode-rev 101101
end

```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, but FortiGate B changes the DSCP field only for reply traffic which passes through its internal interface. Since the example traffic does not pass through the internal interface, FortiGate B does not mark the packets. The binary DSCP values in this configuration map to the following hexadecimal TOS field values:

- DSCP 000000 is TOS field **0x00**
- DSCP 101101 is TOS field **0xb4**, which is configured on FortiGate B but not observed by the sniffer because the example traffic originates from the FortiGate unit itself, and therefore does not match that security policy.
- DSCP 101110 is TOS field **0xb8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field **0xbc**

If you sent HTTPS or DNS traffic from User 1 to FortiGate B, the following would illustrate the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the TOS field, which contains the DSCP value.



Tos and DSCP mapping

The table below lists the mapping of DSCP and ToS hexadecimal values for each service for QoS.

Table 121: ToS to DSCP mappings

Service Class	DSCP Bits	DSCP Value	ToS Value	ToS Hexidecimal
Network Control	111000	56-63	224	0xE0
Internetwork Control	110000	48-55	192	0xC0
Critical - Voice Data (RTP)	101110	46	184	0xB8
	101000	40	160	0xA0
Flash Override Video Data	100010	34	136	0x88
	100100	36	144	0x90
	100110	38	152	0x98
	100000	32	128	0x80
Flash Voice Control	011010	26	104	0x68
	011100	28	112	0x70
	011110	30	120	0x78
	011000	24	96	0x60
Immediate Deterministic (SNA)	010010	18	72	0x48
	010100	20	80	0x50
	010110	22	88	0x58
	010000	16	64	0x40
Priority Controlled Load	001010	10	40	0x28
	001100	12	48	0x30
	001110	14	56	0x38
	001000	8	32	0x20
Routine - Best Effort	000000	0	0	0x00
Routine - Penalty Box	000010	2	8	0x08

Traffic Shaper Monitor

You can view statistical information about traffic shapers and their bandwidth from *Firewall Objects > Monitor > Traffic Shaper Monitor*.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Report By	Select to display dropped packets or current bandwidth. The bar chart changes its name so that you know what current information is being displayed.
Traffic Shaper Usage Dropped packets	The bar chart displays the packets that were dropped by traffic shaper.
Traffic Shaper Usage Current Bandwidth	The bar chart displays the current bandwidth of traffic shapers.



Examples

While it is possible to configure QoS using a combination of security policies and in ToS-based priorities, and to distribute traffic over all six of the possible queues for each physical interface, the results of those configurations can be more difficult to analyze due to their complexity. In those cases, prioritization behavior can vary by several factors, including traffic volume, ToS (type of service) or differentiated services markings, and correlation of session to a security policy.

The following simple examples illustrate QoS configurations using either prioritization by security policy, or prioritization by ToS bit, but not both. The examples also assume you are not configuring traffic shaping for interfaces that receive hardware acceleration from network processing units (NPU).

QoS using priority from security policies

Configurations implementing QoS using the priority values defined in security policies are capable of applying bandwidth limits and guarantees.

In addition to configuring traffic shaping, you may also choose to limit bandwidth accepted by each interface. This can be useful in scenarios where bandwidth being received on source interfaces frequently exceeds the maximum bandwidth limit defined in the security policy. In this case, rather than wasting processing power on packets that will only be dropped later in the processing to enforce those limits, you may choose to preemptively police the traffic.

Note that if you implement QoS using security policies rather than ToS bit, the FortiGate unit applies QoS to all packets controlled by the policy. Control is less granular than prioritization by ToS bit, but has the benefits of correlating quality of service to a security policy, enabling you to distribute traffic over up to four of the possible 6 priority queues (queue 0 to queue 3), not requiring other devices in your network to set or respect the ToS bit, and of enabling you to configure bandwidth limits and guarantees.

In this example, we limit the bandwidth accepted by each source interface, limit the bandwidth used by sessions controlled by the security policy, and then configure prioritized queuing on the destination interface based upon the priority in the security policy, subject to alternative assignment to queue 0 when necessary to achieve the guaranteed packet rate.

To limit bandwidth accepted by an interface

In the CLI, enter the following commands:

```
config system interface
  edit <name_str>
    set inbandwidth <rate_int>
  next
end
```

where <rate_int> is the bandwidth limit in Kb/s. Excess packets will be dropped.

To configure bandwidth guarantees, limits, and priorities

- 1 Go to *Firewall Objects > Traffic Shaper > Shared*, and select *Create New*.
- 2 Enter a name for the shaper.
- 3 Enter the *Guaranteed Bandwidth*, if any.
Bandwidth guarantees affect prioritization. While packet rates are less than this rate, they use priority queue 0. If this is not the effect you intend, consider entering a small guaranteed rate, or enter 0 to effectively disable bandwidth guarantees.
- 4 Enter Maximum Bandwidth.
Packets greater than this rate will be discarded.
- 5 Select the *Traffic Priority*.
High has a priority value of 1, while Low is 3. While the current packet rate is below Guaranteed Bandwidth, the FortiGate unit will disregard this setting, and instead use priority queue 0.
- 6 Select *OK*.

Sample configuration

This sample configuration limits ingress bandwidth to 500 Kb/s. It also applies separate traffic shapers to FTP and HTTP traffic. In addition to the interface bandwidth limit, HTTP traffic is subject to a security policy bandwidth limit of 200 Kb/s.

All egressing FTP traffic greater than 10 Kb/s is subject to a low priority queue (queue 3), while all egressing HTTP traffic greater than 100 Kb/s is subject to a medium priority queue (queue 2). That is, unless FTP traffic rates are lower than their guaranteed rate, and web traffic rates are greater than their guaranteed rate, FTP traffic is lower priority than web traffic.

Traffic less than these guaranteed bandwidth rates use the highest priority queue (queue 0).

Set the inbound bandwidth limits. This setting is only available in the CLI:

```
config system interface
  edit wan1
    set inboundwidth 500
  next
end
```

Create the traffic shapers for FTP and HTTP.

To configure the shapers - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaper > Shared*, and select *Create New*.
- 2 Enter FTP for the name of the shaper.
- 3 Enter the *Guaranteed Bandwidth*, of 10 Kbps.
- 4 Enter Maximum Bandwidth of 500 Kbps.
- 5 Select the *Traffic Priority* of Low.
- 6 Select *OK*.
- 7 Select *Create New*.
- 8 Enter HTTP for the name of the shaper.
- 9 Enter the *Guaranteed Bandwidth*, of 100 Kbps.
- 10 Enter Maximum Bandwidth of 200 Kbps.

11 Select the *Traffic Priority* of *Medium*.

12 Select *OK*.

To configure the shapers - CLI

```
config firewall shaper traffic-shaper
edit FTP
    set maximum-bandwidth 500
    set guaranteed-bandwidth 10
    set per-policy enable
    set priority low
end
next
edit HTTP
    set maximum-bandwidth 200
    set guaranteed-bandwidth 100
    set per-policy enable
    set priority medium
end
```

QoS using priority from ToS or differentiated services

Configurations implementing QoS using the priority values defined in either global or specific ToS bit values are not capable of applying bandwidth limits and guarantees, but are capable of prioritizing traffic at per-packet levels, rather than uniformly to all services matched by the security policy.

In addition to configuring traffic prioritization, you may also choose to limit bandwidth being received by each interface. This can sometimes be useful in scenarios where you want to limit traffic levels, but do not want to configure traffic shaping within a security policy. This has the benefit of policing traffic at a point before the FortiGate unit performs most processing.

Note that if you implement QoS using ToS octet rather than security policies, the FortiGate unit applies QoS on a packet by packet basis, and priorities may be different for packets and services controlled by the same security policy. This is more granular control than prioritization by security policies, but has the drawbacks that quality of service is may not be uniform for multiple services controlled by the same security policy, packets will only use up to three of the six possible queues (queue 0 to queue 2), and bandwidth cannot be guaranteed. Other devices in your network must also be able to set or preserve ToS bits.

In this example, we limit the bandwidth accepted by each source interface, and then configure prioritized queuing on the destination interface based upon the value of the ToS bit located in the IP header of each accepted packet.

To limit bandwidth accepted by an interface, in the CLI, enter the following commands:

```
config system interface
edit <name_str>
    set inbandwidth <rate_int>
next
end
```

where <rate_int> is the bandwidth limit in Kb/s. Excess packets will be dropped.

To configure priorities, in the CLI, configure the global priority value using the following commands:

```
config system global
    set tos-based-priority {high | low | medium}
end
```

where `high` has a priority value of 0 and `low` is 2.

If you want to prioritize some ToS bit values differently than the global ToS-based priority, configure the priority for packets with that ToS bit value using the following commands:

```
config system tos-based-priority
    edit <id_int>
        set tos [0-15]
        set priority {high | low | medium}
    next
end
```

where `id_int` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in this location will override the global ToS-based priority.

Sample configuration

This sample configuration limits ingress bandwidth to 500 Kb/s. It also queues egress traffic based upon the ToS bit in the IP header of ingress packets.

Unless specified for the packet's ToS bit value, packets use the low priority queue (queue 2). For ToS bit values 4 and 15, the priorities are specified as medium (value 1) and high (value 0), respectively.

```
config system interface
    edit wan1
        set inbandwidth 500
    next
end
config system global
    set tos-based-priority low
end
config system tos-based-priority
    edit 4
        set tos 4
        set priority medium
    next
    edit 15
        set tos 15
        set priority high
    next
end
```

Example setup for VoIP

In this example, there are three traffic shaping requirements for a network:

- Voice over IP (VoIP) requires a guaranteed, high-priority for bandwidth for telephone communications.
- FTP bursts must be contained so as not to consume any available bandwidth. As such this traffic needs to be throttled to a smaller amount.
- A consistent bandwidth requirement is needed for all other email and web-based traffic.

To enable this requirement, you need to create three separate shapers and three security policies for each traffic type.



For this example, the actual values are not actual values, they are used for the simplicity of the example.

Creating the traffic shapers

First create the traffic shapers that define the maximum and guaranteed bandwidth. The shared shapers will be used, some with per-policy and some all policies as shown in the table, to better control traffic.

VoIP shaper

The VoIP functionality is a key component to the business as a communication tool and as such requires a guaranteed bandwidth.

To create a VoIP shaper - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaping > Shared*.
- 2 Enter the *Name* `voip`.
- 3 Select *Per Policy*.
- 4 Enter the *Maximum Bandwidth* of 1000 Kb/s
- 5 Enter the *Guaranteed Bandwidth* of 800 Kb/s.
- 6 Select a *Traffic Priority* of *High*.
- 7 Select *OK*.

To create a VoIP shaper - CLI

```
config firewall shaper traffic-shaper
edit voip
set maximum-bandwidth 1000
set guaranteed-bandwidth 800
set per-policy enable
set priority high
end
```

This ensures that whatever number of policies use this shaper, the defined bandwidth will always be the same. At the same time, the bandwidth is continually guaranteed at 800 Kb/s but if available can be as much as 1000 Kb/s. Setting the priority to high ensures that the FortiGate unit always considers VoIP traffic as the most important.

FTP shaper

The FTP shaper sets the maximum bandwidth to use to avoid sudden spikes by sudden uploading or downloading of large files, and interfering with other more important traffic.

To create a FTP shaper - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaping > Shared*.
- 2 Enter the *Name* `ftp`.
- 3 Select *For all Policies Using This Shaper*.
- 4 Enter the *Maximum Bandwidth* of 200 Kb/s

- 5 Enter the *Guaranteed Bandwidth* of 200 Kb/s.
- 6 Select a *Traffic Priority* of Low.
- 7 Select OK.

To create a FTP shaper - CLI

```
config firewall shaper traffic-shaper
edit ftp
set maximum-bandwidth 200
set guaranteed-bandwidth 200
set priority low
end
```

For this shaper, the maximum and guaranteed bandwidth are set low and to the same value. In this case, the bandwidth is restricted to a specific amount. By also setting the traffic priority low ensures more important traffic will be able to pass before FTP traffic.

Regular traffic shaper

The regular shaper sets the maximum bandwidth and guaranteed bandwidth for everyday business traffic such as web and email traffic.

To create a regular shaper - web-based manager

- 1 Go to *Firewall Objects > Traffic Shaping > Shared*.
- 2 Enter the *Name* `daily_traffic`.
- 3 Select *Per Policy*.
- 4 Enter the *Maximum Bandwidth* of 600 Kb/s
- 5 Enter the *Guaranteed Bandwidth* of 600 Kb/s.
- 6 Select a *Traffic Priority* of Medium.
- 7 Select OK.

To create a regular shaper - CLI

```
config firewall shaper traffic-shaper
edit daily_traffic
set maximum-bandwidth 600
set guaranteed-bandwidth 600
set per-policy enable
set priority medium
end
```

For this shaper, the maximum and guaranteed bandwidth are set to a moderate value of 600 Kb/s. It is also set for per policy, which ensures each security policy for day-to-day business traffic has the same distribution of bandwidth.

Creating security policies

To employ the shaper, create security policies that use the shapers within the policies. Create a separate policy for each service and enable traffic shaping. For example, a policy for FTP traffic, a policy for SIP and so on.

For the following steps the VoIP traffic shaper is enabled as well as the reverse direction option. This ensures that return traffic for a VoIP call has the same guaranteed bandwidth as the outgoing call.

To enable traffic shaping in the security policy - web-based manager


- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following and select:

Source interface/Zone	Internal
Source address	All
Destination interface/Zone	WAN1
Destination address	All
Schedule	always
Service	SIP
Action	ACCEPT

- 3 Select *Traffic Shaping*.
- 4 From the drop-down list, select the voip shaper created in the previous steps.
- 5 Select *Reverse Direction Traffic Shaping*.
- 6 Select *OK*.

To enable traffic shaping in the security policy - CLI

```
config firewall policy
  edit 6
    set srcintf internal
    set scraddr all
    set dstintf wan1
    set dstaddr all
    set action accept
    set schedule always
    set service sip
    set traffic-shaper voip
    set reverse-traffic-shaper voip
  end
```

Troubleshooting

This chapter outlines some troubleshooting tips and steps to diagnose the shapers and whether they are working correctly. These diagnose commands include:

- `diagnose system tos-based-priority`
- `diagnose firewall shaper traffic-shaper`
- `diagnose firewall per-ip-shaper`
- `diagnose debug flow`

Interface diagnosis

To optimize traffic shaping performance, first ensure that the network interface's Ethernet statistics are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```

Shaper diagnose commands

There are specific diagnose commands you can use to verify the configuration and flow of traffic, including packet loss due to the employed shaper.

All of these diagnose troubleshooting commands are supported in both IPv4 and IPv6.

TOS command

Use the following command to list command to view information of the TOS lists and traffic.

```
diagnose system tos-based-priority
```

This example displays the priority value currently correlated with each possible TOS bit value. Priority values are displayed in order of their corresponding TOS bit values, which can range between 0 and 15, from lowest TOS bit value to highest.

For example, if you have not configured TOS-based priorities, the following appears...

```
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

...reflecting that all packets are currently using the same default priority, high (value 0).

If you have configured a TOS-based priority of `low` (value 2) for packets with a ToS bit value of 3, the following appears...

```
0 0 0 2 0 0 0 0 0 0 0 0 0 0 0 0
```

...reflecting that most packets are using the default priority value, except those with a ToS bit value of 3.

Shared shaper

To view information for the shared traffic shaper for security policies enter the command

```
diagnose firewall shaper traffic-shaper list
```

The resultant output displays the information on all available shapers. The more shapers available the longer the list. For example:

```
name Throughput
maximum-bandwidth 1200000 Kb/sec
guaranteed-bandwidth 50000 Kb/sec
current-bandwidth 0 B/sec
priority 1
packets dropped 0
```

Additional commands include:

`diagnose firewall shaper traffic-shaper state` - provides the total number of traffic shapers on the FortiGate unit.

`diagnose firewall shaper traffic-shaper stats` - provides summary statistics on the shapers. Sample output looks like the following:

```
shapers 9 ipv4 0 ipv6 0 drops 0
```

Per-IP shaper

To view information for the per-IP shaper for security policies enter the command

```
diagnose firewall shaper per-ip-shaper list
```

The resultant output displays the information on all available per-IP shapers. The more shapers available the longer the list. For example:

```
name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 0
```

Additional commands include:

`diagnose firewall shaper per-ip-shaper state` - provides the total number of per-ip shapers on the FortiGate unit.

`diagnose firewall shaper per-ip-shaper stats` - provides summary statistics on the shapers. Sample output looks like the following:

```
memory allocated 3 packet dropped: 0
```

You can also clear the per-ip statistical data to begin a fresh diagnoses using:

```
diagnose firewall shaper per-ip-shaper clear
```

Packet loss with statistics on shapers

For each shaper there are counters that allow to verify if packets have been discarded. To view this information, in the CLI, enter the command `diagnose firewall shaper`.

The results will look similar to the following output:

```
diagnose firewall shaper traffic-shaper list
```

```
name limit_GB_25_MB_50_LQ
maximum-bandwidth 50 Kb/sec
guaranteed-bandwidth 25 Kb/sec
current-bandwidth 51 Kb/sec
priority 3
dropped 1291985
```

The diagnose command output is different if the shapers are configured either per-policy or shared between policies.

For per-IP the output would be:

```
diagnose firewall shaper per-ip-shaper list

name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 3264220
```

Packet lost with the debug flow

When using the debug flow diagnostic command, there is a specific message information that a packet has exceed the shaper limits and therefor discarded:

```
diagnose debug flow show console enable
diagnose debug flow filter addr 10.143.0.5
diagnose debug flow trace start 1000
```

```
id=20085 trace_id=11 msg="vd-root received a packet(proto=17,
10.141.0.11:3735->10.143.0.5:5001) from port5."
id=20085 trace_id=11 msg="Find an existing session, id=0000eabc,
original direction"
id=20085 trace_id=11 msg="exceeded shaper limit, drop"
```

Session list details with dual traffic shaper

When a Security Policy has a different traffic shaper for each direction, it is reflected in the session list output from the CLI:

```
diagnose system session list

session info: proto=6 proto_state=02 expire=115 timeout=3600
flags=00000000 sock
flag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=Limit_25Mbps prio=1 guarantee 25600/sec max
204800/sec traffic 48/sec
reply-shaper=Limit_100Mbps prio=1 guarantee 102400/sec max
204800/sec traffic 0/sec
ha_id=0 hakey=44020
policy_dir=0 tunnel=/
state=may_dirty rem os rs
statistic(bits/packets/allow_err): org=96/2/1 reply=0/0/0 tuples=2
origin->sink: org pre->post, reply pre->post dev=2->3/3->2
gwy=10.160.0.1/0.0.0.0
hook=pre dir=org act=dnat 192.168.171.243:2538-
>192.168.182.110:80(10.160.0.1:80)
hook=post dir=reply act=snat 10.160.0.1:80-
>192.168.171.243:2538(192.168.182.110:80)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=2 auth_info=0 chk_client_info=0 vd=0
serial=00011e81 tos=ff/ff app=0 dd_type=0 dd_rule_id=0
```

Additional Information

- Packets discarded by the shaper impact flow-control mechanisms like TCP. For more accurate testing results prefer UDP protocol.
- Traffic shaping accuracy is optimum for security policies without a protection profile where no FortiGate content inspection is processed.
- Do not oversubscribe an outbandwidth throughput. For example, $\text{sum}[\text{guaranteed BW}] < \text{outbandwidth}$. For accuracy in bandwidth calculation, it is required to set the “outbandwidth” parameter on the interfaces. For more information see [“Bandwidth guarantee, limit, and priority interactions” on page 2253](#).
- The FortiGate unit is not prioritizing traffic based on the DSCP marking configured in the security policy. However, TOS based prioritizing can be made at ingress. For more information see [“Differentiated Services” on page 2269](#).



Chapter 14 FortiOS Carrier

This FortiOS Handbook chapter contains the following sections:

[Overview of FortiOS Carrier features](#) provides an overview of the three major topics for FortiOS Carrier — Dynamic Profiles, MMS, and GTP.

[Carrier web-based manager settings](#) describes the web-based manager interface of FortiOS Carrier specific features.

[MMS UTM features](#) describes FortiOS UTM features as they apply to MMS including MMS virus scanning, MMS file filtering, MMS content-based Antispam protection, and MMS DLP archiving.

[Message flood protection](#) describes setting thresholds to protect your MMS servers from receiving too many messages from the same sender.

[Duplicate message protection](#) describes setting thresholds to protect your MMS servers from receiving the same message from more than one sender.

[MMS Replacement messages](#) describes customizing MMS replacement messages.

[Configuring GTP on FortiOS Carrier](#) explains configuration of the more basic FortiOS Carrier GTP features.

[GTP message type filtering](#) explains this feature, and how to configure it on FortiOS Carrier.

[GTP identity filtering](#) explains this feature, and how to configure it on FortiOS Carrier.

[Troubleshooting](#) provides answer to common FortiOS Carrier GTP issues.



Overview of FortiOS Carrier features

FortiOS Carrier specific features include Multimedia messaging service (MMS) protection, and GPRS Tunneling Protocol (GTP) protection.

This section includes:

- [Overview](#)
- [MMS background](#)
- [How FortiOS Carrier processes MMS messages](#)
- [MMS protection profiles](#)
- [Bypassing MMS protection profile filtering based on user's carrier end points](#)
- [Applying MMS protection profiles to MMS traffic](#)
- [GTP basic concepts](#)
- [Parts of a GTPv1 network](#)
- [GPRS network common interfaces](#)
- [Packet flow through the GPRS network](#)

Overview

FortiOS Carrier provides all the features found on FortiGate units plus added features specific to carrier networks. These features include:

- [MMS](#)
- [GTP](#)

MMS

MMS is a standard for sending messages that include multimedia content between mobile phones. MMS is also popular as a method of delivering news and entertainment content including videos, pictures, and text. Carrier networks include four different MMS types of messages — MM1, MM3, MM4, and MM7. See [“MMS background” on page 2292](#).

GTP

The GPRS Tunneling Protocol (GTP) runs on GPRS carrier networks. GPRS is a GSM packet radio standard. It provides more efficient usage of the radio interface so that mobile devices can share the same radio channel. FortiOS supports GTPv1 release 7.15.0 and GTPv1 release 8.12.0.

GPRS provides direct connections to the Internet (TCP/IP) and X.25 networks for point-to-point services (connection-less/connection oriented) and point-to-multipoint services (broadcast).

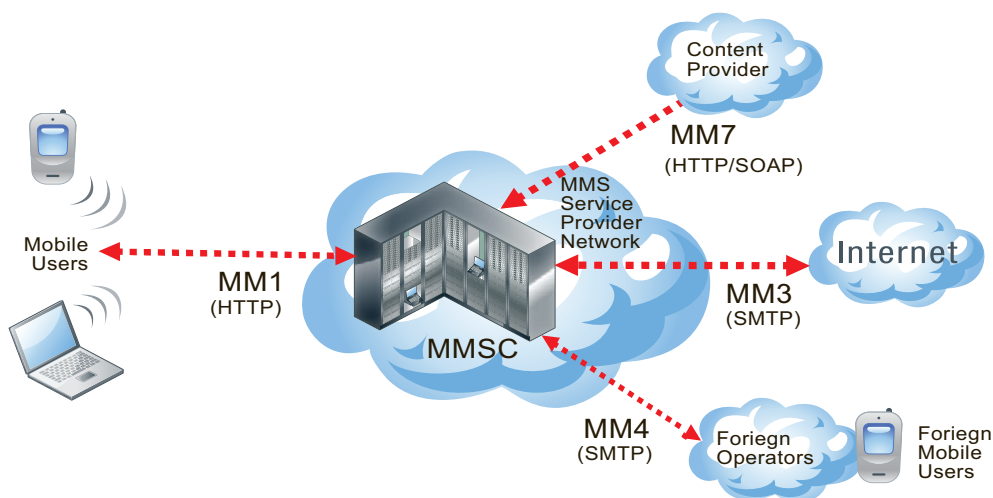
GPRS currently supports data rates from 9.6 kbps to more than 100 kbps, and it is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based) that converts the message from radio to wired, and sends the message to the carrier network and eventually the Internet (wired carrier network). See [“GTP basic concepts” on page 2305](#).

MMS background

MMS is a common method for mobile users to send and receive multimedia content. A Carrier network supports MMS across its network. This makes up the MMS Service Provider Network (MSPN).

Messages can be sent or received between the MMSC and a number of other services including the Internet, content providers, or other carriers. Each of these different service connections uses different MMS formats including MM1 and MM7 MMS messages (essentially HTTP format), and MM3 and MM4 messages (SMTP formatted). These different formats reflect the different purposes and content for each type of MMS message.

Figure 232: MMS content interfaces



MMS content interfaces

MMS messages are sent from devices and servers to other devices and servers using MMS content interfaces

There are eight interfaces defined for the MMS standard, referred to as MM1 through MM8. The most important of these interfaces for the transfer of data is the MM1 interface, as this defines how mobile users communicate from the mobile network to the Multimedia Message Service Center (MMSC). MMS content to be monitored and controlled comes from these mobile users and is going to the provider network.

Other MMS content interfaces that connect a service provider network to other external sources can pose threats as well. MM3 handles communication between the Internet and the MMSC and is a possible source of viruses and other content problems from the Internet. MM4 handles communication between different content provider MMSCs. Filtering MM4 content protects the service provider network from content sent from foreign service providers and their subscribers. Finally MM7 is used for communication between content providers and the MMSC. Filtering MM3 content can also keep harmful content off of the service provider network.

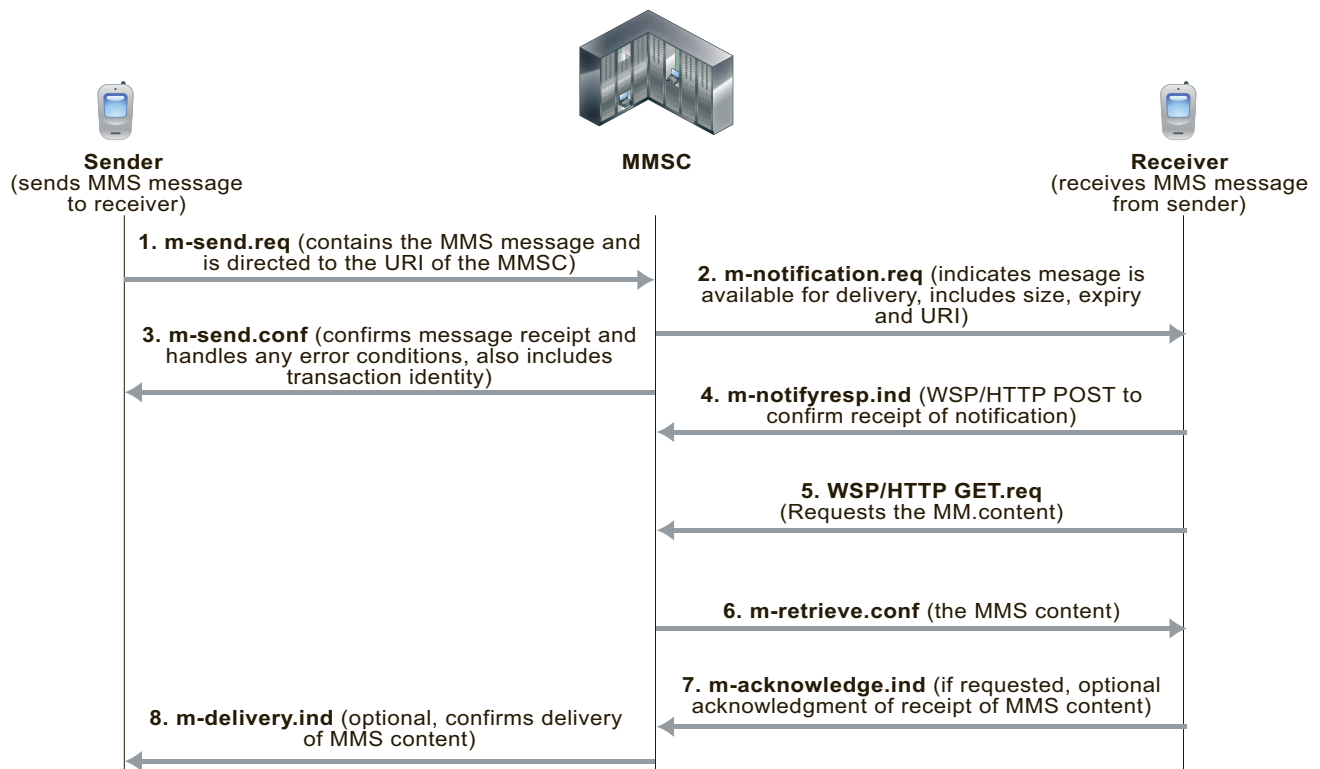
Table 122: MMS content interfaces that

Type	Transaction	Similar to
MM 1	Handset to MMSC	HTTP
MM 3	Between MMSC and Internet	SMTP
MM 4	Between Operator MMSCs	SMTP
MM 7	Content Providers to MMSC	HTTP and SOAP

How MMS content interfaces are applied

As shown in [Figure 233](#), the sender's mobile device encodes the MMS content in a form similar to MIME email message (MMS MIME content formats are defined by the MMS Message Encapsulation specification). The encoded message is then forwarded to the service provider's MMSC. Communication between the sending device and the MMSC uses the MM1 content interface. The MM1 content interface establishes a connection and sends an MM1 send request (`m-send.req`) message that contains the MMS message. The MMSC processes this request and sends back an MM1 send confirmation (`m-send.conf`) HTTP response indicating the status of the message — accepted or an error occurred, for example.

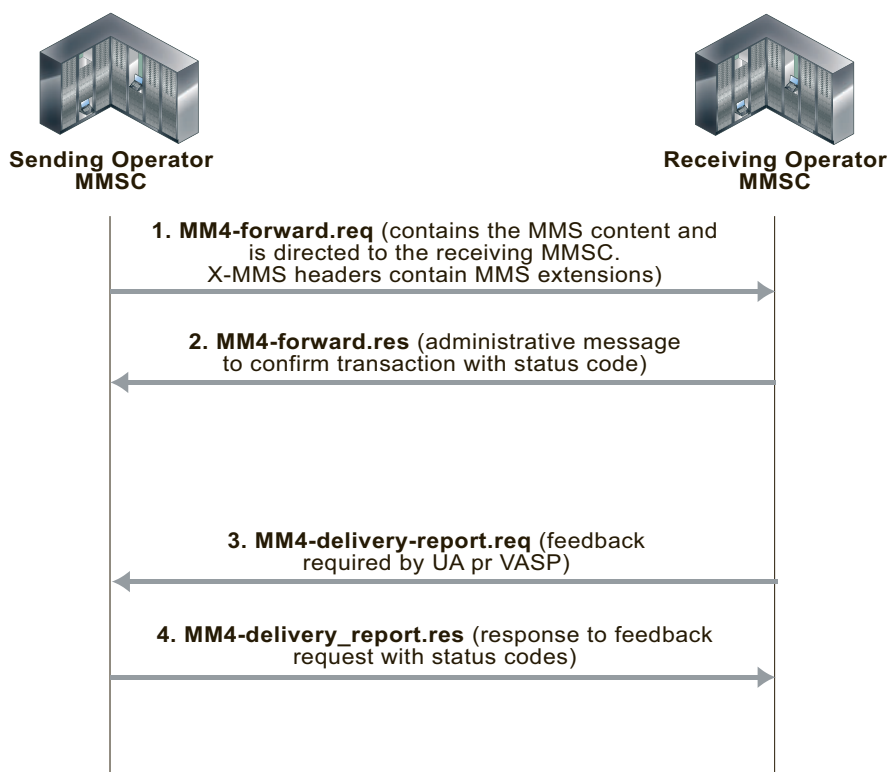
Figure 233: MM1 transactions between senders and receivers and the MMSC



If the recipient is on another carrier, the MMSC forwards the message to the recipient's carrier. This forwarding uses the MM4 content interface for forwarding content between operator MMSCs (see [Figure 234](#)).

Before the MMSC can forward the message to the final recipient, it must first determine if the receiver's handset can receive MMS messages using the MM1 content interface. If the recipient can use the MM1 content interface, the content is extracted and sent to a temporary storage server with an HTTP front-end.

To retrieve the message, the receiver's handset establishes a connection with the MMSC. An HTTP get request is then sent from the recipient to the MMSC. This message contains the URL where the content of the message is stored. The MMSC responds with a retrieve confirmation (**m-retrieve.conf**) HTTP response that contains the message.

Figure 234: MM4 messages sent between operator MMSCs


This causes the receiver's handset to retrieve the content from the embedded URL. Several messages are exchanged to indicate status of the delivery attempt. Before delivering content, some MMSCs also include a content adaptation service that attempts to modify the multimedia content into a format suitable for the recipient's handset.

If the receiver's handset is not MM1 capable, the message can be delivered to a web based service and the receiver can view the content from a normal Internet browser. The URL for the content can be sent to the receiver in an SMS text message. Using this method, non-MM1 capable recipients can still receive MMS content.

Email and web-based gateways from MMSC to the Internet use the MM3 content interface. On the receiving side, the content servers can typically receive service requests both from WAP and normal HTTP browsers, so delivery via the web is simple. For sending from external sources to handsets, most carriers allow MIME encoded message to be sent to the receiver's phone number with a special domain.

How FortiOS Carrier processes MMS messages

MMS messages can be vectors for propagating undesirable content such as spam and viruses. FortiOS Carrier can scan MMS messages sent using the MM1, MM3, MM4, and MM7 content interfaces. You can configure FortiOS Carrier to scan MMS messages for spam and viruses by configuring and adding MMS protection profiles and adding the MMS protection profiles to security policies. You can also use MMS protection profiles to apply content blocking, carrier end point filtering, MMS address translation, sending MMS notifications, DLP archiving of MMS messages, and logging of MMS message activity.

Figure 235: FortiOS Carrier MMS processing

FortiOS Carrier can send MMS messages to senders informing those senders that their devices are infected. FortiOS Carrier can also send MMS notifications to administrators to inform them of suspicious activity on their networks.

For message floods and duplicate messages, FortiOS Carrier does not send notifications to message senders but does send notifications to administrators and sends messages to sender handsets to complete MM1 and MM4 sessions.

Where MMS messaging uses the TCP/IP set of protocols, SMS text messaging uses the Signaling System Number 7 (SS7) set of protocols, which is not supported by FortiOS.

FortiOS Carrier and MMS content scanning

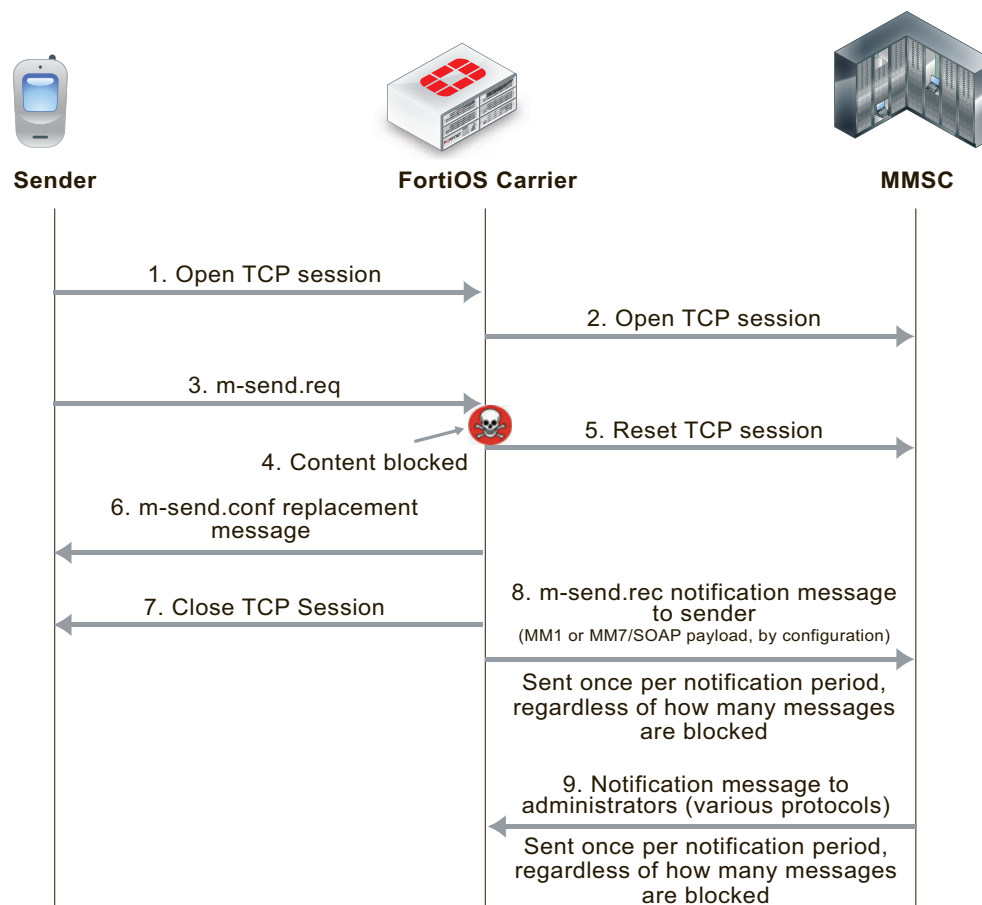
The following applies to MMS content scanning, including virus scanning, file filtering, content spam filtering, carrier end point filtering, and MMS content checksum filtering MMS.

MM1 Content Scanning

During MM1 content scanning a message is first transmitted from the sender, establishing a connection with the MMSC. FortiOS Carrier intercepts this connection and acts as the endpoint. FortiOS Carrier then establishes its own connection to the MMSC. Once connected, the client transmits its `m-send.req` HTTP post request to FortiOS Carrier which scans it according to the MMS protection profile settings. If the content is clean, the message is forwarded to the MMSC. The MMSC returns `m-send.conf` HTTP response through FortiOS Carrier to the sender.

If FortiOS Carrier blocks the message (for example because a virus was found (see [Figure 236](#)), FortiOS Carrier resets the connection to the MMSC and sends `m-send.conf` HTTP response back to the sender. The response message can be customized using replacement messages. Replacement messages are available for the different kinds of MMS scanning that the FortiOS Carrier unit can perform. FortiOS Carrier then terminates the connection. Sending back an `m-send.conf` message prevents the sender from trying to send the message again.

Figure 236: MM1 MMS scanning of message sent by sender (blocking m.send.req messages)



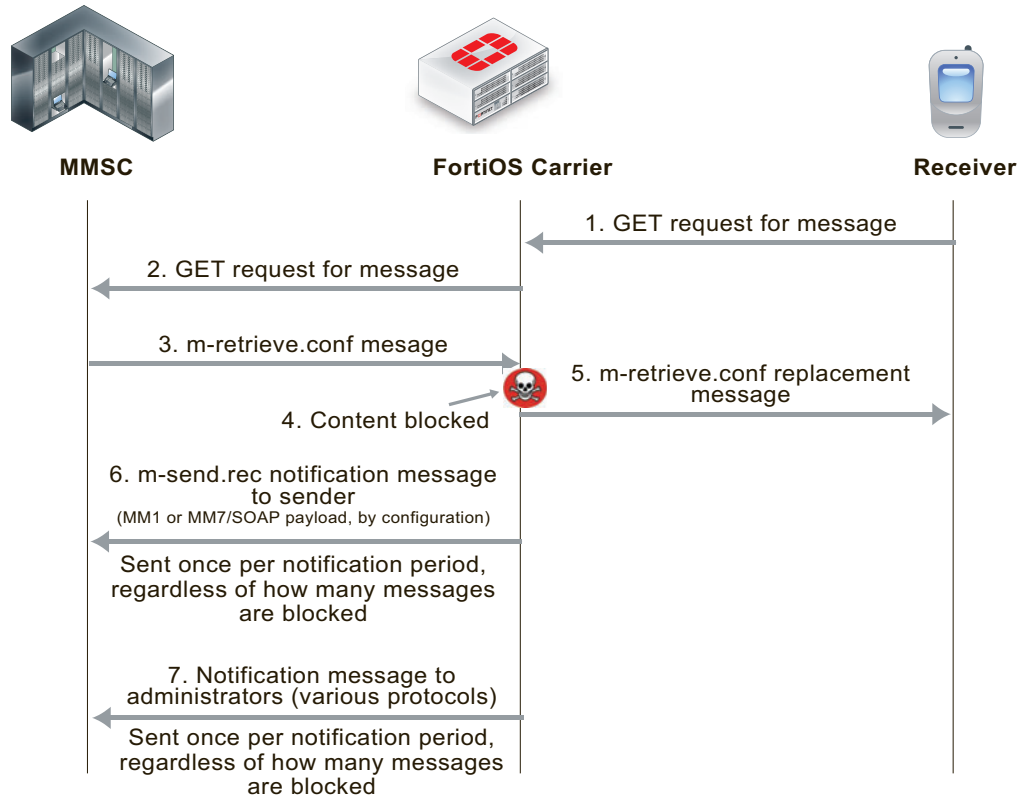
FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the sender to notify them of blocked messages.

Filtering message retrieval

Filtering message retrieval works in a similar way (see [Figure 237](#)). FortiOS Carrier intercepts the connection to the MMSC, and the `m-retrieve.conf` HTTP response from the MMSC is scanned according to the MMS content scanning settings. If the content is clean, the response is forwarded back to the client. If the content is blocked, FortiOS Carrier drops the connection to the MMSC. It then builds an `m-retrieve.conf` message from the associated replacement message and transmits this back to the client.

FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the receiver to notify them of blocked messages.

Figure 237: MM1 MMS scanning of messages received by receiver (blocking m.retrieve.conf messages)



Filtering MM3 and MM4 messages works in an similar way (see [Figure 238](#) and [Figure 239](#)). FortiOS Carrier intercepts connections to the MMSC, and scans messages as configured. When messages are blocked, FortiOS Carrier closes sessions as required, sends confirmation messages to the sender, notifies administrators, and notifies senders and receivers of messages.

Figure 238: MM3 MMS scanning of messages sent from a sender on the Internet to an MMSC

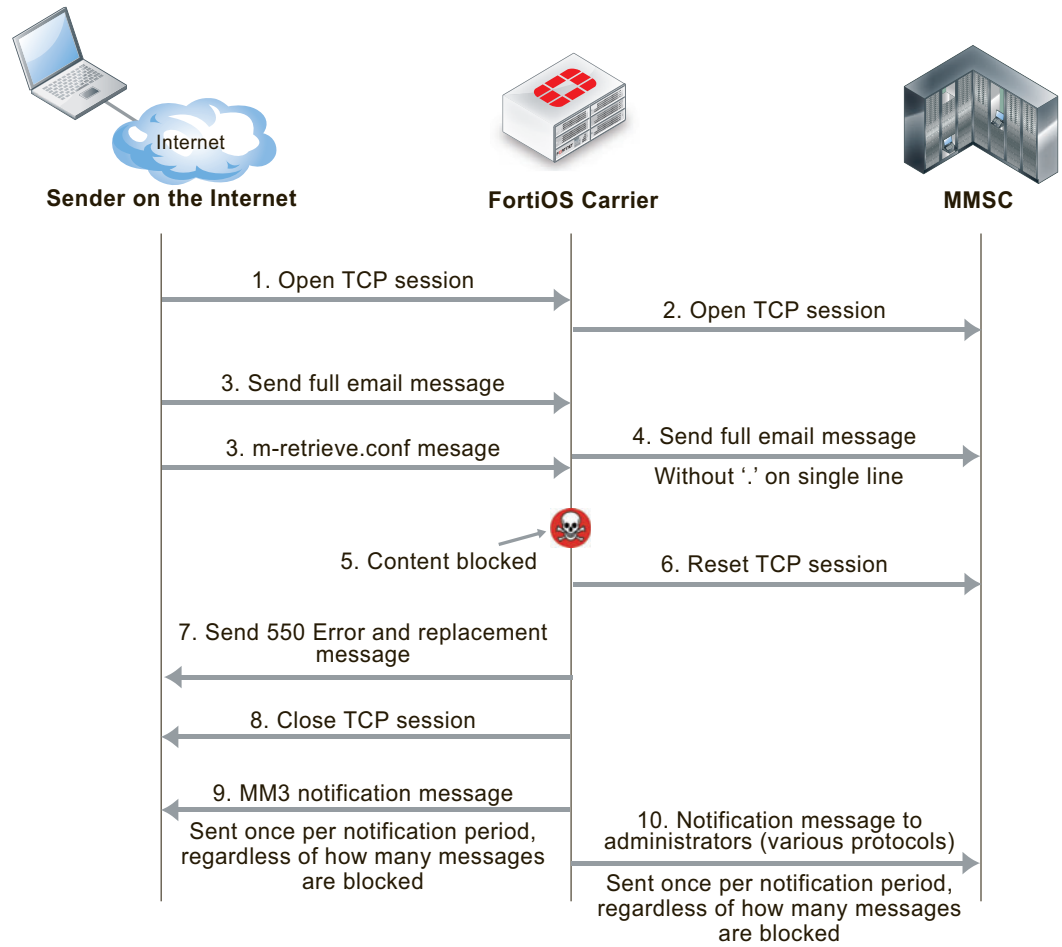


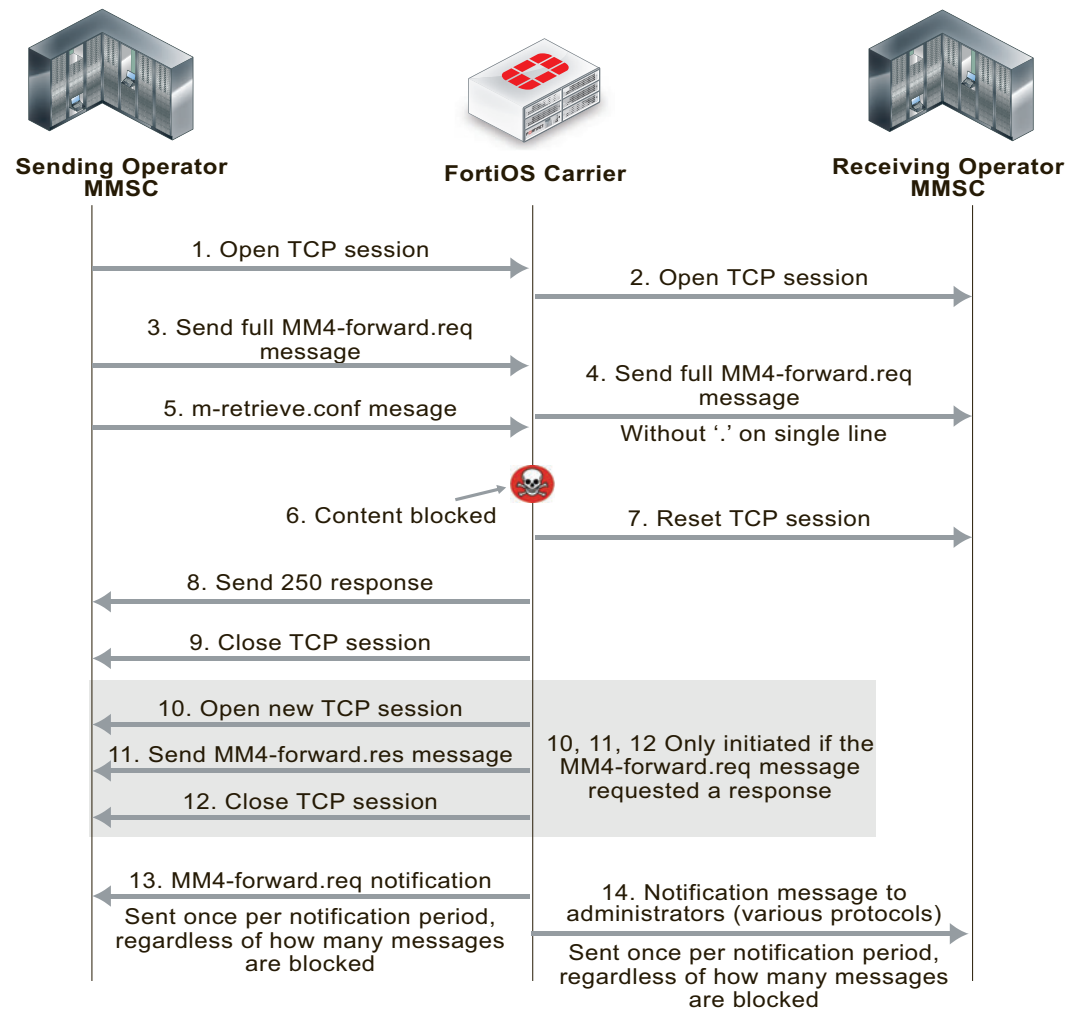
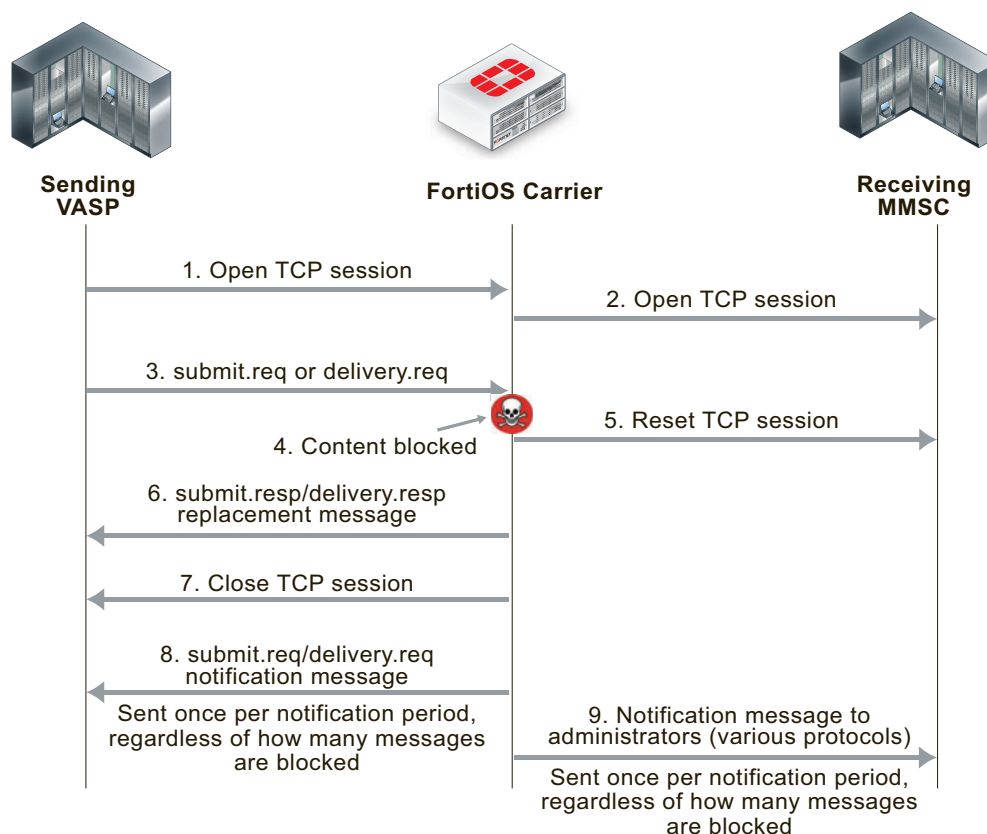
Figure 239: MM4 MMS scanning of messages sent between operator MMSCs

Figure 240: MM7 MMS scanning of messages sent between a VASP and an MMSC

FortiOS Carrier and MMS duplicate messages and message floods

FortiOS Carrier detects duplicate messages and message floods for the MM1 and MM4 interfaces. How FortiOS Carrier detects and responds to duplicate messages and message floods is different from how FortiOS Carrier detects and responds to viruses and other MMS scanning protection measures.

For message floods and duplicate messages, if the sender is an attacker they can gain useful information about message flood and duplicate message thresholds if they receive notifications about floods or duplicate messages. Plus, duplicate messages and message floods are usually a result of a large amount of messaging activity and filtering of these messages is designed to reduce the amount of unwanted messaging traffic. Adding to the traffic by sending notifications to senders and receivers could result in an increase in message traffic.

You can create up to three thresholds for detecting duplicate messages and message floods. For each threshold you can configure the FortiOS Carrier unit to respond by logging the activity, archiving or quarantining the messages, notifying administrators of the activity, and by blocking the messages. In many cases you may only want to configure blocking for higher activity thresholds, and to just monitor and send administrator notifications at lower activity thresholds.

When a block threshold is reached for MM1 messages, FortiOS Carrier sends `m-send.conf` or `m-retrieve.conf` messages to the originator of the activity. These messages are sent to end the MM1 sessions, otherwise the originator would continue to re-send the blocked message. When a block threshold is reached for MM4, FortiOS Carrier sends a `MM4-forward.res` message to close the MM4 session. An MM4 message is sent only if initiated by the originating `MM4-forward.req` message.

Figure 241: MM1 message flood and duplicate message blocking of sent messages

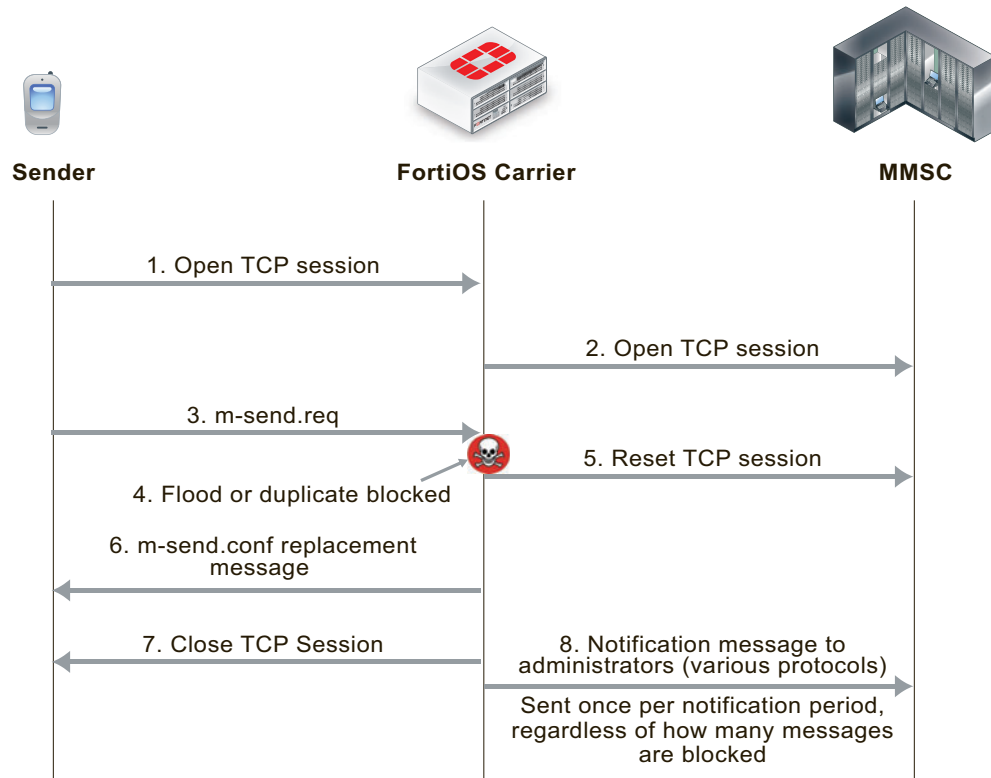


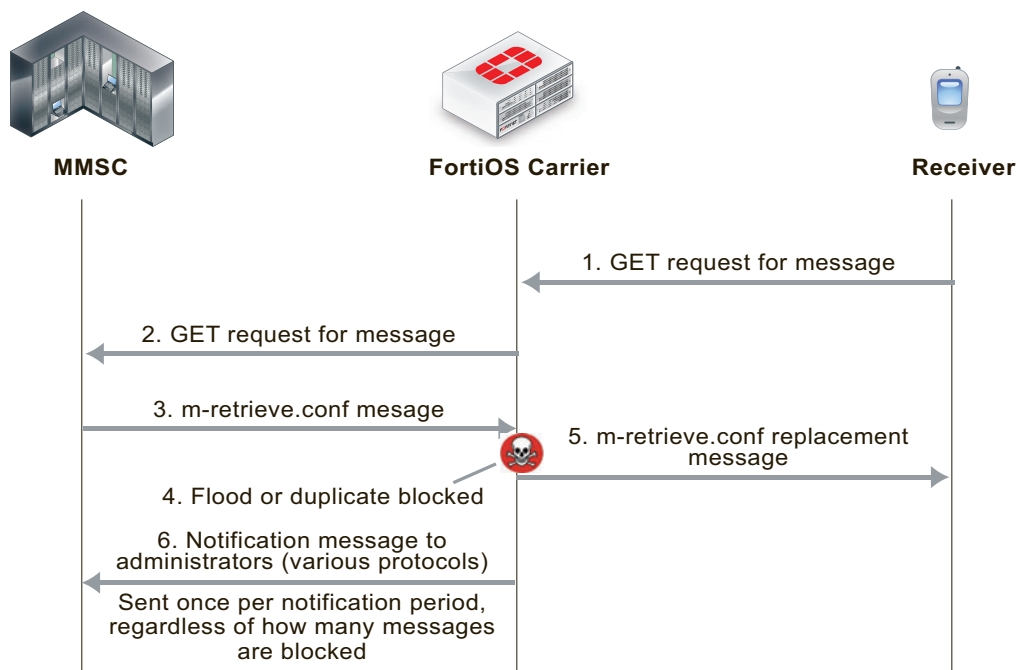
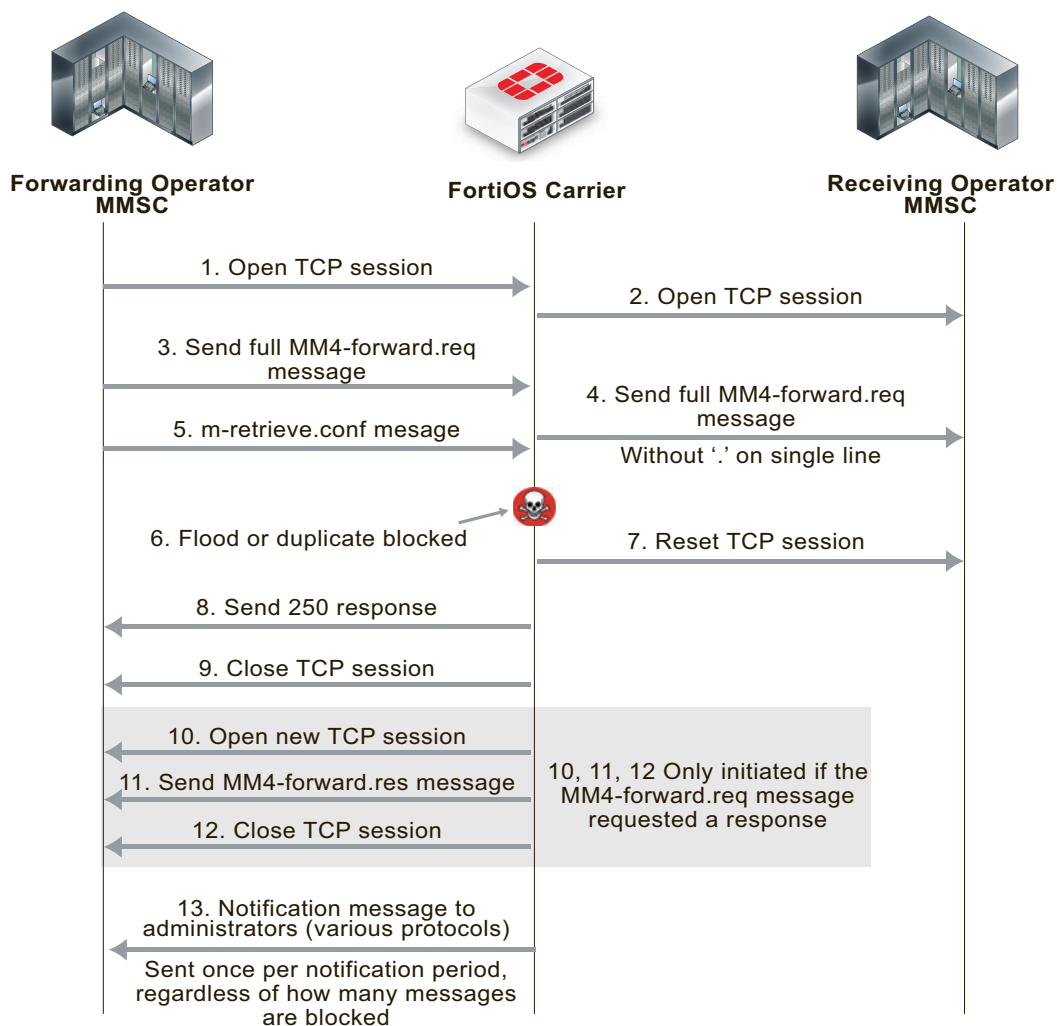
Figure 242: MM1 message flood and duplicate message blocking of received messages

Figure 243: MM4 message flood and duplicate message blocking

MMS protection profiles

An MMS protection profile is a group of settings that you can apply to an MMS session matched by a security policy.

MMS protection profiles are easy to configure and can be used by more than one security policy. You can configure a single MMS protection profile for the different traffic types handled by a set of security policies that require identical protection levels and types. This eliminates the need to repeatedly configure those same MMS protection profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need only moderate protection. You would configure two separate MMS protection profiles to provide the different levels of protection: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS Protection Profile, you need to add it to a security policy to apply the profile to MMS traffic.

See [“MMS UTM features” on page 2353](#).

Bypassing MMS protection profile filtering based on user's carrier end points

You can use carrier end point filtering to exempt MMS sessions from MMS protection profile filtering. Carrier end point filtering matches carrier end points in MMS sessions with carrier end point patterns. If you add a carrier end point pattern to a filter list and set the action to exempt from all scanning, all messages from matching carrier end points bypass MMS protection profile filtering. See [“Bypassing message flood protection based on user's carrier end points”](#) on page 2380.

Applying MMS protection profiles to MMS traffic

To apply an MMS protection profile you must first create the MMS protection profile and then add the MMS protection profile to a security policy by enabling the UTM option. The MMS protection profile then applies to the traffic accepted by that security policy.

MMS protection profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS protection profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS protection profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS protection profile.

To add an MMS protection profile to a security policy

- 1 Go to *UTM Profiles > Carrier > MMS Profile*.
- 2 Select *Create New* to add an MMS protection profile.
- 3 Configure and save the new MMS protection profile.
- 4 Go to *Policy*.
- 5 Select *Create New* to add a security policy, or select an existing policy and *Edit* to add the MMS profile.
- 6 Configure the security policy as required.
- 7 Enable *UTM*.
- 8 Select *Enable MMS Profile*, and select the MMS profile to add to the security policy.
- 9 Select *OK*.

GTP basic concepts

GPRS currently supports data rates from 9.6kbps to more than 100 kbps, and is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based), and the base station unit sends the message to the carrier network and eventually the Internet (wired carrier network).

The network system then either sends the message back to a base station and to the destination mobile unit, or forwards the message to the proper carrier's network where it gets routed to the mobile unit.

This sections includes:

- [PDP Context](#)
- [GPRS security](#)

PDP Context

The packet data protocol (PDP) context is a connection between a mobile station and the end address that goes through the SGSN and GGSN. It includes identifying information about the mobile customer used by each server or device to properly forward the call data to the next hop in the carrier network, typically using a GTP tunnel between the SGSN and GGSN.

When a mobile customer has an active voice or data connection open, both the SGSN and GGSN have the PDP context information for that customer and session.

When a mobile phone wants to communicate with an address on an external packet network, either an IP or X.25 address, the mobile station that phone is connected to opens a PDP context through the SGSN and GGSN to the end address. Before any traffic is sent, the PDP context must first be activated.

The information included in the PDP context includes the customer's IP address, the IMSI number of the mobile handset, and the tunnel endpoint ID for both the SGSN and GGSN. The ID is a unique number, much like a session ID on a TCP/IP firewall. All this information ensures a uniquely identifiable connection is made.

Since one mobile device may have multiple connections open at one time, such as data connections to different Internet services and voice connections to different locations, there may be more than one PDP context with the same IP address making the extra identifying information required.

The end point that the mobile phone is connecting to only knows about the GGSN — the rest of the GPRS connection is masked by the GGSN.

Along the PDP context path, communication is accomplished in using three different protocols.

- The connection between the Mobile Station and SGSN uses the SM protocol.
- Between SGSN and GGSN GTP is used.
- Between GGSN and the end point either IP or X.25 is used.

FortiOS Carrier is concerned with the SGSN to GGSN part of the PDP context — the part that uses GTP.

For more about PDP context, see [“Tunnel Management Messages” on page 2409](#).

Creating a PDP context

While FortiOS Carrier is concerned mostly with the SGSN to GGSN part of the PDP Context, knowing the steps involved in creating a PDP context helps understand the role each device, protocol, and message type plays.

Both mobile stations and GGSNs can create PDP contexts.

A Mobile Station creates a PDP context

- 1 The Mobile Station (MS) sends a `PDP activation request` message to the SGSN including the MS PDP address, and APN.
- 2 Optionally, security functions may be performed to authenticate the MS.
- 3 The SGSN determines the GGSN address by using the APN identifier.
- 4 The SGSN creates a downlink GTP tunnel to send IP packets between the GGSN and SGSN.
- 5 The GGSN creates an entry in its PDP context table to deliver IP packets between the SGSN and the external packet switching network.
- 6 The GGSN creates an uplink GTP tunnel to route IP-PDU from SGSN to GGSN.

- 7 The GGSN then sends back to the SGSN the result of the PDP context creation and if necessary the MS PDP address.
- 8 The SGSN sends an `Activate PDP context accept` message to the MS by returning negotiated the PDP context information and if necessary the MS PDP address.
- 9 Now traffic can pass from the MS to the external network end point.

A GGSN creates a PDP context

- 1 The network receives an IP packet from an external network.
- 2 The GGSN checks if the PDP Context has already been created.
- 3 If not, the GGSN sends a `PDU notification request` to the SGSN in order to initiate a PDP context activation.
- 4 The GGSN retrieves the IP address of the appropriate SGSN address by interrogating the HLR from the IMSI identifier of the MS.
- 5 The SGSN sends to the MS a request to activate the indicated PDP context.
- 6 The PDP context activation procedure follows the one initiated by the MS. See [“A Mobile Station creates a PDP context” on page 2306](#).
- 7 When the PDP context is activated, the IP packet can be sent from the GGSN to the MS.

Terminating a PDP context

A PDP context remains open until it is terminated. To terminate the PDP context an MS sends a `Deactivate PDP context` message to the SGSN, which then sends a `Delete PDP Context` message to the GGSN. When the SGSN receives a PDP context deletion acknowledgment from the GGSN, the SGSN confirms to the MS the PDP context deactivation. The PDP can be terminated by the SGSN or GGSN as well with a slight variation of the order of the messages passed.

When the PDP Context is terminated, the tunnel it was using is deleted as well. If this is not completed in a timely manner, it is possible for someone else to start using the tunnel before it is deleted. This hijacking will result in the original customer being overbilled for the extra usage. Anti-overbilling helps prevent this. See [“Configuring Anti-overbilling in FortiOS Carrier” on page 2405](#).

GPRS security

The GPRS network has some built-in security in the form of GPRS authentication. However this is minimal, and is not sufficient for carrier network security needs. A GTP firewall, such as FortiOS Carrier, is required to secure the Gi, Gn, and Gp interfaces.

GPRS authentication

GPRS authentication is handled by the SGSN to prevent unauthorized GPRS calls from reaching the GSM network beyond the SGSN (the base station system, and mobile station). Authentication is accomplished using some of the customer's information with a random number and uses two algorithms to create ciphers that then allow authentication for that customer.

User identity confidentiality ensures that customer information stays between the mobile station and the SGSN — no identifying information goes past the SGSN. Past that point other numbers are used to identify the customer and their connection on the network.

Periodically the SGSN may request identity information from the mobile station to compare to what is on record. This specifically looks at the IMEI number.

Call confidentiality is achieved through the use of a cipher, similar to the GPRS authentication described earlier. The cipher is applied between the mobile station and the SGSN. Essentially a cipher mask is XORd with each outgoing frame, and the receiving side XORs with its own cipher to result in the original frame and data.

Parts of a GTPv1 network

A sample GTP network consists of the end handset sender, the sender's mobile station, the carrier's network including the SGSN and GGSN, the receiver's mobile station, and the receiver handset.

When a handset moves from one mobile station and SGSN to another, the handset's connection to the Internet is preserved because the tunnel the handset has to the Internet using GTP tracks the user's location and information. For example, the handset could move from one cell to another, or between countries.

The parts of a GPRS network can be separated into the following groups according to the roles of the devices:

- Radio access to the GPRS network is accomplished by mobile phones and mobile stations (MS). See [“Radio access” on page 2309](#).
- Transport the GPRS packets across the GPRS network is accomplished by SGSNs and GGSNs, both local and remote, by delivering packets to the external services. See [“Transport” on page 2309](#).
- Billing and records are handled by CDF, CFR, HLR, and VLR devices. See [“Billing and records” on page 2310](#).

GPRS networks also rely on access points and PDP contexts as central parts of the communication structure. These are not actual devices, but they are still critical .

For more information on APN, see [“Access Point Number \(APN\)” on page 2418](#). For more information on PDP Context, see [“PDP Context” on page 2306](#).

These devices, their roles, neighboring devices, the interfaces and protocols they use are outlined in the following table. These devices and their connections can be viewed in the [“Packet flow through the GPRS network” on page 2312](#).

Figure 244: Carrier network showing the interfaces used (GTPv1)

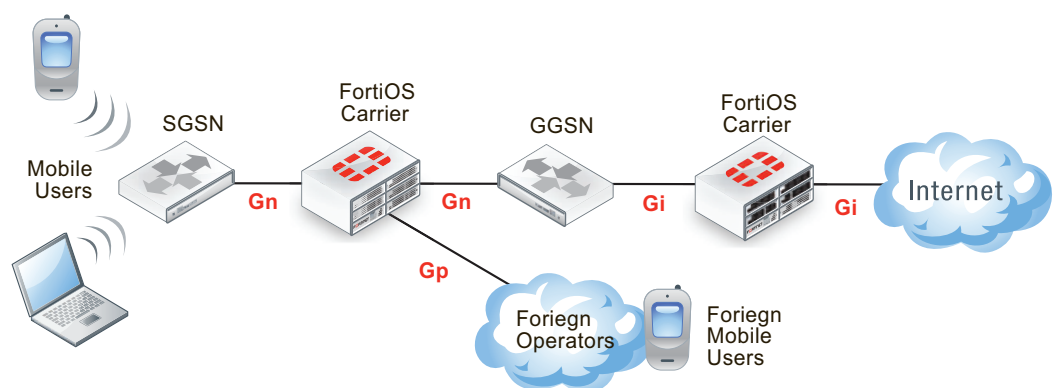


Table 123: Devices on a GTPv1 network

Device role	Neighboring Devices	Interfaces used	Protocols used
-------------	---------------------	-----------------	----------------

Table 123: Devices on a GTPv1 network

Mobile Users	Mobile Stations (MS)	Radio Access Technology (RAT)	
Mobile Stations (MS)	Mobile Users, SGSN	Gb	IP, Frame Relay
SGSN (local)	MS, SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR	Ga, Gb, Gn, Gp, Gz	IP, Frame Relay, GTP, GTP'
SGSN (remote)	SGSN (local)	Gn	GTP
GGSN (local)	SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR	Ga, Gi, Gn, Gp, Gz	IP, GTP, GTP'
GGSN (remote)	SGSN (local), WAP gateway, Internet, other external services	Gi, Gp	IP, GTPv1
CDR, CFR	SGSN (local), GGSN (local)	Ga, Gz	GTP'
HLR, VLR	SGSN (local), GGSN (local)	Ga, Gz	GTP'

Radio access

For a mobile phone to access the GPRS core network, it must first connect to a mobile station. This is a cellular tower that is connected to the carrier network.

How the mobile phone connects to the mobile station (MS) is determined by what Radio Access Technologies (RATs) are supported by the MS.

Transport

Transport protocols move data along the carrier network between radio access and the Internet or other carrier networks.

FortiOS Carrier should be present where information enters the Carrier network, to ensure the information entering is correct and not malicious. This means a FortiOS Carrier unit intercepts the data coming from the SGSN or foreign networks destined for the SSGN or GGSN onto the network, and after the GGSN as the data is leaving the network.

GTP

GPRS Tunnelling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over their network via tunneling. This tunneling allows users to move between SGSNs and still maintain connection to the Internet through the GGSN.

GTP has three versions version 0, 1, and 2. GTP1 and GTP2 are supported by FortiOS Carrier. The only GTP commands that are common to all forms of GTP are the echo request/response commands that allow GSNs to verify up to once every 60 seconds that neighboring GSNs are alive.

GGSN

The Gateway GPRS Support Node (GGSN) connects the GPRS network on one side via the SGSN to outside networks such as the Internet. These outside networks are called packet data networks (PDNs). The GGSN acts as an edge router between the two different networks — the GGSN forwards incoming packets from the external PDN to the addressed SGSN and the GGSN also forwards outgoing packets to the external PDN. The GGSN also converts the packets from the GPRS packets with SGSN to the external packets, such as IP or X.25.

SGSN

The Serving GPRS Support Node (SGSN) connects the GPRS network to GTPv1 compatible mobile stations, and mobile units (such as UTRAN and ETRAN) on one side and to the gateway node (GGSN), which leads to external networks, on the other side. Each SGSN has a geographical area, and mobile phones in that area connect to the GPRS network through this SGSN. The SGSN also maintains a location register that contains customer's location and user profiles until they connect through a different SGSN at which time the customer information is moved to the new SGSN. This information is used for packet routing and transfer, mobility management also known as location management, logical link management, and authentication and billing functions.

MME

MME essentially fills the role of the SGSN in a GTPv1 network — it is how the mobile stations gain access to the Carrier network. GTPv2 supports different mobile stations than GTPv1, so MME handles the GTPv2 MSes and SGSN handles the GTPv1 MSes

Billing and records

A major part of the GPRS network is devoted to billing. Customer billing requires enough information to identify the customer, and then billing specific information such as connection locations and times, as well as amount of data transferred. A modified form of GTP called GTP' is used for billing. The home location records and visitor location records store information about customers that is critical to billing.

GTP' (GTP prime)

GTP is used to handle tunnels of user traffic between SGSNs and GGSNs. However for billing purposes, other devices that are not supported by GTP are required. GTP' (GTP prime) is a modified form of GTP and is used to communicate with these devices such as the Charging Data Function (CDF) that communicates billing information to the Charging Gateway Function (CGF). In most cases, GTP' transports user records from many individual network elements, such as the GGSNs, to a centralised computer which then delivers the charging data more conveniently to the network operator's billing center, often through the CGF. The core network sends charging information to the CGF, typically including PDP context activation times and the quantity of data which the end user has transferred.

GTP' is used by the Ga and Gz interfaces to transfer billing information. GTP' uses registered UDP/TCP port 3386. GTP' defines a different header, additional messages, field values, as well as a synchronisation protocol to avoid losing or duplicating CDRs on CGF or SGSN/GGSN failure. Transferred CDRs are encoded in ASN.1.

HLR

The Home Location Register (HLR) is a central database that contains details of each mobile phone subscriber that is authorized to use the GSM core network. There can be several logical, and physical, HLRs per public land mobile network (PLMN), though one international mobile subscriber identity (IMSI)/MSISDN pair can be associated with only one logical HLR (which can span several physical nodes) at a time. The HLRs store details of every SIM card issued by the mobile phone operator. Each SIM has a unique identifier called an IMSI which is the primary key to each HLR record.

VLR

The Visitor Location Register (VLR) is a database which stores information about all the mobile devices that are currently under the jurisdiction of the Mobile Switching Center which it serves. Of all the information the VLR stores about each Mobile Station, the most important is the current Location Area Identity (LAI). This information is vital in the call setup process.

Whenever an MSC detects a new MS in its network, in addition to creating a new record in the VLR, it also updates the HLR of the mobile subscriber, informing it of the new location of that MS.

For more information on GTP', see [“GTP-U and Charging Management Messages” on page 2410](#).

GPRS network common interfaces

There are interfaces for each connection on the GPRS network. An interface is an established standard form of communication between two devices. Consider a TCP/IP network. In addition to the transport protocol (TCP) there are other protocols on that network that describe how devices can expect communications to be organized, just like GPRS interfaces.

Interfaces between devices on the network

There are a series of interfaces that define how different devices on the carrier network communicate with each other. These interfaces are called G_a to G_z, and each one defines how a specific pair of devices will communicate. For example G_b is the interface between the base station and the SGSN, and G_n is one possible interface between the SGSN and GGSN.

The SGSN and GGSN keep track of the CDR information and forward it to the Charging Data Function (CDF) using the G_r interface between the SGSN and home location register (HLR), G_s interface between the SGSN and MSC (VLR), G_x interface between the GGSN and the Charging Rules Function (CRF), G_y between the GGSN and online charging system (OCS), and finally G_z which is the off-line (CDR-based) charging interface between the GSN and the CG that uses GTP'.

Each of these interfaces on the GPRS network has a name in the format of G_x where _x is a letter of the alphabet that determines what part of the network the interface is used in. It is common for network diagrams of GPRS networks to include the interface name on connections between devices. See [“Packet flow through the GPRS network” on page 2312](#).



Tip: The FortiOS Carrier unit only provides protection on the G_n, G_p, and G_i interfaces.


Table 124: GPRS network interfaces, their roles, and billing

Name	Device connections that use this interface	Traffic Protocol used	Its role or how it affects billing
Ga	CDR and GSN (SGSNs and GGSNs)	GTP' - GTP modified to include CDR role	CDR have the accounting records, that are compiled in the GSN and then sent to the Charging Gateway (CG)
Gb	MS and SGSN	Frame Relay or IP	When an IP address moves to a new MS, the old MS may continue to use and bill that IP address.
Gi	GGSN and public data networks (PDNs)	IP based	This is the connection to the Internet. If the GTP tunnel is deleted without notifying the Gi interface, the connection may remain open incurring additional charges. FortiOS Carrier adds this interface to a firewall. See “Anti-overbilling with FortiOS Carrier” on page 2406 .
Gn	SGSN and external SGSNs and internal GGSNs	GTP	When the GTP tunnel is deleted, need to inform other interfaces immediately to prevent misuse of connections remaining open. FortiOS Carrier adds this interface to a firewall.
Gp	Internal SGSN and external GGSNs	GTP	
Gz	GSN (SGSN and GGSN) and the charging gateway (CG)	GTP'	Used for the offline charging interface. Ga is used for online charging.

Corporate customers may have a direct connection to the Gi interface for higher security. The Gi interface is normally an IP network, though a tunnelling protocol such as GRE or IPsec may be used instead.

Packet flow through the GPRS network

To better understand the GPRS network, we will follow the path data takes for a normal connection. For this example a call placed from a mobile phone involves accessing services on the Internet.

Figure 245: Sample GPRS network topology

- 1 A mobile phone places a call using a mobile station (MS). This connection between the mobile phone and the MS is a radio connection using one of the radio access technologies. See [“Radio Access Technology \(RAT\) type”](#) on page 2418.
- 2 The MS connects to a GPRS System Node (GSN) specifically a Sending GSN. This connection uses the Gb interface and typically uses IP address or Frame Relay.
- 3 The SGSN checks the mobile phone information located in the home location register (HLR) or visitor location register (VLR) to ensure there is subscriber information for that phone. If this mobile phone is from another network, it uses the VLR and updates its home carrier’s information with its current location and information. This connection involves the Ga or Gz interfaces, and uses the GTP’ protocol for communication.
- 4 The SGSN checks to make sure the phone did not transfer this connection from a different MS. If it did, the connection has already been established (along with the billing) and is handed off to this SGSN. If the call is being handed over from another SGSN, it will use the Gn interface between the two SGSNs.
- 5 The SGSN sends GTP messages to the local external Gateway GSN (GGSN) to create a GTP tunnel for this PDP context to access the Internet. It is possible that a remote GGSN has access to a service, such as a WAP gateway, that the local GGSN is missing. In this situation, the local SGSN uses the Gp interface to connect to the remote GGSN. Both the Gn and Gp interfaces use GTP.
- 6 The both the local and remote GGSNs connect to external services outside the GPRS network. These services can include a WAP gateway, a corporate IP network directly connected to the GPRS network, or the Internet. The connection from the GGSN to the external services uses the Gi interface.



Carrier web-based manager settings

The Carrier menu provides settings for configuring FortiOS Carrier features within the UTM menu. These features include MMS and GTP profiles.

In *UTM Profiles > Carrier*, you can configure profiles and settings for MMS and GTP. In the Carrier menu, you can configure an MMS profile and then apply it to a security policy. You can also configure GTP profiles and apply those to security policies as well.

This topic includes the following:

- [MMS profiles](#)
- [MMS Content Checksum](#)
- [Notification List](#)
- [Message Flood](#)
- [Duplicate Message](#)
- [Carrier Endpoint Filter Lists](#)
- [GTP Profile](#)

MMS profiles

Since MMS profiles can be used by more than one security policy, you can configure one profile for the traffic types handled by a set of security policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.



If the security policy requires authentication, do not select the MMS profile in the security policy. This type of profile is specific to the authenticating user group. For details on configuring the profile associated with the user group, see [User Groups in the User Authentication guide](#).

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS profile, you can then apply the profile to MMS traffic by applying it to a security policy.

MMS profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS profile.

The MMS Profile page contains options for each of the following:

- MMS scanning

- MMS Bulk Email Filtering Detection
- MMS Address Translation
- MMS Notifications
- DLP Archive
- Logging

MMS profile configuration settings

The following are MMS profile configuration settings in *UTM Profiles > Carrier > MMS Profile*.

MMS Profile page	
Lists each individual MMS profile that you created. On this page, you can edit, delete or create an MMS profile.	
Create New	Creates a new MMS profile. When you select <i>Create New</i> , you are automatically redirected to the New MMS Profile page.
Edit	Modifies settings within an MMS profile. When you select <i>Edit</i> , you are automatically redirected to the Edit MMS Profile.
Delete	Removes an MMS profile from the list on the MMS Profile page. To remove multiple MMS profiles from within the list, on the MMS Profile page, in each of the rows of the profiles you want removed, select the check box and then select <i>Delete</i> . To remove all MMS profiles from the list, on the MMS Profile page, select the check box in the check box column, and then select <i>Delete</i> .
Name	The name of the MMS profile.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.

New MMS Profile page	
Provides settings for configuring an MMS profile. This page also provides settings for configuring DLP archives and logging.	
Profile Name	Enter a name for the profile.
Comments	Enter a description about the profile. This is optional.
MMS Scanning	Configure MMS Scanning options. See “MMS scanning options” on page 2317 .
MMS Bulk Email Filtering Detection	Configure MMS Bulk Email options. See “MMS bulk email filtering options” on page 2319 .
MMS Address Translation	Configure MMS Address Translation options. See “MMS Address Translation options” on page 2322 .
MMS Notifications	Configure MMS Notification options. “MMS Notifications” on page 2324 .
DLP Archive	Configure DLP archive option. See “DLP Archive options” on page 2326 .
Logging	Configure logging options. See “Logging” on page 2327 .

MMS scanning options

You can configure MMS scanning protection profile options to apply virus scanning, file filtering, content filtering, carrier end point blocking, and other scanning to MMS messages transmitted using the MM1, MM3, MM4 and MM7 protocols.

The following are the MMS Scanning options that are available within an MMS profile. You can create an MMS profile in *UTM Profiles > Carrier > MMS Profile* or edit an existing one. You must expand MMS Scanning to access the following options.

MMS Scanning section of the New MMS Profile page	
Monitor Only	Select to cause the unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Select this option to be able to report on viruses and other problems in MMS traffic without affecting users. Tip: Select <i>Remove Blocked</i> if you want the unit to actually remove content intercepted by MMS scanning options.
Virus Scan	Select to scan attachments in MMS traffic for viruses. Since MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configuration also applies to MM1 and MM7 scanning. MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configuration also applies to MM3 and MM4 scanning.
Scan MM1 message retrieval	Select to scan message retrievals that use MM1. If you enable <i>Virus Scan</i> for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.

File Filter	Select to apply antivirus file filtering to MMS traffic. Select a file filter list to apply.
Quarantine	Select to quarantine the selected MMS traffic
Remove Blocked	<p>Select to remove blocked content from each protocol and replace it with the replacement message.</p> <p>Select <i>Constant</i> if the unit is to preserve the length of the message when removing blocked content, as may occur when billing is affected by the length of the message.</p> <p>Tip: If you only want to monitor blocked content, select <i>Monitor Only</i>.</p>
Content Filter	<p>Select to filter messages based on matching the content of the message with the words or patterns in the selected web content filter list.</p> <p>For information about adding a web content filter list, see the FortiGate CLI Reference.</p>
Carrier Endpoint Block	Select to add <i>Carrier End Point Filtering</i> in this MMS profile. Select the carrier end point filter list to apply it to the profile. For information about carrier end point filtering, see “Carrier Endpoint Filter Lists” on page 2334 .
MMS Content Checksum	Select to add MMS Content Checksum in this MMS profile. Select the MMS content checksum list to apply it to the profile.
Pass Fragmented Messages	Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 message are blocked.
Comfort Clients	<p>Select client comforting for MM1 and MM7 sessions.</p> <p>Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.</p>
Comfort Servers	<p>Select server comforting for each protocol.</p> <p>Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for the unit to buffer and scan large POST requests from slow clients.</p>
Interval (1-900 seconds)	Enter the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.
Amount (1-10240 bytes)	The number of bytes sent by client or server comforting at each interval.

Oversized MMS Message	<p>Select <i>Block</i> or <i>Pass</i> for files and email messages exceeding configured thresholds for each protocol.</p> <p>The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.</p>
Threshold (1KB - 800 MB)	<p>Enter the oversized file threshold and select KB or MB. If a file is larger than the threshold the file is passed or blocked depending on the <i>Oversized MMS Message</i> setting. The web-based manager displays the allowed threshold range. The threshold maximum is 10% of the unit's RAM.</p>

MMS bulk email filtering options

You can use the MMS bulk email filtering options to detect and filter MM1 and MM4 message floods and duplicate messages. You can configure three thresholds that define a flood of message activity and three thresholds that define excessive duplicate messages. The configuration of each threshold includes the response actions for the threshold.

The configurable thresholds for each of the flood and duplicate sensors and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

You can also add MSISDN to the bulk email filtering configuration and select a subset of the bulk email filtering options to applied to these individual MSISDNs.

You must first select MM1 and/or MM4 to detect excessive message duplicates. If excessive message duplicates are detected, the unit will perform the *Duplicate Message Action* for the specified duration.

You can configure three duplicate message thresholds and enable them with separate values and actions. They are labeled Duplicate Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Duplicate Threshold 1 and Duplicate Threshold 2, but you cannot disable Duplicate Threshold 1 and enable Duplicate Threshold 2.

When traffic accepted by a security policy that contains an MMS profile with duplicate message configured receives MM1 or MM4 duplicate messages that match a threshold configured in the MMS protection profile, the unit performs the duplicate message action configured for the matching threshold.

You can configure three message flood thresholds and enable them with separate values and actions. They are labeled Flood Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

When traffic accepted by a security policy that contains an MMS protection profile with message flooding configured experiences MM1 or MM4 message flooding that matches a threshold configured in the MMS profile, the unit performs the message flood action configured for the matching threshold.

MMS Bulk Email Filtering Detection section of the New MMS Profile page
This section of the New MMS Profile page contains numerous sections where you can configure specific settings for flood threshold, duplicate threshold and recipient MSISDNs.

Message Flood section of the new MMS Profile page

The message flood settings for each flood threshold. Expand each to configure settings for a threshold.

Flood Threshold 1		Expand to reveal the flood threshold settings for Flood Threshold 1. The settings for Flood Threshold 1 are the same for Flood Threshold 2 and 3.
	Enable	Select to apply Flood Threshold 1 to the MSISDN exception.
	Message Flood Window	Enter the period of time during which a message flood will be detected if the <i>Message Flood Limit</i> is exceeded. The message flood window can be 1 to 2880 minutes (48 hours).
	Message Flood Limit	Enter the number of messages which signifies a message flood if exceeded within the <i>Message Flood Window</i> .
	Message Flood Block Time	Enter the amount of time during which the unit performs the <i>Message Flood Action</i> after a message flood is detected.
	Message Flood Action	Select one or more actions that the unit is to perform when a message flood is detected.
Flood Threshold 2 Flood Threshold 3		Expand to configure settings for Flood Threshold 2 or 3 respectively.

Duplicate Message section of the new MMS Profile page

The duplicate message threshold settings. Expand each to configure settings for a threshold.

MM1 Retrieve Duplicate Enable		Select to scan MM1 <code>mm1-retr</code> messages for duplicates. By default, <code>mm1-retr</code> messages are not scanned for duplicates as they may often be the same without necessarily being bulk or spam.
	Enable	Select to enable the selected duplicate message threshold and to make the rest of the options available for configuration.
	Duplicate Message Window	Enter the period of time during which excessive message duplicates will be detected if the Duplicate message Limit it exceeded. The duplicate message window can be 1 to 2880 minutes (48 hours).
	Duplicate Message Limit	Enter the number of messages which signifies excessive message duplicates if exceeded within the Duplicate Message Window.

	Duplicate Message Block Time	Enter the amount of time during which the unit will perform the Duplicate Message Action after a message flood is detected.
	Duplicate Message Action	Select one or more actions that the unit is to perform when excessive message duplication is detected.
	Duplicate Threshold 2 Duplicate Threshold 3	Expand to configure settings for Duplicate Threshold 2 or 3 respectively.
Recipient MSISDN section of the New MMS Profile page		
The recipient Mobile Subscriber Integrated Services Digital Network Number (MSISDN) settings for each recipient MSISDN. When you select <i>Create New</i> , you are automatically redirected to the New MSISDN page.		
You need to save the profile before you can add MSISDNs.		
	Recipient MSISDN	The recipient MSISDN.
	Flood Threshold 1	Check to enable Flood Threshold 1 settings for this MSISDN.
	Flood Threshold 2	Check to enable Flood Threshold 2 settings for this MSISDN.
	Flood Threshold 3	Check to enable Flood Threshold 3 settings for this MSISDN.
	Duplicate Threshold 1	Check to enable Duplicate Threshold 1 settings for this MSISDN.
	Duplicate Threshold 2	Check to enable Duplicate Threshold 2 settings for this MSISDN.
	Duplicate Threshold 3	Check to enable Duplicate Threshold 3 settings for this MSISDN.
	Edit	Modifies the settings of a Recipient MSISDN in the Recipient MSISDN list. When you select <i>Edit</i> , you are automatically redirected to the New MSISDN page.
	Delete	Removes a Recipient MSISDN in the Recipient MSISDN list within the Recipient MSISDN section of the page.
New MSISDN page		
	Create New	Creates a new Recipient MSISDN. When you select <i>Create New</i> , you are automatically redirected to the New MSISDN page.
	Recipient MSISDN	Enter a name for the recipient MSISDN.
	Flood Threshold 1	Select to apply Flood Threshold 1 to the MSISDN exception.
	Flood Threshold 2	Select to apply Flood Threshold 2 to the MSISDN exception.
	Flood Threshold 3	Select to apply Flood Threshold 3 to the MSISDN exception.
	Duplicate Threshold 1	Select to apply Duplicate Threshold 1 to the MSISDN exception.

Duplicate Threshold 2	Select to apply Duplicate Threshold 2 to the MSISDN exception.
Duplicate Threshold 3	Select to apply Duplicate Threshold 3 to the MSISDN exception.

MMS Address Translation options

The sender's carrier end point is used to provide logging and reporting details to the mobile operator and to identify the sender of infected content.

When MMS messages are transmitted, the *From* field may or may not contain the sender's address. When the address is not included, the sender information will not be present in the logs and the unit will not be able to notify the user if the message is blocked unless the sender's address is made available elsewhere in the request.

The unit can extract the sender's address from an extended HTTP header field in the HTTP request. This field must be added to the HTTP request before it is received by the unit. If this field is present, it will be used instead of the sender's address in the MMS message for logging and notification. If this header field is present when a message is retrieved, it will be used instead of the *To* address in the message. If this header field is not present the content of the *To* header field is used instead.

Alternatively, the unit can extract the sender's address from a cookie.

You can configure MMS address translation to extract the sender's carrier end point so that it can be added to log and notification messages. You can configure MMS address translation settings to extract carrier end points from HTTP header fields or from cookies. You can also configure MMS address translation to add an end point prefix to the extracted carrier end points. For more information, see Dynamic Profiles and End Points in the [User Authentication guide](#).

MMS Address Translation section of the New MMS Profile page

Sender Address Source	Select to extract the sender's address from the <i>HTTP Header Field</i> or a <i>Cookie</i> . You must also specify the identifier that contains the carrier end point.
------------------------------	---

Sender Address Identifier	<p>Enter the sender address identifier that includes the carrier end point. The default identifier is <code>x-up-calling-line-id</code>.</p> <p>If the <i>Sender Address Source</i> is <i>HTTP Header Field</i>, the address and its identifier in the HTTP request header takes the format:</p> <pre><Sender Address Identifier>: <MSISDN_value></pre> <p>Where the <code><MSISDN_value></code> is the carrier end point. For example, the HTTP header might contain:</p> <pre>x-up-calling-line-id: 6044301297</pre> <p>where <code>x-up-calling-line-id</code> would be the Sender Address Identifier.</p> <p>If the <i>Sender Address Source</i> is <i>Cookie</i>, the address and its identifier in the HTTP request header's <i>Cookie</i> field takes the format of attribute-value pairs:</p> <pre>Cookie: id=<cookie-id>; <Sender Address Identifier>=<MSISDN Value></pre> <p>For example, the HTTP request headers might contain:</p> <pre>Cookie: id=0123jff!a;x-up-calling-line-id=6044301297</pre> <p>where <code>x-up-calling-line-id</code> would be the <i>Sender Address Identifier</i>.</p>
Convert Sender Address From / To HEX	Select to convert the sender address from ASCII to hexadecimal or from hexadecimal to ASCII. This is required by some applications.
Add Carrier Endpoint Prefix for Logging / Notification	Select the following to enable adding endpoint prefixes for logging and notification.
	<p>Enable</p> <p>Select to enable adding the country code to the extracted carrier end point, such as the MSISDN, for logging and notification purposes. You can limit the number length for the test numbers used for internal monitoring without a country code.</p>
	<p>Prefix</p> <p>Enter a carrier end point prefix to be added to all carrier end points. Use the prefix to add extra information to the carrier end point in the log entry.</p>
	<p>Minimum Length</p> <p>Enter the minimum length of the country code information being added. If this and Maximum Length are set to zero (0), length is not limited.</p>
	<p>Maximum Length</p> <p>Enter the maximum length of the country code information being added. If this and Minimum Length are set to zero (0), length is not limited.</p>

MMS Notifications

MMS notifications are messages that a unit sends when an MMS profile matches content in an MM1, MM3, MM4 or MM7 session. For example, the MMS profile detects a virus or uses content blocking to block a web page, text message or email. You can send notifications to the sender of the message using same protocol and the addressing headers in the original message. You can also configure MMS notifications to send notification messages to another destination (such as a system administrator) using the MM1, MM3, MM4 or MM7 protocol.

You need to enable one or more *Notification Types* or you can add an *Antivirus Notification List* to enable sending notifications,.

You can also use MMS notifications options to configure how often notifications are sent. The unit sends notification messages immediately for the first event, then at a configurable interval if events continue to occur. If the interval does not coincide with the window of time during which notices may be sent, the unit waits to send the notice in the next available window. Subsequent notices contain a count of the number of events that have occurred since the previous notification.

There are separate notifications for each notification type, including virus events. Virus event notifications include the virus name. Up to three viruses are tracked for each user at a time. If a fourth virus is found, one of the existing tracked viruses is removed from the list.

The notifications are MM1 `m-send-req` messages sent from the unit directly to the MMSC for delivery to the client. The host name of the MMSC, the URL to which `m-send-req` messages are sent, and the port must be specified.

MMS Notification section of the New MMS Profile page	
Antivirus Notification List	<p>Optionally select an antivirus notification list to select a list of virus names to send notifications for. The unit sends a notification message whenever a virus name or prefix in the antivirus notification list matches the name of a virus detected in a session scanned by the MMS protection profile. Select <i>Disabled</i> if you do not want to use a notification list. To create an antivirus notification list, see “Notification List” on page 2329.</p> <p>Instead of selecting a notification list you can configure the <i>Virus Scan Notification Type</i> to send notifications for all viruses.</p>
Message Protocol	<p>In each column, select the protocol used to send notification messages. You can use a different protocol to send the notification message than the protocol on which the violation was sent. The MMS Notifications options change depending on the message protocol that you select.</p> <p>If you select a different message protocol, you must also enter the User Domain. If selecting MM7 you must also enter the Message Type.</p>
Message Type	<p>Select the MM7 message type to use if sending notifications using MM7. Options include deliver.REQ and submit.REQ</p>

Detect Server Details		<p>Select to use the information in the headers of the original message to set the address of the notification message. If you do not select this option, you can enter the required addressing information manually.</p> <p>You cannot select <i>Detect Server Details</i> if you are sending notification messages using a different message protocol.</p> <p>If you select <i>Detect Server Details</i>, you cannot change the <i>Port</i> where the notification is being sent.</p>
Hostname		Enter the FQDN or the IP address of the server where the notifications will be sent.
URL		<p>Enter the URL of the server. For example if the notification is going to <code>www.example.com/home/alerts</code>, the URL is <code>/home/alerts</code>.</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm1</i> or <i>mm7</i>.</p>
Port		<p>Enter the port number of the server.</p> <p>You cannot change the <i>Port</i> if <i>Detect Server Details</i> is enabled.</p>
Username		<p>Enter the user name required for sending messages using this server (optional).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
Password		<p>Enter the password required for sending messages using this server (optional).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
VASP ID		<p>Enter the value-added-service-provider (VASP) ID to be used when sending a notification message. If a VAS is not offered by the mobile provider, it is offered by a third party or a VAS provider or content provider (CP).</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
VAS ID		<p>Enter the value-added-service (VAS) ID to be used when sending a notification message. A VAS is generally any service beyond voice calls and fax.</p> <p>This option is available only when <i>Message Protocol</i> is <i>mm7</i>.</p>
All Notification Types		<p>In each column, select notification for all MMS event types for that MMS protocol, then enter the amount of time and select the time unit for notice intervals.</p> <p>Alternatively, expand <i>All Notification Types</i>, and then select notification for individual MMS event types for each MMS protocol. Then enter the amount of time and select the time unit for notice intervals.</p> <p>Not all event types are available for all MMS protocols.</p>
	Content Filter	In each column, select to notify when messages are blocked by the content filter, then enter the amount of time and select the time unit for notice intervals.

	File Block	In each column, select to notify when messages are blocked by file block, then enter the amount of time and select the time unit for notice intervals.
	Carrier Endpoint Block	In each column, select to notify when messages are blocked, then enter the amount of time and select the time unit for notice intervals.
	Flood	In each column, select to notify when message flood events occur, then enter the amount of time and select the time unit for notice intervals.
	Duplicate	In each column, select to notify when duplicate message events occur, then enter the amount of time and select the time unit for notice intervals.
	MMS Content Checksum	In each column, select to notify when the content within an MMS message is scanned and banned because of the checksum value that was matched.
	Virus Scan	In each column, select to notify when the content within an MMS message is scanned for viruses.
	Notifications Per Second Limit	For each MMS protocol, enter the number of notifications to send per second. If you enter zero (0), the notification rate is not limited.
	Day of Week	For each MMS protocol, select the days of the week the unit is allowed to send notifications.
	Window Start Time	For each MMS protocol, select the time of day to begin the message alert window. By default, the message window starts at 00:00. You can change this if you want to start the message window later in the day. When configured, notification outside this window will not be sent.
	Window Duration	For each MMS protocol, select the time of day at which to end the message alert window. By default, the message window ends at 00:24. You can change this if you want to end the message window earlier in the day. When configured, notification outside this window will not be sent

DLP Archive options

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. In addition to the MMS profile's DLP archive options, you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select *DLP archiving* for carrier end point patterns in a *Carrier End Point List* and select the *Carrier End Point Block* option in the *MMS Scanning* section of an MMS Profile

The unit only allows one sixteenth of its memory for transferring content archive files. For example, for units with 128MB RAM, only 8MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

DLP Archive section of the New MMS Profile page	
Display DLP meta-information on the system dashboard	Select each required protocol to display the content archive summary in the Log and Archive Statistics dashboard widget on the System Dashboard.
Archive to FortiAnalyzer/FortiGuard	<p>Select the type of archiving that you want for the protocol (MM1, MM3, MM4, and MM7). You can choose from Full, Summary or None.</p> <p>None — Do not send content archives.</p> <p>Summary — Send content archive metadata only. Includes information such as date and time, source and destination, request and response size, and scan result.</p> <p>Full — Send content archive both metadata and copies of files or messages.</p> <p>In some cases, FortiOS Carrier may not archive content, or may make only a partial content archive, regardless of your selected option. This behavior varies by prerequisites for each protocol.</p> <p>This option is available only if a FortiAnalyzer unit or FortiGuard Analysis and Management Service is configured.</p>

Logging

You can enable logging in an MMS profile to write event log messages when the MMS profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS profile logging options to write an event log message every time a virus is detected.

You must first configure how the unit stores log messages so that you can then record these logs messages. For more information, see the [FortiOS Handbook Logging and Reporting chapter](#).

Logging section of the New MMS Profile page	
MMS-Antivirus	If antivirus settings are enabled for this MMS profile, select the following options to record <i>Antivirus Log</i> messages.
Viruses	Record a log message when this MMS profile detects a virus.
Blocked Files	Record a log message when antivirus file filtering enabled in this MMS profile blocks a file.
Intercepted Files	Record a log message when this MMS profile intercepts a file.
Oversized Files/Emails	Record a log message when this MMS profile encounters an oversized file or email message. Oversized files and email messages cannot be scanned for viruses.

MMS Scanning		If MMS scanning settings are enabled for this MMS profile, select the following options to record <i>Email Filter Log</i> messages.
	Notification Messages	Select to log the number of MMS notification messages sent.
	Bulk Messages	Select to log MMS Bulk AntiSpam events. You must also select which protocols to write log messages for in the MMS bulk email filtering part of the MMS profile. For more information, see “MMS bulk email filtering options” on page 2319 .
	Carrier Endpoint Filter Block	Select to log MMS carrier end point filter events, such as MSISDN filtering.
	MMS Content Checksum	Select to log MMS content checksum activity.
	Content Block	Select to log content blocking events.

MMS Content Checksum

The MMS Content Checksum menu allows you to configure content checksum lists. Configure MMS content checksum lists in *UTM Profiles > Carrier > MMS Content Checksum* using the following table.

MMS Content Checksum page	
Lists each individual content checksum list that you created. On this page, you can edit, delete or create a content checksum list.	
Create New	Creates a new MMS content checksum list. When you select <i>Create New</i> , you are automatically redirected to the New List. This page provides a name field and comment field. You must enter a name to go to MMS Content Checksum Settings page.
Edit	Modifies settings to a MMS content checksum. When you select <i>Edit</i> , you are automatically redirected to the MMS Content Checksum Settings page.
Delete	Removes an MMS content checksum from the page. To remove multiple content checksum lists from within the list, on the MMS Content Checksum page, in each of the rows of the content checksum lists you want removed, select the check box and then select <i>Delete</i> . To remove all content checksum lists from list, on the MMS Content Checksum page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the MMS content checksum list that you created.
# Entries	The number of checksums that are included in the content checksum list.

MMS Profiles	The MMS profile or profiles that have the MMS content checksum list applied. For example if two different MMS profiles use this content checksum list, they will both be listed here.
Comments	A description given to the MMS content checksum.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > AntiVirus > Profile</i>), 1 appears in <i>Ref.</i></p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.

Notification List

The Notification List menu allows you to configure a list of viruses. This virus list provides a list for scanning viruses in MMS messages. You can use one virus list in multiple MMS profiles, and configure multiple virus lists.

Notification list configuration settings

The following are notification list configuration settings in *UTM Profiles > Carrier > Notification List*.

Notification List page	
Lists all the notification lists that you created. On this page you can edit, delete or create a new notification list.	
Create New	Creates a new notification list. When you select <i>Create New</i> , you are automatically redirected to the New List page. You must enter a name to go to the Notification List Settings page.
Edit	Modifies settings within the notification list. When you select <i>Edit</i> , you are automatically redirected to the Notification List Settings page.

Delete	<p>Removes a notification list from the list on the Notification List page.</p> <p>To remove multiple notification lists from within the list, on the Notification List page, in each of the rows of the notification lists you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all notification lists from the list, on the Notification List page, select the check box in the check box column and then select <i>Delete</i>.</p>
Name	The name of the MMS content checksum list that you created.
# Entries	The number of checksums that are included in that content checksum list.
MMS Profiles	The MMS profile or profiles that are associated with
Comments	A description given to the MMS notification list.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i></p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
Notification List Settings page Provides settings for configuring a notification list, which is a list of viruses and is used for scanning viruses in MMS messages. This list is called the Antivirus Notification List in an MMS profile.	
Name	If editing the name of a notification list, enter the new name in this field. You must select <i>OK</i> to save the change.
Comments	If you want to enter a comment, enter the comment in the field. You must select <i>OK</i> to save the change.
Create New	Creates a notification entry in the list. When you select <i>Create New</i> , you are automatically redirected to the New Entry page.
Edit	Modifies settings within a notification list. When you select <i>Edit</i> , you are automatically redirected to the Edit Entry page.

Delete	Removes a notification entry from the list on the page. To remove multiple notification entries from within the list, on the Notification List Settings page, in each of the rows of the entries you want removed, select the check box and then select <i>Delete</i> . To remove all notification entries from the list, on the Notification List Settings page, select the check box in the check box column and then select <i>Delete</i> .
Enable	Enables a notification entry that is disabled.
Disable	Disables a notification entry so that it is not active and available for use, but it is not deleted.
Remove All Entries	Removes all notification entries that are listed on the Notification List Settings page.
Enable	Displays whether or not the checksum is enabled.
Virus Name/Profile	The name of the virus that was added to the list.
Entry Type	The type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
New Entry page	
Virus Name/Profile	Enter the virus name.
Entry Type	Select the type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
Enable	Select to enable the virus in the list.

Message Flood

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse. A message flood occurs when a single subscriber sends a volume of messages that exceed the flood threshold that you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected. For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all outgoing messages blocked for 30 minutes.

Action	Description
Log	Add a log entry indicating that a message flood has occurred. You must also enable logging for <i>MMS Scanning > Bulk Messages</i> in the <i>Logging</i> section of the MMS protection profile.
DLP Archive	Save the first message to exceed the flood threshold, or all the messages that exceed the flood threshold, in the DLP archive. DLP archiving flood messages may not always produce useful results. Since different messages can be causing the flood, reviewing the archived messages may not be a good indication of what is causing the problem since the messages could be completely random.

Action		Description
	All messages	All the messages that exceed the flood threshold will be saved in the DLP archive.
	First message only	Save only the first message to exceed the flood threshold in the DLP archive. Other messages in the flood are not saved. For message floods this may not produce much useful information since a legitimate message could trigger the flood threshold.
Intercept		Messages that exceed the flood threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message will also be quarantined for later examination. If the quarantine of intercepted messages is disabled, the <i>Intercept</i> action has no effect.
Block		Messages that exceed the flood threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each message will be quarantined for later examination.
Alert Notification		If the flood threshold is exceeded, the FortiOS Carrier unit will send an MMS flood notification message. In the web-based manager when <i>Alert Notification</i> is selected it displays the fields to configure the notification.

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many messages from the same subscriber to your MMSC.

Message flood configuration settings

The following are message flood configuration settings in *UTM Profiles > Carrier > Message Flood*.

Message Flood page Lists the large amount of messages that are being sent to you, from outside sources.	
Delete	Removes messages from the list. To remove multiple messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select <i>Delete</i> . To remove all messages from the list, on the Message Flood page, select the check box in the check box column and then select <i>Delete</i> .
Remove All Entries	Removes all messages from the list.
Protocol	The protocol used.
MMS Profile	The MMS profile that is used.
Sender	The sender's email address.
Level	The level of severity of the message.

Count	The count column can be up or down and these settings can be turned off by selecting beside the column's name.
Window Size (minutes)	The time in minutes.
Timer (minutes:seconds)	The time in seconds and in minutes. The timer column can be up or down and these settings turned off by selecting beside the column's name.
Page Controls	Use to navigate through the list.

Duplicate Message

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC.

The unit keeps track of the sent messages. If the same message appears more often than the threshold value that you have configured, action is taken. Possible actions are logging the duplicate messages, blocking or intercepting them, archiving, and sending an alert to inform an administrator that duplicate messages are occurring.

Duplicate message configuration settings

View duplicate messages in *UTM Profiles > Carrier > Duplicate Message*.

Duplicate Message page Lists duplicates of messages that were sent to you.	
Delete	Removes a message from the list. To remove multiple duplicate messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select <i>Delete</i> . To remove all duplicate messages from the list, on the Message Flood page, select the check box in the check box column and then select <i>Delete</i> .
Page Controls	Use to navigate through the list.
Remove All Entries	Removes all duplicate messages from the list.
Protocol	Either MM1 or MM4
Profile	The MMS profile that logs the detection.
Checksum	The checksum of the MMS message.
Status	Either flagged or blank. Flagged means that the actions defined in the MMS profile are taken. For more information, see “MMS bulk email filtering options” on page 2319 .
Count	Displays the number of messages in the last window of time.

Window Size (minutes)	The period of time during which a message flood will be detected if the Message Flood Limit is exceeded.
Timer (minutes:seconds)	Either the time left in the window if the message is unflagged, or the time until the message will be unflagged if it is already flagged.

Carrier Endpoint Filter Lists

A carrier end point filter list contains carrier end point patterns. A pattern can match one carrier end point or can use wildcards or regular expressions to match multiple carrier end points. For each pattern, you select the action that the unit takes on a message when the pattern matches a carrier end point in the message. Actions include blocking the message, exempting the message from MMS scanning, and exempting the message from all scanning. You can also configure the pattern to intercept the message and content archive the message to a FortiAnalyzer unit.

Carrier endpoint filter lists configuration settings

The following are Carrier endpoint filter list configuration settings in *UTM Profiles > Carrier > Carrier Endpoint Filter Lists*.

Carrier Endpoint Filter Lists page	
Lists all the endpoint filters that you created. On this page, you can edit, delete or create a new endpoint filter list.	
Create New	Creates a new endpoint filter list. When you select <i>Create New</i> , you are automatically redirected to the New List page. You must enter a name to go to the Carrier Endpoint Filter Lists Settings page.
Edit	Modifies settings within an endpoint filter list in the list.
Delete	Removes an endpoint filter in the list. To remove multiple endpoint filter lists from within the list, on the Carrier Endpoint Filter List page, in each of the rows of the endpoint filter lists you want removed, select the check box and then select <i>Delete</i> . To remove all endpoint filter lists from the list, on the Carrier Endpoint Filter List page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the endpoint filter.
# Entries	The number of carrier end point patterns in each carrier end point filter list.
MMS Profiles	The MMS profile that the carrier end point filter list is added to.
Comments	A description about the endpoint filter.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref.</i></p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
Carrier Endpoint Filter Lists Settings page Provides settings for configuring an endpoint filter.	
Name	The name you entered on the New List page, after selecting <i>Create New</i> on the Carrier Endpoint Filter page.
Comments	A description about the endpoint filter. You can add one here if you did not enter one on the New List page.
Create New	Creates a new endpoint filter list. When you select <i>Create New</i> , you are automatically redirected to the New Entry page.
Edit	Select to modify the settings of a pattern in the list.
Delete	Select to remove a pattern in the list.
Enable	Enables a disabled pattern in the list.
Disable	Disables a pattern in the list.
Remove All Entries	Removes all patterns in the list on the Carrier Endpoint Filter Lists Settings page.
Enable	Indicates whether or not the pattern is enabled.
Pattern	Enter or change the pattern that FortiOS Carrier uses to match with carrier end points. The pattern can be a single carrier end point or consist of wildcards or Perl regular expressions that will match more than one carrier end point. Set <i>Pattern Type</i> to correspond to the pattern that you want to use.

Action		Select the action taken by FortiOS Carrier for messages from a carrier end point that matches the carrier end point pattern:
Pattern Type		The type of pattern chosen.
New Entry page		
Pattern		Enter or change the pattern that FortiOS Carrier uses to match with carrier end points. The pattern can be a single carrier end point or consist of wildcards or Perl regular expressions that will match more than one carrier end point. Set <i>Pattern Type</i> to correspond to the pattern that you want to use.
Action(s)		Select the action taken by FortiOS Carrier for messages from a carrier end point that matches the carrier end point pattern:
	Content Archive	MMS messages from the carrier end point are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
	Intercept	MMS messages from the carrier end point are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.
Pattern Type		Select a pattern type as one of Single Carrier End Point, Wildcard or Regular Expression. Wildcard and Regular Expression will match multiple patterns where Single Carrier End Point matches only one.
Enable		Select to enable this carrier end point filter pattern.

GTP Profile

You can configure multiple GTP profiles within the GTP menu. GTP profiles concern GTP activity flowing through the unit. These GTP profiles are then applied to a security policy.

GTP profile configuration settings

The following are GTP profile configuration settings in *UTM Profiles > Carrier > GTP Profile*.

GTP Profile page		Lists each GTP profile that you have created. On this page, you can edit, delete or create a new GTP profile.
Create New		Creates a new GTP profile. When you select <i>Create New</i> , you are automatically redirected to the New page.
Edit		Modifies settings within a GTP profile in the list. When you select <i>Edit</i> , you are automatically redirected to Edit page.

Delete	<p>Removes a GTP profile from the list.</p> <p>To remove multiple GTP profiles from within the list, on the GTP Profile page, in each of the rows of the profiles you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all GTP profiles from within the list, on the GTP Profile page, select the check box in the check box column and then select <i>Delete</i>.</p>
Name	The name of the GTP profile.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <p>View the list page for these objects – automatically redirects you to the list page where the object is referenced at.</p> <p>Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</p> <p>View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.</p>
New GTP Profile page Provides settings for configuring a GTP profile.	
Name	Enter a name for the GTP profile.
General Settings	Configure general options for the GTP profile. See “General settings options” on page 2338 .
Message Type Filtering	Configure filtering for messages. See “Message type filtering options” on page 2340 .
APN Filtering	Configure filtering options for APN. See “APN filtering options” on page 2340 .
IMSI Filtering	Configure filtering options for IMSI. See “IMSI filtering options” on page 2341 .
Advanced Filtering	Configure advanced filtering options. See “Advanced filtering options” on page 2343 .
IE removal policy	Configure IE removal policy options. See “Information element removal policy options” on page 2346 .
Encapsulated IP Traffic Filtering	Configure filtering options for encapsulated IP traffic. See “Encapsulated IP traffic filtering options” on page 2346 .

Encapsulated Non-IP End User Address Filtering	Configure filtering options for encapsulated non-IP end user addresses. See “Encapsulated non-IP end user traffic filtering options” on page 2347 .
Protocol Anomaly	Configure protocol anomaly options. See “Protocol anomaly prevention options” on page 2348 .
Anti-Overbilling	Configure anti-overbilling options. See “Anti-overbilling options” on page 2349 .
Log	Configure log options. See “Log options” on page 2349 .

General settings options

The following are mostly house keeping options that appear in the General Settings area of the GTP configuration page.

General Settings section of the New GTP Profile page	
Sequence Number Validation	<p>Enable to check that packets are not duplicated or out of order. GTP packets contain a Sequence Number field.</p> <p>This number tells the receiving GGSN the order of the packets it is receiving. Normally the GGSN compares this sequence number in the packets with its own sequence counter — if the two do not match, the packet is dropped. This sequence number validation can be off-loaded to the FortiOS Carrier freeing up resources on the GGSN.</p>
GTP-in-GTP	<p>Select <i>Allow</i> to enable GTP packets to be allowed to contain GTP packets, or a GTP tunnel inside another GTP tunnel.</p> <p>To block all GTP-in-GTP packets, select <i>Deny</i>.</p>
Minimum Message Length	<p>Enter the shortest possible message length in bytes. Normally this is controlled by the protocol, and will vary for different message types. If a packet is smaller than this limit, it is discarded as it is likely malformed and a potential security risk.</p> <p>The default minimum message length is 0 bytes.</p>
Maximum Message Length	<p>Enter the maximum allowed length of a GTP packet in bytes.</p> <p>A GTP packet contains three headers and corresponding parts GTP, UDP, and IP. If a packet is larger than the maximum transmission unit (MTU) size, it is fragmented to be delivered in multiple packets. This is inefficient, resource intensive, and may cause problems with some applications.</p> <p>By default the maximum message length is 1452 bytes.</p>
Tunnel Limit	<p>Enter the maximum number of tunnels allowed open at one time. For additional GTP tunnels to be opened, existing tunnels must first be closed.</p> <p>This feature can help prevent a form of denial of service attack on your network. This attack involves opening more tunnels than the network can handle and consuming all the network resources doing so. By limiting the number of tunnels at any one time, this form of attack will be avoided.</p> <p>The tunnel limiting applies to the Handover Group, and Authorized SGSNs and GGSNs.</p>

Tunnel Timeout	<p>Enter the maximum number of seconds that a GTP tunnel is allowed to remain active. After the timeout the unit deletes GTP tunnels that have stopped processing data. A GTP tunnel may hang for various reasons. For example, during the GTP tunnel tear-down stage, the "delete pdap context response" message may get lost. By setting a timeout value, you can configure the FortiOS Carrier firewall to remove the hanging tunnels.</p> <p>The default is 86400 seconds, or 24 hours.</p>
Control plane message rate limit	<p>Enter the number of packets per second to limit the traffic rate to protect the GSNs from possible Denial of Service (DoS) attacks. The default limit of 0 does not limit the message rate.</p> <p>GTP DoS attacks can include:</p> <ul style="list-style-type: none"> • Border gateway bandwidth saturation: A malicious operator can connect to your GRX and generate high traffic towards your Border Gateway to consume all the bandwidth. • GTP flood: A GSN can be flooded by illegitimate traffic
Handover Group	<p>Select the allowed list of IP addresses allowed to take over a GTP session when the mobile device moves locations.</p> <p>Handover is a fundamental feature of GPRS/UMTS, which enables subscribers to seamlessly move from one area of coverage to another with no interruption of active sessions. Session hijacking can come from the SGSN or the GGSN, where a fraudulent GSN can intercept another GSN and redirect traffic to it. This can be exploited to hijack GTP tunnels or cause a denial of service.</p> <p>When the handover group is defined it acts like a whitelist with an implicit default deny at the end — the GTP address must be in the group or the GTP message will be blocked. This stops handover requests from untrusted GSNs.</p>
Authorized SGSNs	<p>Use <i>Authorized SGSNs</i> to only allow authorized SGSNs to send packets through the unit and to block unauthorized SGSNs. Go to <i>Firewall Objects > Address > Address</i> and add the IP addresses of the authorized SGSNs to a firewall address or address group. Then set <i>Authorized SGSNs</i> to this firewall address or address group.</p> <p>You can use <i>Authorized SGSNs</i> to allow packets from SGSNs that have a roaming agreement with your organization.</p>
Authorized GGSNs	<p>Use <i>Authorized GGSNs</i> to only allow authorized GGSNs to send packets through the unit and to block unauthorized GGSNs. Go to <i>Firewall Objects > Address > Address</i> and add the IP addresses of the authorized GGSNs to a firewall address or address group. Then set <i>Authorized GGSNs</i> to this firewall address or address group.</p> <p>You can use <i>Authorized GGSNs</i> to allow packets from SGSNs that have a roaming agreement with your organization.</p>

Message type filtering options

On the *New GTP Profile* page, you can select to allow or deny the different types of GTP messages, which is referred to as message type filtering. You must expand the Message Type Filtering section to access the settings.

The messages types include Path Management, Tunnel Management, Location Management, Mobility Management, MBMS, and GTP-U and Charging Management messages.



For enhanced security, Fortinet best practices dictate that you set Unknown Message Action to deny. This will block all unknown GTP message types, some of which may be malicious.

To configure message type filter options, expand *Message Type Filtering* in the GTP profile.

APN filtering options

An Access Point Name (APN) is an Information Element (IE) included in the header of a GTP packet. It provides information on how to reach a network.

An APN has the following format:

```
<network_id>[.mnc<mnc_int>.mcc<mcc_int>.gprs]
```

Where:

- `<network_id>` is a network identifier or name that identifies the name of a network, for example, `example.com` or `internet`.
- `[.mnc<mnc_int>.mcc<mcc_int>.gprs]` is the optional operator identifier that uniquely identifies the operator's PLMN, for example `mnc123.mcc456.gprs`.

Combining these two examples results in a complete APN of `internet.mnc123.mcc456.gprs`.

By default, the unit permits all APNs. However, you can configure APN filtering to restrict roaming subscribers' access to external networks.

APN filtering applies only to the GTP *create pdp request* messages. The unit inspects GTP packets for both APN and selected modes. If both parameters match and APN filter entry, the unit applies the filter to the traffic.

Additionally, the unit can filter GTP packets based on the combination of an IMSI prefix and an APN. For more information, see ["IMSI filtering options" on page 2341](#).



You cannot add an APN when creating a new profile.

APN Filtering section on the New GTP Profile page	
Enable APN Filter	Select to enable APN filtering.
Default APN Action	Select the default action for APN filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual APN filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual APN filters.
Value	The APN to be filtered.

Mode	The type of mode chosen that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription:
Action	The type of action that will be taken.
Edit	Modifies the settings within the filter. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the settings of the APN.
Delete	Removes the APN from the list within the table, in the APN Filtering section.
Add APN	Adds a new APN filter to the list. When you select <i>Add APN</i> , the New window appears, which allows you to configure the APN settings.
New APN page	
Value	Enter an APN to be filtered. You can include wild cards to match multiple APNs. For example, the value <i>internet*</i> would match all APNs that begin with <i>internet</i> .
Mode	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Mobile Station provided	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
Subscription Verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
Action	Select <i>Allow</i> or <i>Deny</i> .

IMSI filtering options

The International Mobile Station Identity (IMSI) is used by a GPRS Support Node (GSN) to identify a mobile station. Three elements make up every IMSI:

- the mobile country code (MCC)
- the mobile network code (MNC)
- the mobile subscriber identification number (MSIN).

The subscriber's home network—the public land mobile network (PLMN)—is identified by the IMSI prefix, formed by combining the MCC and MNC.

By default, the unit allows all IMSIs. You can add IMSI prefixes to deny GTP traffic coming from non-roaming partners. Any GTP packets with IMSI prefixes not matching the prefixes you set will be dropped. GTP *Create pdp* request messages are filtered and only IMSI prefixes matching the ones you set are permitted. Each GTP profile can have up to 1000 IMSI prefixes set.

An IMSI prefix and an APN can be used together to filter GTP packets if you set an IMSI filter entry with a non-empty APN.



You cannot add an IMSI when creating a new profile. You must add it after the profile has been created and you are editing the profile.

IMSI Filtering section of the New GTP Profile page

Enable IMSI Filter	Select to enable IMSI filtering.
Default IMSI Action	Select the default action for IMSI filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual IMSI filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual IMSI filters.
APN	The APN that is part of the IMSI that will be filtered.
MCC-MNC	The MCC-MNC part of the IMSI that will be filtered.
Mode	The type of mode that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Action	The type of action that will be taken.
Edit	Modifies settings to an IMSI filter. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the IMSI filter's settings.
Delete	Removes an IMSI filter from within the table, in the IMSI Filtering section.
Add IMSI	Adds a new IMSI filter to the list. When you select <i>Add IMSI</i> , the New window appears, which allows you to configure IMSI filter settings.

New IMSI page

APN	Enter the APN part of the IMSI to be filtered.
MCC-MNC	Enter the MCC-MCC part of the IMSI to be filtered.
Mode	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Mobile Station provided	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
Subscription Verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
Action	Select <i>Allow</i> or <i>Deny</i> .

Advanced filtering options

The FortiOS Carrier firewall supports advanced filtering against the attributes RAT, RAI, ULI, APN restriction, and IMEI-SV in GTP to block specific harmful GPRS traffic and GPRS roaming traffic. The following table shows some of the GTP context requests and responses that the firewall supports.

Table 125: Attributes supported by FortiCarrier firewalls

	GTP Create PDP Context Request	GTP Create PDP Context Response	GTP Update PDP Context Request	GTP Update PDP Context Response
APN	yes	yes	-	
APN Restriction	yes	-	-	yes
IMEI-SV	yes	-	-	-
IMSI	yes	-	yes	-
RAI	yes	-	yes	-
RAT	yes	-	yes	-
ULI	yes	-	yes	-

When editing a GTP profile, select *Advanced Filtering > Add* to create and add a rule. When the rule matches traffic it will either allow or deny that traffic as selected in the rule.

Advanced Filtering section on New GTP Profile page	
Enable	Select to enable advanced filtering.
Default Action	Select the default action for advanced filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual advanced filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual advanced filters.
Messages	The messages, for example, Create PDP Context Request.
APN Restriction	The APN restriction.
RAT Type	The RAT types associated with that filter.
ULI	The ULI pattern.
RAI	The RAI pattern.
IMEI	The IMEI pattern.
Action	The action that will be taken.
Edit	Modifies the filter's settings. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the filter's settings.
Delete	Removes a filter from the list.
Add	Adds a filter to the list. When you select <i>Add</i> , the New window appears, which allows you to configure settings for messages, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI patterns as well as the type of action.
New Filtering page	

Messages		The PDP content messages this profile will match.
	Create PDP Context Request	Select to allow create PDP context requests.
	Create PDP Context Response	Select to allow create PDP context responses.
	Update PDP Context Request	Select to allow update PDP context requests.
	Update PDP Context Response	Select to allow update PDP context responses.
APN		Enter the APN.
	APN Mode	<p>Select an APN mode as one or more of</p> <ul style="list-style-type: none"> • Mobile Station provided • Network provided • Subscription provided <p>This field is only available when an APN has been entered.</p>
	Mobile Station provided	MS-provided PAN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
	Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
	Subscription verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
	APN Restriction	<p>Select the type of restriction that you want. You can choose all of the types, or one of the types. You cannot choose multiple types. Types include:</p> <ul style="list-style-type: none"> • all • Public-1 • Public-2 • Private-1 • Private-2
IMSI		Enter the IMSI.
MSISDN		Enter the MSISDN.

RAT Type	<p>Optionally select the RAT type as any combination of the following:</p> <ul style="list-style-type: none"> • Any • UTRAN • GERAN • Wifi • GAN • HSPA <p>Some RAT types are GTPv1 specific.</p>
ULI pattern	Enter the ULI pattern.
RAI pattern	Enter the RAI pattern.
IMEI pattern	Enter the IMEI pattern.
Action	Select either <i>Allow</i> or <i>Deny</i> .

Adding an advanced filtering rule

When adding a rule, use the following formats:

- Prefix, for example, range 31* for MCC matches MCC from 310 to 319.
- Range, for example, range 310-319 for MCC matches MCC from 310 to 319.
- Mobile Country Code (MCC) consists of three digits. The MCC identifies the country of domicile of the mobile subscriber.
- Mobile Network Code (MNC) consists of two or three digits for GSM/UMTS applications. The MNC identifies the home PLMN of the mobile subscriber. The length of the MNC (two or three digits) depends on the value of the MCC. Best practices dictate not to mix two and three digit MNC codes within a single MCC area.
- Location Area Code (LAC) is a fixed length code (of 2 octets) identifying a location area within a PLMN. This part of the location area identification can be coded using a full hexadecimal representation except for the following reserved hexadecimal values: 0000 and FFFE. These reserved values are used in some special cases when no valid LAI exists in the MS (see 3GPP TS 24.008, 3GPP TS 31.102 and 3GPP TS 51.011).
- Routing Area Code (RAC) of a fixed length code (of 1 octet) identifies a routing area within a location.
- CI or SAC of a fixed length of 2 octets can be coded using a full hexadecimal expression.
- Type Allocation Code (TAC) has a length of 8 digits.
- Serial Number (SNR) is an individual serial number identifying each equipment within each TAC. SNR has a length of 6 digits.
- Software Version Number (SVN) identifies the software version number of the mobile equipment. SVN has a length of 2 digits.



You cannot add an advanced filtering rule when creating a new profile. You must add it after the profile has been created and you are editing the profile.

Information element removal policy options

In some roaming scenarios, the unit is installed on the border of the PLMN and the GRX. In this configuration, the unit supports information element (IE) removal policies to remove any combination of R6 IEs (RAT, RAI, ULI, IMEI-SV and APN restrictions) from the types of messages described in [“Advanced filtering options” on page 2343](#), prior to forwarding the messages to the HGSN (proxy mode).

IE removal policy section of the New GTP Profile page	
Enable	Select to enable this option.
SGSN address of message IE	The firewall address or address group that contains the SGSN addresses.
IEs to be removed	The IE types that will be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.
Add	Adds an IE removal policy. When you select <i>Add</i> , the New window appears, which allows you to configure the IE policy.
Edit	Modifies settings from within the IE removal policy. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the settings within the policy.
Delete	Removes the IE removal policy from the list.
New IE policy page	
SGSN address	Select a firewall address or address group that contains SGSN addresses.
IEs to be removed	Select one or more IE types to be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.

Encapsulated IP traffic filtering options

You can use encapsulated IP traffic filtering to filter GTP sessions based on information contained in the data stream. to control data flows within your infrastructure. You can configure IP filtering rules to filter encapsulated IP traffic from mobile stations by identifying the source and destination policies. For more information, see [“When to use encapsulated IP traffic filtering \(best practices\)” on page 2403](#).

Expand *Encapsulated IP Traffic Filtering* in the GTP profile to reveal the options.

Encapsulated IP Traffic Filtering section of the New GTP Profile page	
Enable IP Filter	Select to enable encapsulated IP traffic filtering options.
Default IP Action	Select the default action for encapsulated IP traffic filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual encapsulated IP traffic filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual encapsulated IP traffic filters.
Source	Select a source IP address from the configured firewall IP address or address group lists. Any encapsulated traffic originating from this IP address will be a match if the destination also matches.

Destination	Select a destination IP address from the configured firewall IP address or address group lists. Any encapsulated traffic being sent to this IP address will be a match if the destination also matches.
Action	The type of action that will be taken. Select to Allow or Deny encapsulated traffic between this source and Destination.
Edit	Modifies the source, destination or action settings.
Add IP Policy	Adds a new encapsulated IP traffic filter. When you select <i>Add IP Policy</i> , the New window appears which allows you to configure IP policy settings.
New (window)	
Source	Select the source firewall address or address group.
Destination	Select the destination firewall address or address group.
Action	Select <i>Allow</i> or <i>Deny</i> .

Encapsulated non-IP end user traffic filtering options

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications list only PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC1700. The PDP types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

Encapsulated Non-IP End User Address Filtering section of the New GTP Profile page	
Enable Non-IP Filter	Select to enable encapsulated non-IP traffic filtering.
Default Non-IP Action	Select the default action for encapsulated non-IP traffic filtering. If you select <i>Allow</i> , all sessions are allowed except those blocked by individual encapsulated non-IP traffic filters. If you select <i>Deny</i> , all sessions are blocked except those allowed by individual encapsulated non-IP traffic filters.
Type	The type chosen, <i>AESTI</i> or <i>IETF</i> .
Start Protocol	The beginning protocol port number range.
End Protocol	The end of the protocol port number range.
Action	The type of action that will be taken.
Edit	Modify a non-IP filter's settings in the list. When you select <i>Edit</i> , the Edit window appears, which allows you to modify the Non-IP policy settings.
Delete	Remove a non-IP policy from the list.

Add Non-IP Policy	Add a new encapsulated non-IP traffic filter. When you select <i>Add Non-IP Policy</i> , you are automatically redirected to the New page.
New (window)	
Type	Select <i>AESTI</i> or <i>IETF</i> .
Start Protocol End Protocol	Select a start and end protocol from the list of protocols in RFC 1700. Allowed range includes 0 to 255 (0x00 to 0xff). Some common protocols include: <ul style="list-style-type: none"> • 33 (0x0021) Internet Protocol • 35 (0x0023) OSI Network Layer • 63 (0x003f) NETBIOS Framing • 65 (0x0041) Cisco Systems • 79 (0x004f) IP6 Header Compression • 83 (0x0053) Encryption
Action	Select <i>Allow</i> or <i>Deny</i> .

Protocol anomaly prevention options

Use protocol anomaly detection options to detect or deny protocol anomalies according to GTP standards and tunnel state. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of the protocol specifications. Packets cannot pass through if they fail the sanity check.

Protocol Anomaly section of the New GTP Profile page	
Invalid Reserved Field	GTP version 0 (GSM 09.60) headers specify a number of fields that are marked as "Spare" and contain all ones (1). GTP packets that have different values in these fields are flagged as anomalies. GTP version 1 (GSM 29.060) makes better use of the header space and only has one, 1-bit, reserved field. In the first octet of the GTP version1 header, bit 4 is set to zero.
Reserved IE	Both versions of GTP allow up to 255 different Information Elements (IE). However, a number of Information Elements values are undefined or reserved. Packets with reserved or undefined values will be filtered.
Miss Mandatory IE	GTP packets with missing mandatory Information Elements (IE) will not be passed to the GGSN.
Out of State Message	<p>The GTP protocol requires a certain level of state to be kept by both the GGSN and SGSN. Some message types can only be sent when in a specific GTP state. Packets that do not make sense in the current state are filtered or rejected</p> <p>Both versions of GTP allow up to 255 different message types. However, a number of message type values are undefined or reserved.</p> <p>Best practices dictate that packets with reserved or undefined values will be filtered.</p>

Out of State IE	GTP Packets with out of order Information Elements are discarded.
Spoofed Source Address	The End User Address Information Element in the PDP Context Create & Response messages contain the address that the mobile station (MS) will use on the remote network. If the MS does not have an address, the SGSN will set the End User Address field to zero when sending the initial PDP Context Create message. The PDP Context Response packet from the GGSN will then contain an address to be assigned to the MS. In environments where static addresses are allowed, the MS will relay its address to the SGSN, which will include the address in the PDP Context Create Message. As the MS address is negotiated within the PDP Context creation handshake, any packets originating from the MS that contain a different source address are detected and dropped.

Anti-overbilling options

You can configure the FortiOS Carrier firewall to prevent overbilling subscribers for traffic over the. To enable anti-overbilling, you must configure both the Gn/Gp firewall and the Gi firewall.

Expand *Anti-Overbilling* in the GTP profile to reveal these settings.

Anti-Overbilling section of the New GTP Profile page	
Gi Firewall IP Address	The IP address of the unit's interface configured as a Gi gateway.
Port	The SG security port number. The default port number is port 21123. Change this number if your system uses a different SG port.
Interface	Select the unit interface configured as a Gi gateway.
Security Context ID	Enter the security context ID. This ID must match the ID entered on the server Gi firewall. The default security context ID is 696.

Log options

All the GTP logs are treated as a subtype of the event logs. To enable GTP logging, you must:

- configure the GTP log settings in a GTP profile
- enable GTP logging when you configure log and report settings.

To enable GTP logging after a GTP profile has been configured

- 1 Go to *Log & Report > Log Config > Log Setting*.
- 2 Select *Event Logging*, and select *GTP service event*.

3 Select *Apply*.

Log section of the New GTP Profile page	
Log Frequency	<p>Enter the number of messages to drop between logged messages.</p> <p>An overflow of log messages can sometimes occur when logging rate-limited GTP packets exceed their defined threshold. To conserve resources on the syslog server and the FortiOS Carrier unit, you can specify that some log messages are dropped. For example, if you want only every twentieth message to be logged, set a logging frequency of 20. This way, 20 messages are skipped and the next logged.</p> <p>Acceptable frequency values range from 0 to 2147483674. When set to '0', no messages are skipped.</p>
Forwarded Log	Select to log forwarded GTP packets.
Denied Log	Select to log GTP packets denied or blocked by this GTP profile.
Rate Limited Log	Select to log rate-limited GTP packets.
State Invalid Log	Select to log GTP packets that have failed stateful inspection.
Tunnel Limit Log	Select to log packets dropped because the maximum limit of GTP tunnels for the destination GSN is reached.

Extension Log	<p>Select to log extended information about GTP packets. When enabled, this additional information will be included in log entries:</p> <ul style="list-style-type: none"> • IMSI • MSISDN • APN • Selection Mode • SGSN address for signaling • SGSN address for user data • GGSN address for signaling • GGSN address for user data
Traffic count Log	<p>Select to log the total number of control and user data messages received from and forwarded to the GGSNs and SGSNs that the unit protects.</p> <p>The unit can report the total number of user data and control messages received from and forwarded to the GGSNs and SGSNs it protects. Alternately, the total size of the user data and control messages can be reported in bytes. The unit differentiates between traffic carried by each GTP tunnel, and also between GTP-User and GTP-Control messages.</p> <p>The number of messages or the number of bytes of data received from and forwarded to the SGSN or GGSN are totaled and logged if a tunnel is deleted.</p> <p>When a tunnel is deleted, the log entry contains:</p> <ul style="list-style-type: none"> • Timestamp • Interface name (if applicable) • SGSN IP address • GGSN IP address • TID • Tunnel duration time in seconds • Number of messages sent to the SGSN • Number of messages sent to the GGSN

Specifying logging types

You can configure the unit to log GTP packets based on their status with GTP traffic logging.

The status of a GTP packet can be any of the following 5 states:

- **Forwarded** - a packet that the unit transmits because the GTP policy allows it
- **Prohibited** - a packet that the unit drops because the GTP policy denies it
- **Rate-limited** - a packet that the unit drops because it exceeds the maximum rate limit of the destination GSN
- **State-invalid** - a packet that the unit drops because it failed stateful inspection
- **Tunnel-limited** - a packet that the unit drops because the maximum limit of GTP tunnels for the destination GSN is reached.

The following information is contained in each log entry:

- Timestamp
- Source IP address
- Destination IP address
- Tunnel Identifier (TID) or Tunnel Endpoint Identifier (TEID)
- Message type
- Packet status: forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited
- Virtual domain ID or name
- Reason to be denied if applicable.



MMS UTM features

FortiOS Carrier includes all the UTM features of FortiOS with extra features specific to MMS carrier networks.

This section includes:

- [Why scan MMS messages for viruses and malware?](#)
- [MMS virus scanning](#)
- [MMS file filtering](#)
- [MMS content-based Antispam protection](#)
- [MMS DLP archiving](#)

Why scan MMS messages for viruses and malware?

The requirement for scanning MMS content comes from the fact that MMS is an increasingly popular technique for propagating malware between mobile devices. See [“MMS virus scanning” on page 2354](#).

Example: COMMWARRIOR

This is a virus for Series 60 type cell phones, such as Nokia, operating Symbian OS version 6 [or higher]. The object of the virus is to spread to other phones using Bluetooth and MMS as transport avenues. The targets are selected from the contact list of the infected phone and also sought via Bluetooth searching for other Bluetooth-enabled devices (phones, printers, gaming devices etc.) in the proximity of the infected phone.

This virus is more than a proof of concept - it has proven successfully its ability to migrate from a zoo collection to being in-the-wild. Currently, this virus is being reported in over 18 different countries around Europe, Asia and North America.

When the virus first infects a cell phone, a prompt is displayed asking the recipient if they want to install “Caribe”. Symptoms of an infected phone may include rapid battery power loss due to constant efforts by the virus to spread to other phones via a Bluetooth seek-and-connect outreach.

The following variants among others are currently scanned by the FortiOS Carrier devices, in addition to more signatures that cover all known threats.

- **SymbOS/COMWAR.V10B!WORM**

Aliases: SymbOS.Commwarrior.B, SymbOS/Commwar.B, SymbOS/Commwar.B!wm, SymbOS/Commwar.B-net, SymbOS/Commwarrior.b!sis, SymbOS/Comwar.B, SymbOS/Comwar.B!wm, SymbOS/Comwar.B-wm, SYMBOS_COMWAR.B, SymbOS/Comwar.1.0.B!wormSYMBOS/COMWAR.V10B.SP!WORM [spanish version]

First Discovered In The Wild: July 04, 2007

Impact Level: 1

Virus Class: Worm

Virus Name Size: 23,320

- **SymbOS/Commwar.A!worm**

Aliases: Commwarrior-A, SymbOS.Commwarrior.A [NAV], SymbOS/Commwar.A-net, SymbOS/Commwar_ezboot.A-ne, SymbOS/Comwar.A, SymbOS/Comwar.A-wm, SYMBOS_COMWAR.A [Trend]

First Discovered In The Wild: May 16 2005

Impact Level: 1

Virus Class: Worm

Virus Name Size: 27,936

- **SymbOS/Commwarriie.C-wm**

Aliases: None

First Discovered In The Wild: Oct 17 2005

Impact Level: 1

Virus Class: File Virus

Virus Name Size: None

For the latest list of threats Fortinet devices detect, go to the [FortiGuard Center Resource Library's Mobile index](#).

MMS virus scanning

You can use MMS virus scanning to scan content contained within MMS messages for viruses. FortiOS Carrier virus scanning can be applied to the MM1, MM3, MM4, and MM7 interfaces to detect and remove content containing viruses at many points in an MMS network. Perhaps the most useful interface to apply virus scanning would be the MM1 interface to block viruses sent by mobile users before they get into the service provider network.

To go to MMS virus scanning, go to *UTM Profiles > Carrier > MMS Profile*, select an existing or create a new profile, and expand *MMS Scanning*. See “[MMS scanning options](#)” on page 2317.

This section includes:

- [MMS virus monitoring](#)
- [MMS virus scanning blocks messages \(not just attachments\)](#)
- [Scanning MM1 retrieval messages](#)
- [Configuring MMS file filtering](#)
- [Removing or replacing blocked messages](#)
- [Carrier Endpoint Block](#)
- [MMS Content Checksum](#)
- [Passing or blocking fragmented messages](#)
- [Client comforting](#)
- [Server comforting](#)
- [Handling oversized MMS messages](#)

MMS virus monitoring

To enable MMS virus monitoring, expand *MMS Scanning* and enable *Monitor only* for the selected MMS types.

This feature causes the FortiOS Carrier unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Selecting this option enables reporting on viruses and other problems in MMS traffic without affecting users.

MMS virus scanning blocks messages (not just attachments)

To enable MMS virus scanning, expand *MMS Scanning* and enable *Virus Scan* for the selected MMS types.

Because MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configurations also apply to MM1 and MM7 scanning. See

MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configurations also apply to MM3 and MM4 scanning.

The message contents will be scanned for viruses, matched against the file extension blocking lists and scanned for banned words. All these items will be configured via the standard GUI interfaces available for the other protocols and will be controlled at the protection profile level with new options specifically for the MM1 messages.

The FortiOS Carrier unit extracts the sender's Mobile Subscriber Integrated Services Digital Network Number (MSISDN) from the HTTP headers if available. The `POST` payload will be sent to the scanunits which will parse the MMS content and scan each message data section. If any part of the data is to be blocked, the proxy will be informed, the connection to the MMSC will be reset and the FortiOS Carrier unit will return an `HTTP 200 OK` message with an `m-send-conf` payload to the client to prevent a retry. Finally the appropriate logging, alert, and replacement message events will be triggered.

For client notification, the `x-mms-response-status` and `x-mms-response-text` fields can also be customized as required.

Scanning MM1 retrieval messages

To scan MM1 retrieval messages, expand *MMS Scanning* and select *Scan MM1 message retrieval*.

Select to scan message retrievals that use MM1. If you enable *Virus Scan* for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.

Configuring MMS file filtering

To configure MMS file filtering, expand *MMS Scanning* and select *File Filter* for the selected MMS types.

When enabling file filtering there are two options available—`builtin_patterns`, and `all_executables`. Filtering built-in patterns allows you to leverage the FortiOS UTM file patterns, providing a much broader level of file protection. `all_executables` will filter any files that can be executed such as `.exe` and `.sis` files.

Removing or replacing blocked messages

To remove blocked messages, expand *MMS Scanning* and select *Remove Blocked* for the selected MMS types.

Select *Remove Blocked* remove blocked content from each protocol and replace it with the replacement message. If FortiOS Carrier is to preserve the length of the message when removing blocked content, as may occur when billing is affected by the length of the message, select *Constant*.

If you only want to monitor blocked content, select *Monitor Only*.

Carrier Endpoint Block

A carrier endpoint defines a specific client on the carrier network. Typically the client IP address is used to identify the client, however on a carrier network this may be impractical when the client is using a mobile device. Other identifying information such as the MSISDN number is used instead.

This information can be used to block a specific endpoint on the network. Reasons for blocking may include clients whose accounts have expired, clients from another carrier, clients who have sent malicious content (phishing, exploits, viruses, etc), or other violations of terms of use.

Enabling carrier endpoint blocking

To enable carrier endpoint blocking you first need to create a carrier endpoint filter list, and then enable it.

To enable carrier end point blocking - web-based manager

- 1 Create a carrier endpoint filter list. See
- 2 Go to *UTM Profiles > Carrier > MMS Profile*.
- 3 Select *Create New*, or select an existing profile to edit and select *Edit*.
- 4 Expand MMS Scanning.
- 5 Select one or more types of MMS messaging to enable end point blocking on.
- 6 Select the carrier endpoint filter list to use in matching the end points to be blocked.



In MMS Profile, endpoints can only be blocked.

Create a carrier endpoint filter list

A carrier endpoint filter list contains one or more carrier endpoints to match. When used in MMS scanning entries in the filter list that are matched are blocked.

You can configure multiple filter lists for different purposes and groups of clients, such as blocking clients, clients with different levels of service agreements, and clients from other carriers. See [“Carrier endpoint filter lists configuration settings” on page 2334](#).

To create a carrier endpoint filter list - web-based manager

- 1 Go to *UTM Profiles > Carrier > Carrier Endpoint Filter Lists*.
- 2 Select *Create New*.
- 3 Enter a descriptive name for the filter list, such as `blocked_clients` or `CountryX_clients`, and select *OK*.
- 4 Select *Create New* to add one or more entries to the list.
- 5 Select *OK* to return to display the list of filter lists.

Configuring end point filter list entries

For each single end point or group of end points have part of their identifying information in common, you create an entry in the end point filter list.

For example a `blocked_clients` filter list may include entries for single end points added as each one needs to be blocked and a group of clients from a country that does not allow certain services.

To configure an end point filter list entry - web-based manager

- 1 Select *Create New*.
- 2 Enter the following information and select OK.

Name	Name of end point filter list. Select this name in an MMS protection profile.
Comments	Optional description of the end point filter list.
Check/Uncheck All	Select the check box to enable all end point patterns in the MMS filter list. Clear the check box to disable all entries on the MMS filter list. You can also select or clear individual check boxes to enable or disable individual end point patterns.
Pattern	The pattern that FortiOS Carrier uses to match with end points. The pattern can be a single end point or consist of wildcards or Perl regular expressions that will match more than one end point. For more on wildcard and regular expressions, see <i>Using wildcards and Perl regular expressions</i> in the UTM handbook chapter.
Action	Select the action taken by FortiOS Carrier for messages from a carrier end point that matches the end point pattern: None - No action is taken. Block - MMS messages from the end point are not delivered and FortiOS Carrier records a log message. Exempt from mass MMS - MMS messages from the end point are delivered and are exempt from mass MMS filtering. Mass MMS filtering is configured in MMS protection profiles and is also called MMS Bulk Email Filtering and includes MMS message flood protection and MMS duplicate message detection. A valid use of mass MMS would be when a service provider notifies customers of a system-wide event such as a shutdown. Exempt from all scanning - MMS messages from the end point are delivered and are exempt from all MMS protection profile scanning. MMS messages are not subject to protection profile filtering, just MMS protection profile filtering.
Content Archive	MMS messages from the end point are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
Intercept	MMS messages from the end point are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.
Pattern Type	The pattern type: <i>Wildcard</i> , <i>Regular Expression</i> , or <i>Single end point</i> .
Enable	Select to enable this end point filter pattern.

Blocking network access based on end points

You can use end point IP filtering to block traffic from source IP addresses associated with end points. You can also configure FortiOS Carrier to record log messages whenever end point IP filtering blocks traffic. End point IP filtering blocks traffic at the IP level, before the traffic is accepted by a security policy.

To configure end point IP filtering, go to *UTM Profiles > Carrier > IP Filter* and add end points to the IP filter list. For each end point you can enable or disable both blocking traffic and logging blocked traffic.



You cannot add end point patterns to the end point IP filter list. You must enter complete and specific end points that are valid for your network.



The only action available is block. You cannot use end point IP filtering to exempt end points from IP filtering or to content archive or quarantine communication sessions.

FortiOS Carrier looks in the current user context list for the end points in the IP filter list and extracts the source IP addresses for these end points. Then any communication session with a source IP address that matches one of these IP addresses is blocked at the IP level, before the communication session is accepted by a security policy.

FortiOS Carrier dynamically updates the list of IP addresses to block as the user context list changes. Only these updated IP addresses are blocked by end point IP filtering.

For information about the carrier end points and the user context list, including how entries are added to and removed from this list, see For more information on carrier end points, see the [FortiOS Handbook User Authentication](#) chapter.

MMS Content Checksum

The MMS content checksum feature attempts to match checksums of known malicious MMS messages, and on a successful match it will be blocked. The checksums are applied to each part of the message—attached files and message body have separate checksums. These checksums are created with CRC-32, the same method as FortiAnalyzer checksums.

For example, if an MMS message contains a browser exploit in the message body, you can add the checksum for that message body to the list, and future occurrences of that exact message will be blocked. Content will be replaced by the content checksum block notification replacement message for that type of MMS message, and if it is enabled the event will be logged.

One possible implementation would to configure all .sis files to be intercepted. When one is found to be infected or malicious it would be added to the MMS content checksum list.

To use this feature a list of one or more malicious checksums must be created and then the feature is enabled using that list. For a detailed list of options, see “[MMS Content Checksum](#)” on page 2328.

To configure an MMS content checksum list

- 1 Go to *UTM Profiles > Carrier > MMS Content Checksum*.
- 2 Select *Create New*.
- 3 Enter a name for the list of checksums, and select OK.

You are taken to the edit screen for that new list.

- 4 Select *Create New* to add a checksum.
- 5 Enter the *Name* and *Checksum*, and select OK.

The checksum is added to the list.

To add more checksums to the list, repeat steps 4 and 5.

To remove a checksum from the list you can either delete the checksum or simply disable it and leave it in the list.

To enable MMS content checksums, expand *MMS Scanning* and select *MMS Content Checksum* for the selected MMS types. Select the checksum list to match.

Passing or blocking fragmented messages

Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 message are blocked.

The *Interval* is the time in seconds before client comforting starts after the download has begun, and the time between sending subsequent data.

The *Amount* is the number of bytes sent by client or server comforting at each interval.

Client comforting

In general, client comforting is available for MM1 and MM7 messaging and provides a visual display of progress for web page loading or HTTP or FTP file downloads. Client comforting does this by sending the first few packets of the file or web page being downloaded to the client at configured time intervals so that the client is not aware that the download has been delayed. The client is the web browser or FTP client. Without client comforting, clients and their users have no indication that the download has started until the FortiOS Carrier unit has completely buffered and scanned the download. During this delay users may cancel or repeatedly retry the transfer, thinking it has failed.

The appearance of a client comforting message (for example, a progress bar) is client-dependent. In some instances, there will be no visual client comforting cue.

During client comforting, if the file being downloaded is found to be infected, then the FortiOS Carrier unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead the download stops, and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, then the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned (and therefore potentially infected) content to the client. Only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

MM1 and MM7 client comforting steps

Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.

The following steps show how client comforting works for a download of a 1 Mbyte file with the client comforting interval set to 20 seconds and the client comforting amount set to 512 bytes.

- 1 The client requests the file.
- 2 The FortiOS Carrier unit buffers the file from the server. The connection is slow, so after 20 seconds about one half of the file has been buffered.
- 3 The FortiOS Carrier unit continues buffering the file from the server, and also sends 512 bytes to the client.
- 4 After 20 more seconds, the FortiGate unit sends the next 512 bytes of the buffered file to the client.
- 5 When the file has been completely buffered, the client has received the following amount of data:

$$ca * (T/ci) \text{ bytes} == 512 * (40/20) == 512 * 2 == 1024 \text{ bytes,}$$
 where *ca* is the client comforting amount, *T* is the buffering time and *ci* is the client comforting interval.
- 6 If the file does not contain a virus, the FortiOS Carrier unit sends the rest of the file to the client. If the file is infected, the FortiOS Carrier unit closes the data connection but cannot send a message to the client.

Server comforting

Server comforting can be selected for each protocol.

Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for FortiOS Carrier to buffer and scan large `POST` requests from slow clients.

The *Interval* is the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.

The *Amount* is the number of bytes sent by client or server comforting at each interval.

Handling oversized MMS messages

Select *Block* or *Pass* for files and email messages exceeding configured thresholds for each protocol.

The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.

MM1 sample messages

```
Internet Protocol, Src Addr: 10.128.206.202 (10.128.206.202), Dst
Addr: 10.129.192.190 (10.129.192.190)
Transmission Control Protocol, Src Port: 34322 (34322), Dst Port:
http (80), Seq: 1, Ack: 1, Len: 1380
  Source port: 34322 (34322)
  Destination port: http (80)
  Header length: 20 bytes
  Flags: 0x0010 (ACK)
  Window size: 24840
  Checksum: 0x63c1 (correct)
```

HTTP proxy

```

Hypertext Transfer Protocol
  POST / HTTP/1.1\r\n
    Request Method: POST
    Request URI: /
    Request Version: HTTP/1.1
  Host: 10.129.192.190\r\n
  Accept: */*, application/vnd.wap.sic,application/vnd.wap.mms-
message,text/x-hdml,image/mng,image/x-mng,video/mng,video/x-
mng,image/bmp\r\n
  Accept-Charset: utf-8,*\r\n
  Accept-Language: en\r\n
  Content-Length: 25902\r\n
  Content-Type: application/vnd.wap.mms-message\r\n
  User-Agent: Nokia7650/1.0 SymbianOS/6.1 Series60/0.9
Profile/MIDP-1.0 Configuration/CLDC-1.0 UP.Link/6.2.1\r\n
  x-up-devcap-charset: utf-8\r\n
  x-up-devcap-max-pdu: 102400\r\n
  x-up-uplink: magh-ip.mi.vas.omnitel.it\r\n
  x-wap-profile: "http://nds.nokia.com/uaprof/N7650r200.xml"\r\n
  x-up-subno: 1046428312-826\r\n
  x-up-calling-line-id: 393475171234\r\n
  x-up-forwarded-for: 10.211.4.12\r\n
  x-forwarded-for: 10.211.4.12\r\n
  Via: 1.1 magh-ip.mi.vas.omnitel.it\r\n
\r\n

```

Scan engine

```

MMS Message Encapsulation, Type: m-send-req
  X-Mms-Message-Type: m-send-req (0x80)
  X-Mms-Transaction-ID: 1458481935
  X-Mms-MMS-Version: 1.0
  From: <insert address>
  To: 3475171234/TYPE=PLMN
  X-Mms-Message-Class: Personal (0x80)
  X-Mms-Expiry: 21600.000000000 seconds
  X-Mms-Priority: Normal (0x81)
  X-Mms-Delivery-Report: No (0x81)
  X-Mms-Read-Report: No (0x81)
  Content-Type: application/vnd.wap.multipart.related;
start=<1822989907>; type=application/smil
    Start: <1822989907>
    Type: application/smil
  Data (Post)
    Multipart body
      Part: 1, content-type: text/plain
        Content-Type: text/plain; charset=iso-10646-ucs-2;
name=Ciao.txt
        Charset: iso-10646-ucs-2
        Name: Ciao.txt
      Headers
        Content-Location: Ciao.txt
        Line-based text data: text/plain

```



```
\377\376C\000i\000a\000o\000  
[Unreassembled Packet: MMSE]
```

Configuring MMS virus scanning

To apply MMS virus scanning you must configure MMS virus scanning in the MMS protection profile, and add the MMS protection profile to a security policy.

The MMS protection profile then applies to the traffic accepted by the security policy.

To apply MMS virus scanning - web-based manager

- 1 Go to *UTM Profiles > Carrier > MMS Profile*.
- 2 Select *Create New* to add an MMS protection profile called `MMS_virus_scan`.
- 3 Configure antivirus settings and save the new MMS protection profile.
- 4 Go to *Policy*.
- 5 Select *Create New* to add a security policy, or select *Edit* for the policy to which you want to add the protection profile.
- 6 Configure the security policy as required.
- 7 Select *UTM Profiles > MMS Profile* and select `MMS_virus_scan`.
- 8 Select *OK*.

Replacement messages

FortiOS Carrier generates replacement messages to notify the sending client that they have sent a virus.

FortiOS Carrier can generate an SMS/SMTP replacement message and an MMS/HTTP POST replacement message. In each case the recipient will be the sender of the initial virus message and a configurable MSISDN parameter will be available to determine the sender (From) – i.e. the FortiOS Carrier unit.

For SMS/SMTP notification a destination email address is configurable and can contain the marker `%%MSISDN%%` which will be replaced with the sender's MSISDN thereby allowing the message to be routed properly.

You need to clarify whether specific headers are required in the SMTP message and whether a predefined format for the message must to be followed. For the MMS message the body will be configurable and this could be specified in WML or SMIL.

For more information on FortiOS Carrier replacement messages, see [“MMS Replacement messages” on page 2393](#). For more information on address translation, see For more information on carrier end points, see the [FortiOS Handbook v3 User Authentication](#) chapter.

Logging and reporting

With each virus infection, or file block, a syslog message is generated. The format of this syslog message is similar to:


```

2005-09-22 19:15:47 device_id=FGT5001ABCDEF1234
log_id=0211060ABC type=virus subtype=infected pri=warning
src=10.1.2.3 dst=10.2.3.4 src_int=port1 dst_int=port2
service=mm1 status=blocked from="<sending MSISDN>"
to="<receiving MSISDN>" file="eicar.com.txt"
virus="EICAR_TEST_FILE" msg="The file eicar.com.txt is
infected with EICAR_TEST_FILE. ref
http://www.fortinet.com/VirusEncyclopedia/search/encyclopediaSearch.do?method=quickSearchDirectly&virusName=EICAR_TEST_FILE.

```

Note that the *from* and *to* fields are samples and not real values.

MMS logging options

You can enable logging in an MMS protection profile to write event log messages when the MMS protection profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS protection profile logging options to write an event log message every time a virus is detected.

To record these log messages you must first configure how the FortiOS Carrier unit stores log messages.

To configure MMS content archiving, go to *UTM Profiles > Carrier > MMS Profile*. Select *Create New* or select the *Edit* icon beside an existing profile. Expand *MMS Bulk AntiSpam Detection > Logging*. Complete the fields as described in the following table and select *OK*. For more a detailed list of options, see “[Logging](#)” on page 2327.

SNMP

A simple SNMP trap will be generated to inform the operators alerting system that a virus has been detected. This SNMP trap could contain the sending and receiving MSISDN however the initial solution would reflect the current behavior, i.e. only the fact that a virus has been detected will be communicated.

MMS file filtering

Use MMS file filtering to apply antivirus file filtering to MMS traffic. Select a file filter list to apply. To configure MMS file filtering, go to *UTM Profiles > Carrier > MMS Profile*, select an existing or create a new profile, and expand *MMS Scanning*.

Configure the FortiGate file filter to block files by:

- **File pattern:** Files can be blocked by name, extension, or any other pattern. File pattern blocking provides the flexibility to block potentially harmful content.
File pattern entries are not case sensitive. For example, adding *.exe to the file pattern list also blocks any files ending in .EXE.
In addition to the built-in patterns, you can specify more file patterns to block.
- **File type:** Files can be blocked by type, without relying on the file name to indicate what type of files they are. When blocking by file type, the FortiGate unit analyzes the file and determines the file type regardless of the file name.

For standard operation, you can choose to disable file filter in the protection profile, and enable it temporarily to block specific threats as they occur.

The FortiGate unit can take either of these actions toward files that match a configured file pattern or type:

- **Allow:** the file is allowed to pass.

- **Block:** the file is blocked and a replacement messages will be sent to the user. If both file filter and virus scan are enabled, the FortiOS Carrier unit blocks files that match the enabled file filter and does not scan these files for viruses.
- **Intercept:** the file will be archived to the local hard disk or the FortiAnalyzer unit.

The FortiOS Carrier unit also writes a message to the virus log and sends an alert email message if configured to do so.

Files are compared to the enabled file patterns and then the file types from top to bottom. If a file does not match any specified patterns or types, it is passed along to antivirus scanning (if enabled). In effect, files are passed if not explicitly blocked.

Using the allow action, this behavior can be reversed with all files being blocked unless explicitly passed. Simply enter all the file patterns or types to be passed with the allow attribute. At the end of the list, add an all-inclusive wildcard (*.*) with a block action. Allowed files continue to antivirus scanning (if enabled) while files not matching any allowed patterns are blocked by the wildcard at the end.

Built-in patterns and supported file types

The FortiGate unit is preconfigured with a default list of file patterns:

- executable files (*.bat, *.com, and *.exe)
- compressed or archive files (*.gz, *.rar, *.tar, *.tgz, and *.zip)
- dynamic link libraries (*.dll)
- HTML application (*.hta)
- Microsoft Office files (*.doc, *.ppt, *.xl?)
- Microsoft Works files (*.wps)
- Visual Basic files (*.vb?)
- screen saver files (*.scr)
- program information files (*.pif)
- control panel files (*.cpl)

The FortiGate unit can take actions against the following file types:

Table 126: Supported file types

arj	activemime	aspack	base64	bat	binhex	bzip	bzip2
cab	class	cod	elf	exe	fsg	gzip	hlp
hta	html	jad	javascript	lzh	mime	msc	msoffice
petite	prc	rar	sis	tar	upx	uue	zip
unknown	ignored						



The “unknown” type is any file type that is not listed in the table. The “ignored” type is the traffic the FortiGate unit typically does not scan. This includes primarily streaming audio and video.

Filtering based on file name

There are filenames that are known to be associated with malware such as viruses and trojans. There are filenames you may associate with other undesirable content in addition to malware. In these situations you want to select specific filenames to filter.

You do not have to match the entire filename. For example if you wanted to block all files with the word `trojan` in them you could use wildcards to accomplish this - `*trojan*`. This allows you to select the entire filename, part of the filename, or just the file type to match.

The following procedure creates a filter list called `filterExampleFiles` that filters two files called `exampleTrojanFile.abc` and `*trojan*.def`. When completed, this file filter list can be included in an MMS profile.

To create a file filter based on file name - web-based manager

- 1 Go to *UTM Profiles > AntiVirus > File Filtering*.
- 2 Select *Create New*, to create a new file filtering list.
- 3 Name the list `filterExampleFiles`.
- 4 Select *Create New* to add a filter to the list.
- 5 Select *File Name Pattern for Filter Type*.
- 6 Enter `exampleTrojanFile.abc`.
- 7 Enter *Block* for the *Action*.
- 8 Select *Enable*, and *OK*.
- 9 Select *Create New* to add a filter to the list.
- 10 Select *File Name Pattern for Filter Type*.
- 11 Enter `*trojan*.def`.
- 12 Enter *Block* for the *Action*.
- 13 Select *Enable*, and *OK*.

Filtering based on file type

When filtering files, it is often useful to filter based on the file type. When malware finds a file type that allows them access to a system, the filename will change but the file type will remain the same. Even for preventing applications that are not malware but simply undesirable, filtering based on file type is often the easiest method.

Simply matching the file type, `.zip` for example, may not be as accurate a method as using the built-in patterns. If users see that `.zip` attachments are blocked, they may simply rename the file so the filters will allow it through. Checking against patterns can help prevent this bypassing.

There are two possible methods available to filter based on file type. If the file type is one of the built-in patterns, you can use them - for example blocking PalmOS files on your network since Palm devices are not supported. Otherwise, you can simply use wildcards to match the file type.

The following example will filter all batch files (`.bat`).

To filter files based on file type using file name pattern - web-based manager

- 1 Go to *UTM Profiles > AntiVirus > File Filter*.
- 2 Select *Create New* and name the list `blockedFileTypes`.
- 3 Select *Create New* to add files to the list.
- 4 Select *File name pattern for Filter Type*.
- 5 Enter `*.bat` for *Pattern*.
- 6 Select an *Action* of *Block*.

- 7 Select *Enable* and *OK*.
- 8 At the file filter list, select *OK*.

The file filter is now available to be used in an MMS profile, and will block all `.bat` files that MMS profile matches.

To filter files based on file type using file type - web-based manager

- 1 Go to *UTM Profiles > AntiVirus > File Filter*.
- 2 Select *Create New* and name the list `blockedFileTypes`.
- 3 Select *Create New* to add files to the list.
- 4 Select *File Type for Filter Type*.
- 5 Select *Batch File (bat)* for *File Type*.
- 6 Select an *Action* of *Block*.
- 7 Select *Enable* and *OK*.
- 8 At the file filter list, select *OK*.

The file filter is now available to be used in an MMS profile, and will block all batch files (that use `.bat` file extension) that the MMS profile matches.

MMS file filtering blocks messages (not just attachments)

When MMS file filtering finds a matching file in an MMS message, the entire message is blocked. This action is more secure, and can reduce the amount of processing required for that message. For example if one MMS message includes three files, and the first one is blocked then the other files won't be scanned or attempted to be matched because the whole message is already being blocked.

Configuring MMS file filtering

To apply MMS file filtering you must begin with a file filter list. You can create your own list or use the built-in patterns list. You then must create the file filter, then add the file filter list to an MMS profile, and then add the MMS profile to a security policy. The MMS profile, and the corresponding file filter then applies to the traffic accepted by the security policy.

The following procedure creates a file filtering list called `MMS_file_filter` that is used in the MMS profile called `filtering_profile`. The filter will be applied to all MMS message types.

To apply MMS file filtering

- 1 Go to *UTM Profiles > Antivirus > File Filter* and create a file filter list called `MMS_file_filter`.
- 2 Select *Create New* to add entries to filter specific file types.
- 3 Select *OK*.
- 4 Go to *UTM Profiles > Carrier > MMS Profile* and create a new profile called `filtering_profile`.
- 5 Expand *MMS Scanning*, and select all the MMS message types.
- 6 Select `MMS_file_filter` from the *Option* drop down menu.
- 7 Set other settings as required in the MMS profile.
- 8 Select *OK*.
- 9 Go to *Policy*, and select *Create New*.

- 10 Within the new policy select UTM Profiles > MMS Profile, and select `filtering_profile`.
- 11 Configure the security policy as required.
- 12 Select OK.

Configuring sender notifications

In most cases you will notify the sender that they are causing problems on the network — either by sending malware content, flooding the network, or some other unwanted activity. The notification assumes the sender is unaware of their activity and will stop or correct it when notified.

However, senders who are notified may use this information to circumvent administration's precautions. For example if flood notification is set to 1000 messages per minute, a notified user may simply reduce their message to 990 messages per minute if this flood is intentional. For this reason, not all problems include sender notifications.

There are two methods of notifying senders:

- [MMS notifications](#)
- [Replacement messages](#)

MMS notifications

MMS notifications enable you to customize notifications for many different situations and differently for all the supported MMS message protocols — MM1, MM3, MM4, and MM7.

MMS notification types include:

- Content Filter
- File Block
- Carrier End Point Block
- Flood
- Duplicate
- MMS Content Checksum
- Virus Scan

Day of Week, *Window start time* and *Window Duration* define what days and what time of day alert notifications will be sent. This allows you to control what alerts are sent on weekends. It also lets you control when to start sending notifications each day. This can be useful if system maintenance is performed at the same time each night — you might want to start alert notifications after maintenance has completed. Another reason to limit the time alert messages are sent could be to limit message traffic to business hours.

Figure 246: Notifications screen for FortiOS Carrier MMS Profile

Edit MMS Profile																
MMS Bulk Email Filtering Detection MMS Address Translation MMS Notifications																
Option																
AntiVirus Notification List -- Disabled --																
	MM1				MM3				MM4				MM7			
Message Protocol	mm1				mm3				mm4				mm7			
Message Type													deliver.REQ			
Detect Server Details	<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>			
Hostname																
URL	/												/			
Port	80				25				25				80			
Username																
Password																
VASP ID																
VAS ID																
All Notification Types	<input checked="" type="checkbox"/> 24 hour(s)				<input checked="" type="checkbox"/> 24 hour(s)				<input checked="" type="checkbox"/> 24 hour(s)				<input checked="" type="checkbox"/> 24 hour(s)			
Notifications Per Second Limit	0				0				0				0			
Day of Week	<input checked="" type="checkbox"/> Sun <input checked="" type="checkbox"/> Mon <input checked="" type="checkbox"/> Tue <input checked="" type="checkbox"/> Wed <input checked="" type="checkbox"/> Thu <input checked="" type="checkbox"/> Fri <input checked="" type="checkbox"/> Sat				<input checked="" type="checkbox"/> Sun <input checked="" type="checkbox"/> Mon <input checked="" type="checkbox"/> Tue <input checked="" type="checkbox"/> Wed <input checked="" type="checkbox"/> Thu <input checked="" type="checkbox"/> Fri <input checked="" type="checkbox"/> Sat				<input checked="" type="checkbox"/> Sun <input checked="" type="checkbox"/> Mon <input checked="" type="checkbox"/> Tue <input checked="" type="checkbox"/> Wed <input checked="" type="checkbox"/> Thu <input checked="" type="checkbox"/> Fri <input checked="" type="checkbox"/> Sat				<input checked="" type="checkbox"/> Sun <input checked="" type="checkbox"/> Mon <input checked="" type="checkbox"/> Tue <input checked="" type="checkbox"/> Wed <input checked="" type="checkbox"/> Thu <input checked="" type="checkbox"/> Fri <input checked="" type="checkbox"/> Sat			
Window Start Time	00 : 00				00 : 00				00 : 00				00 : 00			
Window Duration	24 : 00				24 : 00				24 : 00				24 : 00			
DLP Archive Logging																
<input type="button" value="OK"/> <input type="button" value="Cancel"/>																

For MMS Notification options, see [“MMS Notifications” on page 2324](#).

Replacement messages

FortiOS Carrier units and FortiGate units alike send replacement messages when messages or content is blocked, quarantined, or otherwise diverted from the receiver. In its place a message is sent to notify the receiver what happened.

With FortiOS Carrier MMS replacement messages, send and receive message types are supported separately and receive their own custom replacement messages. This allows the network to potentially notify both the sender and receiver of the problem.

For example the replacement message *MM1 send-req file block message* is sent to the device that sent one or more files that were banned. The default message that is sent is *This device has sent %%NUM_MSG%% messages containing banned files in the last %%DURATION%% hours. The two variables are replaced by the appropriate values.*

Replacement messages are not as detailed or specific as MMS notifications, but they are also not as complicated to configure. They are also useful when content has been removed from an MMS message that was still delivered.

For more information on replacement messages, see [“MMS Replacement messages” on page 2393](#).

MMS content-based Antispam protection

Expand *MMS Scanning* and select *Content Filter* in an MMS protection profile to create content filter black/white lists that block or allow MMS messages based on the content of the message.

Overview

A school computer lab may block age-inappropriate content. A place of business may block unproductive content. A public access internet cafe may block offensive and graphic content. Each installation has its own requirements for what content needs to be blocked, and in what language.

FortiOS Carrier provides the ability to create custom local dictionaries, black lists, and white lists in multiple languages enables you to protect your customers from malicious content around the world.

Content-based protection includes:

- [Configurable dictionary](#)
- [Black listing](#)
- [White listing](#)

Configurable dictionary

You can add a dictionary of configurable terms and phrases by going to *UTM Profiles > Web Filter > Web Content Filter*. The text of MMS messages can be searched for these terms and phrases. Add content filter lists that contain content that you want to match in MMS messages. For every match found, a score is added. If enough matches are found to set the total score above the configured threshold, the MMS message is blocked.

You can add words, phrases, wild cards and Perl regular expressions to create content patterns that match content in MMS messages. For more on wildcard and regular expressions, see *Using wildcards and Perl regular expressions* in the [UTM](#) handbook chapter.

For each pattern you can select *Block* or *Exempt*.

- Block adds an antispam black list pattern. A match with a block pattern blocks a message depending on the score of the pattern and the content filter threshold.
- Exempt adds an antispam white list pattern. A match with an exempt pattern allows the message to proceed through the FortiOS Carrier unit, even if other content patterns in the same content filter list would block it.

If a pattern contains a single word, the FortiOS Carrier unit searches for the word in MMS messages. If the pattern contains a phrase, the FortiOS Carrier unit searches for all of the words in the phrase. If the pattern contains a phrase in quotation marks, the FortiOS Carrier unit searches for the whole phrase.

You can create patterns with Simplified Chinese, Traditional Chinese, Cyrillic, French, Japanese, Korean, Spanish, Thai, or Western character sets.

Black listing

Black listing is the practice of banning entries on the list. For example if an IP address continuously sends viruses, it may be added to the black list. That means any computers that consult that list will not communicate with that IP address.

Sometimes computers or devices can be added to black lists for a temporary problem, such as a virus that is removed when notified. However, as a rule short of contacting the administrator in person to manually be removed from the black list, users have to wait and they generally will be removed after a period without problem.

White listing

White listing is the practice of adding all critical IP addresses to a list, such as company email and web servers. Then if those servers become infected and start sending spam or viruses, those servers are not blocked. This allows the critical traffic through, even if there might be some malicious traffic as well. Blocking all traffic from your company servers would halt company productivity.

Scores and thresholds

Each content pattern includes a score. When a MMS message is matched with a pattern the score is recorded. If a message matches more than one pattern or matches the same pattern more than once, the score for the message increases. When the total score for a message equals or exceeds the threshold the message is blocked.

The default score for a content filter list entry is 10 and the default threshold is 10. This means that by default a message is blocked by a single match. You can change the scores and threshold so that messages can only be blocked if there are multiple matches. For example, you may only want to block messages that contain the phrase "example" if it appears twice. To do this, add the "example" pattern, set action to block and score to 5. Keep the threshold at 10. If "example" is found twice or more in a message the score adds up 10 (or more) and the message is blocked.

Configuring content-based antispam protection

To apply content-based antispam protection - web-based manager

- 1 Go to *UTM Profiles > Web Filter > Web Content Filter* and create or edit a web content filter list.
- 2 Go to *UTM Profiles > Carrier > MMS Profile* and add or edit an MMS protection profile.
- 3 Select *MMS Scanning > Content Filter* and select the web content filter list.
- 4 Optionally change the content filter *Threshold* and save the MMS protection profile.
- 5 Go to *Policy* and create or edit a policy.
- 6 Expand UTM heading and select MMS profile added above.
- 7 Select *OK*.
- 8 Configure the rest of the security policy as required, and select *OK*.

Configuring sender notifications

When someone on the MMS network sends an MMS message that is blocked, in most cases you will notify the sender. Typically an administrator is notified in addition to the sender so action can be taken if required.

There are two types of sender notifications available in FortiOS Carrier:

- [MMS notifications](#)
- [Replacement messages](#)

MMS notifications

MMS notifications to senders are configured in UTM Profiles > Carrier > MMS profile, under MMS Notifications.

In this section you can configure up to four different notification recipients for any combination of MM1/3/4/7 protocol MMS messages. Also for MM7 messages the message type can be `submit.REQ` or `deliver.REQ`.

Useful settings include:

- delay in message based on notification type
- limit on notifications per second to prevent a flood
- schedules for notifications
- log in details for MM7 messages.

For more information on MMS notifications, see [“Notifying message flood senders and receivers” on page 2378](#) and [“MMS Notifications” on page 2324](#).

Replacement messages

Replacement messages are features common to both FortiOS and FortiOS Carrier, however FortiOS Carrier has additional messages for the MMS traffic.

While each MMS protocol has its own different replacement messages, the one common to all MMS protocols is the *MMS blocked content replacement message*. This is the message that the receiver of the message sees when their content is blocked.

For more information on replacement messages, see [“MMS Replacement messages” on page 2393](#).

MMS DLP archiving

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the FortiOS Carrier configuration. The FortiGuard Analysis and Management server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

You can configure full DLP archiving and summary DLP archiving. Full DLP archiving includes all content, for example, full email DLP archiving includes complete email messages and attachments. Summary DLP archiving includes just the meta data about the content, for example, email message summary records include only the email header.

You can archive MM1, MM3, MM4, and MM7 content:

Configuring MMS DLP archiving

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. For each protocol you can archive just session metadata (*Summary*), or metadata and a copy of the associated file or message (*Full*).

In addition to MMS protection profile DLP archive options you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select *DLP archiving* for carrier end point patterns in a *Carrier End Point List* and select the *Carrier End Point Block* option in the *MMS Scanning* section of an MMS Protection Profile

FortiOS Carrier only allows one sixteenth of its memory for transferring content archive files. For example, for FortiOS Carrier units with 128MB RAM, only 8MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

To configure MMS DLP archiving - web-based manager

- 1 Go to *UTM Profiles > Carrier > MMS Profile*.
- 2 Select *Create New* or select the *Edit* icon beside an existing profile.
- 3 Expand *MMS Bulk AntiSpam Detection > Content Archive*.
- 4 Complete the fields as described in “[DLP Archive options](#)” on page 2326.
- 5 Select *OK*.

Viewing DLP archives

You can view DLP archives from the FortiOS Carrier unit web-based manager. Archives are historical logs that are stored on a log device that supports archiving, such as a FortiAnalyzer unit.

These logs are accessed from either *Log&Report > DLP Archive* or if you subscribed to the FortiGuard Analysis and Management Service, you can view log archives from there.


The *DLP Archive* menu is only visible if one of the following is true.

- You have configured the FortiGate unit for remote logging and archiving to a FortiAnalyzer unit.
- You have subscribed to the FortiGuard Analysis and Management Service.

The following tabs are available when you are viewing DLP archives for one of these protocols.

- *E-mail* to view POP3, IMAP, SMTP, POP3S, IMAPS, SMTPS, and spam email archives.
- *Web* to view HTTP and HTTPS archives.
- *FTP* to view FTP archives.
- *IM* to view AIM, ICQ, MSN, and Yahoo! archives.
- *MMS* to view MMS archives.
- *VoIP* to view session control (SIP, SIMPLE and SCCP) archives.

If you need to view log archives in Raw format, select *Raw* beside the *Column Settings* icon.

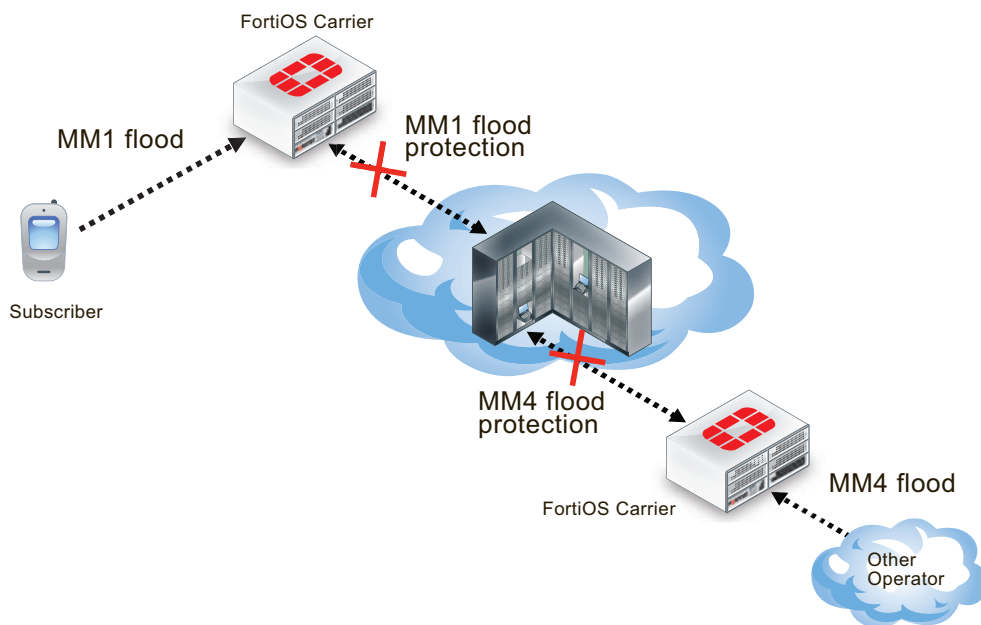


Message flood protection

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse.

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many messages from the same subscriber to your MMSC.

Figure 247: MM1 and MM4 flood protection



The FortiOS Carrier unit keeps track of the number of messages each subscriber sends for the length of time you specify. If the number of messages a subscriber sends exceeds the threshold, a configured action is taken. Possible actions are logging the flood, blocking or intercepting messages in the flood, archiving the flood messages, and sending an alert message to inform the administrator that the flood is occurring.

You can create three different thresholds to take different levels of action at different levels of activity.

With this highly configurable system, you can prevent subscribers from sending more messages than you determine is acceptable, or monitor anyone who exceeds the thresholds.

Setting message flood thresholds

A message flood occurs when a single subscriber sends a volume of messages that exceeds the flood threshold you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected.

If a subscriber exceeds the message flood threshold and is blocked from sending more messages, any further attempts to send messages will re-start the block period. You must also enable logging for *MMS Scanning > Bulk Messages* in the Logging section of the MMS protection profile.



A subscriber is still able to receive messages while they are blocked from sending messages.

Example

For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all messages blocked for half an hour (30 minutes).

Using this example, if the subscriber exceeds the flood threshold, they are blocked from sending message for 30 minutes. If the subscriber tries to send any message after 15 minutes, the message will be blocked and the block period will be reset again to 30 minutes. The block period must expire with no attempts to send a message. Only then will the subscriber be allowed to send more messages.

To configure MM1 message flood threshold - web-based manager

- 1 Go to *UTM Profiles > Carrier > MMS Profile*.
- 2 Select *Create New*.
- 3 Enter `MM1 flood` for *Profile Name*.
- 4 Expand *MMS Bulk Email Filtering Detection*.
- 5 Enter the following information, and select *OK*.

MM1 (first column)	
Enable	Enable
Message Flood Window	60 minutes
Message Flood Limit	100
Message Flood Block Time	30 minutes
Message Flood Action	Block

To configure MM1 message flood threshold - CLI

```
config firewall mms-profile
edit profile_name
config flood mm1
set status1 enable
set window1 60
set limit1 100
set action1 block
```

```
        set block-time1 30
    end
end
```

The threshold values that you set for your network will depend on factors such as how busy your network is and the kinds of problems that your network and your subscribers encounter. For example, if your network is not too busy you may want to set message flood thresholds relatively high so that only an exceptional situation will exceed a flood threshold. Then you can use log messages and archived MMS messages to determine what caused the flood.

If your subscribers are experiencing problems with viruses that send excessive amounts of messages, you may want to set thresholds lower and enable blocking to catch problems as quickly as possible and block access to keep the problem from spreading.

Flood actions

When the FortiOS Carrier unit detects a message flood, it can take any combination of the five actions that you can configure for the flood threshold. For detailed options, see [“Message Flood” on page 2331](#).

Notifying administrators of floods

You can configure alert notifications for message floods by selecting the Alert Notification message flood action.

The FortiOS Carrier unit sends alert notifications to administrators using the MM1, MM3, MM4, or MM7 content interface. To send an alert notification you must configure addresses and other settings required for the content interface.

For example, to send notifications using the MM1 content interface you must configure a source MSISDN, hostname, URL, and port to which to send the notification. You can also configure schedules for when to send the notifications.

Finally you can add multiple MSISDN numbers to the MMS protection profile and set which flood thresholds to send to each MSISDN.

Example — three flood threshold levels with different actions for each threshold

You can set up to three threshold levels to take different actions at different levels of activity.

The first example threshold records log messages when a subscriber’s handset displays erratic behavior by sending multiple messages using MM1 at a relatively low threshold. The erratic behavior could indicate a problem with the subscriber’s handset. For example, you may have determined for your network that if a subscriber sends more the 45 messages in 30 minutes that you want to record log messages as a possible indication or erratic behavior.

From the web-based manager in an MMS profile set message *Flood Threshold 1* to:

Enable	Selected
Message Flood Window	30 minutes

Message Flood Limit	45
Message Flood Action	Log

From the CLI:

```
config firewall mms-profile
edit profile_name
config flood mm1
set status1 enable
set window1 30
set limit1 45
set action1 log
end
end
```

Set a second higher threshold to take additional actions when a subscriber sends more than 100 messages in 30 minutes. Set the actions for this threshold to log the flood, archive the message that triggered the second threshold, and block the sender for 15 minutes.

From the web-based manager in an MMS profile set message *Flood Threshold 2* to:

Enable	Selected
Message Flood Window	30 minutes
Message Flood Limit	100
Message Block Time	15 minutes
Message Flood Action	Log, DLP archive First message only, Block

From the CLI:

```
config firewall mms-profile
edit profile_name
config flood mm1
set status2 enable
set window2 30
set limit2 100
set action2 block log archive-first
set block-time2 15
end
end
```

Set the third and highest threshold to block the subscriber for an extended period and send an administrator alert if the subscriber sends more than 200 messages in 30 minutes. Set the actions for this threshold to block the sender for four hours (240 minutes), log the flood, archive the message that triggered the third threshold, and send an alert to the administrator.

From the web-based manager in an MMS profile set message *Flood Threshold 3* to:

Enable	Selected
Message Flood Window	30 minutes
Message Flood Limit	200
Message Block Time	240 minutes
Message Flood Action	Log, Block, Alert Notification

Because you have selected the *Alert Notification* action you must also configure alert notification settings. For this example, the source MSISDN is 5551234—telephone number 555-1234. When administrators receive MMS messages from this MSISDN they can assume a message flood has been detected.

In this example, alert notifications are sent by the FortiOS Carrier unit to the MMSC using MM1. The host name of the MMSC is `mmscexample`, the MMSC URL is `/`, and the port used by the MMSC is 80. In this example, the alert notification window starts at 8:00am and extends for eight hours on weekdays (Monday-Friday) and the minimum interval between message flood notifications is two hours.

Source MSISDN	5551234
Message Protocol	MM1
Hostname	mmscexample
URL	/
Port	80
Notifications Per Second Limit	0
Window Start Time	8:00
Window Duration	8:00
Day of Week	Mon, Tue, Wed, Thu, Fri, Sat
Interval	2 hours

From the CLI:

```
config firewall mms-profile
edit profile_name
config notification alert-flood-1
set alert-src-msisdn 5551234
set set msg-protocol mm1
set mmsc-hostname mmscexample
set mmsc-url /
set mmsc-port 80
set rate-limit 0
set tod-window-start 8:00
set tod-window-duration 8:00
set days-allowed monday tuesday wednesday thursday friday
set alert-int 2
set alert-int-mode hours
end
end
```

You must also add the MSISDNs of the administrators to be notified of the message flood. In this example, the administrator flood threshold 3 alert notifications are sent to one administrator with MSISDN 5554321.

To add administrator's MSISDNs for flood threshold 3 from the web-based manager when configuring a protection profile, select *MMS Bulk Email Filtering Detection > Recipient MSISDN > Create New*.

MSISDN	5554321
Flood Level 3	Select

From the CLI:

```
config firewall mms-profile
edit profile_name
config notif-msisdn
edit 5554321
set threshold flood-thresh-3
end
end
```

Notifying message flood senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver that cause a message flood. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around flood protection. For this reason, no notification is set to the sender or receiver.

However, FortiOS Carrier does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as message floods. For information about how FortiOS Carrier responds when message flood detection blocks a message, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 2301](#).

Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (`m-send.conf`) to the sender — otherwise the sender’s handset would keep retrying the message. The `m-send.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the MMSC and the MMSC sends the `m-send.conf` message.

You can customize the `m-send.conf` message by editing the *MM1 send-conf flood message* MM1 replacement message (from the CLI the `mm1-send-conf-flood` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted”. To hide the fact that FortiOS Carrier is responding to a flood, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK”:

```
config system replacemsg mm1 mm1-send-conf-flood
set rsp-status ok
set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (`m-retrieve.conf`) to the sender (otherwise the sender’s handset would keep retrying the message). The `m-retrieve.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the receiver, so the MMSC sends the `m-retrieve.conf` message.

You can customize the m-retrieve.conf message by editing the *MM1 retrieve-conf flood message* MM1 replacement message (from the CLI the `mm1-retr-conf-flood` replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mm1 mm1-retr-conf-flood
    set subject "Message blocked"
    set message "Message temporarily blocked by carrier"
end
```

Forward responses for MM4 message floods

When the FortiOS Carrier unit identifies an MM4 message as a flood message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4_forward.res message is sent only when the MM4 message flood action is set to Block and the MM4-forward.req message requested a response. For more information, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 2301](#).

You can customize the MM4_forward.res message by editing the *MM4 flood message* MM4 replacement message (from the CLI the `mm4-flood` replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted” (`err-content-not-accept`). To hide the fact that the FortiOS Carrier unit is responding to a flood, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK” for the MM4 message forward response

```
config system replacemsg mm4 mm4-flood
    set rsp-status ok
    set rsp-text "Message Forwarded OK"
end
```

Viewing DLP archived messages

If *DLP Archive* is a selected message flood action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages, but you can configure the DLP archive setting to save only the first message that exceeds the threshold. This still provides a sample of the offending messages without requiring as much storage.

To select only the first message in a flood for DLP archiving - web-based manager

- 1 Go to *UTM Profiles > Carrier > MMS Profile*.
- 2 Select an existing MMS Profile

Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totalling the number of messages sent by each subscriber regardless of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The FortiOS Carrier unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a *Block* action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

Bypassing message flood protection based on user's carrier end points

You can use carrier end point filtering to exempt MMS sessions from message flood protection. Carrier end point filtering matches carrier end points in MMS sessions with carrier end point patterns.

If you add a carrier end point pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier end points bypass message flood protection. This allows legitimate bulk messages, such as system outage notifications, to be delivered without triggering message flood protection.

For more information on carrier end points, see the [FortiOS Handbook User Authentication](#) chapter.

Configuring message flood detection

To have the FortiOS Carrier unit check for message floods, you must first configure the flood threshold in an MMS profile, select the MMS profile in a protection profile, and select the protection profile in a security policy. All the traffic examined by the security policy will be checked for message floods according to the threshold values you set in the MMS profile.

Configure the MMS profile - web-based manager

- 1 Go to *Firewall Objects > MMS Profile*.
- 2 If you are editing an MMS profile, select the *Edit* icon of the MMS profile.
If you are create a new MMS profile, select *Create New* and enter a profile name.
- 3 Expand *MMS Bulk Email Filtering Detection*.
- 4 Expand *Message Flood*.
- 5 Expand *Flood Threshold 1*.
- 6 Select the *Enable* check box for MM1 messages, MM4 messages, or both.
- 7 In the *Message Flood Window* field, enter the length of time the FortiOS Carrier unit will keep track of the number of messages each subscriber sends.
If the FortiOS Carrier unit detects the quantity of messages specified in the *Message Flood Limit* sent during the number of minutes specified in the *Message Flood Window*, a message flood is in progress.
- 8 In the *Message Flood Limit* field, enter the number of messages required to trigger the flood.
- 9 In the *Message Flood Block Time* field, enter the length of time a user will be blocked from sending messages after causing the message flood.

10 Select the message flood actions the FortiOS Carrier unit will take when the message flood is detected.

11 Select OK.

Configure the protection profile - web-based manager

- 1** Go to *Firewall Objects > Protection Profile*.
- 2** If you are editing a protection profile, select the *Edit* icon of the protection profile.
If you are create a new protection profile, select *Create New* and enter a profile name.
- 3** Expand *MMS Profile*.
- 4** Select the MMS profile from the list.
- 5** Select OK.

Configure the security policy - web-based manager

- 1** Go to *Policy*.
- 2** Select the *Edit* icon of the security policy that controls the traffic in which you want to detect message floods.
- 3** Select the *Protection Profile* check box to enable the use of a protection profile.
- 4** Select the protection profile from the protection profile list.
- 5** Select OK.

Sending administrator alert notifications

When message floods are detected, the FortiOS Carrier unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the message flood action. Each message flood threshold can be configured separately.

This section includes:

- [Configuring how and when to send alert notifications](#)
- [Configuring who to send alert notifications to](#)

Configuring how and when to send alert notifications

You can configure different alert notifications for MM1 and MM4 message floods. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain

For the MM7 content interface you require:

- The message type
 - *submit.REQ* to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
 - *deliver.REQ* to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

For more information, see [“MMS notifications” on page 2367](#).

To configure administrator alert notifications - web-based manager

- 1 Go to *Firewall Objects > MMS Profile* and edit or add a new MMS protection profile.
- 2 Expand *MMS Bulk Email Filtering Detection*.
There are three message flood thresholds.
- 3 Expand the threshold that you want to configure alert notification for.
- 4 For *Message Flood Action*, select the *Alert Notification* check box. Alert notification options appear.
- 5 For the *Source MSISDN*, enter the MSISDN from which the alert notification message will be sent.
- 6 Select the Message Protocol the alert notification will use: *MM1*, *MM3*, *MM4*, or *MM7*.
- 7 Add the information required by FortiOS Carrier to send messages using the selected message protocol:
- 8 For *Notifications Per Second Limit*, enter the number of notifications to send per second.

Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.

- 9 If required, change *Window Start Time* and *Window Duration* configure when the FortiOS Carrier unit sends alert notifications.

By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.

For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.

You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.

- 10 For *Day of Week*, select the days of the week to send notifications.

For example, you may only want to send alert notifications on weekends for higher thresholds.

- 11 In the *Interval field*, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.

All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

Configuring who to send alert notifications to


In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the message flood threshold that triggers sending notifications to this MSISDN.

To configure the alert notification recipients - web-based manager

- 1 Go to *Firewall Objects > MMS Profile*.
- 2 Select the *Edit* icon of the MMS profile in which you want to configure the alert notification recipients.
- 3 Expand *MMS Bulk Email Filtering Detection*.
- 4 Expand *Recipient MSISDN*.
- 5 Select *Create New*.
- 6 In the *New MSISDN* window, enter the MSISDN to use for flood threshold alert notification.
- 7 Select the duplicate thresholds at which to send alert notifications to the MSISDN.



For the flood threshold to be able to send an alert notification to the MSISDN, the alert notification action must be enabled and configured within the flood threshold.




Duplicate message protection

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or other unwanted messages. Often, the same message will be sent by multiple subscribers. The message can be spam, viral marketing, or worm-generated messages. MMS duplicate prevention can help prevent this type of abuse by keeping track of the messages being sent.

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC. This can help prevent a potential flood that would otherwise become widespread between carriers.

Figure 248: MM1 and MM4 duplicate message protection



The FortiOS Carrier unit keeps track of the sent messages. If the same message appears more often than the threshold value you configure, then action is taken. Possible actions are logging the duplicates, blocking or intercepting duplicate messages, archiving the duplicate messages, and sending an alert to inform an administrator that duplicates are occurring.

With this highly configurable system, you can prevent the transmission of duplicate messages when there are more than you determine is acceptable.

For detailed configuration options, see [“Duplicate Message” on page 2333](#).

Using message fingerprints to identify duplicate messages

The FortiOS Carrier unit detects duplicates by keeping a record of all the messages travelling on the network and comparing new messages to those that have already been sent.

Rather than save the messages, the FortiOS carrier creates a checksum using the message body and subject. This serves as a fingerprint to identify the message. If another message with the same message body and subject appears, the fingerprint will also be the same and the FortiOS Carrier unit will recognize it as a duplicate.

By creating and saving message fingerprints instead of saving the messages, the FortiOS Carrier unit can save resources and time.

Messages from any sender to any recipient

Duplicate message detection will detect duplicate messages regardless of the sender or recipient. To do this, message fingerprints are generated using only the message body and subject. The sender, recipient, and other header information is not included.

If multiple messages appear with the same subject and message body, the FortiOS Carrier unit will recognize them as being the same.

Setting duplicate message thresholds

The FortiOS Carrier recognizes all duplicate messages, but it will take action when it detects a volume of duplicate messages that exceed the duplicate threshold you set. The threshold defines the maximum number of duplicate messages allowed, the period during which the messages are considered, and the length of time the duplicate message can not be sent by anyone.

For example, you may determine that once a duplicate message is sent more than 300 times in an hour, any attempt to send the same duplicate message will be blocked for 30 minutes.

If a particular duplicate message exceeds the duplicate message threshold and is blocked, any further attempts to send the same message will re-start the block period.

Using the example above, if the duplicate message count exceeds the duplicate threshold, any attempt to send a copy of the duplicate message will be blocked for 30 minutes. If a subscriber tries to send a copy of the message after waiting 15 minutes, the message will be blocked and the block period will be reset to 30 minutes. The block period must expire with no attempts to send a duplicate message. Only then will a subscriber be allowed to send the message. Non-duplicate messages will not reset the block period.

Duplicate message actions

When the FortiOS Carrier unit detects that a duplicate message has exceeded duplicate threshold, it can take any combination of the five actions you configure for the duplicate threshold.

Action		Description
Log		Add a log entry indicating that a duplicate message event has occurred. You must also enable logging for <i>MMS Scanning > Bulk Messages</i> in the <i>Logging</i> section of the MMS protection profile.
DLP Archive		
	All messages	Save all the messages that exceed the duplicate threshold in the DLP archive.
	First message only	Save the first message to exceed the duplicate threshold in the DLP archive. Subsequent messages that exceed the duplicate threshold will not be saved.
Intercept		Messages that exceed the duplicate threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message is also quarantined for later examination. If the quarantine of intercepted messages is disabled, the <i>Intercept</i> action has no effect.
Block		Messages that exceed the duplicate threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each blocked message is quarantined for later examination.
Alert Notification		If the duplicate threshold is exceeded, the FortiOS Carrier unit will send an MMS duplicate message notification message.

Notifying duplicate message senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver of duplicate messages. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around duplicate message protection. For this reason, no notification is set to the sender or receiver.

However, the FortiOS Carrier unit does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as duplicate messages. For information about how FortiOS Carrier responds when message flood detection blocks a message, see [“FortiOS Carrier and MMS duplicate messages and message floods” on page 2301](#).

Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (m-send.conf) to the sender (otherwise the sender's handset would keep retrying the message). The m-send.conf message is sent only when the MM1 duplicate message action is set to Block. For other duplicate message actions the message is actually delivered to the MMSC and the MMSC sends the m-send.conf message.

You can customize the m-send.conf message by editing the *MM1 send-conf duplicate message* MM1 replacement message (from the CLI the `mm1-send-conf-dupe` replacement message). You can customize the response status and message text for this message. The default response status is "Content not accepted". To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to "Success". The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to "Success" and changes the message text to "Message Sent OK":

```
config system replacemsg mm1 mm1-send-conf-dupe
  set rsp-status ok
  set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (m-retrieve.conf) to the sender (otherwise the sender's handset would keep retrying). The m-retrieve.conf message is sent only when the MM1 duplicate message action is set to Block. For other message flood actions the message is actually received by the receiver, so the MMSC sends the m-retrieve.conf message.

You can customize the m-retrieve.conf message by editing the *MM1 retrieve-conf duplicate message* MM1 replacement message (from the CLI the `mm1-retr-conf-dupe` replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mm1 mm1-retr-conf-dupe
  set subject "Message blocked"
  set message "Message temporarily blocked by carrier"
end
```

Forward responses for duplicate MM4 messages

When the FortiOS Carrier unit identifies an MM4 message as a duplicate message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4_forward.res message is sent only when the MM4 duplicate message action is set to Block and the MM4-forward.req message requested a response. For more information, see ["FortiOS Carrier and MMS duplicate messages and message floods" on page 2301](#).

You can customize the MM4_forward.res message by editing the *MM4 duplicate message* MM4 replacement message (from the CLI the mm4-dupe replacement message). You can customize the response status and message text for this message. The default response status is “Content not accepted” (err-content-not-accept). To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to “Success”. The default message text informs the sender that the message was blocked. You could change this to something more generic. For example, the following command sets the submission confirmation response status to “Success” and changes the message text to “Message Sent OK”:

```
config system replacemsg mm4 mm4-dupe
    set rsp-status ok
    set rsp-text "Message Forwarded OK"
end
```

Viewing DLP archived messages

If *DLP Archive* is a selected duplicate message action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages but you can configure the DLP archive setting to save only the first message that exceeds the threshold. See [“Viewing DLP archived messages” on page 2379](#).

Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totalling the number of messages sent by each subscriber regardless of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The FortiOS Carrier unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a *Block* action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.

Bypassing duplicate message detection based on user’s carrier end points

You can use carrier end point filtering to exempt MMS sessions from duplicate message detection. Carrier end point filtering matches carrier end points in MMS sessions with carrier end point patterns. If you add a carrier end point pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier end points bypass duplicate message detection. For more information about end points, see [FortiOS Handbook User Authentication chapter](#).

Configuring duplicate message detection

To have the FortiOS Carrier unit check for duplicate messages, configure the duplicate threshold in an MMS profile, select the MMS profile in a protection profile, and select the protection profile in a security policy.

- 1 Create an MMS profile and configure one or more of the duplicate thresholds as required.

- 2 Select the MMS profile in a protection profile.
- 3 Select the protection profile in a security policy.

All traffic matching the security policy will be checked for duplicate messages according to the settings in the MMS profile.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

The modular nature of the profiles allows you great flexibility in how you configure the scanning options. MMS profiles can be used in any number of protection profiles. Similarly, protection profiles can be used in any number of security policies.

In a complex configuration, there may be many security policies, each with a different protection profile and MMS profile. For a simpler network, you may have many security policies all using the same protection profile and MMS profile.

Sending administrator alert notifications

When duplicate messages are detected, the FortiOS Carrier unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the duplicate message action. Each duplicate message threshold can be configured separately.

This section includes:

- [Configuring how and when to send alert notifications](#)
- [Configuring who to send alert notifications to](#)

Configuring how and when to send alert notifications

You can configure different alert notifications for MM1 and MM4 duplicate messages. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually “/”)
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain

For the MM7 content interface you require:

- The message type
 - *submit.REQ* to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
 - *deliver.REQ* to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server

- The URL of the server (usually “/”)
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

To configure administrator alert notifications - web-based manager

- 1 Go to *Firewall Objects > MMS Profile* and edit or add a new MMS protection profile.
- 2 Expand *MMS Bulk Email Filtering Detection*.
There are three duplicate message thresholds.
- 3 Expand the threshold that you want to configure alert notification for.
- 4 For *Duplicate Message Action*, select the *Alert Notification* check box. Alert notification options appear.
- 5 For the *Source MSISDN*, enter the MSISDN from which the alert notification message will be sent.
- 6 Select the Message Protocol the alert notification will use: *MM1*, *MM3*, *MM4*, or *MM7*.
- 7 Add the information required by FortiOS Carrier to send messages using the selected message protocol:
- 8 For *Notifications Per Second Limit*, enter the number of notifications to send per second.
Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.
- 9 If required, change *Window Start Time* and *Window Duration* configure when the FortiOS Carrier unit sends alert notifications.
By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.
For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.
You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.
- 10 For *Day of Week*, select the days of the week to send notifications.
For example, you may only want to send alert notifications on weekends for higher thresholds.
- 11 In the *Interval field*, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.
All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

Configuring who to send alert notifications to

In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the duplicate threshold that triggers sending notifications to this MSISDN.

To configure the alert notification recipients - web-based manager

- 1 Go to *Firewall Objects > MMS Profile*.
- 2 Select the *Edit* icon of the MMS profile in which you want to configure the alert notification recipients.
- 3 Expand *MMS Bulk Email Filtering Detection*.
- 4 Expand *Recipient MSISDN*.
- 5 Select *Create New*.
- 6 In the *New MSISDN* window, enter the MSISDN to use for duplicate threshold alert notification.
- 7 Select the duplicate thresholds at which to send alert notifications to the MSISDN.



For the duplicate threshold to be able to send an alert notification to the MSISDN, the duplicate message threshold alert notification action must be enabled and configured.



MMS Replacement messages

Go to *System > Config > Replacement Message* to change replacement messages and customize notifications that the FortiOS Carrier unit adds to MMS content streams and for administrator alert notifications.

The replacement messages configured here are the default replacement message group selected in a protection profile. To add a replacement message to a protection profile go to *Firewall Objects > Protection Profile*, add or edit a protection profile and in the and under *Replacement Messages* select a replacement message group. The *default* replacement message group is selected by default.



Disclaimer replacement messages provided by Fortinet are examples only.

This section includes:

- [Changing replacement messages](#)
- [Multimedia content for MMS replacement messages](#)
- [MMS replacement message types](#)
- [Replacement message tags](#)
- [Replacement message groups](#)

Changing replacement messages

To change a replacement message list go to *System > Config > Replacement Message*. Use the expand arrows to view the replacement message that you want to change. You can change the content of the replacement message by editing the text and HTML codes and by working with replacement message tags.

Replacement messages can be text or HTML messages. You can add HTML code to HTML messages. Allowed Formats shows you which format to use in the replacement message. There is a limit of 8192 characters for each replacement message. The following fields and options are available when editing a replacement message. Different replacement messages have different sets of fields and options.

Message Setup	The name of the replacement message.
Allowed Formats	<p>The type of content that can be included in the replacement message. Allowed formats can be either Text or HTML.</p> <p>You can include replacement message tags in text and HTML messages.</p> <p>Do not use HTML code in Text messages — it will be incorrectly displayed.</p>

Size	<p>The number of characters allowed in the replacement message. A typical size is 8192 characters. This is the combined total size of the messages text, any SMIL content, and image.</p> <p>Each part of the MMS replacement message is limited to smaller sizes — 1023 characters for message text, 1023 characters for SMIL content, and 6000 bytes for an image if included.</p>
Response Status	<p>Select a response status for the replacement message. Many options are available including <i>Content not accepted</i>, <i>Success</i>, and <i>Unspecified error</i>.</p>
Priority	<p>Set the priority used by the protocol for sending the message. This priority is not used by FortiOS Carrier — it is added to the message. Options include <i>Not Specified</i>, <i>Low</i>, <i>Normal</i>, and <i>High</i>.</p>
Class	<p>Select the classification used by the protocol for the message. The classification is not used by FortiOS Carrier but is added to the message. Select <i>Not Specified</i>, <i>Personal</i>, <i>Advertisement</i>, <i>Information</i>, or <i>Automatic</i>.</p> <p>Note that not all MMS replacement messages include this field.</p>
Sender Visibility	<p>Select whether to show or hide the message sender. You can also select <i>Not Specified</i>.</p>
Use Sender MSISDN	<p>Select to include the sender's MSISDN in the replacement message.</p> <p>If the <i>From</i> field is used, the <i>Use Sender MSISDN</i> field is disabled.</p>
From	<p>Enter the name to appear as the sender of the replacement message.</p> <p>You cannot include replacement message tags in the <i>From</i> field.</p> <p>If the <i>Use Sender MSISDN</i> field is selected, the <i>From</i> field is disabled.</p>
Subject	<p>Enter or edit the subject for the replacement message. You cannot include replacement message tags in the subject field.</p>
Character Set	<p>Select the character set to use for the replacement message. You can select UTF-8 or US ASCII.</p>
Add SMIL Part	<p>Select to include Synchronized Multimedia Integration Language (SMIL) code in the message. Enter SMIL code into the SMIL Contents part of the replacement message.</p> <p>The <i>Image</i> and <i>SMIL Contents</i> fields are disabled unless <i>Add SMIL Part</i> is selected.</p>
Image	<p>Select a replacement message image to include in the replacement message. Use the <code>%%IMAGE_CID%%</code> tag to include the image in the SMIL contents.</p> <p>Any image you select must be 6000 bytes or less in size.</p>

SMIL Contents	Enter SMIL instructions for the replacement message to define multimedia content presentation. The size limit on the SMIL content is limited to 1023 characters.
Message Text	The body or editable text of the replacement message. The message text can include plain text, HTML instructions (if HTML is the allowed format) and replacement message tags. The size limit on this message is 1023 characters.

Multimedia content for MMS replacement messages

One of the main differences MMS replacement messages have is the addition of the Synchronized Multimedia Integration Language (SMIL) message portion. SMIL is a markup language, based on XML, that controls the presentation of media items such as text, images, audio, video, or even links to other SMIL presentations. This allows you to create a multimedia replacement message that includes your company colors and logo, an animated company logo, or other more advanced personations as appropriate. For more information on SMIL, see the W3C website <http://www.w3.org/AudioVideo/>.

The most important limit to be aware of is the size limit. The SMIL code portion can only be 1023 characters long which is the same maximum size for the regular message portion of the replacement message. If you include an image, its maximum size is 6000 bytes and valid formats are GIF, JPEG, PNG, or TIFF.

The following procedure will create an entry for a PNG image file called `test_image1`. It will then be included in the `MM1 send-req virus message` replacement message.

To upload an image for a SMIL message - web-based manager

- 1 Go to *System > Config > Replacement Message Image*.
- 2 Select *Create New*.
- 3 Enter `test_image1` for *Name*.
- 4 For *Content Type* select PNG.
- 5 Browse to the file location on your local computer using *Browse* next to *Upload*.
- 6 Select OK.

At this point, you will see your image displayed with a tag of `test_image1`.



You can not edit a replacement message image. To change an entry, you must delete it, and create a new entry.

To use an image in an MMS replacement message - web-based manager

- 1 Go to *System > Config > Replacement Messages*.
- 2 Expand *MM1*.
- 3 Edit the `MM1 send-req virus message` replacement message.
- 4 Enable *Add SMIL Part*.
- 5 For *Image*, select `test_image1` from the drop down list.
- 6 Enter SMIL code to display the image as required.
- 7 Enter message text including any replacement message tags required.
- 8 Enter rest of information as required.

- 9 Select OK.

MMS replacement message types

There are three types of replacement message used with MMS. The three types are:

Table 127: MMS replacement messages

m-send-conf	These messages are sent in response to an m-send-req message initiated by the mobile client. This message can only contain a response code and a plain-text response message to the client.
m-retrieve-conf	These messages are sent in response to a GET request resulting in the return of an MMS message. A new message is built using the options specified in these replacement messages, including a WML format message, and sent back to the user to inform them of what has occurred.
m-send-req	These messages are sent to notify the user of how many messages have been sent from their phone that violate the message content rules. A message is built from the options specified in these replacement messages, including a WML format message, and sent to the MMSC for delivery to the client.

Replacement message tags

Replacement messages can include replacement message tags. When users receive the replacement message, the replacement message tag is replaced with content relevant to the message. [Table 128](#) lists the replacement message tags that you can add.

Table 128: Replacement message tags

Tag	Description
%%AUTH_LOGOUT%%	The URL that will immediately delete the current policy and close the session. Used on the auth-keepalive page.
%%AUTH_REDIR_URL%%	The auth-keepalive page can prompt the user to open a new window which links to this tag.
%%CATEGORY%%	The name of the content category of the web site.
%%DEST_IP%%	The IP address of the request destination from which a virus was received. For email this is the IP address of the email server that sent the email containing the virus. For HTTP this is the IP address of web page that sent the virus.
%%DURATION%%	The amount of time in the reporting period. This is user defined in the protection profile.
%%EMAIL_FROM%%	The email address of the sender of the message from which the file was removed.
%%EMAIL_TO%%	The email address of the intended receiver of the message from which the file was removed.
%%FAILED_MESSAGE%%	The failed to login message displayed on the auth-login-failed page.

Table 128: Replacement message tags (Continued)

Tag	Description
%%FILE%%	The name of a file that has been removed from a content stream. This could be a file that contained a virus or was blocked by antivirus file blocking. %%FILE%% can be used in virus and file block messages.
%%FORTIGUARD_WF%	The FortiGuard - Web Filtering logo.
%%FORTINET%%	The Fortinet logo.
%%IMAGE_CID%%	The reference name of an image you have uploaded to the FortiOS Carrier unit. Use this to display the image in the message. The message with the image is generated as a MIME multipart message For example if you upload a file called <code>example.jpg</code> , and call it <code>example_logo</code> , then %%IMAGE_CID%% would resolve to <code>example_logo</code> .
%%LINK%%	The link to the FortiClient Host Security installs download for the Endpoint Control feature.
%%HTTP_ERR_CODE%	The HTTP error code. "404" for example.
%%HTTP_ERR_DESC%	The HTTP error description.
%%KEEPALIVEURL%%	<code>auth-keepalive-page</code> automatically connects to this URL every %%TIMEOUT%% seconds to renew the connection policy.
%%MMS_SENDER%%	Senders MSISDN from message header.
%%MMS_RECIPIENT%	Recipients MSISDN from message header.
%%MMS_SUBJECT%%	MMS Subject line to help with message identity.
%%MMS_HASH_CHECKSUM%%	Value derived from hash calculation - will only be shown on duplicate message alerts.
%%MMS_THRESH%%	Mass MMS alert threshold that triggered this alert.
%%NIDSEVENT%%	The IPS attack message. %%NIDSEVENT%% is added to alert email intrusion messages.
%%NUM_MSG%%	The number of times the device tried to send the message with banned content within the reporting period.
%%OVERRIDE%%	The link to the FortiGuard Web Filtering override form. This is visible only if the user belongs to a group that is permitted to create FortiGuard web filtering overrides.
%%OVRD_FORM%%	The FortiGuard web filter block override form. This tag must be present in the FortiGuard Web Filtering override form. It is not to be used in other replacement messages.
%%PROTOCOL%%	The protocol (http, ftp, pop3, imap, or smtp) in which a virus was detected. %%PROTOCOL%% is added to alert email virus messages.

Table 128: Replacement message tags (Continued)

Tag	Description
%%QUARFILENAME%%	The name of a file that has been removed from a content stream and added to the quarantine. This could be a file that contained a virus or was blocked by antivirus file blocking. %%QUARFILENAME%% can be used in virus and file block messages. Quarantining is only available on FortiOS Carrier units with a local disk.
%%QUOTA_INFO%%	Display information about the traffic shaping quota setting that is blocking the user. Used in traffic quota control replacement messages.
%%QUESTION%%	Authentication challenge question on auth-challenge page. Prompt to enter username and password on auth-login page.
%%SERVICE%%	The name of the web filtering service.
%%SOURCE_IP%%	The IP address of the request originator who would have received the blocked file. For email this is the IP address of the user's computer that attempted to download the message from which the file was removed.
%%TIMEOUT%%	Configured number of seconds between authentication keepalive connections. Used on the auth-keepalive page.
%%URL%%	The URL of a web page. This can be a web page that is blocked by web filter content or URL blocking. %%URL%% can also be used in http virus and file block messages to be the URL of the web page from which a user attempted to download a file that is blocked.
%%VIRUS%%	The name of a virus that was found in a file by the antivirus system. %%VIRUS%% can be used in virus messages

Replacement message groups

You can add up to five replacement message groups that can be applied to specific protection profiles allowing the customizing of messages for specific groups of users.

For example if your network has residential, corporate, and administrator users each group could have their own set of customized replacement messages with different information, graphics, and design. Another example could be if you provide services to five different companies, you could customize the replacement message groups for each company with their logo, colors, and so on.

You configure the default replacement message group from *System > Config > Replacement Message*. This replacement message group is the default replacement message group selected in a protection profile. All new replacement message groups that you add inherit their configuration from the default group.



Modifying messages in the default group automatically changes any messages that are unmodified in the other groups.

If you enable virtual domains (VDOMs) on the FortiOS Carrier unit, replacement message groups are configured separately for each virtual domain. Each virtual domain has its own default replacement message group, configured from *System > Config > Replacement Message*. When you modify a message in a replacement message group, a Reset icon appears beside the message in the group. You can select this Reset icon to reset the message in the replacement message group to the default version.

All MM1, MM4, MM7 notification messages (and MM1 retrieve-conf messages) can contain an SMIL layer and all MM4 notification messages can contain an HTML layer in the message. These layers can be used to brand messages by using logos uploaded to the unit via the *Manage Images* link found on the replacement message group configuration page. See [“Multimedia content for MMS replacement messages” on page 2395](#).

Replacement message group example

In this example, the message group is for a customer company called Example.com. Your company is called MyCarrier. Their group will be named `example_group`. Their logo is in a file called `example_logo.jpg`. Their employees do not want excessive information in the messages, so three replacement messages will be changed (mm1 send-req, send-conf, and retrieve-conf virus) to just the barebones information as part of this example.

To upload the logo image - web-based manager

- 1 Go to *System > Config > Replacement Message Group > Manage Images*.
- 2 Select *Create New*.
- 3 Enter `example_logo` for *Name*, and select *JPEG* for *Content Type*.
- 4 Browse to the file location of the file `example_logo.jpg` on your computer.
- 5 Select *OK*.
- 6 Select *Return* to return to the *Replacement Message Group* list.

To create the message group - web-based manager

- 1 Go to *System > Config > Replacement Message Group*.
- 2 Select *Create New*.
- 3 Enter `example_group` for *Name*.
- 4 Select *OK*.
- 5 Select `example_group`, and select *Edit*.
- 6 Expand *MM1*.
- 7 Select *MM1 send-req virus message*.
- 8 Enter the following information.


Priority	Normal
Class	Information
Sender Visibility	Show
From	%%MSISDN%% MyCarrier
Subject	Virus infected message(s) detected
Add SMIL Part	Enable
Image	example_logo

9 Enter the following SMIL code:

```
<smil><head><meta name="author" content="MyCarrier" /></head>
<body></body>
</smil>
```

10 Enter the following replacement message code:

```
This device has sent %%NUM_MSG%% virus infected messages in the
last %%DURATION%% hours. Contact MyCarrier customer support
for farther instructions.
```



Configuring GTP on FortiOS Carrier

Configuring GTP support on FortiOS Carrier involves configuring a number of areas of features. Some features require longer explanations, and have their own chapters. The other features are addressed here.

This section includes:

- [GTP support on the FortiOS Carrier unit](#)
- [Configuring General Settings on the FortiOS Carrier unit](#)
- [Configuring Encapsulated Filtering in FortiOS Carrier](#)
- [Configuring the Protocol Anomaly feature in FortiOS Carrier](#)
- [Configuring Anti-overbilling in FortiOS Carrier](#)
- [Logging events on the FortiOS Carrier unit](#)

GTP support on the FortiOS Carrier unit

The FortiCarrier unit needs to have access to all traffic entering and exiting the carrier network for scanning, filtering, and logging purposes. This promotes one of two configurations — hub and spoke, or bookend.

A hub and spoke configuration with the FortiOS Carrier unit at the hub and the other GPRS devices on the spokes is possible for smaller networks where a lower bandwidth allows you to divide one unit into multiple virtual domains to fill multiple roles on the carrier network. It can be difficult with a single FortiOS Carrier as the hub to ensure all possible entry points to the carrier network are properly protected from potential attacks such as [“Relayed network attacks” on page 2404](#).

A bookend configuration uses two FortiOS Carrier units to protect the carrier network between them with high bandwidth traffic. One unit handles traffic from mobile stations, SGSNs, and foreign carriers. The other handles GGSN and data network traffic. Together they ensure the network is secure.

The FortiOS Carrier unit can access all traffic on the network. It can also verify traffic between devices, and verify that the proper GPRS interface is being used. For example there is no reason for a Gn interface to be used to communicate with a mobile station — the mobile station will not know what to do with the data — so that traffic is blocked.



When you are configuring your FortiOS Carrier unit's GTP profile, you must first configure the APN. It is critical to GTP communications — no traffic will flow without the APN.

The FortiOS Carrier unit does more than just forward and route GTP packets over the network. It also performs:

- [Packet sanity checking](#)
- [GTP stateful inspection](#)
- [Protocol anomaly detection and prevention](#)
- [HA](#)

- [Virtual domain support](#)

Packet sanity checking

The FortiOS Carrier firewall checks the following items to determine if a packet conforms to the UDP and GTP standards:

- GTP release version number — must be 0, 1, or 2
- Settings of predefined bits
- Protocol type
- UDP packet length

If the packet in question does not confirm to the standards, the FortiOS Carrier firewall drops the packet, so that the malformed or forged traffic will not be processed.

GTP stateful inspection

Apart from the static inspection (checking the packet header), the FortiOS Carrier firewall performs stateful inspection.

Stateful inspection provides enhanced security by keeping track of communications sessions and packets over a period of time. Both incoming and outgoing packets are examined. Outgoing packets that request specific types of incoming packets are tracked; only those incoming packets constituting a proper response are allowed through the firewall.

The FortiOS Carrier firewall can also index the GTP tunnels to keep track of them.

Using the enhanced Carrier traffic policy, the FortiOS Carrier firewall can block unwanted encapsulated traffic in GTP tunnels, such as infrastructure attacks. Infrastructure attacks involve attempts by an attacker to connect to restricted machines, such as GSN devices, network management systems, or mobile stations. If these attempts to connect are detected, they are to be flagged immediately by the firewall.

Protocol anomaly detection and prevention

The FortiOS Carrier firewall detects and optionally drops protocol anomalies according to GTP standards and specific tunnel states. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of protocol specifications. These packets are not seen on a production network. Protocol anomaly attacks exploit poor programming practices when decoding packets, and are typically used to maliciously impair system performance or elevate privileges.

FortiOS Carrier also detects IP address spoofing inside GTP data channel.

See [“Configuring the Protocol Anomaly feature in FortiOS Carrier”](#) on page 2405.

HA

FortiOS Carrier active-passive HA provides failover protection for the GTP tunnels. This means that an active-passive cluster can provide FortiOS Carrier firewall services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiOS Carrier firewall. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially for mission-critical environments.

FortiOS HA synchs TCP sessions by default, but UDP sessions are not synchronized by default. However synchronizing a session is only part of the solution if the goal is to continue GTP processing on a synchronized session after a HA switch. For that to be successful we also need to synch the GTP tunnel state. So, once the master completes tunnel setup then the GTP tunnel is synchronized to the slave.

For more information on HA in FortiOS, see [“High Availability” on page 1983](#).

Virtual domain support

FortiOS Carrier is suited to both large and smaller carriers. A single FortiOS Carrier unit can serve either one large carrier, or several smaller ones through virtual domains. As with any FortiGate unit, FortiOS Carrier units have the ability to split their resources into multiple virtual units. This allows smaller carriers to use just the resources that they need without wasting the extra. For more information on virtual domains in FortiOS, see [“Virtual Domains” on page 1875](#).

Configuring General Settings on the FortiOS Carrier unit

To configure the GTP General Settings, go to *UTM Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *General Settings* to configure settings. See [“General settings options” on page 2338](#).

Configuring Encapsulated Filtering in FortiOS Carrier

Encapsulated traffic on the GPRS network can come in a number of forms as it includes traffic that is “wrapped up” in another protocol. This detail is important for firewalls because it requires “unwrapping” to properly scan the data inside. If encapsulated packets are treated as regular packets, that inside layer will never be scanned and may allow malicious data into your network.

On FortiOS Carrier units, GTP related encapsulated filtering falls under encapsulated IP traffic filtering, and encapsulated non-IP end user address filtering.

Configuring Encapsulated IP Traffic Filtering

Generally there are a very limited number of IP addresses that are allowed to encapsulate GPRS traffic. For example GTP tunnels are a valid type of encapsulation when used properly. This is the GTP tunnel which uses the Gp or Gn interfaces between SGSNs and GGSNs. However, a GTP tunnel within a GTP tunnel is not accessible — FortiOS Carrier will either block or forward the traffic, but is not able to open it for inspection.

The ability to filter GTP sessions is based on information contained in the data stream and provides operators with a powerful mechanism to control data flows within their infrastructure. You can also configure IP filtering rules to filter encapsulated IP traffic from Mobile Stations.

To configure the Encapsulated IP Traffic Filtering, go to *UTM Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *Encapsulated IP Traffic Filtering* to configure settings. See [“Encapsulated IP traffic filtering options” on page 2346](#).

When to use encapsulated IP traffic filtering (best practices)

The following are the typical cases that need encapsulated IP traffic filtering:

Mobile station IP pools

In a well-designed network, best practices dictate that the mobile station address pool is to be completely separate from the GPRS network infrastructure range of addresses. Encapsulated IP packets originating from a mobile station will not contain source or destination addresses that fall within the address range of GPRS infrastructures. In addition, traffic originating from the users handset will not have destination/source IP addresses that fall within any Network Management System (NMS) or Charging Gateway (CG) networks.

Communication between mobile stations

Mobile stations on the same GPRS network are not be able to communicate with other mobile stations. Best practices dictate that packets containing both source and destination addresses within the mobile station's range of addresses are to be dropped.

Direct mobile device or internet attacks

It may be possible for attackers to wrap attack traffic in GTP protocols and submit the resulting GTP traffic directly to a GPRS network element from their mobile stations or a node on the Internet. It is possible that the receiving SGSN or GGSN would then strip off the GTP header and attempt to route the underlying attack. This underlying attack could have any destination address and would probably have a source address spoofed as if it were valid from that PLMN.



You cannot add an IE removal policy when you are creating a new profile.

Relayed network attacks

Depending on the destination the attack could be directly routed, such as to another node of the PLMN, or rewrapped in GTP for transmission to any destination on the Internet outside the PLMN depending on the routing table of the GSN enlisted as the unwitting relay.

The relayed attack could have any source or destination addresses and could be any of numerous IP network attacks, such as an attack to hijack a PDP context, or a direct attack against a management interface of a GSN or other device within the PLMN. Best practices dictate that any IP traffic originating on the Internet or from an MS with a destination address within the PLMN is to be filtered.

Configuring Encapsulated Non-IP End User Address Filtering

Much of the traffic on the GPRS network is in the form of IP traffic. However some parts of the network do not used IP based addressing, so the FortiOS Carrier unit is unable to perform Encapsulated IP Traffic Filtering.

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications only list PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC 1700. The PDP types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

To configure the Encapsulated Non-IP End User Address Filtering, go to *UTM Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand *Encapsulated Non-IP End User Address Filtering* to configure settings. See [“Encapsulated non-IP end user traffic filtering options” on page 2347](#).

Configuring the Protocol Anomaly feature in FortiOS Carrier

When anomalies do happen, it is possible for the anomaly to interrupt network traffic or consume network resources — if precautions were not taken. Anomalies can be generated by accident or maliciously, but both methods can have the same results — degrading the performance of the carrier network, or worse.

To configure GTP protocol anomalies, go to *UTM Profiles > Carrier > GTP Profile*, and edit a GTP profile. Expand the *Protocol Anomaly* option. See [“Protocol anomaly prevention options” on page 2348](#).

The following are some examples:

- The GTP header specifies the length of the packet excluding the mandatory GTP header. In GTP version 0 (GSM 09.60), the mandatory GTP header size is 20 bytes, whereas GTP version 1 (GSM 29.060) specifies that the minimum length of the GTP header is 8 bytes. The GTP packet is composed of the header, followed by Information Elements typically presented in a Type-Length-Value format. It is possible for an attacker to create a GTP packet with a GTP header field length that is incompatible with the length of the necessary information elements.
- The same concepts are true for GTP version 2 headers even though there are different fields in them.
- It is similarly possible for an attacker to create a packet with an invalid IE length. Invalid lengths may cause protocol stacks to allocate incorrect amounts of memory, and thereby cause crashes or buffer overflows.

By default the FortiOS Carrier firewall detects these problems as well as other protocol anomalies and drops the packets. However, you can change the policy to allow them.

All protocol anomaly options are set to *Deny* by default.

Configuring Anti-overbilling in FortiOS Carrier

This section includes:

- [Overbilling in GPRS networks](#)
- [Anti-overbilling with FortiOS Carrier](#)

Overbilling in GPRS networks

GPRS overbilling attacks can be prevented with a properly configured FortiOS Carrier unit.

Overbilling can occur when a subscriber returns his IP address to the IP pool. Before the billing server closes it, the subscriber's session is still open and vulnerable. If an attacker takes control of the subscriber's IP address, he can send or receive data and the subscriber will be billed for the traffic.

Overbilling can also occur when an available IP address is reassigned to a new mobile station (MS). Subsequent traffic by the previous MS may be forwarded to the new MS. The new MS would then be billed for traffic it did not initiate.

Anti-overbilling with FortiOS Carrier

The FortiOS Carrier unit can be configured to assist with anti-overbilling measures. These measures ensure that the customer is only billed for connection time and data transfer that they actually use.

Anti-overbilling on the FortiOS Carrier unit involves:

- the FortiOS Carrier unit configures the overbilling settings in the GTP profile to notify the Gi firewall when a GTP tunnel is deleted
- the unit clears the sessions when the Gi firewall receives a notification from the Gn/Gp firewall about a GTP tunnel being deleted. This way, the Gi firewall prevents overbilling by blocking traffic initiated by other users.

The three locations to configure anti-overbilling options include:

- *System > Network > Interface > Gi Gatekeeper* — enable to monitor Gi anti-overbilling traffic on this interface
- *System > Admin > Settings > Gi Gatekeeper Settings* — set the context ID and port that anti-overbilling will take place on.
- *UTM Profiles > Carrier > GTP Profile > Anti-Overbilling* — the IP address, port, interface and context ID to use for anti-overbilling measures.

For detailed options, see [“Anti-overbilling options” on page 2349](#).

Logging events on the FortiOS Carrier unit

Logging on the FortiOS Carrier unit is just like logging on any other FortiOS unit. The only difference with FortiOS Carrier is that there are a few additional events that you can log beyond the regular ones. These additional events are covered here. For more information on other logging issues, see the [Logging and Reporting Guide](#) and [FortiOS CLI Reference](#).

To enable FortiOS Carrier logging, go to *Log&Report > Event Log*, and ensure *GTP service event* is enabled. Once this option is selected, the logging options under *UTM Profiles > Carrier > GTP Profile* will be active.

To change FortiOS Carrier specific logging event settings, go to *UTM Profiles > Carrier > GTP Profile* and edit a GTP profile. Expand the *Log* section to change the settings. For detailed options, see [“Log options” on page 2349](#).

The following information is contained in each log entry:

Timestamp	The time and date when the log entry was recorded
Source IP address	The sender's IP address.
Destination IP address	The receiver's IP address. The sender-receiver pair includes a mobile phone on the GPRS local network, and a device on a network external to the GPRS network, such as the Internet.

Tunnel Identifier (TID) Tunnel Endpoint Identifier (TEID)	An identifier for the start and end points of a GTP tunnel. This information uniquely defines all tunnels. It is important for billing information based on the length of time the tunnel was active and how much data passed over the tunnel.
Message type	For available message types, see “Common message types on carrier networks” on page 2409 .
Packet status	What action was performed on the packet. This field matches the logging options while you are configuring GTP logging. See “Anti-overbilling with FortiOS Carrier” on page 2406 . The status can be one of forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited
Virtual domain ID or name	A FortiOS Carrier unit can be divided into multiple virtual units, each being a complete and self-contained virtual FortiCarrier unit. This field indicates which virtual domain (VDOM) was responsible for the log entry. If VDOMs are not enabled on your unit, this field will be <code>root</code> .
Reason to be denied if applicable	If the packet that generated this log entry was denied or blocked, this field will include what part of FortiOS denied or blocked that packet. Such as firewall, antivirus, webfilter, or spamfilter.

An example of the above log message format is for a Tunnel deleted log entry. When a tunnel is deleted, the log entry contains the following information:

- Timestamp
- Interface name (if applicable)
- SGSN IP address (source IP)
- GGSN IP address (destination IP)
- Tunnel ID
- Tunnel duration time in seconds
- Number of messages sent to the SGSN
- Number of messages sent to the GGSN



GTP message type filtering

FortiOS Carrier supports message filtering in GTP by the type of message.

This section includes:

- [Common message types on carrier networks](#)
- [Configuring message type filtering in FortiOS Carrier](#)

Common message types on carrier networks

Carrier networks include many types of messages — some concern the network itself, others are content moving across the network, and still others deal with handshaking, billing, or other administration based issues.

GTP contains two major parts GTP for the control plane (GTP-C) and GTP for user data tunnelling (GTP-U). Outside of those areas there are only unknown message types.

GTP-C messages

GTP-C contains the networking layer messages. These address routing, versioning, and other similar low level issues.

When a subscriber requests a Packet Data Protocol (PDP) context, the SGSN will send a create PDP context request GTP-C message to the GGSN giving details of the subscriber's request. The GGSN will then respond with a create PDP context response GTP-C message which will either give details of the PDP context actually activated or will indicate a failure and give a reason for that failure. This is a UDP message on port 212.

GTP-C message types include Path Management Messages, Location Management Messages, and Mobility Management Messages.

Path Management Messages

Path management is used by one GSN to detect if another GSN is alive, or if it has restarted after a failure.

The path management procedure checks if a given GSN is alive or has been restarted after a failure. In case of SGSN restart, all MM and PDP contexts are deleted in the SGSN, since the associated data is stored in a volatile memory. In the case of GGSN restart, all PDP contexts are deleted in the GGSN.

Tunnel Management Messages

The tunnel management procedures are used to create, update, and delete GTP tunnels in order to route IP PDUs between an MS and an external PDN via the GSNs.

The PDP context contains the subscriber's session information when the subscriber has an active session. When a mobile wants to use GPRS, it must first attach and then activate a PDP context. This allocates a PDP context data structure in the SGSN that the subscriber is currently visiting and the GGSN serving the subscriber's access point.

Tunnel management procedures are defined to create, update, and delete tunnels within the GPRS backbone network. A GTP tunnel is used to deliver packets between an SGSN and a GGSN. A GTP tunnel is identified in each GSN node by a TEID, an IP address, and a UDP port number.

Location Management Messages

The location-management procedure is performed during the network-requested PDP context activation procedure if the GGSN does not have an SS7 MAP interface (i.e., Gc interface). It is used to transfer location messages between the GGSN and a GTP-MAP protocol-converting GSN in the GPRS backbone network.

Location management subprocedures are used between a GGSN that does not support an SS7 MAP interface (i.e., Gc interface) and a GTP-MAP protocol-converting GSN. This GSN supports both Gn and Gc interfaces and is able to perform a protocol converting between GTP and MAP.

Mobility Management Messages

The MM procedures are used by a new SGSN in order to retrieve the IMSI and the authentication information or MM and PDP context information in an old SGSN. They are performed during the GPRS attach and the inter-SGSN routing update procedures.

The MM procedures are used between SGSNs at the GPRS-attach and inter-SGSN routing update procedures. An identity procedure has been defined to retrieve the IMSI and the authentication information in an old SGSN. This procedure may be performed at the GPRS attach. A recovery procedure enables information related to MM and PDP contexts in an old SGSN to be retrieved. This procedure is started by a new SGSN during an inter-SGSN RA update procedure.

GTP-U messages

GTP-U is focused on user related issues including tunneling, and billing. GTP-U message types include MBMS messages, and GTP-U and Charging Management Messages

MBMS messages

Multimedia Broadcast and Multicast Services (MBMS) have recently begun to be offered over GSM and UMTS networks on UTRAN and GERAN radio access technologies. MBMS is mainly used for mobile TV, using up to four GSM timeslots for one MBMS connection. One MBMS packet flow is replicated by GGSN, SGSN and RNCs.

MBMS is split into the MBMS Bearer Service and the MBMS User Service. The MBMS User Service is basically the MBMS Service Layer and offers a Streaming- and a Download Delivery Method. The Streaming Delivery method can be used for continuous transmissions like Mobile TV services. The Download Method is intended for "Download and Play" services.

GTP-U and Charging Management Messages

SGSNs and GGSNs listen for GTP-U messages on UDP port 2152.

GTP' (GTP prime) is used for billing messages. It uses the common GTP messages (GTP Version Not Supported, Echo Request and Echo Response) and adds additional messages related to billing procedures.

Unknown Action messages

If the system doesn't know what type of message it is, it falls into this category. This is an important category of message because malformed messages may appear and need to be handled with security in mind.



Fortinet best practices dictate that you set *Unknown Action messages* to deny for security reasons.

Configuring message type filtering in FortiOS Carrier

GPRS Tunneling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) traffic within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over their network via tunneling.

In the CLI, there is a keyword for each type of GTP message for both message filtering, and for message rate limiting.



GTP message rate limiting is only accessible from the CLI using the command `configure firewall gtp`.

To configure GTP message type filtering - web-based manager

- 1 Go to *UTM Profiles > Carrier > GTP Profile*.
- 2 Select *Create New*.
- 3 Enter a name for this profile such as `msg_type_filtering`.
- 4 Select *Message Type Filtering* to expand it.
- 5 For each type of message in the list, select Allow or Deny. All messages are set to Allow by default.



Fortinet best practices dictate that the unknown message action should be set to Deny for security reasons as this will block malformed messages.

- 6 Optionally select and configure any other GTP features for this profile, such as logging.
- 7 Select OK to save the profile.
- 8 Apply the `msg_type_filtering` profile a security policy configured for GTP tunnel traffic.

To configure GTP message filtering and block Unknown Message Action messages- CLI

```
config firewall gtp
  edit msg_type_filtering
    config message-filter
      set unknown-message-action deny
    next
  end
end
```

Message Type Fields

Each of the following message types can be allowed or denied by your FortiOS Carrier unit depending on your carrier network and GTP traffic.

The message types include:

- [Unknown Message Action](#)
- [Path Management Messages](#)
- [Tunnel Management Messages](#)
- [Location Management Messages](#)
- [Mobility Management Messages](#)
- [MBMS messages](#)
- [GTP-U and Charging Management Messages](#)

Unknown Message Action

Set this message type to deny.

Many attempts to hack into a carrier network will result in this unknown message type and therefore it is denied for security reasons.

Path Management Messages

Message Type	Used by	Description
Echo Request/Response	GTP-C, GTP-U, GTP'	Echo Request is sent on a path to another GSN to determine if the other node is alive. Echo Response is the reply.
Version not Supported	GTP-C, GTP-U, GTP'	There are multiple versions of GTP. Both devices communicating must use the same version of GTP, or this message will be the response.
Support Extension Headers Notification		Extensions are optional parts that a device can choose to support or not. If a device includes these extensions, it must include headers for the extensions to sure ensure proper formatting.

Tunnel Management Messages

Message Type	Used by	Description
Create PDP Context Request/ Response	GTP-C	Sent from an SGSN to a GGSN node as part of a GPRS PDP Context Activation procedure or the Network-Requested PDP Context Activation procedure. A valid request initiates the creation of a tunnel.
Update PDP Context Request/ Response	GTP-C	Used when PDP Context information changes, such as when a mobile device changes location.
Delete PDP Context Request/ Response	GTP-C	Used to terminate a PDP Context, and confirm the context has been deleted.

Create AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS Anonymous Access PDP Context Activation. It is used to create a tunnel between a context in the SGSN and a context in the GGSN.
Delete AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS PDP Anonymous Access Context Deactivation procedure to deactivate an activated PDP Context. It contains Cause and Private Extension Information Elements
Error Indication	GTP-U	Sent to the GGSN when a tunnel PDU is received for the following conditions: <ul style="list-style-type: none"> — No PDP context exists — PDP context is inactive — No MM context exists — GGSN deletes its PDP context when the message is received.
PDU Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	When receiving a Tunneled PDU (T-PDU), the GGSN checks if a PDP context is established for the given PDP address. If no PDP context has been established, the GGSN may initiate the Network-requested PDP Context Activation procedure by sending a PDU Notification Request to the SGSN. Reject Request - Sent when the PDP context requested by the GGSN cannot be established.

Location Management Messages

Message Type	Used By	Description
Send Routing Information for GPRS Request/ Response	GTP-C	Sent by the GGSN to obtain location information for the MS. This message type contains the IMSI of the MS and Private Extension.
Failure Report Request/ Response	GTP-C	Sent by the GGSN to the HLR when a PDU reject message is received. The GGSN requests the HLR to set the flag and add the GGSN to the list of nodes to report to when activity from the subscriber that owns the PDP address is detected. The message contains the subscriber IMSI and Private Extension
Note MS GPRS Present Request/ Response	GTP-C	When the HLR receives a message from a mobile with MDFG set, it clears the MDFG and sends the Note MS Present message to all GGSN's in the subscriber's list. This message type contains subscriber IMSI, GSN Address and Private Extension

Mobility Management Messages


Message Type	Used By	Description
Identification Request/Response	GTP-C	Sent by the new SGSN to the old SGSN to request the IMSI for a MS when a GPRS Attach is done with a P-TMSI and the MS has changed SGSNs since the GPRS Detach was done.
SGSN context Request/ Response/ Acknowledge	GTP-C	Sent by the new SGSN to the old SGSN to request the MM and PDP Contexts for the MS.
Forward Relocation Request/ Response/ Complete/ Complete Acknowledge	GTP-C	Indicates mobile activation/deactivation within a Routing Area. This prevents paging of a mobile that is not active (visited VLR rejects calls from the HLR or applies Call Forwarding). Note that the mobile station does not maintain an attach/detach state. SRNS contexts contain for each concerned RAB the sequence numbers of the GTP-PDUs next to be transmitted in uplink and downlink directions.
Relocation Cancel Request/ Response	GTP-C	Send to cancel the relocation of a connection.
Forward SRNS Context/ Context Acknowledge	GTP-C	This procedure may be used to trigger the transfer of SRNS contexts from RNC to CN (PS domain) in case of inter system forward handover.
RAN Information Relay	GTP-C	Forward the Routing Area Network (RAN) information. A Routing Area (RA) is a subset of a GSM Location Area (LA). A RA is served by only one SGSN. Ensures that regular radio contact is maintained by the mobile

MBMS messages

Message Type	Used By	Description
MBMS Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	Notification of the radio access devices.
Create MBMS Context Request/ Response	GTP-C	Request to create an active MBMS context. The context will be pending until the response is received. Once active, the MBMS context allows the MS to receive data from a specific MBMS source
Update MBMS Context Request/ Response	GTP-C	
Delete MBMS Context Request/ Response	GTP-C	Request to deactivate the MBMS context. When the response is received, the MBMS context will be inactive.

GTP-U and Charging Management Messages

Message Type	Used By	Description
G-PDU	GTP-C, GTP-U	GPRS Packet data unit delivery message.
Node Alive Request/Response	GTP-C, GTP-U	Used to inform rest of network when a node starts service.
Redirection Request/Response	GTP-C, GTP-U	Used to divert the flow of CDRs from the CDFs to another CGF when the sender is being removed, or they are used when the CGF has lost its connection to a downstream system.
Data Record Transfer Request/Response	GTP-C, GTP-U	Used to reliably transport CDRs from the point of generation (SGSN/GGSN) to non-volatile storage in the CGF



GTP identity filtering

FortiOS Carrier supports a number of filtering methods based on subscriber identity such as APN filtering, IMSI filtering, and advanced filtering.

This section includes:

- [IMSI on carrier networks](#)
- [Other identity and location based information elements](#)
- [Configuring APN filtering in FortiOS Carrier](#)
- [Configuring IMSI filtering in FortiOS Carrier](#)
- [Configuring advanced filtering in FortiOS Carrier](#)


IMSI on carrier networks

The International Mobile Subscriber Identity (IMSI) number is central to identifying users on a carrier network. It is a unique number that is assigned to a cell phone or mobile device to identify it on the GSM or UTM network.

Typical the IMSI number is stored on the SIM card of the mobile device and is sent to the network as required.

An IMSI number is 15 digits long, and includes the Mobile Country Code (MCC), Mobile Network Code (MNC), and Mobile Station Identification Number (MSIN).

Figure 249: IMSI codes



The Home Network Identity (HNI) is made up of the MCC and MNC. The HNI is used to fully identify a user's home network. This is important because some large countries have more than one country code for a single carrier. For example a customer with a mobile carrier on the East Coast of the United States would have a different MCC than a customer on the West Coast with the same carrier because even though the MNC would be the same the MCC would be different — the United States uses MCCs 310 to 316 due to its size.

If an IMSI number is not from the local carrier's network, IMSI analysis is performed to resolve the number into a Global Title which is used to access the user's information remotely on their home carrier's network for things like billing and international roaming.

See [“Configuring IMSI filtering in FortiOS Carrier” on page 2421](#).

Other identity and location based information elements

IMSI focuses on the user, their location, and carrier network. There are other numbers used to identify different user related Information Elements (IE).

These identity and location based elements include:

- [Access Point Number \(APN\)](#)
- [Mobile Subscriber Integrated Services Digital Network \(MSISDN\)](#)
- [Radio Access Technology \(RAT\) type](#)
- [User Location Information \(ULI\)](#)
- [Routing Area Identifier \(RAI\)](#)
- [International Mobile Equipment Identity \(IMEI\)](#)

Access Point Number (APN)

The Access Point Number (APN) is used in GPRS networks to identify an IP packet data network that a user wants to communicate with. The Network Identifier describes the network and optionally the service on that network that the GGSN is connected to. The APN also includes the MCC and MCN, which together locate the network the GGSN belongs to. An example of an APN in the Barbados using Digicel as the carrier that is connecting to the Internet is `internet.mcc342.mnc750.gprs`.

When you are configuring your FortiOS Carrier unit's GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

The access point can then be used in a DNS query to a private DNS network. This process (called APN resolution) gives the IP address of the GGSN which serves the access point. At this point a PDP context can be activated. See [“Configuring APN filtering in FortiOS Carrier” on page 2420](#).

Mobile Subscriber Integrated Services Digital Network (MSISDN)

This is a 15-digit number that, along with the IMSI, uniquely identifies a mobile user. Normally this number includes a 2-digit country code, a 3-digit national destination code, and a 10-digit subscriber number or the phone number of the mobile device, and because of that may change over time if the user changes their phone number. The MSISDN number follows the ITU-T E.164 numbering plan.

Radio Access Technology (RAT) type

The RAT type represents the radio technology used by the mobile device. This can be useful in determining what services or content can be sent to a specific mobile device. FortiOS Carrier supports:

- **UMTS Terrestrial Radio Access Network (UTRAN)**, commonly referred to as 3G, routes many types of traffic including IP traffic. This is one of the faster types.
- **GSM EDGE Radio Access Network (GERAN)** is a key part of the GSM network which routes both phone calls and data.
- **Wireless LAN (WLAN)** is used but not as widely as the other types. It is possible for the mobile device to move from one WLAN to another such as from an internal WLAN to a commercial hot spot.
- **Generic Access Network (GAN)** can also be called unlicensed mobile access (UMA). It routes voice, data, and SIP over IP networks. GAN is commonly used for mobile devices that have a dual-mode and can hand-off between GSM and WLANs.

- **High Speed Packet Access (HSPA)** includes two other protocols High Speed Downlink and Uplink Packet Access protocols (HSDPA and HSUPA respectively). It improves on the older WCDMA protocols by better using the radio bandwidth between the mobile device and the radio tower. This results in an increased data transfer rate for the user.

RAT type is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 2422](#).

User Location Information (ULI)

Gives Cell Global Identity/Service Area Identity (CGI/SAI) of where the mobile station is currently located. The ULI and the RAI are commonly used together to identify the location of the mobile device.

ULI is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 2422](#).

Routing Area Identifier (RAI)

Routing Areas (RAs) divide the carrier network and each has its own identifier (RAI). When a mobile device moves from one routing area to another, the connection is handled by a different part of the network. There are normally multiple cells in a routing area. There is only one SSGN per routing area. The RAI and ULI are commonly used to determine a user's location.

RAI is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 2422](#).

International Mobile Equipment Identity (IMEI)

IMEI is a unique 15-digit number used to identify mobile devices on mobile networks. It is very much like the MAC address of a TCP/IP network card for a computer. It can be used to prevent network access by a stolen phone — the carrier knows the mobile phone's IMEI, and when it is reported stolen that IMEI is blocked from accessing the carrier network no matter if it has the same SIM card as before or not. It is important to note that the IMEI stays with the mobile phone or device where the other information is either location based or stored on the removable SIM card.

IMEI type is part of advanced filtering configuration. See [“Configuring advanced filtering in FortiOS Carrier” on page 2422](#).

When to use APN, IMSI, or advanced filtering

At first glance APN, IMSI, and advanced filtering have parts in common. For example two can filter on APN, and another two can filter on IMSI. The difficulty is knowing when to use which type of filtering.

Figure 250: Identity filtering comparison

Filtering type	Filter on the following data:	When to use this type of filtering
APN	APN	Filter based on GTP tunnel start or destination
IMSI	IMSI, MCC-MNC	Filter based on subscriber information
Advanced	PDP context, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI	When you want to filter based on: user phone number (MSISDN) what wireless technology the user employed to get on the network (RAT type) user location (ULI and RAI) handset ID, such as for stolen phones (IMEI)

APN filtering is very specific — the only identifying information that is used to filter is the APN itself. This will always be present in GTP tunnel traffic, so all GTP traffic can be filtered using this value. See [“Configuring APN filtering in FortiOS Carrier” on page 2420](#).

IMSI filtering can use a combination of the APN and MCC-MNC numbers. The MCC and MNC are part of the APN, however filtering on MCC-MNC separately allows you to filter based on country and carrier instead of just the destination of the GTP Tunnel. See [“Configuring IMSI filtering in FortiOS Carrier” on page 2421](#).

Advanced filtering can go into much deeper detail covering PDP contexts, MSISDN, IMEI, and more not to mention APN, and IMSI as well. If you can’t find the information in APN or IMSI that you need to filter on, then use Advanced filtering. See [“Configuring advanced filtering in FortiOS Carrier” on page 2422](#).

Configuring APN filtering in FortiOS Carrier

To configure APN filtering go to *UTM Profiles > Carrier > GTP Profile*. Select a profile or create a new one, and expand *APN filtering*.



When you are configuring your FortiOS Carrier unit’s GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

For more information on APN, see [“Access Point Number \(APN\)” on page 2418](#).

Enable APN Filter	Select to enable filtering based on APN value.
Default APN Action	Select either Allow or Deny for all APNs that are not found in the list. The default is Allow.
Value	Displays the APN value for this entry. Partial matches are allowed using wildcard. For example *.mcc333.mcn111.gprs would match all APNs from country 333 and carrier 111 on the gprs network.

Mode	Select one or more of the methods used to obtain APN values. Mobile Station provided - The APN comes from the mobile station where the mobile device connected. This is the point of entry into the carrier network for the user's connection. Network provided - The APN comes from the carrier network. Subscription Verified - The user's subscription has been verified for this APN. This is the most secure option.
Action	One of allow or deny to allow or block traffic associated with this APN.
Delete icon	Select to remove this APN entry from the list.
Edit icon	Select to change the information for this APN entry.
Add APN	Select to add an APN to the list. Not active while creating GTP profile, only when editing an existing GTP profile. Save all changes before adding APNs. A warning to this effect will be displayed when you select the <i>Add APN</i> button.

The Add APN button is not activated until you save the new GTP profile. When you edit that GTP profile, you will be able to add new APNs.

Configuring IMSI filtering in FortiOS Carrier

In many ways the IMSI on a GPRS network is similar to an IP address on a TCP/IP network. Different parts of the number provide different pieces of information. This concept is used in IMSI filtering on FortiOS Carrier.

To configure IMSI filtering go to *UTM Profiles > Carrier > GTP Profile* and expand *IMSI filtering*.

While both the APN and MCC-MCN fields are optional, without using one of these fields the IMSI entry will not be useful as there is no information for the filter to match.

Enable IMSI Filter	Select to turn on IMSI filtering.
Default IMSI Action	Select Allow or Deny. This action will be applied to all IMSI numbers except as indicated in the IMSI list that is displayed. The default value is Allow.
APN	The Access Point Number (APN) to filter on. This field is optional.
MCC-MNC	The Mobile Country Code (MCC) and Mobile Network Code (MNC) to filter on. Together these numbers uniquely identify the carrier and network of the GGSN being used. This field is optional.

Mode	<p>Select the source of the IMSI information as one or more of the following:</p> <p>Mobile Station provided - the IMSI number comes from the mobile station the mobile device is connecting to.</p> <p>Network provided - the IMSI number comes from the GPRS network which could be a number of sources such as the SGSN, or HLR.</p> <p>Subscription Verified - the IMSI number comes from the user's home network which has verified the information.</p> <p>While Subscription Verified is the most secure option, it may not always be available. Selecting all three options will ensure the most complete coverage.</p>
Action	Select the action to take when this IMSI information is encountered. Select one of Allow or Deny.
Delete Icon	Select the delete icon to remove this IMSI entry.
Edit Icon	Select the edit icon to change information for this IMSI entry.
Add IMSI	<p>Select to add an IMSI to the list. Not active while creating GTP profile, only when editing an existing GTP profile.</p> <p>Save all changes before adding IMSIs. A warning to this effect will be displayed when you select the <i>Add IMSI</i> button.</p>

Also see ["IMSI filtering options"](#) on page 2341.

Configuring advanced filtering in FortiOS Carrier

Compared to ADN or IMSI filtering, advanced filtering is well named. Advanced filtering can be viewed as a catch-all filtering option — if ADN or IMSI filtering doesn't do what you want, then advanced filtering will. The advanced filtering can use more information elements to provide considerably more granularity for your filtering.

Enable	Select to turn on advanced filtering.
Default Action	Select Allow or Deny as the default action to take when traffic does not match an entry in the advanced filter list .
Messages	<p>Optionally select one or more types of messages this filter applies to:</p> <p>Create PDP Context Request, Create PDP Context Response, Update PDP Context Request, or Update PDP Context Response. Selecting <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> limits RAT type to only GAN and HSPA, and disables the APN, APN Mode, IMSI, MSISDN, ULI, RAI, and IMEI fields.</p> <p>To select <i>Update PDP Context Request</i>, APN Restriction must be set to <i>all</i>. Selecting <i>Update PDP Context Request</i> disables the APN, MSISDN, and IMEI fields.</p> <p>if all message types are selected, only the RAT Types of GAN and HSPA are available to select.</p>

APN Restriction	APN Restriction either allows all APNs or restricts the APNs to one of four categories — Public-1, Public-2, Private-1, or Private-2. This can also be combined with a specific APN or partial APN as well as specifying the APN mode. See “Access Point Number (APN)” on page 2418 .
RAT Type	Select one or more of the Radio Access Technology Types listed. These fields control how a user accesses the carrier’s network. You can select one or more of UTRAN, GERAN, WLAN, GAN, HSPA, or any. See “Radio Access Technology (RAT) type” on page 2418 .
ULI	The user location identifier. Often the ULI is used with the RAI to locate a user geographically on the carrier’s network. The ULI is disabled when <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> messages are selected. See “User Location Information (ULI)” on page 2419 .
RAI	The router area identifier. There is only one SGSN per routing area on a carrier network. This is often used with ULI to locate a user geographically on a carrier network. The RAI is disabled when <i>Create PDP Context Response</i> or <i>Update PDP Context Response</i> messages are selected. See “Routing Area Identifier (RAI)” on page 2419 .
IMEI	The International Mobile Equipment Identity. The IMEI uniquely identifies mobile hardware, and can be used to block stolen equipment. The IMEI is only available when <i>Create PDP Context Request</i> or no messages are selected. See “International Mobile Equipment Identity (IMEI)” on page 2419 .
Action	Select Allow or Deny as the action when this filter matches traffic. The default is Allow.
Delete Icon	Select to delete this entry from the list.
Edit Icon	Select to edit this entry.
Add	Select to add an advanced filter to the list. Not active while creating GTP profile, only when editing an existing GTP profile. Save all changes before adding advanced filters. A warning to this effect will be displayed when you select the <i>Add</i> button.

Also see [“Advanced filtering options” on page 2343](#).



Troubleshooting

This section highlights troubleshooting for Carrier related issues.

This section includes:

- [FortiOS Carrier diagnose commands](#)
- [Applying Intrusion and Prevention System \(IPS\) signatures to IP packets within GTP-U tunnels](#)
- [GTP packets are not moving along your network](#)

FortiOS Carrier diagnose commands

This section includes diagnose commands specific to FortiOS Carrier features such as GTP.

GTP related diagnose commands

This CLI command allows you to gain information on GTP packets, logs, statistics, and other information.

```
diag firewall gtp <command>
```

apn list <gtp_profile>	The APN list entries in the specified GTP profile
auth-ggsns show <gtp_profile>	The authorized GGSNs entries for the specified GTP profile. Any GGSNs not on this list will not be recognized.
auth-sgsns show <gtp_profile>	The authorized SGSNs list entries for the specified GTP profile. Any SGSNs not on this list will not be recognized.
handover-grp show <gtp_profile>	The handover group showing the range of allowed handover group IP addresses. The handover group acts like a whitelist of allowed GTP addresses with a default deny at the end — if the GTP address is not on the list, it is denied.
ie-remove-policy list <gtp_profile>	List of IE policies in the IE removal policy for this GTP profile. The information displayed includes the message count for this policy, the length of the SGSN, the list of IEs, and list of SGSN IP addresses.
imsi list <gtp_profile>	IMSI filter entries for this GTP profile. The information displayed includes the message count for this filter, length of the IMSI, the length of the APN and IMSI, and of course the IMSI and APN values.
invalid-sgsns-to-long list <gtp_profile>	List of SGSNs that do not match the filter criteria. These SGSNs will be logged.
ip-policy list <gtp_profile>	List the IP policies including message count for each policy, the action to take, the source and destination IP addresses or ranges, and masks.

noip-policy <gtp_profile>	List the non-IP policies including the message count, which mode, the action to take, and the start and end protocols to be used by decimal number.
path {list flush}	Select list or flush. List the GTP related paths in FortiOS Carrier memory. Flush the GTP related paths from memory.
policy list <gtp_policy>	The GTP advanced filter policy information for this GTP profile. The information displayed for each entry includes a count for messages matching this filter, a hexadecimal mask of which message types to match, the associated flags, action to take on a match, APN selection mode, MSISDN, RAT types, RAI, ULI, and IMEI.
profile list	Displays information about the configured GTP profiles. You will not be able to see the bulk of the information if you do not log the output to a file.
runtime-stat flush	Select to flush the GTP runtime statistics from memory.
stat	Display the GTP runtime statistics — details on current GTP activity. This information includes how many tunnels are active, how many GTP profiles exist, how many IMSI filter entries, how many APN filter entries, advanced policy filter entries, IE remove policy filter entries, IP policy filter entries, clashes, and dropped packets.
tunnel {list flush}	Select one of list or flush. List lists all the GTP tunnels currently active. Flush clears the list of active GTP tunnels.

Applying Intrusion and Prevention System (IPS) signatures to IP packets within GTP-U tunnels

GTP-U (GTP user data tunnelling) tunnels carry user data packets, signalling messages and error information. GTP-U uses UDP port 2152. FortiOS Carrier units can apply IPS intrusion protection and detection to GTP-U user data sessions.

To apply IPS to GTP-U user data sessions, add an IPS Sensor to a profile and add the profile to a security policy that accepts GTP-U tunnels. The security policy Service field must be set to GTP or ANY to accept GTP-U packets.

The FortiOS Carrier unit intercepts packets with destination port 2152, removes the GTP header and handles the packets as regular IP packets. Applying an IPS sensor to the IP packets, the FortiOS Carrier unit can log attacks and pass or drop packets depending on the configuration of the sensor.

If the packet is GTP-in-GTP, or a nested tunnel, the packets are passed or blocked without being inspected.

To apply an IPS sensor to GTP-U tunnels

- 1 Go to *UTM Profiles > Intrusion Protection > IPS Sensor* and select *Create New* to add an IPS Sensor.

- 2 Configure the IPS Sensor to detect attacks and log, drop, or pass attack packets.
See the Intrusion Protection section of the [FortiOS UTM Guide](#).
- 3 Go to *Policy > Policy* and apply the IPS sensor to the security policy.
- 4 Go to *Policy > Policy* and select Create New to add a security policy or select a security policy.
- 5 Configure the security policy to accept GTP traffic.
In the security policy configure the source and destination settings to match the GTP traffic. Service to GTP or ANY so that the security policy accepts GTP traffic.
- 6 Select the GTP profile within the security policy.
- 7 Configure any other required security policy settings.
- 8 Select OK to save the security policy.

GTP packets are not moving along your network

When GTP packets are not getting to their destination, this could be caused by any one of a number of issues. General troubleshooting principals apply here.

The following sections provide some suggestions on how to troubleshoot this issue:

- [Attempt to identify the section of your network with the problem](#)
- [Ensure you have an APN configured](#)
- [Check the logs and adjust their settings if required](#)
- [Check the routing table](#)
- [Perform a sniffer trace](#)
- [Generate specific packets to test the network](#)

Attempt to identify the section of your network with the problem

The first step is to determine how widespread this problem is. Does it affect the whole GPRS network, or just one or two devices?

If the entire network is has this problem, the solution is likely a more general one such as ensuring the security policies allow GTP traffic to pass, the GTP profile specifies SSGNs and GSGNs, or ensuring the GTP general settings are not overly limiting.

If one part of the network is affected, the problem is more likely centered around configurations with those network devices specified such as the handover group, or authorized SGSNs/GSGNs. It is also possible that small portions of the network may have hardware related issues such as cabling or faulty hardware. This section does not address those issues, and assumes hardware is not the problem.

The handover group is a whitelist of GTP addresses allowed to handle GTP messages. If a device's address is not on this list, it will be denied.

Ensure you have an APN configured

When you configure your GTP profile, ensure you first configure the APN. Without it, there will be no flow of traffic. The APN is used in nearly all GTP communications and without it, the FortiOS Carrier unit doesn't have the information it needs.

Check the logs and adjust their settings if required

During normal operation, the log settings will show any problems on the network but may not provide the level of details required to fully troubleshoot the problem. The reason for this is that the level of detail required for troubleshooting would quickly overwhelm the daily logs without any real benefit.

GTP related events in the event log will have message IDs in the range 41216 to 41222. For more information on GTP log messages, see the [Log Message Reference](#). For more information on logging in general, see the [Logging and Reporting handbook chapter](#).

Once there is a problem to troubleshoot, check the logs to trace the traffic patterns and narrow down the possible sources of the problem. There may be enough detail for you to locate and fix the problem without changing the log settings.



Remember to set any changes you made to the log settings back to their original values when you are done troubleshooting. Otherwise, the amount of detail will overwhelm your logging.

However, if more detail is required you can change settings such as:

- Lower the Log Frequency number in GTP Profiles so fewer or no log messages are dropped. This will allow a more accurate picture of everything happening on the network, where you may have had only a partial picture before.
- Ensure all the GTP log events are enabled to provide you with a complete picture.
- Increase the minimum log level to Information or Debug to ensure you are seeing all possible log entries. This is found if you go to *Log&Report > Log Config > Log Setting > Local Logging & Archiving > Minimum log level*.
- Ensure that all relevant event types are enabled under *Log&Report > Log Config > Event Log*.

For more information on GTP related logging, see “[Logging events on the FortiOS Carrier unit](#)” on page 2406. See the log and report chapters of [Logging and Reporting Guide](#) and [FortiOS CLI Reference](#).

General information to look for in the logs includes:

- Are all packets having problems or just certain types?
- Are all devices on the network having problem, or just certain devices?
- Is it just GTP traffic that is having problems or are all types of traffic having the same problem?

Check the routing table

On any network, the routing table determines how packets reach their destination. This is also true on a carrier network.

If the FortiOS Carrier unit is running in NAT mode, verify that all desired routes are in the routing table — local subnets, default routes, specific static routes, and dynamic routing protocols. For complete information, it is best to check the routing table in the CLI. This method provides more complete information.



If VDOMs are enabled on your FortiOS Carrier unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

To check the routing table using the CLI

```
# get router info routing-table all
```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
 O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS
 inter area
 * - candidate default

```
S*      0.0.0.0/0 [10/0] via 192.168.183.254, port2
S       1.0.0.0/8 [10/0] via 192.168.183.254, port2
S       2.0.0.0/8 [10/0] via 192.168.183.254, port2
C       10.142.0.0/23 is directly connected, port3
B       10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C       192.168.182.0/23 is directly connected, port2
```

Examining an entry from the routing table above:

```
B       10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route including netmask.
[20/0]	20 indicates and administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
10.142.0.74	The gateway, or next hop.
port3	The interface used by this route.
2d18h02m	How old this route is, in this case almost three days old.

Perform a sniffer trace

When troubleshooting network traffic, it helps to look inside the headers of packets to determine if they are traveling along the route you expect. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiOS Carrier unit has NP2 interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP2 interfaces, disable offloading on those interfaces.

What can sniffing packets tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiOS Carrier unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the FortiOS Carrier unit is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing. This prevents an IP packet from being forwarded if its source IP address either does not belong to a locally attached subnet (local interface), or be a hop on the routing between the FortiOS Carrier and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the FortiOS Carrier unit and consequently cause many features to be turned off.



If you configure virtual IP addresses on your FortiOS Carrier unit, the unit will use those addresses in preference to the physical IP addresses. If not configured properly, secondary IP addresses can cause a broadcast storm. You will notice the secondary address being preferred when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

How to sniff packets

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose>
<count>
```

To stop the sniffer, type `CTRL+C`.

<interface_name>	The name of the interface to sniff, such as <code>port1</code> or <code>internal</code> . This can also be <code>any</code> to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. <code>none</code> indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <code><CTRL C></code> .

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the `port1` interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as `ack`), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh
3177924955 ack 1854307757
```

```
0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh
1854307757 ack 3177925808
```

```
0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh
4225311614 ack 3314279933
```

Generate specific packets to test the network

If some packets are being delivered as expected while others are not, or after you believe you have fixed the problem, it is a good idea to generate specific traffic to test your network.

For example if you discover through log messages and packet sniffing that Create PDP Context Request messages are not being delivered between two SGSNs, you can generate those specific messages on your network to confirm they are the problem, and later that you have solved the problem and they are now being delivered as expected.

This step requires a third party traffic generation tool, either hardware or software. This is not supported by Fortinet.



Chapter 15 Deploying Wireless Networks

This FortiOS Handbook chapter contains the following sections:

[Introduction to wireless networking](#) explains the basic concepts of wireless networking and how to plan your wireless network.

[Configuring a WiFi LAN](#) explains how to set up a basic wireless network, prior to deploying access point hardware.

[Access point deployment](#) explains how to deploy access point hardware and add it to your wireless network configuration.

[Wireless network monitoring](#) explains how to monitor your wireless clients and how to monitor other wireless access points, potentially rogues, in your coverage area.

[Configuring wireless network clients](#) explains how to configure typical wireless clients to work with a WPA-Enterprise protected network.

[Wireless network examples](#) provides two examples. The first is a simple WiFi network using automatic configuration. The second is a more complex example of a business with two WiFi networks, one for employees and another for guests or customers.

[Using a FortiWiFi unit as a client](#) explains how to use a FortiWiFi unit as a wireless client to connect to other WiFi networks. This connection can take the place of an Ethernet connection where wired access to a network or to the Internet is not available.

[WiFi Reference](#) provides information about WiFi radio channels.

[WiFi Controller Reference](#) details the web-based manager pages that configure the WiFi controller, manage access points, and monitor your WiFi network.



Introduction to wireless networking

This chapter introduces some concepts you should understand before working with wireless networks, describes Fortinet's wireless equipment, and then describes the factors you need to consider in planning deployment of a wireless network.

The following topics are included in this section:

- [Wireless concepts](#)
- [Security](#)
- [Authentication](#)
- [Wireless networking equipment](#)
- [Deployment considerations](#)
- [Automatic Radio Resource Provisioning](#)

Wireless concepts

Wireless networking is radio technology, subject to the same characteristics and limitations as the familiar audio and video radio communications. Various techniques are used to modulate the radio signal with a data stream.

Bands and channels

Depending on the wireless protocol selected, you have specific channels available to you, depending on what region of the world you are in.

- IEEE 802.11a,b, and g protocols provide up to 14 channels in the 2.400-2.500 GHz Industrial, Scientific and Medical (ISM) band.
- IEEE 802.11a,n (5.150-5.250, 5.250-5.350, 5.725-5.875 GHz, up to 16 channels) in portions of Unlicensed National Information Infrastructure (U-NII) band

Note that the width of these channels exceeds the spacing between the channels. This means that there is some overlap, creating the possibility of interference from adjacent channels, although less severe than interference on the same channel. Truly non-overlapping operation requires the use of every fourth or fifth channel, for example ISM channels 1, 6 and 11.

The capabilities of your wireless clients is the deciding factor in your choice of wireless protocol. If your clients support it, 5GHz protocols have some advantages. The 5GHz band is less used than 2.4GHz and its shorter wavelengths have a shorter range and penetrate obstacles less. All of these factors mean less interference from other access points, including your own.

When configuring your WAP, be sure to correctly select the Geography setting to ensure that you have access only to the channels permitted for WiFi use in your part of the world.

For detailed information about the channel assignments for wireless networks for each supported wireless protocol, see [“Wireless radio channels”](#) on page 2503.

Power

Wireless LANs operate on frequencies that require no license but are limited by regulations to low power. As with other unlicensed radio operations, the regulations provide no protection against interference from other users who are in compliance with the regulations.

Power is often quoted in dBm. This is the power level in decibels compared to one milliwatt. 0dBm is one milliwatt, 10dBm is 10 milliwatts, 27dBm, the maximum power on Fortinet FortiAP equipment, is 500 milliwatts. The FortiGate unit limits the actual power available to the maximum permitted in your region as selected by the WiFi controller country setting.

Received signal strength is almost always quoted in dBm because the received power is very small. The numbers are negative because they are less than the one milliwatt reference. A received signal strength of -60dBm is one millionth of a milliwatt or one nanowatt.

Antennas

Transmitted signal strength is a function of transmitter power and antenna gain. Directional antennas concentrate the signal in one direction, providing a stronger signal in that direction than would an omnidirectional antenna.

FortiWiFi units have detachable antennas. However, these units receive regulatory approvals based on the supplied antenna. Changing the antenna might cause your unit to violate radio regulations.

Security

There are several security issues to consider when setting up a wireless network.

Whether to broadcast SSID

Users who want to use a wireless network must configure their computers with the wireless service set identifier (SSID) or network name. Broadcasting the SSID makes connection to a wireless network easier because most wireless client applications present the user with a list of network SSIDs currently being received. This is desirable for a public network.

To obscure the presence of a wireless network, do not broadcast the SSID. This does not prevent attempts at unauthorized access, however, because the network is still detectable with wireless network “sniffer” software.

Encryption

Wireless networking supports the following security modes for protecting wireless communication, listed in order of increasing security.

None — Open system. Any wireless user can connect to the wireless network.

WEP64 — 64-bit Web Equivalent Privacy (WEP). This encryption requires a key containing 10 hexadecimal digits.

WEP128 — 128-bit WEP. This encryption requires a key containing 26 hexadecimal digits.

WPA — 256-bit Wi-Fi Protected Access (WPA) security. This encryption can use either the TKIP or AES encryption algorithm and requires a key of either 64 hexadecimal digits or a text phrase of 8 to 63 characters. It is also possible to use a RADIUS server to store a separate key for each user.

WPA2 — WPA with security improvements fully meeting the requirements of the IEEE 802.11i standard. Configuration requirements are the same as for WPA.

For best security use the WPA2 with AES encryption and a RADIUS server to verify individual credentials for each user. WEP, while better than no security at all, is an older algorithm that is easily compromised. With either WEP or WAP, changing encryption passphrases on a regular basis further enhances security.

Separate access for employees and guests

Wireless access for guests or customers should be separate from wireless access for your employees. This does not require additional hardware. Both FortiWiFi units and FortiAP units support multiple wireless LANs on the same access point. Each of the two networks can have its own SSID, security settings, firewall policies, and user authentication.

A good practice is to broadcast the SSID for the guest network to make it easily visible to users, but not to broadcast the SSID for the employee network.

Two separate wireless networks are possible because multiple virtual APs can be associated with an AP profile. The same physical APs can provide two or more virtual WLANs.

Captive portal

As part of authenticating your users, you might want them to view a web page containing your acceptable use policy or other information. This is called a captive portal. No matter what URL the user initially requested, the portal page is returned. Only after authenticating and agreeing to usage terms can the user access other web resources.

For information about setting up a captive portal, see [“Captive Portal security” on page 2452](#).

Power

Reducing power reduces unwanted coverage and potential interference to other WLANs. Areas of unwanted coverage are a potential security risk. There are people who look for wireless networks and attempt to access them. If your office WLAN is receivable out on the public street, you have created an opportunity for this sort of activity.

Monitoring for rogue APs

It is likely that there are APs available in your location that are not part of your network. Most of these APs belong to neighboring businesses or homes. They may cause some interference, but they are not a security threat. There is a risk that people in your organization could connect unsecured WiFi-equipped devices to your wired network, inadvertently providing access to unauthorized parties. The optional On-Wire Rogue AP Detection Technique compares MAC addresses in the traffic of suspected rogues with the MAC addresses on your network. If wireless traffic to non-Fortinet APs is also seen on the wired network, the AP is a rogue, not an unrelated AP.

Decisions about which APs are rogues are made manually on the Rogue AP monitor page. For detailed information about monitoring rogue APs, see [“Monitoring rogue APs” on page 2470](#).

Suppressing rogue APs

When you have declared an AP to be a rogue, you have the option of suppressing it. To suppress an AP, the FortiGate WiFi controller sends reset packets to the rogue AP. Also, the MAC address of the rogue AP is blocked in the firewall policy. You select the suppression action on the Rogue AP monitor page. For more information, see [“Suppressing rogue APs” on page 2473](#).



Rogue suppression is available only when there is a radio dedicated to scanning. It will not function during background scanning.

Authentication

Wireless networks usually require authenticated access. FortiOS authentication methods apply to wireless networks the same as they do to wired networks because authentication is applied in the firewall policy.

The types of authentication that you might consider include:

- user accounts stored on the FortiGate unit
- user accounts managed and verified on an external RADIUS, LDAP or TACACS+ server
- Windows Active Directory authentication, in which users logged on to a Windows network are transparently authenticated to use the wireless network.

This Wireless chapter of the FortiOS Handbook will provide some information about each type of authentication, but more detailed information is available in the Authentication chapter.

What all of these types of authentication have in common is the use of user groups to specify who is authorized. For each wireless LAN, you will create a user group and add to it the users who can use the WLAN. In the identity-based firewall policies that you create for your wireless LAN, you will specify this user group.

Some access points, including FortiWiFi units, support MAC address filtering. You should not rely on this alone for authentication. MAC addresses can be “sniffed” from wireless traffic and used to impersonate legitimate clients.

Wireless networking equipment

Fortinet produces two types of wireless networking equipment:

- FortiWiFi units, which are FortiGate units with a built-in wireless access point/client
- FortiAP units, which are wireless access points compliant with the CAPWAP standard that you can control from any FortiGate unit that supports the WiFi Controller feature.

FortiWiFi units

A FortiWiFi unit can:

- Provide an access point for clients with wireless network cards. This is called Access Point mode, which is the default mode.

or

- Connect the FortiWiFi unit to another wireless network. This is called Client mode. A FortiWiFi unit operating in client mode can only have one wireless interface.

or

- Monitor access points within radio range. This is called Monitoring mode. You can designate the detected access points as Accepted or Rogue for tracking purposes. No access point or client operation is possible in this mode. But, you can enable monitoring as a background activity while the unit is in Access Point mode.

FortiWiFi unit capabilities differ by model as follows:

Table 129: FortiWiFi model capabilities

Model	Radio	Simultaneous SSIDs
20C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
30B	802.11 b/g 2.4GHz	7 for AP, 1 for monitoring
40C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
50B	802.11 b/g 2.4GHz	7 for AP, 1 for monitoring
60B	802.11 b/g 2.4GHz 802.11 a 5GHz	7 for AP, 1 for monitoring
60C	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring
80/81CM	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	7 for AP, 1 for monitoring

Using a FortiWiFi unit as a managed AP

A FortiWiFi unit can also be used much like a FortiAP unit to provide an access point managed by another FortiGate unit. To use a FortiWiFi unit as a managed WAP, you need to switch it to wireless terminal mode by using the CLI as follows:

```
config system global
    set wireless-mode wtp
end
```



FortiWiFi-80CM supports WTP mode only in FortiOS 4.3 patch 2 or later.

The wireless functionality of a FortiWiFi unit in wireless terminal mode cannot be controlled from the unit itself.

If there are firewall devices between the WiFi controller FortiGate unit and the managed FortiWiFi units, make sure that ports 5246 and 5247 are open. These ports carry, respectively, the encrypted control channel data and the wireless network data. If needed, you can change these ports in the CLI:

```
config system global
    set wireless-controller-port <port_int>
end
```

This command sets the control channel port. The data channel port is always the control port plus one. The port setting must match on the access controller and all access points.

FortiAP units

FortiAP series wireless access points are controlled by a FortiGate unit over Ethernet. Capabilities differ by model as follows:

Table 130: FortiAP model capabilities

Model	Radio 1	Radio 2	Simultaneous SSIDs
210B (indoor)	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	N/A	7 for AP, 1 for monitoring
220A (indoor)	802.11 b/g/n 2.4GHz	802.11 a/n 5GHz	14 for AP, 2 for monitoring
220B (indoor)	802.11 b/g/n 2.4GHz 802.11 a/n 5GHz	802.11 b/g/n 2.4GHz	14 for AP, 2 for monitoring
222B (outdoor)	802.11 b/g/n 2.4GHz	802.11 a/n 5GHz	14 for AP, 2 for monitoring

Dual-band radios can function as an AP on either band or as a dual-band monitor. The monitoring function is also available during AP operation if Background Scan is enabled in the custom AP profile for the device.

Third-party WAPs

FortiOS implements the CAPWAP standard.

Deployment considerations

Several factors need to be considered when planning a wireless deployment.

Types of wireless deployment

This Handbook chapter describes two main types of wireless deployment: single WAP and multiple WAP. You will know which type of deployment you need after you have evaluated the coverage area environment.

Deployment methodology

- 1 Evaluate the coverage area environment.
- 2 Position access point(s).
- 3 Select access point hardware.
- 4 Install and configure the equipment.
- 5 Test and tune the network.

Evaluating the coverage area environment

Consider the following factors:

- **Size of coverage area** — Even under ideal conditions, reliable wireless service is unlikely beyond 100 metres outdoors or 30 metres indoors. Indoor range can be further diminished by the presence of large metal objects that absorb or reflect radio frequency energy. If wireless users are located on more than one floor of a building, a minimum of one WAP for each floor will be needed.

- **Bandwidth required** — Wireless interface data rates are between 11 and 150 Mb/s, depending on the 802.11 protocol that is used. This bandwidth is shared amongst all users of the wireless data stream. If wireless clients run network-intensive applications, fewer of them can be served satisfactorily by a single WAP.

Note that on some FortiWiFi units you can define up to four wireless interfaces, increasing the available total bandwidth.

- **Client wireless capabilities** — The 802.11n protocol provides the highest data rates and has channels in the less interference-prone 5GHz band, but it is supported only on the latest consumer devices. The 802.11g protocol is more common but offers lower bandwidth. Some older wireless client equipment supports only 802.11b with a maximum data rate of 11Mb/s. WAP radios support the protocol that you select with backward compatibility to older modes. For example, if you select 802.11n, clients can also connect using 802.11g or 802.11b.

The most important conclusion from these considerations is whether more than one WAP is required.


Positioning access points

When placing the access point, your main concern is providing a strong signal to all users. A strong signal ensures a fast connection and efficient data transfer. A weaker signal means a greater chance of data transmission errors and the need to re-send information, slowing down data transfer.

Consider the following guidelines when placing access points:

- Physical barriers can impede the radio signals. Solid objects such as walls, furniture and people absorb radio waves, weakening the signal. Be aware of the physical barriers in your office space that may reduce a signal. If there is enough physical interference, you may encounter dead spots that receive no signal.
- Ensure the access point is located in a prominent location within a room for maximum coverage, rather than in a corner.
- Construction materials used in a building can also weaken radio signals. Rooms with walls of concrete or metal can affect the signal strength.

If you cannot avoid some of these impediments due to the shape of the office or building materials used, you may need to use multiple access points to help distribute the radio signal around the room. [Figure 251](#) shows how positioning two FortiAP-220A units within a uniquely shaped office space helps to distribute signals around the area.

Figure 251: Using multiple APs to provide a constant strong signal.

This sample office has washrooms, a stairwell and an elevator shaft in the center of the building, making it impossible to use a single access point effectively. The elevator shaft and multiple metal stalls in the washrooms can cause signal degradation. However, placing access points in diagonally opposite areas of the office provides maximum coverage.

When using multiple access points, set each access point to a different channel to avoid interference in areas where signals from both access points can be received.

Selecting access point hardware

For a single WAP installation, you could deploy a single FortiWiFi unit. If the site already has a FortiGate unit that supports the WiFi controller feature, adding a FortiAP unit is the most economical solution.

For a multiple WAP deployment you need a FortiGate unit as a WiFi controller and multiple FortiAP units. A FortiWiFi unit can be used as a managed WAP, but it is more expensive.

The FortiAP unit offers more flexible placement. FortiWiFi units either sit on a shelf or are rack mounted. FortiAP units can be attached to any wall or ceiling, enabling you to locate them where they will provide the best coverage.

Single access point networks

A single access point is appropriate for a limited number of users in a small area. For example, you might want to provide wireless access for a group of employees in one area on one floor of an office building.

A good rule of thumb is that one access point can serve 3000 to 4000 square feet of space, with no user more than 60 feet from the access point. Walls and floors reduce the coverage further, depending on the materials from which they are made.

Multiple access point networks

To cover a larger area, such as multiple floors of a building, or multiple buildings, multiple access points are required.

In the WiFi controller, you configure a single virtual access point, but the controller manages multiple physical access points that share the same configuration. A feature known as “fast roaming” enables users to move from one physical access point coverage area to another while retaining their authentication.

Fast Roaming

Users in a multi-AP network, especially with mobile devices, can move from one AP coverage area to another. But, the process of re-authentication can often take seconds to complete and this can impair wireless voice traffic and time sensitive applications. The FortiAP fast roaming feature solves this problem and is available only when moving between FortiAP units managed by the same FortiGate unit.

Fast roaming uses two standards-based techniques:

- Pairwise Master Key (PMK) Caching enables a RADIUS-authenticated user to roam away from an AP and then roam back without having to re-authenticate. To accomplish this, the FortiGate unit stores in a cache a master key negotiated with the first AP. This enables the 802.11i-specified method of “fast roam-back.”
- Pre-authentication or “fast-associate in advance” enables an 802.11 AP associated to a client to bridge to other APs over the wired network and pre-authenticate the client to the “next” AP to which the client might roam. This enables the PMK to be derived in advance of a roam and cached. When the client does roam, it will already have negotiated authentication in advance and will use its cached PMK to quickly associate to the next AP. This capability will ensure that wireless clients that support Pre-authentication to continue the data transfer without noticeable connection issues.

Automatic Radio Resource Provisioning

To prevent interference between APs, the FortiOS WiFi Controller includes the Automatic Radio Resource Provisioning (ARRP) feature. When enabled in an access point profile, this feature measures utilization and interference on the available channels and selects the clearest channel at each access point. The measurement can be repeated periodically to respond to changing conditions.



Configuring a WiFi LAN

When working with a FortiGate WiFi controller, you can configure your wireless network before you install any access points. If you are working with a standalone FortiWiFi unit, the access point hardware is already present but the configuration is quite similar. Both are covered in this section.


The following topics are included in this section:

- [Overview of WiFi controller configuration](#)
- [Setting your geographic location](#)
- [Creating a custom AP Profile](#)
- [Defining a wireless network interface \(SSID\)](#)
- [Configuring user authentication](#)
- [Configuring firewall policies for the SSID](#)
- [Customizing captive portal pages](#)
- [Configuring the built-in access point on a FortiWiFi unit](#)

Overview of WiFi controller configuration

The FortiGate WiFi controller configuration is composed of three types of object, the SSID, the AP Profile and the physical Access Point.

- An **SSID** defines a virtual wireless network interface, including security settings. One SSID is sufficient for a wireless network, regardless how many physical access points are provided. You might, however, want to create multiple SSIDs to provide different services or privileges to different groups of users. Each SSID has separate firewall policies and authentication. Each radio in an access point can support up to 8 SSIDs.
A more common use of the term SSID is for the identifier that clients must use to connect to the wireless network. Each SSID (wireless interface) that you configure will have an SSID field for this identifier. In Managed Access Point configurations you choose wireless networks by SSID values. In firewall policies you choose wireless interfaces by their SSID name.
- An **AP Profile** defines the radio settings, such as band (802.11g for example) and channel selection. The AP Profile names the SSIDs to which it applies. Managed APs can use automatic profile settings or you can create custom AP profiles.
- **Managed Access Points** represent local wireless APs on FortiWiFi units and FortiAP units that the FortiGate unit has discovered. There is one managed access point definition for each AP device. An access point definition can use automatic AP profile settings or select a custom AP Profile. When automatic profile settings are used, the managed AP definition also selects the SSIDs to be carried on the AP.

Figure 252: Conceptual view of FortiGate WiFi controller configuration

About SSIDs on FortiWiFi units

FortiWiFi units have a default SSID (wireless interface) named *wlan*. You can modify or delete this SSID as needed. As with external APs, the built-in wireless AP can be configured to carry any SSID.

The AP settings for the built-in wireless access point are located at *WiFi Controller > Managed Access Points > Local WiFi Radio*. The available operational settings are the same as those for external access points which are configured at *WiFi Controller > Managed Access Points > Managed FortiAP*.

About automatic AP profile settings

FortiOS simplifies wireless network configuration by providing an automatic setting for the access point profile. You can enable wireless AP operation and Rogue AP scanning with the radios in the AP automatically allocated as follows:

Table 131: Radio functions in automatic profile

No. of Radios	Wireless Access enabled	Rogue AP Scan enabled	Wireless Access and Rogue AP Scan enabled
1	Radio 1 - AP	Radio 1 - scan	Radio 1 - AP + background scan
2	Radio 1 - AP Radio 2 - disabled	Radio 1 - disabled Radio 2 - scan	Radio 1 - AP Radio 2 - scan

You can select which SSIDs (wireless networks) will be available through the access point and adjust the wireless power level of the AP.

Process to create a wireless network

To set up your wireless network, you will need to perform the following steps.

- Make sure the FortiGate wireless controller is configured for your geographic location. This ensures that the available radio channels and radio power are in compliance with the regulations in your region.
- Optionally, if you don't want to use automatic AP profile settings, configure a custom Access Point (AP) profile, specifying the radio settings and the SSIDs to which they apply.
- Configure one or more SSIDs for your wireless network. The SSID configuration includes DHCP and DNS settings.
- Configure the user group and users for authentication on the WLAN.
- Configure the firewall policy for the WLAN.
- Optionally, customize the captive portal.
- Configure access points.

Configuration of the built-in AP on FortiWiFi units is described in this chapter. Connection and configuration of FortiAP units is described in the next chapter, [“Access point deployment”](#).

Setting your geographic location

The maximum allowed transmitter power and permitted radio channels for Wi-Fi networks depend on the region in which the network is located. By default, the WiFi controller is configured for the United States. If you are located in any other region, you need to set your location before you begin configuring wireless networks.

To change the location setting - CLI

To change the country to France, for example, enter

```
config wireless-controller setting
  set country FR
end
```

To see the list of country codes, enter a question mark (?) instead of a country code.



Before changing the country setting, you must remove all Custom AP profiles. To do this, go to *WiFi Controller > Managed Access Points > Custom AP Profile*.

Creating a custom AP Profile

If the automatic AP profile settings don't meet your needs, you can define a custom AP Profile. For information about the automatic profile settings, see [“About automatic AP profile settings” on page 2446](#).

An AP Profile configures radio settings and selects the Virtual APs to which the settings apply. FortiAP units contain two radio transceivers, making it possible, for example, to provide both 2.4GHz 802.11b/g/n and 5GHz 802.11a/n service from the same access point.

FortiAP units also provide a monitoring function for the Rogue AP feature.

To configure an AP Profile - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Custom AP Profile* and select *Create New*.
- 2 Enter a *Name* for the AP Profile.
- 3 In *Platform*, select the FortiWiFi or FortiAP model to which this profile applies.
- 4 In *Mode*, select *Access Point*.
- 5 Optionally, enable *Background Scan* to support the Rogue AP feature.
For more information see [“Wireless network monitoring” on page 2469](#).
- 6 Optionally, select *Radio Resource Provision* to enable the ARRP feature.
For more information see [“Automatic Radio Resource Provisioning” on page 2443](#).
- 7 In *Band*, select the 802.11 wireless protocol that you want to support.
Note that there are two choices for 802.11n. Select *802.11n* for 2.4GHz operation or *802.11n-5G* for 5GHz operation. The available choices depend on the radio’s capabilities.
- 8 In *Channel*, select the channels that the AP is permitted to use. By default, all channels are selected.
- 9 Leave the *TX Power* at its default setting. You can adjust this later.
- 10 In *SSID*, use the arrow buttons to move the SSIDs (wireless LANs) to which these settings apply into the *Selected* list.
- 11 Repeat steps 4 through 10 for Radio 2, if required.
Note that on the FortiAP-220 unit Radio 1 is 2.4GHz and Radio 2 is 5GHz.
Radio 2 also supports 40MHz wide channels on the 5GHz band on 802.11n.
- 12 Select *OK*.

To configure an AP Profile - CLI

This example configures a FortiAP-220A to use only Radio 1 for 802.11g operation applied to SSID `example_wlan`.

```
config wireless-controller wtp-profile
  edit guest_prof
    config platform
      set type 220A
    end
    config radio-1
      set mode ap
      set band 802.11g
      set vaps example_wlan
    end
  end
```

Defining a wireless network interface (SSID)

You begin configuring your wireless network by defining one or more SSIDs to which your users will connect.

A virtual AP defines the SSID and security settings that can be applied to one or more physical APs. On the FortiGate unit, this creates a virtual network interface with the virtual AP’s name. With this interface you can define the DHCP services, firewall policies, and other settings for your WiFi LAN.

To configure an SSID - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
- 2 Enter the *Interface Name* that will identify the wireless interface.
- 3 In the *Addressing Mode* section, enter the *IP/Netmask* for the interface.
If IPv6 is enabled, you can also enter an *IPv6 Address*.
- 4 In *Administrative Access*, select *Ping*.
Ping is useful for testing. For security it is better not to enable administrative access on wireless interfaces.
- 5 Enter the *SSID* for your WLAN and choose whether to enable *SSID Broadcast* or not.
For more information, see [“Whether to broadcast SSID” on page 2436](#).
- 6 If you want to provide DHCP service to your clients, select *Enable DHCP* and enter the range of IP addresses to assign.
For more information, see [“Configuring DHCP for WiFi clients” on page 2450](#).
- 7 Select the *Security Mode* and enter the required settings.
For more information, see [“Configuring security” on page 2450](#).
- 8 If you want to prevent direct communication between your wireless clients, enable *Block Intra-SSID Traffic*.
- 9 Optionally, set the *Maximum Clients* limit.
The default of 0 sets no limit on the number of clients.
- 10 If you want to restrict access to the wireless network by MAC address, select *Enable MAC Filter*.
For more information, see [“Adding a MAC filter” on page 2453](#).
- 11 Select OK.

Each Virtual AP that you create is a wireless interface that establishes a wireless LAN. Go to *System > Network > Interface* to configure its IP address.

To configure a virtual access point - CLI

This example creates an access point with SSID “example” and WPA2-Personal security. The wireless interface is named `example_wlan`.

```
config wireless-controller vap
  edit example_wlan
    set ssid "example"
    set broadcast-ssid enable
    set security wpa2-only-personal
    set passphrase "hardtoguess"
    set vdom root
  end
config system interface
  edit example_wlan
    set ip 10.10.120.1 255.255.255.0
  end
```

Configuring DHCP for WiFi clients

Wireless clients need to have IP addresses. If you use RADIUS authentication, each user's IP address can be stored in the Framed-IP-Address attribute. Otherwise, you need to configure a DHCP server on the WLAN interface to assign IP addresses to wireless clients.

To configure a DHCP server for WiFi clients - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
- 2 In the *WiFi Settings* section, select *Enable DHCP*.
- 3 In the *Address Start* and *Address End* fields, enter the IP address range to assign.
The address range needs to be in the same subnet as the wireless interface IP address, but not include that address.
- 4 Set the *Default Gateway* to the wireless interface IP address.
- 5 Set the *Netmask* to an appropriate value, such as 255.255.255.0.
- 6 Enter the IP address of the *DNS Server* that your users will access.
- 7 Select *OK*.

The DHCP server automatically configures itself to serve only FortiAP units.

You can also configure DHCP through *System > Network > DHCP Server*, but that page offers additional options that might not be suitable for a wireless network.

To configure a DHCP server for WiFi clients - CLI

In this example, WiFi clients on the `example_wlan` interface are assigned addresses in the 10.10.120.2-9 range to connect with the WiFi access point on 10.10.120.1.

```
config system dhcp server
  edit 0
    set default-gateway 10.10.120.1
    set dns-service default
    set interface example_wlan
    set netmask 255.255.255.0
    config ip-range
      edit 1
        set end-ip 10.10.120.9
        set start-ip 10.10.120.2
      end
    end
  end
```



You cannot delete an SSID (wireless interface) that has DHCP enabled on it.

Configuring security

The FortiGate WiFi controller supports both Wireless Equivalent Privacy (WEP) and Wi-Fi Protected Access (WPA) security. WPA support includes WPA2, which has additional security improvements.

WEP security uses an encryption key between the wireless device and the access point. WEP64 uses a key of ten hexadecimal digits. WEP128 keys are 26 digits long. WEP security is relatively easy to break. Wherever possible, use WPA security. WEP can be enabled only through the CLI.

WPA security offers more robust encryption that is much more difficult to break. WPA provides two methods of authentication: through RADIUS (802.1X) authentication or by pre-shared key.

WPA security with a preshared key for authentication is called WPA-Personal. This can work well for one person a small group of trusted people. But, as the number of users increases, it is difficult to distribute new keys securely and there is increased risk that the key could fall into the wrong hands.

A more secure form of WPA security is WPA-Enterprise. Users each have their own authentication credentials, verified through an authentication server, usually RADIUS. FortiOS can also authenticate WPA-Enterprise users through its built-in user group functionality. FortiGate user groups can include RADIUS servers and can select users by RADIUS user group. This makes possible Role-Based Access Control (RBAC).

WPA security can encrypt communication with either Temporal Key Integrity Protocol (TKIP) or Advanced Encryption Standard (AES). AES is the preferred encryption, but some older wireless clients do not support it. You can select the encryption during setup.

Captive Portal security connects users to an open web portal defined in replacement messages. To navigate to any location beyond the web portal, the user must pass FortiGate user authentication.

WPA-Personal security

WPA-Personal security setup requires only the preshared key that you will provide to your clients.

To configure WPA-Personal security - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
- 2 In *Security Mode*, select *WPA/WPA2-Personal*.
- 3 In *Data Encryption*, select *AES*.
If some of your wireless clients do not support AES, select TKIP.
- 4 In *Pre-shared Key*, enter a key between 8 and 63 characters long.
- 5 Select *OK*.

To configure WPA-Personal security - CLI

```
config wireless-controller vap
edit example_wlan
set security wpa-personal
set passphrase "hardtoguess"
set encrypt AES
end
```

WPA-Enterprise security

If you will use FortiOS user groups for authentication, go to *User > User Group* and create those groups first. The groups should be Firewall groups.

If you will use a RADIUS server to authenticate wireless clients, you must first configure the FortiGate unit to access the RADIUS server.

To configure FortiGate unit access to the RADIUS server - web-based manager

- 1 Go to *User > Remote > RADIUS* and select *Create New*.
- 2 Enter a *Name* for the server.
- 3 In *Primary Server Name/IP*, enter the network name or IP address for the server.
- 4 In *Primary Server Secret*, enter the shared secret used to access the server.
- 5 Optionally, enter the information for a secondary or backup RADIUS server.
- 6 Select *OK*.

To configure the FortiGate unit to access the RADIUS server - CLI

```
config user radius
  edit exampleRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret aoewmntiasf
  end
```

To configure WPA-Enterprise security - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
- 2 In *Security Mode*, select *WPA/WPA2-Enterprise*.
- 3 In *Data Encryption*, select *AES*.
If some of your wireless clients do not support AES, select *TKIP*.
- 4 In *Authentication*, do one of the following:
 - If you will use a RADIUS server for authentication, select *RADIUS Server* and then select the RADIUS server.
 - If you will use a local user group for authentication, select *Usergroup* and then select the user group that is permitted to use the wireless network.
- 5 Select *OK*.

To configure WPA-Enterprise security - CLI

```
config wireless-controller vap
  edit example_wlan
    set security wpa-enterprise
    set encrypt AES
    set auth radius
    set radius-server exampleRADIUS
  end
```

Captive Portal security

Captive Portal security provides an access point that initially appears open. The wireless client can connect to the AP with no security credentials. The AP responds to the client's first HTTP request with a web page requesting user name and password. Until the user enters valid credentials, no communication beyond the AP is permitted.

The wireless controller authenticates users through the FortiGate user accounts. In the SSID configuration, you select the user groups that are permitted access through the captive portal.

The captive portal contains the following web pages:

- **Login page**—requests user credentials

- **Login failed page**—reports that the entered credentials were incorrect and enables the user to try again.
- **Disclaimer page**—is statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding.
- **Declined disclaimer page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting *Customize Portal Messages*. Each SSID can have its own unique portal content.

To configure Captive Portal security - web-based manager

- 1 Configure user groups as needed in *User > User Group*.
- 2 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID entry.
- 3 In *Security Mode*, select *Captive Portal*.
- 4 Optionally, select *Customize Portal Messages* and modify the portal pages that users of this SSID will see.
- 5 In *User Groups*, select the group(s) that are allowed to use the wireless network and move them to the *Selected* list.
- 6 Select *OK*.

Adding a MAC filter

On each SSID, you can create a MAC address filter list to either permit or exclude a list of clients identified by their MAC addresses.

This is actually not as secure as it appears. Someone seeking unauthorized access to your network can obtain MAC addresses from wireless traffic and use them to impersonate legitimate users. A MAC filter list should only be used in conjunction with other security measures such as encryption.

To configure a MAC filter list, you must use the CLI.

To configure a MAC filter - CLI

In this example, the MAC addresses 11:11:11:11:11:11 and 12:12:12:12:12:12 will be excluded from the example_wlan wireless interface.

```
config wireless-controller vap
edit example_wlan
config mac-filter-list
edit 1
set mac 11:11:11:11:11:11
set mac-filter-policy deny
edit 2
set mac 12:12:12:12:12:12
set mac-filter-policy deny
end
end
```

Configuring user authentication

You can perform user authentication when the wireless client joins the wireless network and when the wireless user communicates with another network through a firewall policy. WEP and WPA-Personal security rely on legitimate users knowing the correct key or passphrase for the wireless network. The more users you have, the more likely it is that the key or passphrase will become known to unauthorized people. WPA-Enterprise and captive portal security provide separate credentials for each user. User accounts can be managed through FortiGate user groups or an external RADIUS authentication server.

WPA-Enterprise authentication

If your WiFi network uses WPA-Enterprise authentication verified by a RADIUS server, you need to configure the FortiGate unit to connect to that RADIUS server.

Configuring connection to a RADIUS server - web-based manager

- 1 Go to *User > Remote > RADIUS* and select *Create New*.
- 2 Enter a *Name* for the server.
This name is used in FortiGate configurations. It is not the actual name of the server.
- 3 In *Primary Server Name/IP*, enter the network name or IP address for the server.
- 4 In *Primary Server Secret*, enter the shared secret used to access the server.
- 5 Optionally, enter the information for a secondary or backup RADIUS server.
- 6 Select *OK*.

To configure the FortiGate unit to access the RADIUS server - CLI

```
config user radius
  edit exampleRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret aoewmntiasf
  end
```

To implement WPA-Enterprise security, you select this server in the SSID security settings. See [“Configuring security” on page 2450](#).

To use the RADIUS server for authentication, you can create individual FortiGate user accounts that specify the authentication server instead of a password, and you then add those accounts to a user group. Or, you can add the authentication server to a FortiGate user group, making all accounts on that server members of the user group.

Creating a wireless user group

Most wireless networks require authenticated access. To enable creation of identity-based firewall policies, you should create at least one user group for your wireless users. You can add or remove users later. There are two types of user group to consider:

- A Firewall user group can contain user accounts stored on the FortiGate unit or external authentication servers such as RADIUS that contain and verify user credentials.
- A Directory Services user group is used for integration with Windows Active Directory or Novell eDirectory. The group can contain Windows or Novell user groups who will be permitted access to the wireless LAN. Fortinet Single Sign On (FSSO) agent must be installed on the network.

Configuring firewall policies for the SSID

For users on the WiFi LAN to communicate with other networks, firewall policies are required. Before you create firewall policies, you need to define any firewall addresses you will need. This section describes creating a WiFi network to Internet policy.

To create a firewall address for WiFi users - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information and select *OK*.

Address Name	Enter a name for the address, wifi_net for example.
Type	Select <i>Subnet / IP Range</i>
Subnet / IP Range	Enter the subnet address, 10.10.110.0/24 for example.
Interface	Select the interface where this address is used, e.g., example_wifi

To create a firewall address for WiFi users - CLI

```
config firewall address
  edit "wifi_net"
    set associated-interface "example_wifi"
    set subnet 10.10.110.0 255.255.255.0
  end
```

To create a firewall policy - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.
- 2 In *Source Interface/Zone*, select the wireless interface.
- 3 In *Source Address*, select the address of your WiFi network, wifi_net for example.
- 4 In *Destination Interface/Zone*, select the Internet interface, for example, port1.
- 5 In *Destination Address*, select *All*.
- 6 In *Service*, select *ANY*, or select the particular services that you want to allow, and then select the right arrow button to move the service to the *Selected Services* list.
- 7 In *Schedule*, select *Always*, unless you want to define a schedule for limited hours.
- 8 In *Action*, select *ACCEPT*.
- 9 Select *Enable NAT*.
- 10 Optionally, select *UTM* and set up UTM features for wireless users.
- 11 Select *OK*.

To create a firewall policy - CLI

```
config firewall policy
  edit 0
    set srcintf "example_wifi"
    set dstintf "port1"
    set srcaddr "wifi_net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
```

```
set nat enable
end
```

Customizing captive portal pages

If you select Captive Portal authentication in the SSID, the wireless controller presents to the user pages defined in Captive Portal Default replacement pages.

The captive portal contains the following web pages:

- **Login page**—requests user credentials
- **Login failed page**—reports that the entered credentials were incorrect and enables the user to try again.
- **Disclaimer page**—is statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding.
- **Declined disclaimer page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting *Customize Portal Messages*. Each SSID can have its own unique portal content.



Only the Login page and Login failed page are available in FortiOS 4.3.0. If you require a disclaimer, consider using the User Authentication disclaimer that can be enabled in the firewall policy.

Modifying the login page

The login page requests the user's credentials. Typical modifications for this page would be to change the logo and modify some of the text.

Figure 253: Default captive portal login page

FORTINET®

SSID "example_guest" Authentication Required

Please enter your username and password to continue.

Username:

Password:

Changing the logo

You can replace the default Fortinet logo with your organization's logo. First, import the logo file into the FortiGate unit and then modify the Login page code to reference your file.

To import a logo file

- 1 Go to *System > Config > Replacement Message* and select *Manage Images*.
- 2 Select *Create New*.
- 3 Enter a *Name* for the logo and select the appropriate *Content Type*.
The file must not exceed 6000 bytes.
- 4 Select *Browse*, find your logo file and then select *Open*.
- 5 Select *OK*.

To specify the new logo in the replacement message

- 1 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID.
The SSID *Security Mode* must be *Captive Portal*.
- 2 Make sure that *Customize Portal Messages* is selected and then select the adjacent *Edit* icon.
- 3 In the *Edit Message* window, select the *Login page* message.
- 4 In the Message HTML, find the `%%IMAGE` tag.
By default it specifies the Fortinet logo: `%%IMAGE:logo_fw_auth%%`
- 5 Change the image name to the one you provided for your logo.
The tag should now read, for example, `%%IMAGE:mylogo%%`
- 6 Select *OK*.

Modifying text

You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. There are two exceptions to this rule:

- The line “Please enter your username and password to continue” is provided by the `%%QUESTION%%` tag. You can replace this tag with text of your choice.
- The line “SSID ... Authentication Required” includes the name of the SSID, provided by the `%%CPAUTH_SSID%%` tag. You can remove or change the position of this tag.

Except for these items, you should not remove any tags because they may carry information that the FortiGate unit needs.

To modify login page text

- 1 Go to *WiFi Controller > WiFi Network > SSID* and edit your SSID.
The SSID *Security Mode* must be *Captive Portal*.
- 2 Make sure that *Customize Portal Messages* is selected and then select the adjacent *Edit* icon.
- 3 In the *Edit Message* window, select the *Login page* message.
- 4 In the Message HTML box, edit the text, then select *OK*.
- 5 Select *OK*.

Modifying the login failed page

The Login failed page is similar to the Login page. It even contains the same login form. You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. There are two exceptions to this rule:

- The line “Firewall authentication failed. Please try again.” is provided by the `%%FAILED_MESSAGE%%` tag. You can replace this tag with text of your choice.

- The line “SSID ... Authentication Required” includes the name of the SSID, provided by the %%CPAUTH_SSID%% tag. You can remove or change the position of this tag.

Except for these items, you should not remove any tags because they may carry information that the FortiGate unit needs.

Figure 254: Default login failed page

Configuring the built-in access point on a FortiWiFi unit

Both FortiGate and FortiWiFi units have the WiFi controller feature. If you configure a WiFi network on a FortiWiFi unit, you can also use the built-in wireless capabilities in your WiFi network as one of the access points.

If Virtual Domains are enabled, you must select the VDOM to which the built-in access point belongs. You do this in the CLI. For example:

```
config wireless-controller global
  set local-radio-vdom vdom1
end
```

To configure the FortiWiFi unit's built-in WiFi access point

- 1 Go to *WiFi Controller > Managed Access Points > Local WiFi Radio*.
- 2 Make sure that *AP Profile* is *Automatic*.
- 3 Make sure that *Enable WiFi Radio* is selected.
- 4 In *SSID*, if you do not want this AP to carry all SSIDs, select *Select SSIDs* and then select the required SSIDs.
- 5 Optionally, adjust the *TX Power* slider.
If you have selected your location correctly (see [“Setting your geographic location” on page 2447](#)), the 100% setting corresponds to the maximum power allowed in your region.
- 6 If you do not want the built-in WiFi radio to be used for rogue scanning, select *Do not participate in Rogue AP scanning*.
- 7 Select *OK*.

If you want to connect external APs, such as FortiAP units, see the next chapter, [“Access point deployment”](#).



Access point deployment

This chapter describes how to configure access points for your wireless network. The following topics are included in this section:

- [Overview](#)
- [Network topology for managed APs](#)
- [Discovering and authorizing APs](#)
- [Advanced WiFi controller discovery](#)

Overview

FortiAP units discover WiFi controllers. The administrator of the WiFi controller authorizes the FortiAP units that the controller will manage.

In most cases, FortiAP units can find WiFi controllers through the wired Ethernet without any special configuration. Review the following section, “[Network topology for managed APs](#)”, to make sure that your method of connecting the FortiAP unit to the WiFi controller is valid. Then, you are ready to follow the procedures in “[Discovering and authorizing APs](#)” on page 2460.

If your FortiAP units are unable to find the WiFi controller, refer to “[Advanced WiFi controller discovery](#)” on page 2464 for detailed information about the FortiAP unit’s controller discovery methods and how you can configure them.

Network topology for managed APs

The FortiAP unit can be connected to the FortiGate unit in any of the following ways:

Direct connection: The FortiAP unit is directly connected to the FortiGate unit with no switches between them. This configuration is common for locations where the number of FortiAP’s matches up with the number of ‘internal’ ports available on the FortiGate. In this configuration the FortiAP unit requests an IP address from the FortiGate unit, enters discovery mode and should quickly find the FortiGate WiFi controller. This is also known as a wirecloset deployment. See [Figure 255](#), below.

Switched Connection: The FortiAP unit is connected to the FortiGate WiFi controller by an Ethernet switch operating in L2 switching mode or L3 routing mode. There must be a routable path between the FortiAP unit and the FortiGate unit and ports 5246 and 5247 must be open. This is also known as a gateway deployment. See [Figure 255](#), below

Connection over WAN: The FortiGate WiFi controller is off-premises and connected by a VPN tunnel to a local FortiGate. In this method of connectivity its best to configure each FortiAP with the static IP address of the WiFi controller. Each FortiAP can be configured with three WiFi controller IP addresses for redundant failover. This is also known as a datacenter remote management deployment. See [Figure 256](#), below.

Figure 255: Wirecloset and Gateway deployments

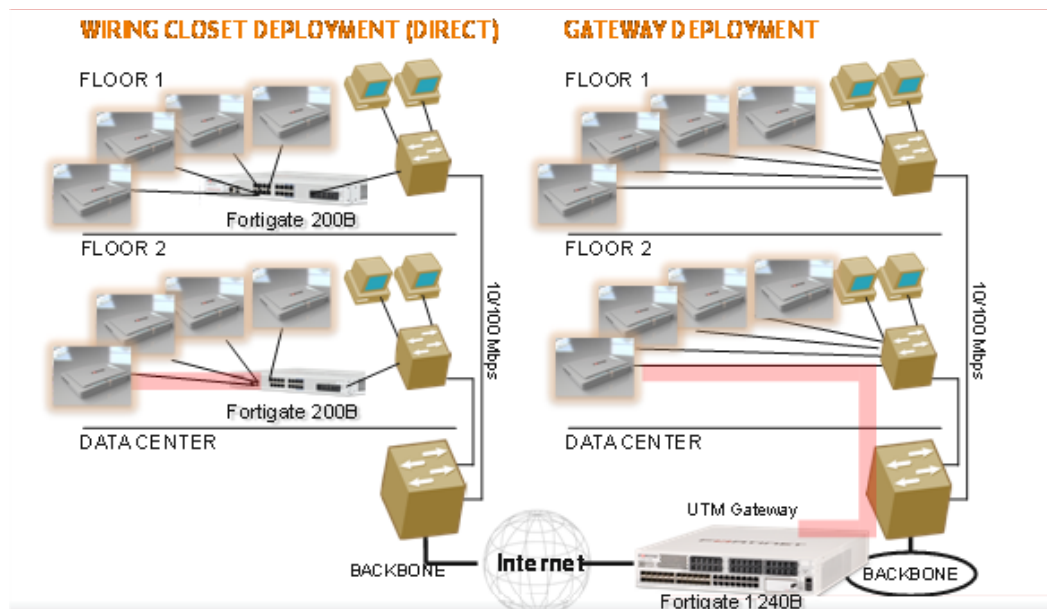
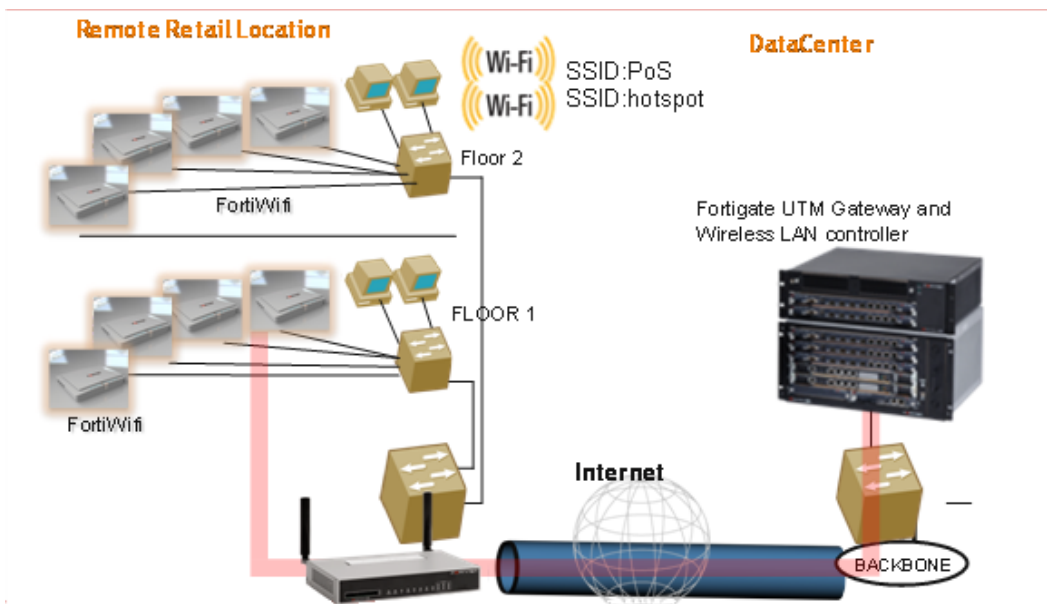


Figure 256: Remote deployment



Discovering and authorizing APs

After you prepare your FortiGate unit, you can connect your APs to discover them using the discovery methods described earlier. To prepare the FortiGate unit, you need to

- Configure the network interface to which the AP will connect.
- Configure DHCP service on the interface to which the AP will connect.
- Connect the AP units and let the FortiGate unit discover them.
- Enable each discovered AP and configure it or assign it to an AP profile.

Configuring the network interface for the AP unit

The interface to which you connect your wireless access point needs an IP address. No administrative access, DNS Query service or authentication should be enabled.

To configure the interface for the AP unit - web-based manager

- 1 Go to *System > Network > Interface* and edit the interface to which the AP unit connects.
- 2 Set *Addressing Mode* to *Manual* and enter the IP address and netmask to use.
- 3 Enable *Dedicate this interface to FortiAP connection* and set *Reserve IP addresses for FortiAP* to the range of addresses that you want to use for FortiAP units.

The address range needs to be in the same subnet as the interface IP address, but not include that address. This step automatically configures a DHCP server for the AP units.

- 4 Select OK.

To configure the interface for the AP unit - CLI

```
config system interface
  edit port3
    set mode static
    set ip 192.168.8.1 255.255.255.0
  end
```

To configure the DHCP server for AP unit - CLI

```
config system dhcp server
  edit 0
    set default-gateway 192.168.8.1
    set interface wan2
    config ip-range
      edit 1
        set end-ip 192.168.8.9
        set start-ip 192.168.8.2
      end
    set netmask 255.255.255.0
  end
```

Enabling a discovered AP

Within two minutes of connecting the AP unit to the FortiGate unit, the discovered unit should be listed on *WiFi Controller > Managed Access Points > Managed FortiAP* page.

Figure 257: Discovered access point unit

Create New Edit Delete Refresh Reset All						
<input type="checkbox"/>	Admin	Name	AP Profile	Clients	Join Time	Reset
<input type="checkbox"/>		FAP22A3U10600118		0	06/24/10 14:54	

To add the discovered AP unit - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Select the FortiAP unit from the list and edit it.
- 3 Optionally, enter a *Name*. Otherwise, the unit will be identified by serial number.
- 4 Select *Authorize*.

5 Select OK.

The physical access point is now added to the system. If the rest of the configuration is complete, it should be possible to connect to the wireless network through the AP.

To add the discovered AP unit - CLI

First get a list of the discovered access point unit serial numbers:

```
get wireless-controller wtp
```

Add a discovered unit and associate it with AP-profile1, for example:

```
config wireless-controller wtp
edit FAP22A3U10600118
set admin enable
set wtp-profile AP-profile1
end
```

To use the automatic profile, leave the `wtp-profile` field unset.

To view the status of the added AP unit

```
config wireless-controller wtp
edit FAP22A3U10600118
get
```

The `join-time` field should show a time, not “N/A”. See the preceding web-based manager procedure for more information.

Configuring a managed AP

When you add a FortiAP unit, it is configured by default to

- use the Automatic profile
- operate at the maximum radio power permitted in your region
- carry all SSIDs

You can change the radio power and selection of SSIDs or assign the unit to a custom AP profile which defines the entire configuration for the AP.

To modify settings within Automatic profile - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Select the FortiAP unit from the list and edit it.
AP Profile should be *Automatic*.
- 3 Make sure that *Enable WiFi Radio* is selected.
- 4 In *SSID*, if you do not want this AP to carry all SSIDs, select *Select SSIDs* and then select the required SSIDs.
- 5 Optionally, adjust the *TX Power* slider.

If you have selected your location correctly (see [“Setting your geographic location” on page 2447](#)), the 100% setting corresponds to the maximum power allowed in your region.

- 6 Select OK.

To modify settings within Automatic profile - CLI

When `wtp-profile` is unset (null value), the Automatic profile is in use and some of its settings can be adjusted. This example sets the AP to carry only the employee and guest SSIDs and operate at 80% of maximum power.

```
config wireless-controller wtp
edit FAP22A3U10600118
set radio-enable enable
set vap-all disable
set vaps employee guest
set power-level 80
end
```

To select a custom AP profile - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Select the FortiAP unit from the list and edit it.
- 3 In *AP Profile*, select the custom AP Profile to use, and then select *Apply*.
Only AP Profiles that are appropriate for this AP unit are available.
- 4 Select *OK*.

To select a custom AP profile - CLI

```
config wireless-controller wtp
edit FAP22A3U10600118
set wtp-profile AP-profile1
end
```

To select automatic AP profile - CLI

```
config wireless-controller wtp
edit FAP22A3U10600118
unset wtp-profile
end
```

Updating FortiAP unit firmware

You can update the FortiAP unit's firmware from the FortiGate unit that acts as its WiFi controller.

Updating FortiAP firmware from the FortiGate unit

You can update the FortiAP firmware using either the web-based manager or the CLI. Only the CLI method can update all FortiAP units at once.

To update FortiAP unit firmware - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Select the FortiAP unit from the list and edit it.
- 3 In *FortiAP OS Version*, select *[Upgrade]*.
- 4 Select *Browse* and locate the firmware upgrade file.
- 5 Select *OK*.
- 6 When the upgrade process completes, select *OK*.
The FortiAP unit restarts.

To update FortiAP unit firmware - CLI

- 1 Upload the FortiAP image to the FortiGate unit.

For example, the Firmware file is `FAP_22A_v4.3.0_b0212_fortinet.out` and the server IP address is `192.168.0.100`.

```
execute wireless-controller upload-wtp-image tftp
FAP_22A_v4.3.0_b0212_fortinet.out 192.168.0.100
```

If your server is FTP, change `tftp` to `ftp`, and if necessary add your user name and password at the end of the command.

- 2 Verify that the image is uploaded:

```
execute wireless-controller list-wtp-image
```

- 3 Upgrade the FortiAP units:

```
exec wireless-controller reset-wtp all
```

If you want to upgrade only one FortiAP unit, enter its serial number instead of `all`.

Updating FortiAP firmware from the FortiAP unit

You can connect to a FortiAP unit's internal CLI to update its firmware from a TFTP server on the same network. This method does not require access to the wireless controller.

- 1 Place the FortiAP firmware image on a TFTP server on your computer.
- 2 Connect the FortiAP unit to a separate private switch or hub or directly connect to your computer via a cross-over cable.
- 3 Change your computer's IP address to 192.168.1.3.
- 4 Telnet to IP address 192.168.1.2.
This IP address is overwritten if the FortiAP is connected to a DHCP environment. Ensure that the FortiAP unit is in a private network with no DHCP server.
- 5 Login with the username "admin" and no password.
- 6 Enter the following command.

For example, the FortiAP image file name is `FAP_22A_v4.3.0_b0212_fortinet.out`.

```
restore FAP_22A_v4.3.0_b0212_fortinet.out 192.168.1.3
```

Advanced WiFi controller discovery

A FortiAP unit can use any of four methods to locate a controller. By default, FortiAP units cycle through all four of the discovery methods. In most cases there is no need to make configuration changes on the FortiAP unit.

There are exceptions. The following section describes the WiFi controller discovery methods in more detail and provides information about configuration changes you might need to make so that discovery will work.

You can also configure a FortiWiFi unit to act as an AP. But in this case you must choose which discovery method it will use. See ["Configuring a FortiWiFi unit as a WiFi AP" on page 2466](#).

Controller discovery methods

There are four methods that a FortiAP unit can use to discover a WiFi controller.

Static IP configuration

If FortiAP and the controller are not in the same subnet, broadcast and multicast packets cannot reach the controller. The admin can specify the controller's static IP on the AP unit. The AP unit sends a discovery request message in unicast to the controller. Routing must be properly configured in both directions.

To specify the controller's IP address on a FortiAP unit

```
cfg -a AC_IPADDR_1="192.168.0.1"
```

By default, the FortiAP unit receives its IP address by DHCP. If you prefer, you can assign the AP unit a static IP address.

To assign a static IP address to the FortiAP unit

```
cfg -a ADDR_MODE=STATIC
cfg -a AP_IPADDR="192.168.0.100"
cfg -a AP_NETMASK="255.255.255.0"
```

For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI” on page 2466](#).

Broadcast request

The AP unit broadcasts a discovery request message to the network and the controller replies. The AP and the controller must be in the same broadcast domain. No configuration adjustments are required.

Multicast request

The AP unit sends a multicast discovery request and the controller replies with a unicast discovery response message. The AP and the controller do not need to be in the same broadcast domain if multicast routing is properly configured.

The default multicast destination address is 224.0.1.140. It can be changed through the CLI. The address must be same on the controller and AP. For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI” on page 2466](#).

To change the multicast address on the controller

```
config wireless-controller global
set discovery-mc-addr 224.0.1.250
end
```

To change the multicast address on a FortiAP unit

```
cfg -a AC_DISCOVERY_MC_ADDR="224.0.1.250"
```

For information about connecting to the FortiAP CLI, see [“Connecting to the FortiAP CLI” on page 2466](#).

DHCP

If you use DHCP to assign an IP address to your FortiAP unit, you can also provide the WiFi controller IP address at the same time. This is useful if the AP is located remotely from the WiFi controller and other discovery techniques will not work.

When you configure the DHCP server, configure Option 138 to specify the WiFi controller IP address. You need to convert the address into hexadecimal. Convert each octet value separately from left to right and concatenate them. For example, 192.168.0.1 converts to C0A80001.

If Option 138 is used for some other purpose on your network, you can use a different option number if you configure the AP units to match.

To change the FortiAP DHCP option code

To use option code 139 for example, enter

```
cfg -a AC_DISCOVERY_DHCP_OPTION_CODE=139
```

For information about connecting to the FortiAP CLI, see “[Connecting to the FortiAP CLI](#)” below.

Connecting to the FortiAP CLI

The FortiAP unit has a CLI through which some configuration options can be set.

To access the FortiAP unit CLI

- 1 Connect your computer to the FortiAP directly with a cross-over cable or through a separate switch or hub.
- 2 Change your computer's IP address to 192.168.1.3
- 3 Telnet to IP address 192.168.1.2.
Ensure that FortiAP is in a private network with no DHCP server for the static IP address to be accessible.
- 4 Login with user name admin and no password.
- 5 Enter commands as needed.
- 6 Optionally, use the `passwd` command to assign an administrative password for better security.
- 7 Save the configuration by entering the following command:

```
cfg -c .
```
- 8 Unplug the FortiAP and then plug it back in, in order for the configuration to take effect.



When a WiFi controller has taken control of the FortiAP unit, Telnet access to the FortiAP unit's CLI is no longer available.

Configuring a FortiWiFi unit as a WiFi AP

FortiWiFi units can also be deployed as managed APs controlled by a FortiGate unit wireless controller.

In the CLI, enter

```
config system global
  set wireless-mode wtp
end
```

Setting the discovery mode

Unlike FortiAP units, a FortiWiFi unit deployed as an AP does not cycle through the discovery methods. You must select one discovery method to use.

To select DHCP discovery

```
config wireless-controller global
  set ac-discovery-type dhcp
end
```

The DHCP discovery method is the simplest to use and will work when the AP is connected directly to the WiFi controller unit.

To select multicast discovery


In this example, the FortiWiFi AP is configured for multicast discovery and its multicast address is changed:

```
config wireless-controller global
  set ac-discovery-type multicast
  set discovery-mc-addr 224.0.1.250
end
```

Discovery by multicast will work even when the FortiWiFi AP is not in the same domain as the WiFi controller.

Completing configuration

The rest of the configuration is located in `config wireless-controller` and is similar to the FortiGate WiFi controller configuration.



Wireless network monitoring

You can monitor both your wireless clients and other wireless networks that are available in your coverage area.

The following topics are included in this section:

- [Monitoring wireless clients](#)
- [Monitoring rogue APs](#)
- [Suppressing rogue APs](#)

Monitoring wireless clients

To view connected clients on a FortiWiFi unit

- Go to *WiFi Controller > Monitor > Client Monitor*.

The following information can be displayed, depending on the *Column Settings* you have selected.

Association Time	How long the client has been connected to this access point.
Auth	The type of authentication used.
Bandwidth Rx	Received bandwidth used by the client, in Kbps.
Bandwidth Tx	Transmit bandwidth used by the client, in Kbps.
Bandwidth Tx/Rx	<i>Bandwidth Rx + Bandwidth Tx</i> .
FortiAP	The serial number of the FortiAP unit to which the client connected.
Idle Time	The total time this session that the client was idle.
IP	The IP address assigned to the wireless client.
MAC	The MAC address of the wireless client.
Manufacturer	Manufacturer of the client wireless device.
Physical AP	The name of the physical access point with which the client is associated.
Rate	The data rate that the wireless connection can support.
Signal Strength / Noise	The signal-to-noise ratio in deciBels calculated from signal strength and noise level.
SSID	The SSID that the client connected to.
Virtual AP	The name of the virtual access point with which the client is associated.

Monitoring rogue APs

The access point radio equipment can scan for other available access points, either as a dedicated monitor or as a background scan performed while the access point is idle.

Discovered access points are listed in the *Rogue AP Monitor* list. You can then mark them as either Accepted or Rogue access points. This designation helps you to track access points. It does not affect anyone's ability to use these access points.

It is also possible to suppress rogue APs. See [“Suppressing rogue APs” on page 2473](#).

On-wire rogue AP detection technique

Other APs that are available in the same area as your own APs are not necessarily rogues. A neighboring AP that has no connection to your network might cause interference, but it is not a security threat. A rogue AP is an unauthorized AP connected to your wired network. This can enable unauthorized access. When rogue AP detection is enabled, the *On-wire* column in the *Rogue AP Monitor* list shows a green up-arrow on detected rogues.

Rogue AP monitoring of WiFi client traffic builds a table of WiFi clients and the Access Points that they are communicating through. The FortiGate unit also builds a table of MAC addresses that it sees on the LAN. The FortiGate unit's on-wire correlation engine constantly compares the MAC addresses seen on the LAN to the MAC addresses seen on the WiFi network.

There are two methods of Rogue AP on-wire detection operating simultaneously: Exact MAC address match and MAC adjacency.

Exact MAC address match

If the same MAC address is seen on the LAN and on the WiFi network, this means that the wireless client is connected to the LAN. If the AP that the client is using is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue. This scheme works for non-NAT rogue APs.

MAC adjacency

If an access point is also a router, it applies NAT to WiFi packets. This can make rogue detection more difficult. However, an AP's WiFi interface MAC address is usually in the same range as its wired MAC address. So, the MAC adjacency rogue detection method matches LAN and WiFi network MAC addresses that are within a defined numerical distance of each other. By default, the MAC adjacency value is 7. If the AP for these matching MAC addresses is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue.

Limitations

On-wire rogue detection has some limitations. There must be at least one WiFi client connected to the suspect AP and continuously sending traffic. If the suspect AP is a router, its WiFi MAC address must be very similar to its Ethernet port MAC address.

Logging

Information about detected rogue APs is logged and uploaded to your FortiAnalyzer unit, if you have one. By default, rogue APs generate an alert level log, unknown APs generate a warning level log. This log information can help you with PCI-DSS compliance requirements.

Rogue AP scanning as a background activity

Each WiFi radio can perform monitoring of radio channels in its operating band while acting as an AP. It does this by briefly switching from AP to monitoring mode. By default, a scan period starts every 300 seconds. Each second a different channel is monitored for 20ms until all channels have been checked.

During heavy AP traffic, it is possible for background scanning to cause lost packets when the radio switches to monitoring. To reduce the probability of lost packets, you can set the CLI `ap-bgscan-idle` field to delay the switch to monitoring until the AP has been idle for a specified period. This means that heavy AP traffic may slow background scanning.

The following CLI example configures default background rogue scanning operation except that it sets `ap-bgscan-idle` to require 100ms of AP inactivity before scanning the next channel.

```
config wireless-controller wtp-profile
edit ourprofile
config radio-1
set ap-bgscan enable
set rogue-scan enable
set ap-bgscan-period 300
set ap-bgscan-intv 1
set ap-bgscan-duration 20
set ap-bgscan-idle 100
end
end
```

Configuring rogue scanning

Rogue scanning is easily enabled for all of your APs.

To enable the rogue AP scanning feature - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > Rogue AP Settings*.
- 2 Select *Enable Rogue AP Detection*.
- 3 Select *Enable On-wire Rogue AP Detection Technique* if you want to use that method of distinguishing rogues from neighbors.
- 4 Select *Apply*.

To enable the rogue AP scanning feature - CLI

```
config wireless-controller setting
set ap-scan enable
set on-wire-scan enable
end
```

To adjust MAC adjacency

You can adjust the maximum WiFi to Ethernet MAC difference used when determining whether an suspect AP is a rogue. For example, to change the adjacency to 8, enter

```
config wireless-controller global
set rogue-scan-mac-adjacency 8
end
```

Exempting an AP from rogue scanning

By default, if Rogue AP Detection is enabled, it is enabled on all managed FortiAP units. Optionally, you can exempt an AP from scanning. You should be careful about doing this if your organization must perform scanning to meet PCI-DSS requirements.

To exempt an AP from rogue scanning - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Select which AP to edit.
- 3 Select *Do not participate in Rogue AP Scanning* and then select *OK*.

To exempt an AP from rogue scanning - CLI

This example shows how to exempt access point AP1 from rogue scanning.




```
config wireless-controller wtp
  edit AP1
    set ap-scan disable
  end
```

Using the Rogue AP Monitor

Go to *WiFi Controller > Monitor > Rogue AP Monitor* to view the list of other wireless access points that are receivable at your location. Available information about the APs includes:

SSID	channel
security type	signal strength
detected by which of your APs	MAC address
AP equipment vendor	on-wire status
time first seen	time last seen

The status of newly detected APs is Unclassified. You can manually change the status using the *Mark* menu.

	Rogue AP. Use this status for unauthorized APs attached to your wired networks. The On-wire detection technique determines which unknown APs are rogues.
	Accepted AP. Use this status for APs that are an authorized part of your network or are neighboring APs that are not a security threat.
	Unclassified AP. This is the initial status of a discovered AP. You can change an AP back to unclassified if you have mistakenly marked it as Rogue or Accepted.

Suppressing rogue APs

In addition to monitoring rogue APs, you can actively prevent your users from connecting to them. When suppression is activated against an AP, the FortiGate WiFi controller sends deauthentication messages to the rogue AP's clients, posing as the rogue AP, and also sends deauthentication messages to the rogue AP, posing as its clients. This is done using the monitoring radio.


To enable rogue AP suppression, you must enable monitoring of rogue APs with the on-wire detection technique. See [“Monitoring rogue APs” on page 2470](#). The monitoring radio must be in the Dedicated Monitor mode.

To activate AP suppression against a rogue AP

- 1 Go to *WiFi Controller > Monitor > Rogue AP Monitor*.
- 2 When you see an AP listed that is a rogue detected “on-wire”, select it and then select *Mark > Mark Rogue*.
- 3 To suppress an AP that is marked as a rogue, select it and then select *Suppress AP*.

To deactivate AP suppression

- 1 Go to *WiFi Controller > Monitor > Rogue AP Monitor*.
- 2 Select the suppressed rogue AP and then select *Suppress AP > Unsuppress AP*.



Configuring wireless network clients

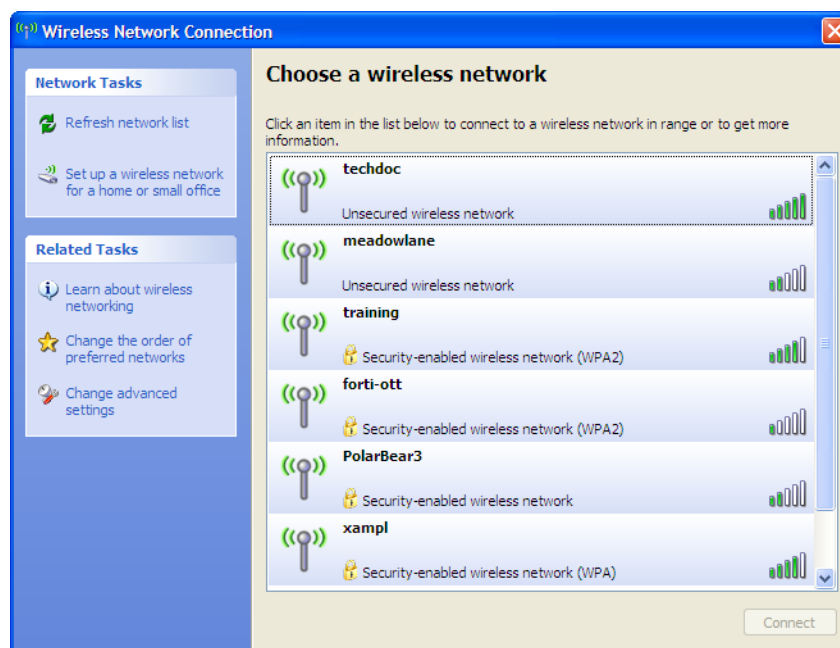
This chapter shows how to configure typical wireless network clients to connect to a wireless network with WPA-Enterprise security. The following topics are included in this section:

- [Windows XP client](#)
- [Windows 7 client](#)
- [Mac OS client](#)
- [Linux client](#)
- [Troubleshooting](#)

Windows XP client

To configure the WPA-Enterprise network connection

- 1 In the Windows Start menu, go to *Control Panel > Network Connections > Wireless Network Connection* or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.

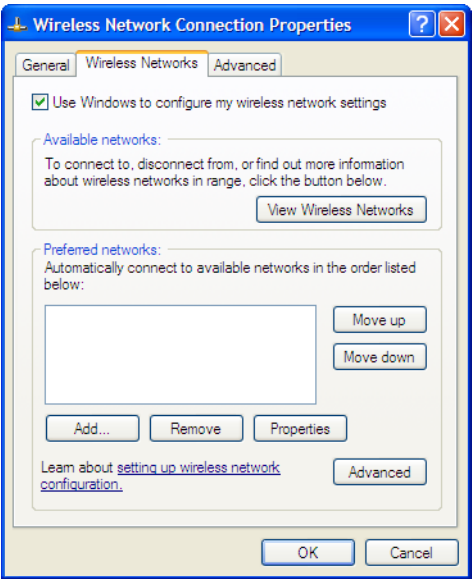


If you are already connected to another wireless network, the Connection Status window displays. Select *View Wireless Networks* on the *General* tab to view the list.

If the network broadcasts its SSID, it is listed. But do not try to connect until you have completed the configuration step below. Because the network doesn't use the Windows XP default security configuration, configure the client's network settings manually before trying to connect.

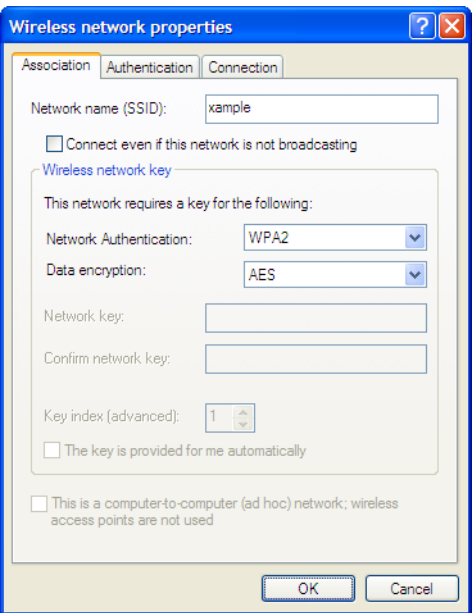
- 2 You can configure the WPA-Enterprise network to be accessible from the *View Wireless Networks* window even if it does not broadcast its SSID.

- 3 Select *Change Advanced Settings* and then select the *Wireless Networks* tab.



Any existing networks that you have already configured are listed in the *Preferred Networks* list.

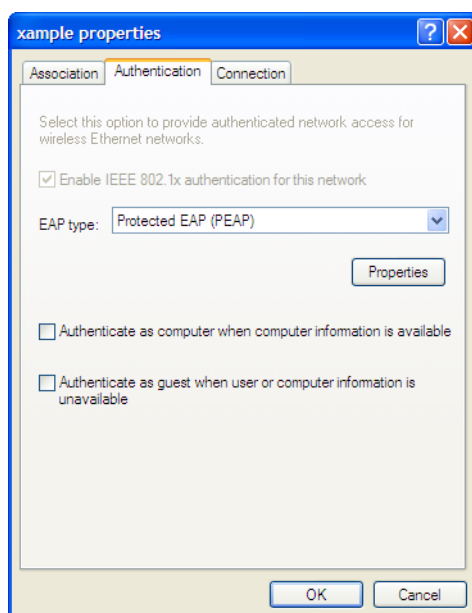
- 4 Select *Add* and enter the following information:



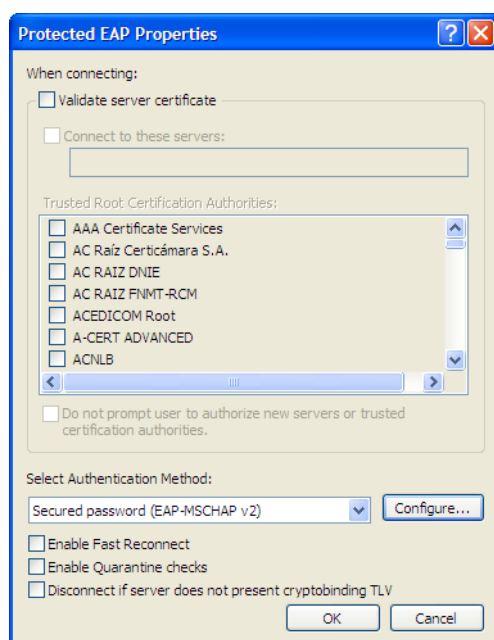
Network Name (SSID)	The SSID for your wireless network
Network Authentication	WPA2
Data Encryption	AES

- 5 If this wireless network does not broadcast its SSID, select *Connect even if this network is not broadcasting* so that the network will appear in the *View Wireless Networks* list.


- 6 Select the *Authentication* tab.



- 7 In *EAP Type*, select *Protected EAP (PEAP)*.
 8 Make sure that the other two authentication options are not selected.
 9 Select *Properties*.



- 10 Make sure that *Validate server_certificate* is selected.
 11 Select the server certificate *UTN-USERFirst-Hardware*.
 12 In *Select Authentication Method*, select *Secured Password (EAP-MSCHAPv2)*.
 13 Ensure that the remaining options are not selected.

14 Select Configure.

15 If your wireless network credentials are the same as your Windows logon credentials, select *Automatically use my Windows logon name and password*. Otherwise, make sure that this option is not selected.


16 Select OK. Repeat until you have closed all of the *Wireless Network Connection Properties* windows.

To connect to the WPA-Enterprise wireless network

- 1** Select the wireless network icon in the Notification area of the Taskbar.
- 2** In the *View Wireless Networks* list, select the network you just added and then select *Connect*.

You might need to log off of your current wireless network and refresh the list.

- 3** When the following popup displays, click on it.




- 4** In the *Enter Credentials* window, enter your wireless network *User name*, *Password*, and *Logon domain* (if applicable). Then, select OK.



In future, Windows will automatically send your credentials when you log on to this network.

Windows 7 client

- 1 In the Windows Start menu, go to *Control Panel > Network and Internet > Network and Sharing Center > Manage Wireless Networks* or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.




- 2 Do one of the following:
 - If the wireless network is listed (it broadcasts its SSID), select it from the list.
 - Select *Add > Manually create a network profile*.
- 3 Enter the following information and select *Next*.

Network name	Enter the SSID of the wireless network. (Required only if you selected <i>Add</i> .)
Security type	WPA2-Enterprise
Encryption type	AES
Start this connection automatically	Select
Connect even if the network is not broadcasting.	Select

The Wireless Network icon will display a popup requesting that you click to enter credentials for the network. Click on the popup notification.

- 4 In the *Enter Credentials* window, enter your wireless network *User name*, *Password*, and *Logon domain* (if applicable). Then, select *OK*.
- 5 Select *Change connection settings*.

- 6 On the *Connection* tab, select *Connect automatically when this network is in range*.
- 7 On the *Security* tab, select the Microsoft PEAP authentication method and then select *Settings*.
- 8 Make sure that *Validate server_certificate* is selected.
- 9 Select the server certificate *UTN-USERFirst-Hardware*.
- 10 In *Select Authentication Method*, select *Secured Password (EAP-MSCHAPv2)*.
- 11 Select *Configure*.

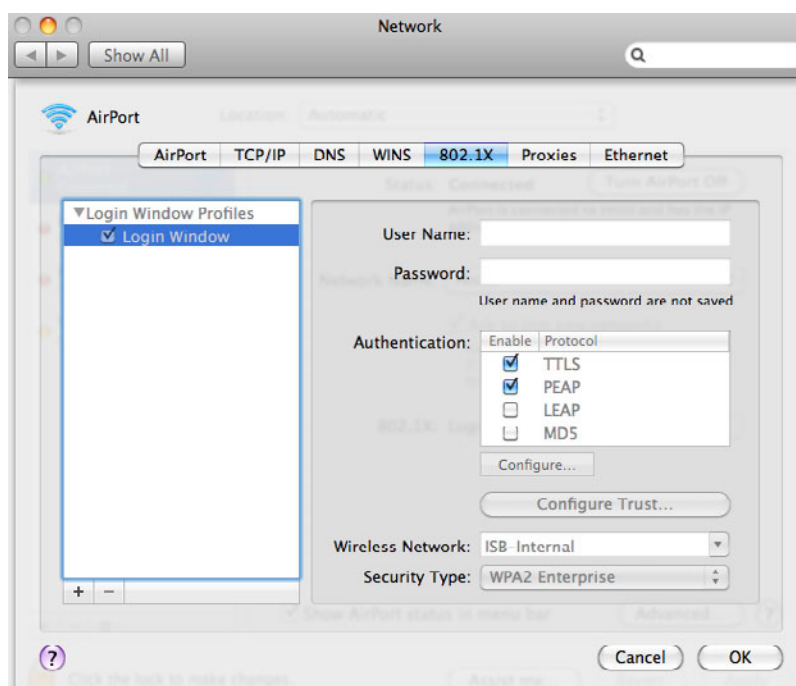


- 12 If your wireless network credentials are the same as your Windows logon credentials, select *Automatically use my Windows logon name and password*. Otherwise, make sure that this option is not selected.
- 13 Ensure that the remaining options are not selected.
- 14 Select OK. Repeat until you have closed all of the *Wireless Network Properties* windows.

Mac OS client

To configure network preferences

- 1 Right-click the *AirPort* icon in the toolbar and select *Open Network Preferences*.
- 2 Select *Advanced* and then select the *802.1X* tab.



- 3 If there are no Login Window Profiles in the left column, select the + button and then select *Add Login Window Profile*.
- 4 Select the Login Window Profile and then make sure that both TTLS and PEAP are selected in *Authentication*.

To configure the WPA-Enterprise network connection

- 1 Select the *AirPort* icon in the toolbar.
- 2 Do one of the following:
 - If the network is listed, select the network from the list.
 - Select *Connect to Other Network*.

One of the following windows opens, depending on your selection.

- 3 Enter the following information and select *OK* or *Join*:

Network name	Enter the SSID of your wireless network. (Other network only)
Wireless Security	WPA Enterprise
802.1X	Automatic
Username Password	Enter your logon credentials for the wireless network.
Remember this network	Select.

You are connected to the wireless network.




Mac OS supports only PEAP with MSCHAPv2 authentication and therefore can authenticate only to a RADIUS server, not an LDAP or TACACS+ server.

Linux client

This example is based on the Ubuntu 10.04 Linux wireless client.

To connect to a WPA-Enterprise network

- 1 Select the Network Manager icon to view the Wireless Networks menu.



Wireless networks that broadcast their SSID are listed in the *Available* section of the menu. If the list is long, it is continued in the *More Networks* submenu.

- 2 Do one of the following:
 - Select the network from the list (also check *More Networks*).
 - Select *Connect to Hidden Wireless Network*.

One of the following windows opens, depending on your selection.



- 3 Enter the following information:

Connection	Leave as <i>New</i> . (Hidden network only)
Network name	Enter the SSID of your wireless network. (Hidden network only)
Wireless Security	WPA & WPA2 Enterprise
Authentication	Protected EAP (PEAP) for RADIUS-based authentication Tunneled TLS for TACACS+ or LDAP-based authentication
Anonymous identity	This is not required.
CA Certificate	If you want to validate the AP's certificate, select the UTN-USERFirst-Hardware root certificate. The default location for the certificate is /usr/share/ca-certificates/mozilla/.
PEAP version	Automatic (applies only to PEAP)
Inner authentication	MSCHAPv2 for RADIUS-based authentication PAP or CHAP for TACACS+ or LDAP-based authentication
Username Password	Enter your logon credentials for the wireless network.

- 4 If you did not select a CA Certificate above, you are asked to do so. Select *Ignore*.



- 5 Select *Connect*. You are connected to the wireless network.

To connect to a WPA-Enterprise network

- 1 Select the Network Manager icon to view the Wireless Networks menu.
- 2 Select the network from the list (also check *More Networks*).

If your network is not listed (but was configured), select *Connect to Hidden Wireless Network*, select your network from the Connection drop-down list, and then select *Connect*.

Troubleshooting

Using tools provided in your operating system, you can find the source of common wireless networking problems.


Checking that the client has received IP address and DNS server information

Windows XP

- 1 Double-click the network icon in the taskbar to display the *Wireless Network Connection Status* window. Check that the correct network is listed in the *Connection* section.
- 2 Select the *Support* tab.
Check that the *Address Type* is *Assigned by DHCP*. Check that the *IP Address*, *Subnet Mask*, and *Default Gateway* values are valid.
- 3 Select *Details* to view the DNS server addresses.
The listed address should be the DNS servers that were assigned to the WAP. Usually a wireless network that provides access to the private LAN is assigned the same DNS servers as the wired private LAN. A wireless network that provides guest or customer users access to the Internet is usually assigned public DNS servers.
- 4 If any of the addresses are missing, select *Repair*.
If the repair procedure doesn't correct the problem, check your network settings.

Mac OS

- 1 From the Apple menu, open *System Preferences > Network*.
- 2 Select *AirPort* and then select *Configure*.
- 3 On the *Network* page, select the *TCP/IP* tab.



- 4 If there is no IP address or the IP address starts with 169, select *Renew DHCP Lease*.
- 5 To check DNS server addresses, open a terminal window and enter the following command:

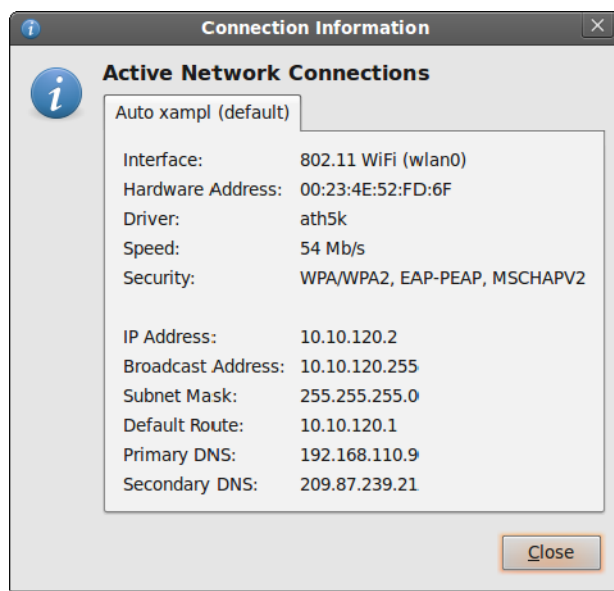
```
cat /etc/resolv.conf
```

Check the listed nameserver addresses. A network for employees should use the wired private LAN DNS server. A network for guests should specify a public DNS server.

Linux

This example is based on the Ubuntu 10.04 Linux wireless client.

- 1 Right-click the Network Manager icon and select *Connection Information*.



- 2 Check the IP address, and DNS settings. If they are incorrect, check your network settings.



Wireless network examples

This chapter provides an example wireless network configuration. The following topics are included in this section:

- [Basic wireless network](#)
- [A more complex example](#)

Basic wireless network

This example uses automatic configuration to set up a basic wireless network.

To configure this wireless network, you must:

- Configure authentication for wireless users
- Configure the SSID (WiFi network interface)
- Configure the firewall policy
- Configure and connect FortiAP units

Configuring authentication for wireless users

You need to configure user accounts and add the users to a user group. This example shows only one account, but multiple accounts can be added as user group members.

To configure a WiFi user - web-based manager

- 1 Go to *User > User* and select *Create New*.
- 2 Enter a *User Name* and *Password* and then select *OK*.

To configure the WiFi user group - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter the following information and then select *OK*:

Name	wlan_users
Type	Firewall
Allow SSL-VPN Access	Not selected.
Available Users/Groups / Members	Move users to the <i>Members</i> list.

To configure a WiFi user and the WiFi user group - CLI

```
config user user
  edit "user01"
    set type password
    set passwd "asdf12ghjk"
  end
```

```

config user group
  edit "wlan_users"
    set member "user01"
  end

```

Configuring the SSID

First, establish the SSID (network interface) for the network. This is independent of the number of physical access points that will be deployed. The network assigns IP addresses using DHCP.

To configure the SSID - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
- 2 Enter the following information and select OK:

Name		example_wifi
IP/Netmask		10.10.110.1/24
Administrative Access		Ping (to assist with testing)
SSID		example_wifi
Enable DHCP		Enable
	Address Range	10.10.110.2 - 10.10.110.199
	Netmask	255.255.255.0
	Default Gateway	Same As Interface IP
	DNS Server	Same as System DNS
Security Mode		WPA/WPA2-Enterprise
Data Encryption		AES
Authentication		Usergroup, select <i>wlan_users</i> .
Leave other settings at their default values.		

To configure the SSID - CLI

```

config wireless-controller vap
  edit example_wifi
    set ssid "example_wifi"
    set broadcast-ssid enable
    set security wpa-enterprise
    set auth usergroup
    set usergroup wlan_users
  end
config system interface
  edit example_wifi
    set ip 10.10.110.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set default-gateway 10.10.110.1
    set dns-service default
    set interface "example_wifi"

```

```

config ip-range
  edit 1
    set end-ip 10.10.110.199
    set start-ip 10.10.110.2
  end
  set netmask 255.255.255.0
end

```

Configuring firewall policies

A firewall policy is needed to enable WiFi users to access the Internet on port1. First you create firewall address for the WiFi network, then you create the example_wifi to port1 policy.

To create a firewall address for WiFi users - web-based manager

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information and select *OK*.

Address Name	wlan_user_net
Type	Subnet / IP Range
Subnet / IP Range	10.10.110.0/24
Interface	example_wifi

To create a firewall address for WiFi users - CLI

```

config firewall address
  edit "wlan_user_net"
    set associated-interface "example_wifi"
    set subnet 10.10.110.0 255.255.255.0
  end

```

To create a firewall policy for WiFi users - web-based manager

- 1 Go to *Firewall Objects > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*:

Source Interface/Zone	example_wifi
Source Address	wlan_user_net
Destination Interface/Zone	port1
Destination Address	All
Action	ACCEPT
NAT	Enable NAT. Select <i>Use Destination Interface Address</i> (default).
Leave other settings at their default values.	

To create a firewall policy for WiFi users - CLI

```
config firewall policy
edit 0
    set srcintf "example_wifi"
    set dstintf "port1"
    set srcaddr "wlan_user_net"
    set dstaddr "all"
    set action accept
    set nat enable
end
```

Connecting the FortiAP units

You need to connect each FortiAP unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

To configure the interface for the AP unit - web-based manager

- 1 Go to *System > Network > Interface* and edit the port3 interface.
- 2 Set the *Addressing mode* to *Manual* and set the *IP/Netmask* to 192.168.8.1.
- 3 Enable *Connect FortiAP to this interface* and set *Reserve IP addresses for FortiAP* to 192.168.8.2 - 192.168.8.9.

This step automatically configures a DHCP server for the AP units. You can see this configuration in *System > Network > DHCP Server*.

- 4 Select OK.

To configure the interface for the AP unit - CLI

```
config system interface
edit port3
    set mode static
    set ip 192.168.8.1 255.255.255.0
end
```

To configure the DHCP server for AP units - CLI

```
config system dhcp server
edit 0
    set interface port3
    config ip-range
    edit 1
        set end-ip 192.168.8.9
        set start-ip 192.168.8.2
    end
    set netmask 255.255.255.0
    set vci-match enable
    set vci-string "FortiAP"
end
```


To connect a FortiAP unit - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Connect the FortiAP unit to port 3.
- 3 Periodically select *Refresh* while waiting for the FortiAP unit to be listed.
Recognition of the FortiAP unit can take up to two minutes.
If FortiAP units are connected but cannot be recognized, try disabling VCI-Match in the DHCP server settings.
- 4 When the FortiAP unit is listed, select the entry to edit it.
The *Edit Managed Access Point* window opens.
- 5 In *State*, select *Authorize*.
- 6 Make sure that AP Profile is set to *Automatic*.
- 7 In *SSID*, select *Automatically Inherit all SSIDs*.
- 8 Select *OK*.
- 9 Repeat Steps 2 through 8 for each FortiAP unit.

To connect a FortiAP unit - CLI

- 1 Connect the FortiAP unit to port 3.
- 2 Enter

```
config wireless-controller wtp
```
- 3 Wait 30 seconds, then enter *get*.
Retry the *get* command every 15 seconds or so until the unit is listed, like this:

```
== [ FAP22A3U10600118 ]  
wtp-id: FAP22A3U10600118
```
- 4 Edit the discovered FortiAP unit like this:

```
edit FAP22A3U10600118  
    set admin enable  
end
```
- 5 Repeat Steps 2 through 4 for each FortiAP unit.

A more complex example

This example creates multiple networks and uses custom AP profiles.

Scenario

In this example, Example Co. provides two wireless networks, one for its employees and the other for customers or other guests of its business. Guest users have access only to the Internet, not to the company's private network. The equipment for these WiFi networks consists of FortiAP-220A units controlled by a FortiGate unit.

The employee network operates in 802.11n mode on both the 2.4GHz and 5GHz bands. Client IP addresses are in the 10.10.120.0/24 subnet, with 10.10.120.1 the IP address of the WAP. The guest network also operates in 802.11n mode, but only on the 2.4GHz band. Client IP addresses are on the 10.10.115.0/24 subnet, with 10.10.115.1 the IP address of the WAP.

On FortiAP-220A units, the 802.11n mode also supports 802.11g and 802.11b clients on the 2.4GHz band and 802.11a clients on the 5GHz band.

The guest network WAP broadcasts its SSID, the employee network WAP does not.

The employees network uses WPA-Enterprise authentication through a FortiGate user group. The guest network features a captive portal. When a guest first tries to connect to the Internet, a login page requests logon credentials. Guests use numbered guest accounts authenticated by RADIUS. The captive portal for the guests includes a disclaimer page.

In this example, the FortiAP units connect to port 3 and are assigned addresses on the 192.168.8.0/24 subnet.

Configuration

To configure these wireless networks, you must:

- Configure authentication for wireless users
- Configure the SSIDs (network interfaces)
- Configure the AP profile
- Configure the WiFi LAN interface and a DHCP server
- Configure firewall policies

Configuring authentication for employee wireless users

Employees have user accounts on the FortiGate unit. This example shows creation of one user account, but you can create multiple accounts and add them as members to the user group.

To configure the user group for employee access - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter the following information and then select OK:

Name	employee-group
Type	Firewall
Allow SSL-VPN Access	Disabled
Available Users/Groups / Members	Move appropriate user accounts to the <i>Members</i> list.

To configure the user group for employee access - CLI

```
config user group
  edit "employee-group"
    set member "user01"
  end
```

The user authentication setup will be complete when you select the employee-group in the SSID configuration.

Configuring authentication for guest wireless users

Guests are assigned temporary user accounts created on a RADIUS server. The RADIUS server stores each user's group name in the Fortinet-Group-Name attribute. Wireless users are in the group named "wireless".

The FortiGate unit must be configured to access the RADIUS server.

To configure the FortiGate unit to access the guest RADIUS server - web-based manager

- 1 Go to *User > Remote > RADIUS* and select *Create New*.
- 2 Enter the following information and select OK:

Name	guestRADIUS
Primary Server Name / IP	10.11.102.100
Primary Server Secret	grikfwpdfg
Secondary Server Name / IP	Optional
Secondary Server Secret	Optional
Authentication Scheme	Use default, unless server requires otherwise.
Leave other settings at their default values.	

To configure the FortiGate unit to access the guest RADIUS server - CLI

```
config user radius
  edit guestRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret grikfwpdfg
  end
```

To configure the user group for guest access - web-based manager

- 1 Go to *User > User Group* and select *Create New*.
- 2 Enter the following information and then select OK:

Name	guest-group
Type	Firewall
Allow SSL-VPN Access	Disabled
Available Users/Groups / Members	Move <i>guestRADIUS</i> to the <i>Members</i> list.
Match one of these group names	Select Add and fill in the following fields:
Remote Server	Select <i>guestRADIUS</i> .
Group Name	Enter <i>wireless</i>

To configure the user group for guest access - CLI

```
config user group
  edit "guest-group"
    set member "guestRADIUS"
  config match
    edit 0
      set server-name "guestRADIUS"
      set group-name "wireless"
    end
  end
```

The user authentication setup will be complete when you select the guest-group user group in the SSID configuration.

Configuring the SSIDs

First, establish the SSIDs (network interfaces) for the employee and guest networks. This is independent of the number of physical access points that will be deployed. Both networks assign IP addresses using DHCP.

To configure the employee SSID - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
- 2 Enter the following information and select OK:

Interface Name	example_inc
IP/Netmask	10.10.120.1/24
Administrative Access	Ping (to assist with testing)
SSID	example_inc
Enable DHCP	Enable
Address Range	10.10.120.2 - 10.10.120.199
Netmask	255.255.255.0
Default Gateway	Same As Interface IP
DNS Server	Same as System DNS
Security Mode	WPA/WPA2-Enterprise
Data Encryption	AES
Authentication	Select <i>Usergroup</i> , then select <i>employee-group</i> .
Leave other settings at their default values.	

To configure the employee SSID - CLI

```

config wireless-controller vap
  edit example_inc
    set ssid "example_inc"
    set security wpa-enterprise
    set auth usergroup
    set usergroup employee-group
  end
config system interface
  edit example_inc
    set ip 10.10.120.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set default-gateway 10.10.120.1
    set dns-service default
    set interface example_inc
    config ip-range
      edit 1
        set end-ip 10.10.120.199
        set start-ip 10.10.120.2
      end
    end
  end

```

```

end
set lease-time 7200
set netmask 255.255.255.0
end

```

To configure the example_guest SSID - web-based manager

- 1 Go to *WiFi Controller > WiFi Network > SSID* and select *Create New*.
- 2 Enter the following information and select OK:

Name	example_guest
IP/Netmask	10.10.115.1/24
Administrative Access	Ping (to assist with testing)
SSID	example_guest
Enable DHCP	Enable
Address Range	10.10.115.2 - 10.10.115.50
Netmask	255.255.255.0
Default Gateway	Same as Interface IP
DNS Server	Same as System DNS
Security Mode	Captive Portal
Customize Portal Messages	Select
User Groups	Select <i>guest-group</i>
Leave other settings at their default values.	

To configure the example_guest SSID - CLI

```

config wireless-controller vap
  edit example_guest
    set ssid "example_guest"
    set security captive-portal
    set selected-usergroups guest-group
  end
config system interface
  edit example_guest
    set ip 10.10.115.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set default-gateway 10.10.115.1
    set dns-service default
    set interface "example_guest"
    config ip-range
      edit 1
        set end-ip 10.10.115.50
        set start-ip 10.10.115.2
      end
    set lease-time 7200
    set netmask 255.255.255.0
  end
end

```

Configuring the custom AP profile

The custom AP Profile defines the radio settings for the networks. The profile provides access to both Radio 1 (2.4GHz) and Radio 2 (5GHz) for the employee virtual AP, but provides access only to Radio 1 for the guest virtual AP.

To configure the AP Profile - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Custom AP Profile* and select *Create New*.
- 2 Enter the following information and select OK:

Name	example_AP
Platform	FAP220A
Radio 1	
Mode	Access Point
Background Scan	Enable
Rogue AP On-wire Scan	Enabled.
Radio Resource Provision	Not enabled.
Band	802.11n
Short Guard Interval	Not enabled.
Channel	Select 1, 6, and 11.
Tx Power	100%
SSID	Select <i>example_inc</i> and <i>example_guest</i> .
Radio 2	
Mode	Access Point
Background Scan	Enable
Rogue AP On-wire Scan	Enabled.
Radio Resource Provision	Enabled.
Band	802.11n_5G
Short Guard Interval	Not enabled.
20/40 MHz Channel Width	Not enabled.
Channel	Select all.
Tx Power	100%
SSID	Select <i>example_inc</i> .

To configure the AP Profile - CLI

```
config wireless-controller wtp-profile
edit "example_AP"
  config platform
    set type 220A
  end
  config radio-1
    set ap-bgscan enable
    set band 802.11n
```

```

set channel "1" "6" "11"
set rogue-scan enable
set vaps "example_inc" "example_guest"
end
config radio-2
set ap-bgscan enable
set band 802.11n-5G
set channel "36" "40" "44" "48" "149" "153" "157" "161"
    "165"
set rogue-scan enable
set vaps "example_inc"
end

```

Configuring firewall policies

Identity-based firewall policies are needed to enable the WLAN users to access the Internet on Port1. First you create firewall addresses for employee and guest users, then you create the firewall policies.

To create firewall addresses for employee and guest WiFi users

- 1 Go to *Firewall Objects > Address > Address*.
- 2 Select *Create New*, enter the following information and select *OK*.

Address Name	employee-wifi-net
Type	Subnet / IP Range
Subnet / IP Range	10.10.120.0/24
Interface	example_inc

- 3 Select *Create New*, enter the following information and select *OK*.

Address Name	guest-wifi-net
Type	Subnet / IP Range
Subnet / IP Range	10.10.115.0/24
Interface	example_guest

To create firewall policies for employee WiFi users - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*:

Source Interface/Zone	example_inc
Source Address	employee-wifi-net
Destination Interface/Zone	port1
Destination Address	All
Schedule	always
Service	Any
Action	ACCEPT
NAT	Enable NAT

- 3 Optionally, select *UTM* and set up UTM features for wireless users.
- 4 Select *OK*.
- 5 Repeat steps 1 through 4 but select Internal as the Destination Interface/Zone to provides access to the ExampleCo private network.

To create firewall policies for employee WiFi users - CLI

```
config firewall policy
edit 0
set srcintf "employee_inc"
set dstintf "port1"
set srcaddr "employee-wifi-net"
set dstaddr "all"
set action accept
set schedule "always"
set service "ANY"
set nat enable
set schedule "always"
set service "ANY"
next
edit 0
set srcintf "employee_inc"
set dstintf "internal"
set srcaddr "employee-wifi-net"
set dstaddr "all"
set action accept
set schedule "always"
set service "ANY"
set nat enable
set schedule "always"
set service "ANY"
end
```

To create a firewall policy for guest WiFi users - web-based manager

- 1 Go to *Policy > Policy* and select *Create New*.
- 2 Enter the following information and select *OK*:

Source Interface/Zone	example_guest
Source Address	guest-wifi-net
Destination Interface/Zone	port1
Destination Address	All
Schedule	always
Service	Any
Action	ACCEPT
NAT	Enable NAT

- 3 Optionally, select *UTM* and set up UTM features for wireless users.
- 4 Select *OK*.

To create a firewall policy for guest WiFi users - CLI

```
config firewall policy
edit 0
    set srcintf "example_guest"
    set dstintf "port1"
    set srcaddr "guest-wifi-net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
end
```

Connecting the FortiAP units

You need to connect each FortiAP-220A unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

To configure the interface for the AP unit - web-based manager

- 1 Go to *System > Network > Interface* and edit the port3 interface.
- 2 Set the *Addressing mode* to *Manual* and set the *IP/Netmask* to 192.168.8.1.
- 3 Enable *Connect FortiAP to this interface* and set *Reserve IP addresses for FortiAP* to 192.168.8.2 - 192.168.8.9.

This step automatically configures a DHCP server for the AP units.

- 4 Select *OK*.

To configure the interface for the AP unit - CLI

```
config system interface
edit port3
    set mode static
    set ip 192.168.8.1 255.255.255.0
end
```

To configure the DHCP server for AP units - CLI

```
config system dhcp server
edit 0
    set interface port3
    config ip-range
    edit 1
        set end-ip 192.168.8.9
        set start-ip 192.168.8.2
    end
    set netmask 255.255.255.0
    set vci-match enable
    set vci-string "FortiAP"
end
```

To connect a FortiAP-220A unit - web-based manager

- 1 Go to *WiFi Controller > Managed Access Points > Managed FortiAP*.
- 2 Connect the FortiAP unit to port 3.
- 3 Periodically select *Refresh* while waiting for the FortiAP unit to be listed.
Recognition of the FortiAP unit can take up to two minutes.
If there is persistent difficulty recognizing FortiAP units, try disabling VCI-Match in the DHCP server settings.
- 4 When the FortiAP unit is listed, select the entry to edit it.
The *Edit Managed Access Point* window opens.
- 5 In *State*, select *Authorize*.
- 6 In the *AP Profile*, select *[Change]* and then select the *example_AP* profile.
- 7 Select *OK*.
- 8 Repeat Steps 2 through 8 for each FortiAP unit.


To connect a FortiAP-220A unit - CLI

- 1 Connect the FortiAP unit to port 3.
- 2 Enter

```
config wireless-controller wtp
```
- 3 Wait 30 seconds, then enter *get*.
Retry the *get* command every 15 seconds or so until the unit is listed, like this:

```
== [ FAP22A3U10600118 ]  
wtp-id: FAP22A3U10600118
```
- 4 Edit the discovered FortiAP unit like this:

```
edit FAP22A3U10600118  
  set admin enable  
  set wtp-profile example_AP  
end
```
- 5 Repeat Steps 2 through 4 for each FortiAP unit.



Using a FortiWiFi unit as a client

A FortiWiFi unit by default operates as a wireless access point. But a FortiWiFi unit can also operate as a wireless client, connecting the FortiGate unit to another wireless network.

This section includes the following topics:

- [Use of client mode](#)
- [Configuring client mode](#)


Use of client mode

In client mode, the FortiWiFi unit connects to a remote WiFi access point to access other networks or the Internet. This is most useful when the FortiWiFi unit is in a location that does not have a wired infrastructure.

For example, in a warehouse where shipping and receiving are on opposite sides of the building, running cables might not be an option due to the warehouse environment. The FortiWiFi unit can support wired users using its Ethernet ports and can connect to another access point wirelessly as a client. This connects the wired users to the network using the 802.11 WiFi standard as a backbone.

Note that in client mode the FortiWiFi unit cannot operate as an AP. WiFi clients cannot see or connect to the FortiWiFi unit in Client mode.

Figure 258: FortiGate unit in Client mode



Configuring client mode

To set up the FortiAP unit as a WiFi client, you must use the CLI. Before you do this, be sure to remove any AP WiFi configurations such as SSIDs, DHCP servers, policies, and so on.

To configure wireless client mode

- 1 Change the WiFi mode to client.

In the CLI, enter the following commands:

```
config system global
  set wireless-mode client
end
```

Respond “y” when asked if you want to continue. The FortiWiFi unit will reboot.

- 2 Create a WiFi interface and configure the appropriate WiFi client settings.

For example, to configure the client for WPA-Personal authentication on the *our_wifi* SSID with passphrase *justforus*, enter the following in the CLI:


```
config system interface
  edit client_wifi
    set type wireless
    set mode dhcp
    set wifi-ssid our_wifi
    set wifi-security wpa-personal
    set wifi-passphrase "justforus"
  end
```

The WiFi interface *client_wifi* will receive an IP address using DHCP.

- 3 Configure a *client_wifi* to port1 policy.

You can use either CLI or web-based manager to do this. The important settings are:

Source Interface/Zone	client_wifi
Source Address	All
Destination Interface/Zone	port1
Destination Address	All
Action	ACCEPT
NAT	Enable NAT. Select <i>Use Destination Interface Address</i> (default).
Leave other settings at their default values.	



WiFi Reference

This chapter provides some reference information pertaining to wireless networks. The following topics are included in this section:

- [Wireless radio channels](#)

Wireless radio channels

IEEE 802.11a/n channels

[Table 132](#) lists the channels supported on FortiWiFi products that support the IEEE 802.11a and 802.11n wireless standards. 802.11a is available on FortiWiFi models 60B and higher. 802.11n is available on FortiWiFi models 80CM and higher.

All channels are restricted to indoor usage except in the Americas, where both indoor and outdoor use is permitted on channels 52 through 64 in the United States.

Table 132: IEEE 802.11a/n (5-GHz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas				
		Americas	Europe	Taiwan	Singapore	Japan
34	5170					•
36	5180	•	•		•	
38	5190					
40	5200	•	•		•	•
42	5210					
44	5220	•	•		•	•
46	5230					
48	5240	•	•		•	•
149	5745	•		•	•	
153	5765	•		•	•	
157	5785	•		•	•	
161	5805	•		•	•	
165	5825	•			•	

IEEE 802.11b/g/n channel numbers

[Table 133](#) lists IEEE 802.11b/g/n channels. All FortiWiFi units support 802.11b and 802.11g. Newer models also support 802.11n.

Mexico is included in the Americas regulatory domain. Channels 1 through 8 are for indoor use only. Channels 9 through 11 can be used indoors and outdoors. You must make sure that the channel number complies with the regulatory standards of Mexico.

Table 133: IEEE 802.11b/g/n (2.4-Ghz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas			
		Americas	EMEA	Israel	Japan
1	2412	•	•	•?	•
2	2417	•	•	•?	•
3	2422	•	•	•?	•
4	2427	•	•	•	•
5	2432	•	•	•	•
6	2437	•	•	•	•
7	2442	•	•	•	•
8	2447	•	•	•	•
9	2452	•	•	•	•
10	2457	•	•	•	•
11	2462	•	•	•?	•
12	2467		•	•?	•
13	2472		•	•?	•
14	2484				b only



WiFi Controller Reference

This section introduces you to the web-based manager *WiFi Controller* menu.

The following topics are included in this section:

- [WiFi Controller overview](#)
- [WiFi Network](#)
- [Managed access points](#)
- [Monitor](#)



The word “unit” refers to the FortiGate unit. The words “FortiGate unit” are used when talking about different Fortinet products in one sentence. For example, “The Central Management menu provides the option of remotely managing your FortiGate unit by a FortiManager unit.”

WiFi Controller overview

The WiFi Controller menu configures WiFi networks on your FortiWiFi or FortiGate unit. Your WiFi networks can use any of the following WiFi networking equipment:

- your FortiWiFi unit’s built-in wireless access point/client (see [“FortiWiFi units” on page 2438](#))
- FortiAP units—wireless access points compliant with the CAPWAP standard (see [“FortiAP units” on page 2440](#))
- the built-in wireless access point/client of a FortiWiFi unit connected to your unit (see [“Using a FortiWiFi unit as a managed AP” on page 2439](#))

Each of these pieces of WiFi networking equipment is an access point. Each access point can carry multiple networks to which clients can connect.



The WiFi Controller feature is available on all models running FortiOS or FortiOS Carrier, **except** model 30B.

The wireless controller feature can also:

- monitor activity on your WiFi networks
- monitor neighboring access points that might cause interference
- detect rogue (unauthorized) access points connected to your wired networks
- suppress access points that you have designated as rogues

WiFi Network

In the WiFi Network menu, you can configure SSID and rogue AP detection settings.

An SSID defines the security settings for a wireless LAN. For each SSID, the FortiGate unit creates a virtual network interface. You create firewall policies to control traffic between the SSID interface and other networks. Users need the correct security settings to connect to the access point, and they can also be required to authenticate to use a firewall policy.

A Rogue AP is an unauthorized AP connected to your network. This can be a security issue. Other APs may be receivable in your area. These APs belong to neighboring businesses or homes. They can cause interference but are not a security threat. The on-wire detection technique can distinguish between neighbors and rogues.

This topic includes the following:

- [SSID list](#)
- [SSID configuration settings](#)
- [Rogue AP Settings](#)

SSID list

The list of SSIDs (WiFi networks) at *WiFi Controller > WiFi Network > SSID* contains the following columns:

Create New	Creates a new SSID. When you select <i>Create New</i> , you are automatically redirected to the New SSID page. See SSID configuration settings .
Edit	Modifies an SSID's settings. When you select <i>Edit</i> , you are automatically redirected to the Edit SSID page. See SSID configuration settings .
Delete	Removes an SSID from the list on the SSID page. To remove multiple SSIDs from within the list, on the SSID page, in each of the rows of the SSIDs you want removed, select the check box and then select <i>Delete</i> . To remove all SSIDs from the list, on the SSID page, select the check box in the check box column and then select <i>Delete</i> .
SSID	The SSID or network name for the wireless interface.
Administrative Status	Indicates whether the SSID's administrative status is up or down. A green up arrow indicates that it is up; a red down arrow indicates that it is not.
Security mode	The type of security for the wireless interface: WPA/WPA2 Personal — user must know pre-shared key value to connect. WPA/WPA2 Enterprise — user must know user name and password to connect. Captive Portal — user connects to the open access point and then must authenticate to use the network

Data Encryption	<p>The type of encryption for the wireless interface in WPA/WPA2 modes.</p> <p>AES is the most secure, but some older clients support only TKIP.</p>
Clients	<p>The maximum number of clients permitted to connect simultaneously.</p>
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i>.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the <i>Object Usage</i> window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <p>View the list page for these objects – automatically redirects you to the list page where the object is referenced at.</p> <p>Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page.</p> <p>View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting than the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.</p>

SSID configuration settings

When you edit an SSID or create a new one in *WiFi Controller > WiFi Network > SSID*, the following configuration settings are available:

Interface Name	Enter a name for the SSID.
Administrative Status	Select to have the SSID's status either up or down. If the SSID's status is down, the SSID is not being used. If the SSID's status is up, the SSID is being used.
Addressing Mode	
IP/Netmask	Enter the IP address and netmask for the SSID.
IPv6 Address	Enter the IPv6 address. This is available only when IPv6 has been enabled on the unit.
Administrative Access	Select the administrative access for the SSID.
IPv6 Administrative Access	If you have IPv6 addresses, select the administrative access for IPv6 for SSID.
Enable Explicit Web Proxy	Select to enable explicit web proxy for the SSID.

WiFi Settings	
SSID	Enter the SSID. By default, this field contains fortinet.
Enable DHCP	Select to enable a DHCP server and configure basic DHCP server settings. The <i>Address Start</i> and <i>Address End</i> settings are used to create an appropriate DHCP server in the DHCP server list. If the unit is in transparent mode, the DHCP server settings will be unavailable.
Address Range	Enter the starting IP address of the DHCP server.
Netmask	Enter the netmask of the DHCP server.
Default Gateway	Enter the default gateway for the DHCP server.
DNS Server	Enter the DNS server.
Security Mode	Select the security mode for the wireless interface. Wireless users must use the same security mode to be able to connect to this wireless interface. Additional security mode options are available in the CLI.
	<i>WPA/WPA2-Personal</i> – WPA or WPA2 security. WPA is WiFi protected access. WPA2 is WPA with additional security features. There is one shared key (password) that all users use.
	<i>WPA/WPA2-Enterprise</i> – similar to WPA/WPA2-Personal, but is best used for enterprise networks. Each user is separately authenticated by user name and password.
	<i>Captive Portal</i> – authenticates users through a customizable web page.
Customize Portal Messages	Available only when <i>Security Mode</i> is <i>Captive Portal</i> . Select to customize the endpoint replacement messages. When you select <i>Edit</i> , the Edit Message window appears. Within the window, you can modify each one of the endpoint replacement messages.
User Groups	Available only when <i>Security Mode</i> is <i>Captive Portal</i> . Select the user groups that can authenticate. To select a user group, select the group in <i>Available</i> and then use the -> arrow to move that group to <i>Selected</i> . To remove a user group from <i>Selected</i> , select the group and then use the <- arrow to move the group back to <i>Available</i> .
Data Encryption	Available only when <i>Security Mode</i> is <i>WPA/WPA2-Enterprise</i> . Select <i>TKIP</i> or <i>AES</i> encryption as appropriate for the capabilities of your wireless clients. This is available for WPA/WPA2 security modes.
Pre-shared Key	Available only when <i>Security Mode</i> is <i>WPA/WPA2-Personal</i> . Enter the encryption key that the clients must use.

Authentication	Available only when <i>Security Mode</i> is <i>WPA/WPA2-Enterprise</i> . Select one of the following: <i>RADIUS Server</i> — Select the RADIUS server that will authenticate the clients. <i>Usergroup</i> – Select the user group that can authenticate.
Block Intra-SSID Traffic	Select to enable the unit to block intra-SSID traffic.
Maximum Clients	Select to limit the number of clients permitted to connect simultaneously. Enter the limit value.
Comments	Enter a description or comment for the SSID.

Rogue AP Settings

From the Rogue AP Settings page, you can enable rogue AP detection and the on-wire rogue AP detection technique. Rogue APs are APs that are not known to the WiFi controller and these unknown APs can be monitored by using these two features.

The feature, *Enable On-Wire Rogue AP Detection Technique*, determines which unknown APs are actually connected to your network. The unknown APs are considered rogues.

You can enable or disable these settings in *WiFi Controller > WiFi Network > Rogue AP Settings*.

Managed access points

The WiFi controller needs to be configured to manage each physical access point and configure its radio settings for the wireless LAN.

From the Managed Access Points menu, you can configure managed FortiAP and local WiFi radio settings and create custom AP profiles.

This topic contains the following:

- [Local WiFi Radio configuration settings](#)
- [Managed FortiAP list](#)
- [Managed FortiAP configuration settings](#)
- [Custom AP Profiles](#)
- [Custom AP Profile Settings](#)



The Local WiFi Radio submenu is available only on FortiWiFi units.

Local WiFi Radio configuration settings

Go to *WiFi Controller > Managed Access Points > Local WiFi Radio* to configure the WiFi radio facility of your FortiWiFi unit. FortiGate units do not have this page.

Local WiFi Radio page	
Displays the local WiFi radio settings. From this page you can change the local WiFi settings, such as changing the AP profile. You must select <i>Apply</i> to save the changes.	
AP Profile	Select <i>Change</i> to change the profile. A drop-down list appears when you select <i>Change</i> ; select the profile from the list and then select <i>Apply</i> .
Enable Wireless Radio	Select to enable the wireless radio settings on the unit.
Automatically Inherit All SSIDs	Select to have the unit automatically inherit all SSID broadcasts.
Select SSIDs	Select to manually choose which SSID broadcasts.
TX Power	Displays the transmission power in percent. Use the slider to change the power.
Channel	The channel that the unit is broadcasting on. For example, channel 2.
Band	The IEEE wireless protocol that the unit is using.

Managed FortiAP list

Go to *WiFi Controller > Managed Access Points > Managed FortiAP* to view the list of managed APs that have discovered the WiFi Controller. On this page, you can edit, delete, authorize, ignore or restart access points.

Edit	Modifies a managed physical AP's settings. When you select <i>Edit</i> , you are automatically redirected to the Managed FortiAP configuration settings .
Delete	Removes a managed physical AP in the list on the Managed Physical AP page. To remove multiple managed APs from within the list, on the Managed Physical AP page, in each of the rows of the APs you want removed, select the check box and then select <i>Delete</i> . To remove all APs in the list, on the Managed Physical AP page, select the check box in the check box column, and then select <i>Delete</i> .
Column Settings	Not all columns are shown by default. Select <i>Column Settings</i> to choose which columns to display.
Refresh	Select to refresh the current information on the page.
Restart	Select to restart an AP.
Show Ignored	Select to show the FortiAP units that the unit currently ignores.
State	The state of the FortiAP unit. The state is indicated by an icon, for example, if there is a blue question mark, this indicates that the unit is connecting.
Status	The status of the FortiAP unit.

Name	The name or serial number of the physical AP.
Clients	The maximum number of clients that are permitted to connect simultaneously. When you click on the number, you are automatically redirected to <i>WiFi Controller > Monitor > Client Monitor</i> .
SSIDs	The SSIDs of the FortiAP unit.
Connecting From	The wired IP address of the FortiAP unit.

Managed FortiAP configuration settings

You can select a managed AP on the *WiFi Controller > Managed Access Points > Managed FortiAP* page and modify the following settings:

Serial Number	The serial number of the unit (read-only).
Name	Enter a name for the access point. Otherwise, the serial number is displayed as the AP name.
Description	Select <i>Change</i> to change the existing description. If there is no description, "N/A" displays.
Managed AP status section	
Status	Indicates the connection status of the access point. For example, if the access point is connecting, <i>Connecting</i> displays.
Connecting From	The IP address of the unit.
State	The type of state the unit is in. Select <i>Authorize</i> to authorize the managed access point. If you want to deauthorize the managed access point, select <i>Deauthorize</i> .
Wireless Settings	
AP Profile	The name of the AP Profile or Automatic if a custom profile is not used. Select <i>Change</i> to select a different profile or Automatic settings, then select <i>Apply</i> .
Wireless Settings with Automatic AP Profile	
Enable WiFi Radio	Select to enable operation of this AP.
SSID	Automatically Inherit all SSIDs — AP will carry all WiFi networks. Select SSIDs — selects individual SSIDs for this AP to carry.
Tx Power	Adjust AP transmitter power. The 100% setting is the maximum permitted in your country.
Band	The WiFi radio band to be used. 802.11n, for example.
Channel	The radio channel currently in use.
Do not participate in Rogue AP Scanning	Select if AP performance is poor due to heavy traffic. The scanning function can affect performance.
Wireless Settings with Custom AP Profile	
Radio 1	Specific information about the AP's first radio.
Radio 2	Specific information about the AP's second radio.

Custom AP Profiles

The following are profile configuration settings in *WiFi Controller > Managed Access Points > Custom AP Profile*.

Custom AP Profile page Lists each individual physical FortiAP that are currently on your network. On this page, you can edit, delete, authorize, ignore or restart.	
Create New	Creates a new AP profile. When you select <i>Create New</i> , you are automatically redirected to the Custom AP Profile Settings page.
Edit	Modifies a managed physical AP settings. When you select <i>Edit</i> , you are automatically redirected to the Custom AP Profile Settings page.
Delete	Removes a managed physical AP in the list on the Managed Physical AP page. To remove multiple managed APs from within the list, on the Managed Physical AP page, in each of the rows of the APs you want removed, select the check box and then select <i>Delete</i> . To remove all APs in the list, on the Managed Physical AP page, select the check box in the check box column, and then select <i>Delete</i> .
Name	The name of the AP profile.
Comments	A description or comment about the AP profile.
Platform	The type of model that is associated with the AP profile.
Radio 1	The selected radio band and channels for the first (or only) radio in the managed access point.
Radio 2	The selected radio band and channels for the second radio in the managed access point. This is for FortiAP units only.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM > Antivirus > Profile</i>), 1 appears in <i>Ref.</i> . To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object. To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window: <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when the icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.

Custom AP Profile Settings

You can edit or create new custom AP profiles on the *WiFi Controller > Managed Access Points > Custom AP Profile* page.

Name	Enter a name for the AP profile.
Comments	Enter a description or comment about the AP profile. This is optional.
Platform	Select the type of Fortinet platform that will be using the AP profile. For example, FortiWiFi-60C.
Radio 1 settings / Radio 2 settings	
Radio 1 settings are the same as Radio 2 settings except for the options for Channel.	
Note: Radio 2 settings are available only for FortiAP models with dual radios.	
Mode	Select the type of mode. <ul style="list-style-type: none"> • <i>Disable</i> – no mode is set. • <i>Access Point</i> – allows for the platform to be an access point • <i>Dedicated Monitor</i> – allows for the platform to be a dedicated monitor
Background Scan	Select to enable a background scan, which monitors for rogue APs. This is for the Rogue AP feature. By default, a background scan is disabled.
Radio Resource Provision	Select to enable the radio resource provision feature.
Band	Select the IEEE wireless protocol that is available to the region.
Short Guard Interval	Select to enable the short guard interval feature for 802.11n.
20/40 Mhz Channel Width	Select to enable the channel width to have 20/40 megahertz for 802.11n-5G.
Channel	Select the channel or channels to include. These channels change with regards to what IEEE wireless protocol you selected in <i>Band</i> .
TX Power	By default, the TX power is set to 100% of the maximum power permitted in your region. To change the level, drag the slider.
SSID	Choose the SSIDs (WiFi networks) that APs using this profile will carry. Select the required SSIDs in the <i>Available</i> list and use the -> arrow to move them to the <i>Selected</i> list. To remove an SSID from the <i>Selected</i> list, select the SSID and then use the <- arrow to move it back to the <i>Available</i> list.

Monitor

The Monitor menu allows you to view monitored wireless activity.

This topic contains the following:

- [Client Monitor](#)
- [Rogue AP Monitor](#)

Client Monitor

In *WiFi Controller > Monitor > Client Monitor*, you can view information about wireless clients of your managed access points.


Refresh	Update the information in the table.
Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <ul style="list-style-type: none"> • To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings. • To modify settings, select <i>Change</i> beside the setting and edit the settings. • To clear all filter settings, select the icon beside <i>Clear all filters</i>. • To use a filter icon to filter settings within a column, select the filter icon in the column; <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column. <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p>
Column Settings	Select the columns to display in the list. You can also determine the order in which they appear.
Page Controls	Use to navigate through the list.
Information columns	
Actual columns displayed depends on <i>Column Settings</i> .	
MAC	The MAC address of the wireless client.
Auth	The authentication type.
IP	The IP address assigned to the wireless client.
FortiAP	The name of the physical access point with which the client is associated.
SSID	The SSID for the managed access point.
Bandwidth Tx/Rx	<i>The current bandwidth.</i>
Signal Strength/Noise	The signal-to-noise ratio in deciBels calculated from signal strength and noise level.
Association Time	The time period that the client has been connected to this access point.
Bandwidth Rx	Received bandwidth used by the client, in Kbps.

Bandwidth Tx	Transmit bandwidth used by the client, in Kbps.
Idle Time	The total time this session that the client was idle.
Rate	The data rate of the client connection.
Manufacturer	The manufacturer of the client wireless device.

Rogue AP Monitor

View information about detected APs in *WiFi Controller > Monitor > Rogue AP Monitor*. You can also mark and suppress rogue APs.

Mark	Select the down arrow to mark the AP as accepted, rogue or unclassified.
Suppress AP	Select the down arrow to suppress the AP or unsuppress the AP. Available only if the AP is marked as a rogue AP.
Column Settings	Select the columns to display in the list. You can also determine the order in which they appear.
Refresh	Select to update the information. <i>none</i> means no updates.
Show Accepted	Select to show only the accepted APs.
Total detected APs	Displays the total number of APs that are detected by the FortiGate unit.
Information Columns Actual columns displayed depends on <i>Column Settings</i> .	
State	The state of the rogue AP.
Online Status	A green check mark indicates an active access point. A grey X indicates that the access point is inactive.
SSID	The wireless service set identifier (SSID) or network name for the wireless interface.
Security Type	The type of security currently being used.
Channel	The wireless radio channel that the access point uses.
MAC Address	The MAC address of the Wireless interface.
Vendor Info	The name of the vendor.
Signal Strength	The relative signal strength of the AP. Mouse over the symbol to view the signal-to-noise ratio.
Detected By	The name or serial number of the AP unit that detected the signal.
On-wire	A green up-arrow indicates a suspected rogue, based on the on-wire detection technique. A red down-arrow indicates AP is not a suspected rogue.



Chapter 16 VoIP Solutions: SIP & FortiGate Voice

This FortiOS Handbook chapter contains the following sections:

[FortiGate VoIP solutions: SIP](#) describes FortiGate SIP support.

[Example FortiGate Voice branch office configuration](#) describes how to configure a FortiGate Voice-80C unit to operate in NAT/Route mode and provide basic UTM and SIP services for an example branch office network.

[FortiGate Voice web-based manager configuration reference](#) describes FortiGate Voice web-based manager configuration settings.

[Using the PBX user web portal](#) describes how to log into and use the FortiGate Voice PBX portal.

[FortiGate Voice VoIP, PBX, and PSTN CLI Reference](#) describes FortiGate Voice VoIP, PBX, and PSTN CLI commands.



FortiGate VoIP solutions: SIP

This chapter includes the following sections:

- [SIP overview](#)
- [Common SIP VoIP configurations](#)
- [SIP messages and media protocols](#)
- [The SIP session helper](#)
- [The SIP ALG](#)
- [How the SIP ALG performs NAT](#)
- [Enhancing SIP pinhole security](#)
- [Hosted NAT traversal](#)
- [SIP over IPv6](#)
- [Deep SIP message inspection](#)
- [Blocking SIP request messages](#)
- [SIP rate limiting](#)
- [SIP logging and DLP archiving](#)
- [SIP and HA: session failover and geographic redundancy](#)
- [SIP and IPS](#)
- [SIP debugging](#)
- [VoIP Profile options](#)

SIP overview

The Session Initiation Protocol (SIP) is an IETF application layer signaling protocol used for establishing, conducting, and terminating multiuser multimedia sessions over TCP/IP networks using any media. SIP is often used for Voice over IP (VoIP) calls but can be used for establishing streaming communication between end points.

SIP employs a request and response transaction model similar to HTTP for communicating between endpoints. SIP sessions begin with a SIP client sending a SIP request message to another client to initiate a multimedia session. The other client responds with a SIP response message. Using these request and response messages, the clients engage in a SIP dialog to negotiate how to communicate and then start, maintain, and end the communication session.

SIP commonly uses TCP or UDP port 5060 and/or 5061. Port 5060 is used for non-encrypted SIP signaling sessions and port 5061 is typically used for SIP sessions encrypted with Transport Layer Security (TLS).

Devices involved in SIP communications are called SIP User Agents (UAs) (also sometimes called a User Element (UE)). UAs include User Agent Clients (UACs) that communicate with each other and User Agent Servers (UASs) that facilitate communication between UACs. For a VoIP application, an example of a UAC would be a SIP phone and an example of a UAS would be a SIP proxy server.

A SIP message contains headers that include client and server names and addresses required for the communication sessions. The body of a SIP message contains Session Description Protocol (SDP) statements that establish the media communication (port numbers, protocols and codecs) that the SIP UAs use. SIP VoIP most commonly uses the Real Time Protocol (RTP) and the Real Time Control Protocol (RTCP) for voice communication. Once the SIP dialog establishes the SIP call the VoIP stream can run independently, although SIP messages can affect the VoIP stream by changing port numbers or addresses and by ending it.

Once SIP communication and media settings are established, the UAs communicate with each using the established media settings. When the communication session is completed, one of the UAs ends the session by sending a final SIP request message and the other UA sends a SIP response message and both UAs end the SIP call and stop the media stream.

FortiGate units provide security for SIP communications using the SIP session helper and the SIP ALG:

- The SIP session-helper provides basic high-performance support for SIP calls passing through the FortiGate unit by opening SIP and RTP pinholes and performing source and destination IP address and port translation for SIP and RTP packets and for the IP addresses and port numbers in the SIP headers and the SDP body of the SIP messages. For more about the SIP session helper, see [“The SIP session helper” on page 2536](#).
- The SIP Application Layer Gateway (ALG) provides the same features as the session helper plus additional advanced features such as deep SIP message inspection, SIP logging, SIP IPv6 support, SIP message checking, HA failover of SIP sessions, and SIP rate limiting. For more about the SIP ALG, see [“The SIP ALG” on page 2541](#).


There are a large number of SIP-related Internet Engineering Task Force (IETF) documents (Request for Comments) that define behavior of SIP and related applications. FortiGate units provide complete support of [RFC 3261](#) for SIP and [RFC 4566](#) for SDP. FortiGate units also provide support for other SIP and SIP-related RFCs and performs [“Deep SIP message inspection” on page 2586](#) for SIP statements defined in other SIP RFCs.

Common SIP VoIP configurations

This section describes some common SIP VoIP configurations and simplified SIP dialogs for these configurations. This section also shows some examples of how adding a FortiGate unit affects SIP processing.

Peer to peer configuration

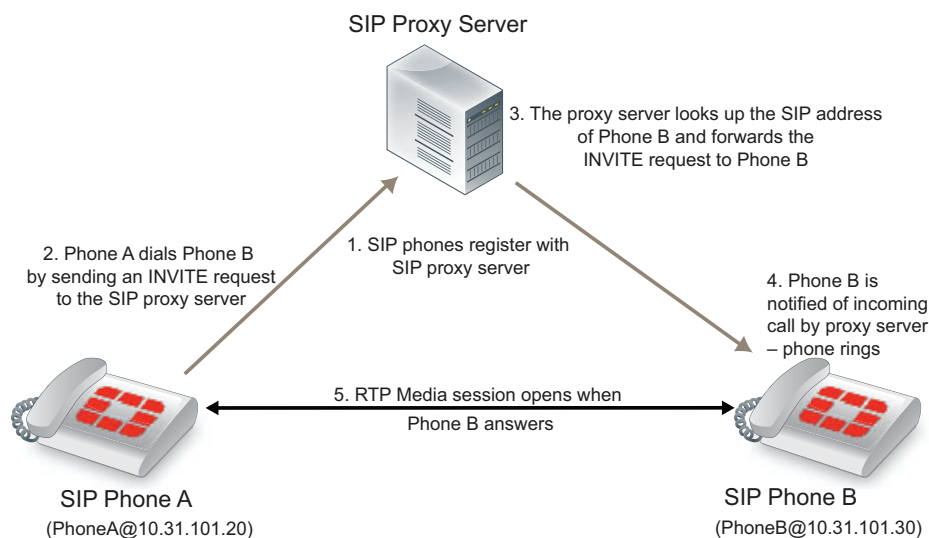
In the peer to peer configuration shown in [Figure 259](#), two SIP phones (in the example, FortiPhones) communicate directly with each other. The phones send SIP request and response messages back and forth between each other to establish the SIP session.

Figure 259: SIP peer to peer configuration

Peer to peer configurations are not very common because they require the SIP phones to keep track of the names and addresses of all of the other SIP phones that they can communicate with. In most cases a SIP proxy or re-direct server maintains addresses of a large number of SIP phones and a SIP phone starts a call by contacting the SIP proxy server.

SIP proxy server configuration

A SIP proxy server act as intermediary between SIP phones and between SIP phones (for example, two FortiFones) and other SIP servers. As shown in [Figure 260](#), SIP phones send request and response messages the SIP proxy server. The proxy server forwards the messages to other clients or to other SIP proxy servers. Proxy servers can hide SIP phones by proxying the signaling messages. To the other users on the VoIP network, the signaling invitations look as if they come from the SIP proxy server.

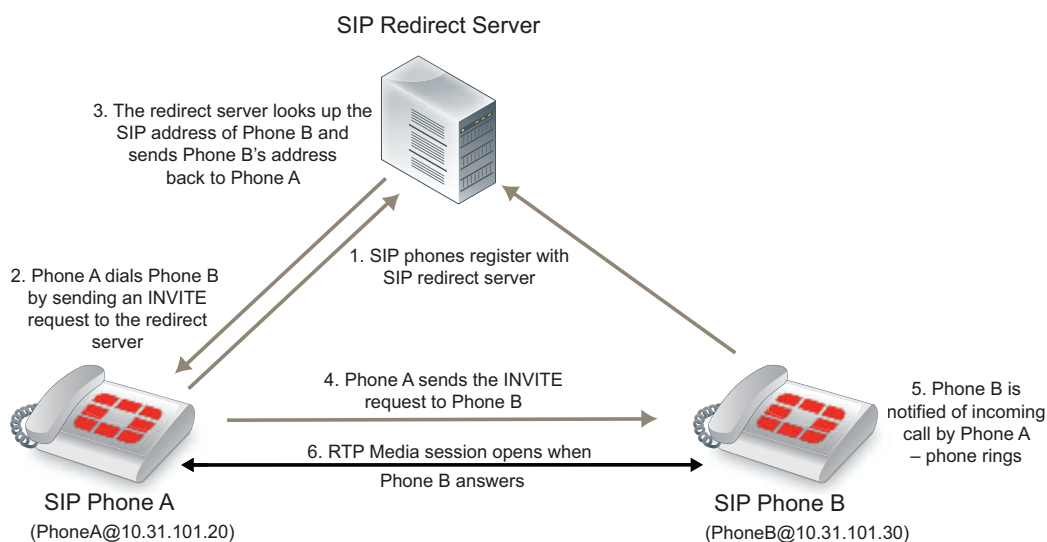
Figure 260: SIP in proxy mode

A common SIP configuration would include multiple networks of SIP phones. Each of the networks would have its own SIP server. Each SIP server would proxy the communication between phones on its own network and between phones in different networks.

SIP redirect server configuration


A SIP redirect server accepts SIP requests, maps the addresses in the request into zero or more new addresses and returns those addresses to the client. The redirect server does not initiate SIP requests or accept calls. As shown in [Figure 261](#), SIP clients send INVITE requests to the redirect server, which then looks up the destination address. The redirect server returns the destination address to the client. The client uses this address to send the INVITE request directly to the destination SIP client.

Figure 261: SIP in redirect mode



SIP registrar configuration

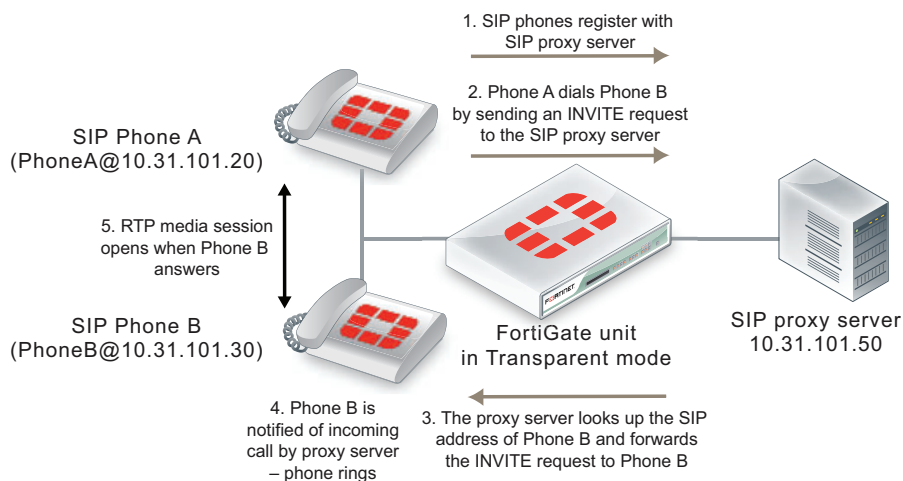
A SIP registrar accepts SIP REGISTER requests from SIP phones for the purpose of updating a location database with this contact information. This database can then become a SIP location service that can be used by SIP proxy servers and redirect servers to locate SIP clients. As shown in [Figure 262](#), SIP clients send REGISTER requests to the SIP registrar.

Figure 262: SIP registrar and proxy servers

SIP with a FortiGate unit

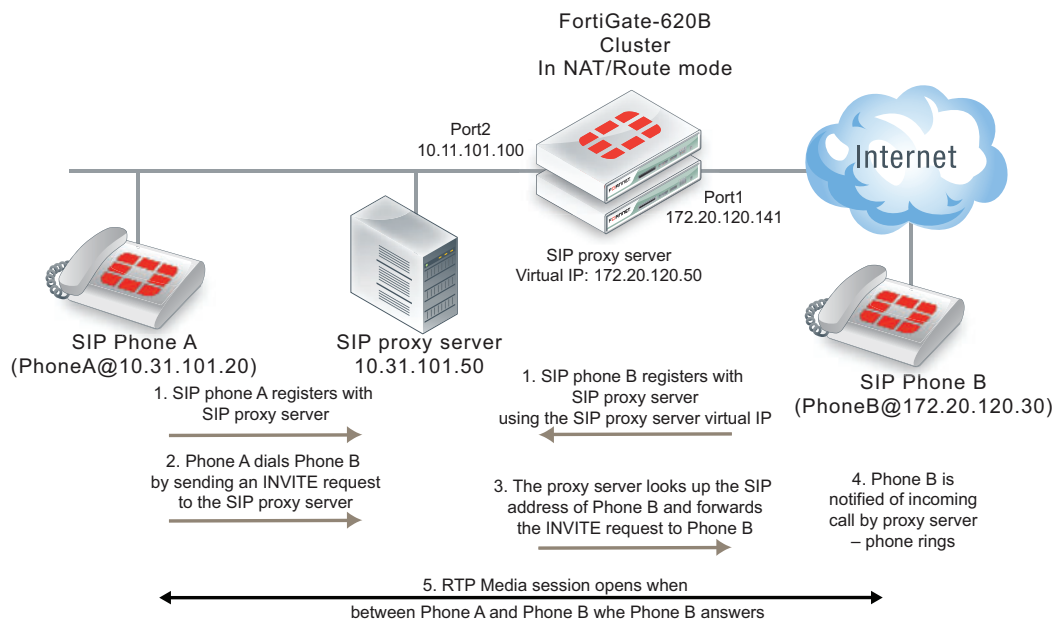
Depending on your security requirements and network configuration FortiGate units may be in many different places in a SIP configuration. This section shows a few examples.

Figure 263 shows a FortiGate unit installed between a SIP proxy server and SIP phones on the same network. The FortiGate unit is operating in Transparent mode so both the proxy server and the phones are on the same subnet. In this configuration, called SIP inspection without address translation, the FortiGate unit could be protecting the SIP proxy server on the private network by implementing SIP security features for SIP sessions between the SIP phones and the SIP proxy server.

Figure 263: SIP network with FortiGate unit in Transparent mode

The phones and server use the same SIP dialogs as they would if the FortiGate unit was not present. However, the FortiGate unit can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities.

Figure 264 shows a FortiGate unit operating in NAT/Route mode and installed between a private network and the Internet. Some SIP phones and the SIP proxy server are connected to the private network and some SIP phones are connected to the Internet. The SIP phones on the Internet can connect to the SIP proxy server through the FortiGate unit and communication between SIP phones on the private network and SIP phones on the Internet must pass through the FortiGate unit.

Figure 264: SIP network with FortiGate unit in NAT/Route mode

The phones and server use the same SIP dialog as they would if the FortiGate unit was not present. However, the FortiGate unit can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities. In addition, the FortiGate unit has a firewall virtual IP that forwards packets sent to the SIP proxy server Internet IP address (172.20.120.50) to the SIP proxy server internal network IP address (10.31.101.30).

Since the FortiGate unit is operating in NAT/Route mode it must translate packet source and destination IP addresses (and optionally ports) as the sessions pass through the FortiGate unit. Also, the FortiGate unit must translate the addresses contained in the SIP headers and SDP body of the SIP messages. As well the FortiGate unit must open SIP and RTP pinholes through the FortiGate unit. SIP pinholes allow SIP signalling sessions to pass through the FortiGate between phones and between phones and SIP servers. RTP pinholes allow direct RTP communication between the SIP phones once the SIP dialog has established the SIP call. Pinholes are opened automatically by the FortiGate unit. Administrators do not add security policies for pinholes or for RTP sessions. All that is required is a security policy that accepts SIP traffic.

Opening an RTP pinhole means opening a port on a FortiGate interface to allow RTP traffic to use that port to pass through the FortiGate unit between the SIP phones on the Internet and SIP phones on the internal network. A pinhole only accepts packets from one RTP session. Since a SIP call involves at least two media streams (one from Phone A to Phone B and one from Phone B to Phone A) the FortiGate unit opens two RTP pinholes. Phone A sends RTP packets through a pinhole in port2 and Phone B sends RTP packets through a pinhole in port1. The FortiGate unit opens the pinholes when required by the SIP dialog and closes the pinholes when the SIP call is completed. The FortiGate unit opens new pinholes for each SIP call.

Each RTP pinhole actually includes two port numbers. The RTP port number as defined in the SIP message and an RTCP port number, which is the RTP port number plus 1. For example, if the SIP call used RTP port 3346 the FortiGate unit would create a pinhole for ports 3346 and 3347.

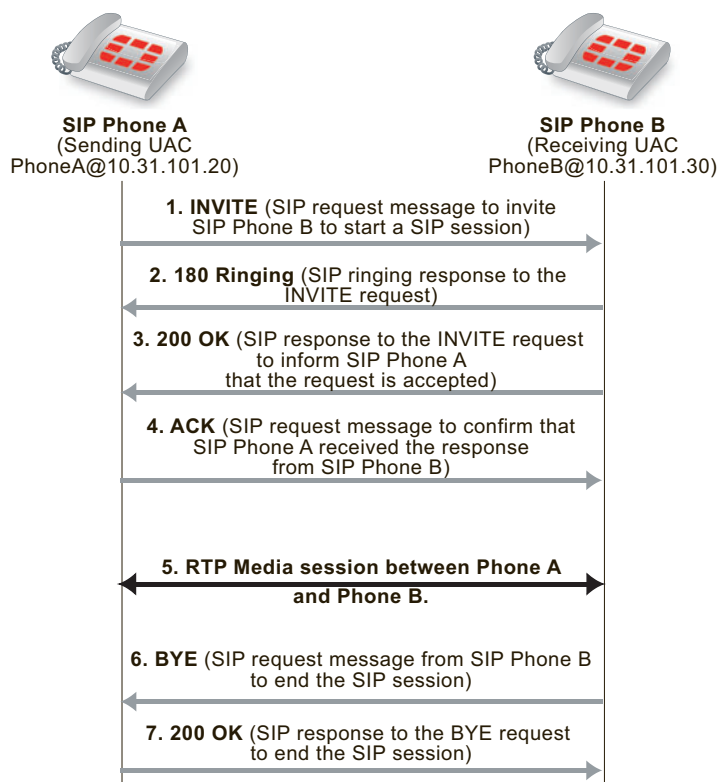
SIP messages and media protocols

This section provides an overview of SIP messages and how they communicate information about SIP sessions and how SDP, RTP, and RTCP fits in with SIP communications.

SIP uses clear text messages to start, maintain, and end media sessions between SIP user agent clients (UACs) and user agent servers (UASs). These messages form a SIP dialog. A typical SIP dialog begins with an INVITE request message sent from a UAC to another UAC or to a UAS. The first INVITE request message attempts to start a SIP call and includes information about the sending UAC and the receiving UAC as well as information about the communication session.

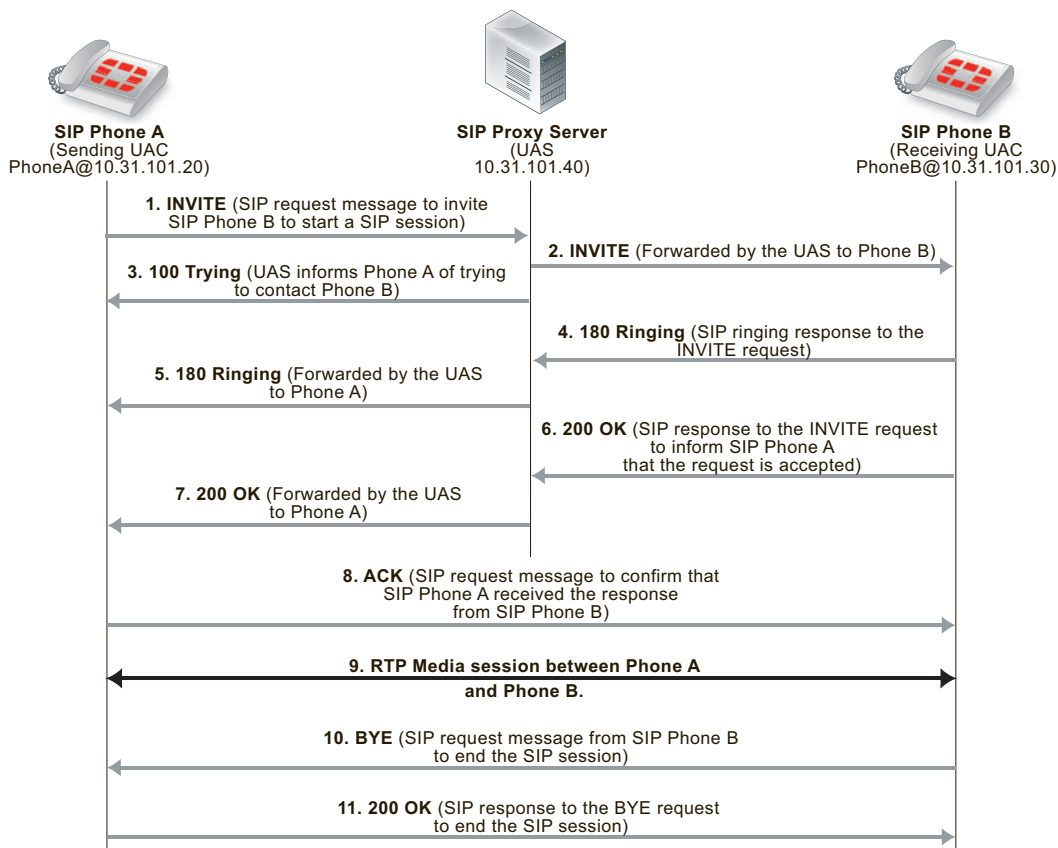
If only two UACs are involved as shown in [Figure 265](#), the receiving UAC (Phone B) responds with a 180 Ringing and then a 200 OK SIP response message that informs Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. Phone A and Phone B can then participate in the RTP media session set up by the SIP messages.

When the phone call is complete, one of the UACs (in the example Phone B) hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

Figure 265: Basic SIP dialog between two UACs

If a UAS in the form of a SIP proxy server is involved, similar messages are sent and received, but the proxy server participates as an intermediary in the initial call setup. In the example in [Figure 266](#) the SIP proxy server receives the INVITE request from Phone A and forwards it to Phone B. The proxy server then sends a 100 Trying response to Phone A. Phone B receives the INVITE request and responds with a 180 Ringing and then a 200 OK SIP response message. These messages are received by the proxy server and forwarded to Phone A to notify Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. This response is received by the proxy server and forwarded to Phone B. Phone A and Phone B can then participate in the media session independently of the proxy server.

When the phone call is complete Phone B hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

Figure 266: Basic SIP dialog between UACs with a SIP proxy server UAS

The SIP messages include SIP headers that contain names and addresses of Phone A, Phone B and the proxy server. This addressing information is used by the UACs and the proxy server during the call set up.

The SIP message body includes Session Description Protocol (SDP) statements that Phone A and Phone B use to establish the media session. The SDP statements specify the type of media stream to use for the session (for example, audio for SIP phone calls) and the protocol to use for the media stream (usually the Real Time Protocol (RTP) media streaming protocol).

Phone A includes the media session settings that it would like to use for the session in the INVITE message. Phone B includes its response to these media settings in the 200 OK response. Phone A's ACK response confirms the settings that Phone A and Phone B then use for the media session.

Hardware accelerated RTP processing

FortiGate units can offload RTP packet processing to network processor (NP) interfaces. This acceleration greatly enhance the overall throughput and resulting in near speed RTP performance.

SIP request messages

SIP sessions always start with a SIP request message (also just called a SIP request). SIP request messages also establish, maintain, and terminate SIP communication sessions. [Table 134](#) lists some common SIP request message types.

Table 134: Common SIP request message types

Message Type	Description
INVITE	A client sends an INVITE request to invite another client to participate in a multimedia session. The INVITE request body usually contains the description of the session.
ACK	The originator of an INVITE message sends an ACK request to confirm that the final response to an INVITE request was received. If the INVITE request did not contain the session description, it must be included in the ACK request.
PRACK	In some cases, SIP uses provisional response messages to report on the progress of the response to a SIP request message. The provisional response messages are sent before the final SIP response message. Similar to an ACK request message, a PRACK request message is sent to acknowledge that a provisional response message has been received.
OPTIONS	The UA uses OPTIONS messages to get information about the capabilities of a SIP proxy. The SIP proxy server replies with a description of the SIP methods, session description protocols, and message encoding that are supported.
BYE	A client sends a BYE request to end a session. A BYE request from either end of the SIP session terminates the session.
CANCEL	A client sends a CANCEL request to cancel a previous INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE sends a final response to the INVITE before receiving the CANCEL.
REGISTER	A client sends a REGISTER request to a SIP registrar server with information about the current location (IP address and so on) of the client. A SIP registrar server saves the information it receives in REGISTER requests and makes this information available to any SIP client or server attempting to locate the client.
Info	For distributing mid-session signaling information along the signaling path for a SIP call. I
Subscribe	For requesting the current state and state updates of a remote node.
Notify	Informs clients and servers of changes in state in the SIP network.
Refer	Refers the recipient (identified by the Request-URI) to a third party according to the contact information in the request.
Update	Opens a pinhole for new or updated SDP information.
Response codes (1xx, 202, 2xx, 3xx, 4xx, 5xx, 6xx)	Indicates the status of a transaction. For example: 200 OK, 202 Accepted, or 400 Bad Request.

SIP response messages

SIP response messages (often just called SIP responses) provide status information in response to SIP request messages. All SIP response messages include a response code and a reason phrase. There are five SIP response message classes. They are described below.

There are also two types of SIP response messages, provisional and final. Final response messages convey the result of the request processing, and are sent reliably. Provisional responses provide information on the progress of the request processing, but may not be sent reliably. Provisional response messages start with 1xx and are also called informational response messages.

Informational (or provisional)

Informational or provisional responses indicate that a request message was received and imply that the endpoint is going to process the request. Information messages may not be sent reliably and may not require an acknowledgement.

If the SIP implementation uses Provisional Response Acknowledgement (PRACK) ([RFC 3262](#)) then informational or provisional messages are sent reliably and require a PRACK message to acknowledge that they have been received.

Informational responses can contain the following reason codes and reason phrases:

```
100 Trying
180 Ringing
181 Call is being forwarded
182 Queued
183 Session progress
```

Success

Success responses indicate that a request message was received, understood, and accepted. Success responses can contain the following reason codes and reason phrases:

```
200 OK
202 Accepted
```

Redirection

Redirection responses indicate that more information is required for the endpoint to respond to a request message. Redirection responses can contain the following reason codes and reason phrases:

```
300 Multiple choices
301 Moved permanently
302 Moved temporarily
305 Use proxy
380 Alternative service
```

Client error

Client error responses indicate that a request message was received by a server that contains syntax that the server cannot understand (i.e. contains a syntax error) or cannot comply with. Client error responses include the following reason codes and reason phrases:

400 Bad request	401 Unauthorized
402 Payment required	403 Forbidden
404 Not found	405 Method not allowed

406 Not acceptable	407 Proxy authentication required
408 Request time-out	409 Conflict
410 Gone	411 Length required
413 Request entity too large	414 Request-URL too large
415 Unsupported media type	420 Bad extension
480 Temporarily not available	
481 Call leg/transaction does not exist	
482 Loop detected	483 Too many hops
484 Address incomplete	485 Ambiguous
486 Busy here	487 Request canceled
488 Not acceptable here	

Server error

Server error responses indicate that a server was unable to respond to a valid request message. Server error responses include the following reason codes and reason phrases:

```

500 Server internal error
501 Not implemented
502 Bad gateway
502 Service unavailable
504 Gateway time-out
505 SIP version not supported

```

Global failure

Global failure responses indicate that there are no servers available that can respond to a request message. Global failure responses include the following reason codes and reason phrases:

```

600 Busy everywhere
603 Decline
604 Does not exist anywhere
606 Not acceptable

```

SIP message start line

The first line in a SIP message is called the start line. The start line in a request message is called the request-line and the start line in a response message is called the status-line.

Request-line	<p>The first line of a SIP request message. The request-line includes the SIP message type, the SIP protocol version, and a Request URI that indicates the user or service to which this request is being addressed. The following example request-line specifies the INVITE message type, the address of the sender of the message (inviter@example.com), and the SIP version:</p> <pre>INVITE sip:inviter@example.com SIP/2.0</pre>
Status-line	<p>The first line of a SIP response message. The status-line includes the SIP protocol version, the response code, and the reason phrase. The example status-line includes the SIP version, the response code (200) and the reason phrase (OK).</p> <pre>SIP/2.0 200 OK</pre>

SIP headers

Following the start line, SIP messages contain SIP headers (also called SIP fields) that convey message attributes and to modify message meaning. SIP headers are similar to HTTP header fields and always have the following format:

```
<header_name>:<value>
```

SIP messages can include the SIP headers listed in [Table 135](#):

Table 135: SIP headers

SIP Header	Description
Allow	Lists the set of SIP methods supported by the UA generating the message. All methods, including ACK and CANCEL, understood by the UA MUST be included in the list of methods in the Allow header field, when present. For example: <code>Allow: INVITE, ACK, OPTIONS, CANCEL, BYE</code>
Call-ID	A globally unique identifier for the call, generated by the combination of a random string and the sender's host name or IP address. The combination of the To, From, and Call-ID headers completely defines a peer-to-peer SIP relationship between the sender and the receiver. This relationship is called a SIP dialog. <code>Call-ID: ddeg45e793@10.31.101.30</code>
Contact	Included in SIP request messages, the Contact header contains the SIP URI of the sender of the SIP request message. The receiver uses this URI to contact the sender. For example: <code>Contact: Sender <sip:sender@10.31.100.20>t</code>
Content-Length	The number of bytes in the message body (in bytes). <code>Content-Length: 126</code>
Content-Type	In addition to SIP headers, SIP messages include a message body that contains information about the content or communication being managed by the SIP session. The Content-Type header specifies what the content of the SIP message is. For example, if you are using SIP with SDP, the content of the SIP message is SDP code. <code>Content-Type: application/sdp</code>
CSeq	The command sequence header contains a sequence integer that is increased for each new SIP request message (but is not incremented in the response message). This header also includes the request name found in the request message request-line. For example: <code>CSeq: 1 INVITE</code>
Expires	Gives the relative time after which the message (or content) expires. The actual time and how the header is used depends on the SIP method. For example: <code>Expires: 5</code>
From	Identifies the sender of the message. Responses to a message are sent to the address of the sender. The following example includes the sender's name (Sender) and the sender's SIP address (sender@10.31.101.20): <code>From: Sender <sip:sender@10.31.101.20></code>

Table 135: SIP headers (Continued)

SIP Header	Description
Max-forwards	An integer in the range 0-255 that limits the number of proxies or gateways that can forward the request message to the next downstream server. Also called the number of hops, this value is decreased every time the message is forwarded. This can also be useful when the client is attempting to trace a request chain that appears to be failing or looping in mid-chain. For example: <code>Max-Forwards: 30</code>
P-Asserted-Identity	The P-Asserted-Identity header is used among trusted SIP entities to carry the identity of the user sending a SIP message as it was verified by authentication. See RFC 3325 . The header contains a SIP URI and an optional display-name, for example: <code>P-Asserted-Identity: "Example Person" <sip:10.31.101.50></code>
RAck	Sent in a PRACK request to support reliability of information or provisional response messages. It contains two numbers and a method tag. For example: <code>RAck: 776656 1 INVITE</code>
Record-Route	Inserted into request messages by a SIP proxy to force future requests to be routed through the proxy. In the following example, the host at IP address 10.31.101.50 is a SIP proxy. The <code>lr</code> parameter indicates the URI of a SIP proxy in Record-Route headers. <code>Record-Route: <sip:10.31.101.50;lr></code>
Route	Forces routing for a request message through one or more SIP proxies. The following example includes two SIP proxies: <code>Route: <sip:172.20.120.10;lr>,
<sip:10.31.101.50;lr></code>
RSeq	The RSeq header is used in information or provisional response messages to support reliability of informational response messages. The header contains a single numeric value. For example: <code>RSeq: 33456</code>
To	Identifies the receiver of the message. The address in this field is used to send the message to the receiver. The following example includes the receiver's name (Receiver) and the receiver's SIP address (receiver@10.31.101.30.): <code>To: Receiver <sip:receiver@10.31.101.30></code>
Via	Indicates the SIP version and protocol to be used for the SIP session and the address to which to send the response to the message that contains the Via field. The following example Via field indicates to use SIP version 2, UDP for media communications, and to send the response to 10.31.101.20 using port 5060. <code>Via: SIP/2.0/UDP 10.31.101.20:5060</code>

The SIP message body and SDP session profiles

The SIP message body describes the session to be initiated. For example, in a SIP phone call the body usually includes audio codec types, sampling rates, server IP addresses and so on. For other types of SIP session the body could contain text or binary data of any type which relates in some way to the session. The message body is included in request and response messages.

Two possible SIP message body types:

- Session Description Protocol (SDP), most commonly used for SIP VoIP.
- Multipurpose Internet Mail Extensions (MIME)

SDP is most often used for VoIP and FortiGate units support SDP content in SIP message bodies. SDP is a text-based protocol used by SIP to control media sessions. SDP does not deliver media but provides a session profile that contains media details, transport addresses, parameter negotiation, and other session description metadata for the participants in a media session. The participants use the information in the session profile to negotiate how to communicate and to manage the media session. SDP is described by [RFC 4566](#).

An SDP session profile always contains session information and may contain media information. Session information appears at the start of the session profile and media information (using the `m=` attribute) follows.

SDP session profiles can include the attributes listed in [Table 136](#).

Table 136: SDP session profile attributes

Attribute	Description
a=	Attributes to extend SDP in the form <code>a=<attribute></code> or <code>a=<attribute>:<value></code> .
b=	Contains information about the bandwidth required for the session or media in the form <code>b=<bandwidth_type>:<bandwidth></code> .
c=	Connection data about the session including the network type (usually IN for Internet), address type (IPv4 or IPv6), the connection source address, and other optional information. For example: <code>c=IN IPv4 10.31.101.20</code>
i=	A text string that contains information about the session. For example: <code>i=A audio presentation about SIP</code>
k=	Can be used to convey encryption keys over a secure and trusted channel. For example: <code>k=clear:444gdduudjffdee</code>

Table 136: SDP session profile attributes (Continued)

Attribute	Description
m=	<p>Media information, consisting of one or more lines all starting with m= and containing details about the media including the media type, the destination port or ports used by the media, the protocol used by the media, and a media format description.</p> <pre>m=audio 49170 RTP 0 3 m=video 3345/2 udp 34 m=video 2910/2 RTP/AVP 3 56</pre> <p>Multiple media lines are needed if SIP is managing multiple types of media in one session (for example, separate audio and video streams).</p> <p>Multiple ports for a media stream are indicated using a slash. 3345/2 udp means UDP ports 3345 and 3346. Usually RTP uses even-numbered ports for data with the corresponding one-higher odd ports used for the RTCP session belonging to the RTP session. So 2910/2 RTP/AVP means ports 2910 and 2912 are used for RTP and 2911 and 2913 are used for RTCP.</p> <p>Media types include udp for an unspecified protocol that uses UDP, RTP or RTP/AVP for standard RTP and RTP/SAVP for secure RTP.</p>
o=	<p>The sender's username, a session identifier, a session version number, the network type (usually IN for Internet), the address type (for example, IPv4 or IPv6), and the sending device's IP address. The o= field becomes a universal identifier for this version of this session description. For example:</p> <pre>o=PhoneA 5462346 332134 IN IP4 10.31.101.20</pre>
r=	<p>Repeat times for a session. Used if a session will be repeated at one or more timed intervals. Not normally used for VoIP calls. The times can be in different formats. For example.</p> <pre>r=7d 1h 0 25h r=604800 3600 0 90000</pre>
s=	<p>Any text that describes the session or s= followed by a space. For example:</p> <pre>s=Call from inviter</pre>
t=	<p>The start and stop time of the session. Sessions with no time restrictions (most VoIP calls) have a start and stop time of 0.</p> <pre>t=0 0</pre>
v=	<p>SDP protocol version. The current SDP version is 0 so the v= field is always:</p> <pre>v=0</pre>
z=	<p>Time zone adjustments. Used for scheduling repeated sessions that span the time between changing from standard to daylight savings time.</p> <pre>z=2882844526 -1h 2898848070 0</pre>

Example SIP messages

The following example SIP INVITE request message was sent by PhoneA to PhoneB. The first nine lines are the SIP headers. The SDP profile starts with **v=0** and the media part of the session profile is the last line, starting with **m=**.

```
INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
```

```

To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3

```

The following example shows a possible 200 OK SIP response message in response to the previous INVITE request message. The response includes 200 OK which indicates success, followed by an echo of the original SIP INVITE request followed by PhoneB's SDP profile.

```

SIP/2.0 200 OK
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneB@10.31.101.30
Content-Type: application/sdp
Content-Length: 107
v=0
o=PhoneB 124333 67895 IN IP4 172.20.120.30
s=Hello!
t=0 0
c=IN IP4 172.20.120.30
m=audio 3456 RTP 0

```

SIP can support multiple media streams for a single SIP session. Each media stream will have its own c= and m= lines in the body of the message. For example, the following message includes three media streams:

```

INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.20:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49172 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49174 RTP 0 3

```

The SIP session helper

The SIP session-helper is a high-performance solution that provides basic support for SIP calls passing through the FortiGate unit by opening SIP and RTP pinholes and by performing NAT of the addresses in SIP messages.

The SIP session helper:

- Understands SIP dialog messages.
- Keeps the states of the SIP transactions between SIP UAs and SIP servers.
- Translates SIP header and SDP information to account for NAT operations performed by the FortiGate unit.
- Opens up and closes dynamic SIP pinholes for SIP signalling traffic.
- Opens up and closes dynamic RTP and RTSP pinholes for RTP and RTSP media traffic.
- Provides basic SIP security as an access control device.
- Uses the intrusion protection (IPS) engine to perform basic SIP protocol checks.

SIP session helper configuration overview

The SIP session helper is enabled by default and set to listen for SIP traffic on TCP or UDP port 5060. SIP sessions using port 5060 accepted by a security policy that does not include a VoIP profile are processed by the SIP session helper.

You can enable and disable the SIP session helper, change the TCP or UDP port that the session helper listens on for SIP traffic, and enable or disable SIP NAT tracing. If the FortiGate unit is operating with multiple VDOMs, each VDOM can have a different SIP session helper configuration.

To have the SIP session helper process SIP sessions you need to add a security policy that accepts SIP sessions on the configured SIP UDP or TCP ports. The security policies can have service set to ANY, or to the SIP pre-defined firewall service, or a custom firewall service. The SIP pre-defined firewall service restricts the security policy to only accepting sessions on UDP port 5060.

If NAT is enabled for security policies that accept SIP traffic, the SIP session helper translates addresses in SIP headers and in the RDP profile and opens up pinholes as required for the SIP traffic. This includes security policies that perform source NAT and security policies that contain virtual IPs that perform destination NAT and port forwarding. No special SIP configuration is required for this address translation to occur, it is all handled automatically by the SIP session helper according to the NAT configuration of the security policy that accepts the SIP session.

To use the SIP session helper you must not add a VoIP profile to the security policy. If you add a VoIP profile, SIP traffic bypasses the SIP session helper and is processed by the SIP ALG.



In most cases you would want to use the SIP ALG since the SIP session helper provides limited functionality. However, the SIP session helper is available and can be useful for high-performance solutions where a high level of SIP security is not a requirement.

Disabling and enable the SIP session helper

You can use the following steps to disable the SIP session helper. You might want to disable the SIP session helper if you don't want the FortiGate unit to apply NAT or other SIP session help features to SIP traffic. With the SIP session helper disabled, the FortiGate unit can still accept SIP sessions if they are allowed by a security policy, but the FortiGate unit will not be able to open pinholes or NAT the addresses in the SIP messages.

To disable the sip session helper

- 1 Enter the following command to find the sip session helper entry in the session-helper list:

```
show system session-helper
.
.
.
edit 13
    set name sip
    set port 5060
    set protocol 17
next
.
.
.
```

This command output shows that the sip session helper listens in UDP port 5060 for SIP sessions.

- 2 Enter the following command to delete session-helper list entry number 13 to disable the sip session helper:

```
config system session-helper
    delete 13
end
```

If you want to use the SIP session helper you can verify whether it is enabled or disabled using the `show system session-helper` command.



You do not have to disable the SIP session helper to use the SIP ALG.

If the SIP session helper has been disabled by being removed from the session-helper list you can use the following command to enable the SIP session helper by adding it back to the session helper list:

```
config system session-helper
    edit 0
        set name sip
        set port 5060
        set protocol 17
    end
```

Changing the port numbers that the SIP session helper listens on

You can use the following command to change the port number that the SIP session helper listens on for SIP traffic to 5061. The SIP session helper listens on the same port number for UDP and TCP SIP sessions. In this example, the SIP session helper is session helper 13:

```
config system session-helper
edit 13
set port 5061
end
```



The `config system settings` options `sip-tcp-port` and `sip-udp-port` control the ports that the SIP ALG listens on for SIP sessions. See [“Changing the port numbers that the SIP ALG listens on” on page 2545](#).

Your FortiGate unit may use a different session helper number for SIP. Enter the following command to view the session helpers:


```
show system session-helper
.
.
.
edit 13
set name sip
set port 5060
set protocol 17
end
.
.
.
```

Configuration example: SIP session helper in Transparent Mode

[Figure 267](#) shows an example SIP network consisting of a FortiGate unit operating in Transparent mode between two SIP phones. Since the FortiGate unit is operating in Transparent mode both phones are on the same network and the FortiGate unit and the SIP session helper does not perform NAT. Even though the SIP session helper is not performing NAT you can use this configuration to apply SIP session helper security features to the SIP traffic.

The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

Figure 267: SIP network with FortiGate unit in Transparent mode



General configuration steps

The following general configuration steps are required for this SIP configuration. This example includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

- 1 Add firewall addresses for Phone A and Phone B.
- 2 Add a security policy that accepts SIP sessions initiated by Phone A.
- 3 Add a security policy that accepts SIP sessions initiated by Phone B.

Configuration steps - web-based manager

To add firewall addresses for the SIP phones

- 1 Go to *Firewall Objects > Address*.
- 2 Add the following addresses for Phone A and Phone B:

Address Name	Phone_A
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	port1

Address Name	Phone_B
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.30/255.255.255.255
Interface	port2

To add security policies to accept SIP sessions

- 1 Go to *Policy > Policy > Policy*.
- 2 Select Create New to add a security policy.
- 3 Add a security policy to allow Phone A to send SIP request messages to Phone B:

Source Interface/Zone	port1
Source Address	Phone_A
Destination Interface/Zone	port2
Destination Address	Phone_B
Schedule	always
Service	SIP
Action	ACCEPT

- 4 Select OK.
- 5 Add a security policy to allow Phone B to send SIP request messages to Phone A:

Source Interface/Zone	port2
Source Address	Phone_B

Destination Interface/Zone	port1
Destination Address	Phone_A
Schedule	always
Service	SIP
Action	ACCEPT

- 6 Select OK.

Configuration steps - CLI

To add firewall addresses for Phone A and Phone B and security policies to accept SIP sessions

- 1 Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
  edit Phone_A
    set associated interface port1
    set type ipmask
    set subnet 10.31.101.20 255.255.255.255
  next
  edit Phone_B
    set associated interface port2
    set type ipmask
    set subnet 10.31.101.30 255.255.255.255
  end
```

- 2 Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr Phone_A
    set dstaddr Phone_B
    set action accept
    set schedule always
    set service SIP
  next
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr Phone_B
    set dstaddr Phone_A
    set action accept
    set schedule always
    set service SIP
    set utm-status enable
  end
```

SIP session helper diagnose commands

You can use the `diagnose sys sip` commands to display diagnostic information for the SIP session helper.

Use the following command to set the debug level for the SIP session helper. Different debug masks display different levels of detail about SIP session helper activity.

```
diagnose sys sip debug-mask <debug_mask_int>
```

Use the following command to display the current list of SIP dialogs being processed by the SIP session help. You can also use the `clear` option to delete all active SIP dialogs being processed by the SIP session helper.

```
diagnose sys sip dialog {clear | list}
```

Use the following command to display the current list of SIP NAT address mapping tables being used by the SIP session helper.

```
diagnose sys sip mapping list
```

Use the following command to display the current SIP session helper activity including information about the SIP dialogs, mappings, and other SIP session help counts. This command can be useful to get an overview of what the SIP session helper is currently doing.

```
diagnose sys sip status
```

The SIP ALG

In most cases you should use the SIP Application Layer Gateway (ALG) for processing SIP sessions. The SIP ALG provides the same basic SIP support as the SIP session helper. Additionally, the SIP ALG provides a wide range of features that protect your network from SIP attacks, can apply rate limiting to SIP sessions, can check the syntax of SIP and SDP content of SIP messages, and provide detailed logging and reporting of SIP activity.

You apply the SIP ALG to SIP traffic by adding a VoIP profile with SIP enabled to a security policy that accepts SIP traffic. The SIP session helper is automatically bypassed by traffic accepted by a security policy that includes a VoIP profile.

As shown in [Figure 268](#), the FortiGate SIP ALG intercepts SIP packets after they have been routed by the routing module, accepted by a security policy and passed through DoS and IPS Sensors (if DoS and IPS are enabled). The ALG raises SIP packets to the application layer, analyzes the SIP and SDP addressing information in the SIP messages, makes adjustments (for example, NAT) to this addressing if required, and then sends the packets out the egress interface to their destination.

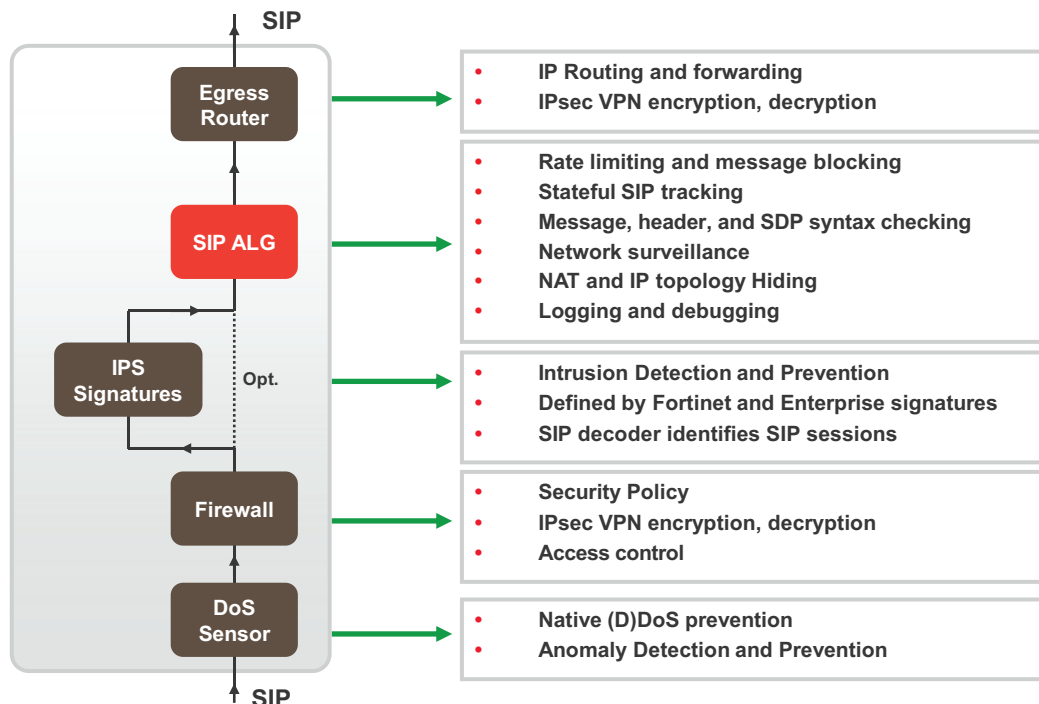
The SIP ALG provides:

- All the same features as the SIP session help including NAT and SIP and RTP Pinholes.

In addition for the ALG you can enable or disable RTP pinholing, SIP register pinholing and SIP contact pinholing. In a signalling only environment where the RTP stream bypasses the FortiGate unit, you can disable RTP pinholing to improve performance.

- SIP TCP and UDP support
- SIP Message order checking
- Configurable Header line length maximums

Figure 268: The SIP ALG works at the application level after ingress packets are accepted by a security policy



- Message fragment assembly (TCP)

If SIP messages are fragmented across multiple packets, the FortiGate unit assembles the fragments, does inspection and pass the message in its entirety to the SIP server as one packet. This offloads the server from doing all the TCP processing of fragments.

- L4 Protocol Translation

- Message Flood Protection

Protects a SIP server from intentional or unintentional DoS of flooding INVITE, REGISTER, and other SIP methods by allowing control of the rate that these messages pass through the FortiGate unit.

- SIP message type filtering

The FortiGate unit can prevent specified SIP message types from passing through the FortiGate unit to a SIP server. For example In a voice only SIP implementation, there may be no need to permit a SUBSCRIBE message to ever make it's way to the SIP call processor. Also, if a SIP server cannot process some SIP message types you can use SIP message type filtering to block them. For example, a SIP server could have a bug that prevents it from processing certain SIP messages. In this case you can temporarily block these message types until problem with the SIP server has been fixed.

- SIP statistics and logging

- SIP over IPv6

- Deep SIP message syntax checking (also called deep SIP header inspection or SIP fuzzing protection). Prevents attacks that use malformed SIP messages. Can check many SIP headers and SDP statements. Configurable bypass and modification options.

- Hosted NAT traversal, Resolves IP address issue in SIP and SDP lines due to NAT-PT in far end firewall. Important feature for VoIP access networks.
- SIP High Availability (HA), including active-passive clustering and session pickup (session failover) for SIP sessions.
- Geographical Redundancy. In an HA configuration, if the active SIP server fails (missing SIP heartbeat messages or SIP traffic) SIP sessions can be redirected to a secondary SIP server in another location.
- SIP per request method message rate limitation with configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.
- RTP Bypass, Supports configurations with and without RTP pinholing. May inspect and protect SIP signaling only.
- SIP NAT with IP address conservation. Performs SIP and RTP aware IP Network Address translation. Preserves the lost IP address information in the SDP profile i= line for later processing/debugging in the SIP server. See [“NAT with IP address conservation” on page 2573](#).
- IP topology hiding

The IP topology of a network can be hidden through NAT and NAPT manipulation of IP and SIP level addressing. For example, see [“SIP NAT configuration example: destination address translation \(destination NAT\)” on page 2567](#).

- SIP inspection without address translation

The SIP ALG inspects SIP messages but addresses in the messages are not translated. This feature can be applied to a FortiGate unit operating in Transparent mode or in NAT/Route mode. In Transparent mode you add normal Transparent mode security policies that enable the SIP ALG and include a VoIP profile that causes the SIP ALG to inspect SIP traffic as required. For an example configuration, see [“Configuration example: SIP in Transparent Mode” on page 2550](#).

For a FortiGate unit operating in NAT/Route mode, if SIP traffic can pass between different networks without requiring NAT because is supported by the routing configuration, you can add security policies that accept SIP traffic without enabling NAT. In the VoIP profile you can configure the SIP ALG to inspect SIP traffic as required.

SIP ALG configuration overview

To apply the SIP ALG, you add a SIP VoIP profile to a security policy that accepts SIP sessions. All SIP sessions accepted by the security policy will be processed by the SIP ALG using the settings in the VoIP profile. The VoIP profile contains settings that are applied to SIP, Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE) and Skinny Call Control Protocol (SCCP) sessions. You configure SIP and SCCP settings separately. SIP settings also apply to SIMPLE sessions.

Enabling VoIP support on the web-based manager

Before you begin adding VoIP profiles to security policies you may have to enable VoIP support on the web-based manager. To do this, on the web-based manager go to *System > Admin > Settings* and make sure that the VoIP Support on GUI checkbox is selected.

From the CLI you can also enter the following command enable VoIP support on the GUI:

```
config system global
    set gui-voip-profile enable
```

end

VoIP profiles

To add a new VoIP profile from the web-based manager go to *UTM Profiles > VoIP > Profile* and select *Create New*.

For SIP, from the web-based manager you can configure the VoIP profile to limit the number of SIP REGISTER and INVITE requests and enable logging of SIP sessions and SIP violations. Many additional options for configuring how the ALG processes SIP sessions are available from the CLI.

Use the following command to add a VoIP profile named VoIP_Pro_1 from the CLI:

```
config voip profile
  edit VoIP_Pro_1
end
```

FortiGate units include two pre-defined VoIP profiles. On the web-based manager these profiles look identical. However, the CLI-only settings result in the following functionality.

default	<p>The most commonly used VoIP profile. This profile enables both SIP and SCCP and places the minimum restrictions on what calls will be allowed to negotiate. This profile allows normal SCCP, SIP and RTP sessions and enables the following security settings:</p> <ul style="list-style-type: none"> • <code>block-long-lines</code> to block SIP messages with lines that exceed maximum line lengths. • <code>block-unknown</code> to block unrecognized SIP request messages. • <code>log-call-summary</code> to write log messages that record SIP call progress (similar to DLP archiving). • <code>nat-trace</code> (see “NAT with IP address conservation” on page 2573). • <code>contact-fixup</code> perform NAT on the IP addresses and port numbers in SIP headers in SIP CONTACT messages even if they don't match the session's IP address and port numbers.
strict	<p>This profile is available for users who want to validate SIP messages and to only allow SIP sessions that are compliant with RFC 3261. In addition to the settings in the default VoIP profile, the strict profile sets all SIP deep message inspection header checking to block and drop SIP messages that contain malformed SIP or SDP lines that can be detected by the ALG. For more information about SIP deep header inspection, see “Deep SIP message inspection” on page 2586.</p>

Neither of the default profiles applies SIP rate limiting or message blocking. To apply more ALG features to SIP sessions you can clone (copy) the pre-defined VoIP profiles and make your own modifications to them. For example, to clone the default profile and configure the limit for SIP NOTIFY request messages to 1000 messages per second per security policy and block SIP INFO request messages.

```
config voip profile
  clone default to my_voip_pro
  edit my_voip_pro
    config sip
      set notify-rate 1000
      set block-info enable
    end
  end
end
```

Changing the port numbers that the SIP ALG listens on

Most SIP configurations use TCP or UDP port 5060 for SIP sessions. If your SIP network uses different ports for SIP sessions you can use the following command to configure the SIP ALG to listen on a different TCP or UDP ports. For example, to change the TCP port to 5061 and the UDP port to 5065.

```
config system settings
  set sip-tcp-port 5061
  set sip-udp-port 5065
end
```

Disabling the SIP ALG in a VoIP profile

SIP is enabled by default in a VoIP profile. Usually you would want SIP to be enabled in a VoIP profile. But in some cases if you are just using the VoIP profile for SCCP you can use the following command to disable SIP in a VoIP profile.

```
config voip profile
  edit VoIP_Pro_2
    config sip
      set status disable
    end
  end
```

SIP ALG get and diagnose commands

You can use the following commands to display diagnostic information for the SIP ALG.

Use the following commands to enter a test level to display information about the SIP ALG.

```
get test sip <test_level_int>
diagnose test application sip <test_level_int>
```

Use the following command to list all active SIP calls being processed by the SIP ALG. You can also use the `clear` option to delete all active SIP calls being processed by the SIP ALG.

```
diagnose sys sip-proxy calls {clear | list}
```

Use the following commands to use filters to display specific information about the SIP ALG and the session that it is processing.

```
diagnose sys sip-proxy filter <filter_options>
diagnose sys sip-proxy log-filter <filter_options>
```

Use the following command to display the active SIP rate limiting meters and their current settings.

```
diagnose sys sip-proxy meters list
```

Use the following command to display status information about the SIP sessions being processed by the SIP ALG. You can also clear all SIP ALG statistics.

```
diagnose sys sip-proxy stats {clear | list}
```

Conflicts between the SIP ALG and the session helper

Even if the SIP session helper is enabled, if a security policy with a VoIP profile that has SIP enabled accepts a SIP session on the TCP or UDP port that the SIP ALG listens on the ALG is used. You don't need to turn off the session helper to use the ALG.

You may find that the session helper is being used for some SIP sessions even when you only want to use the ALG. This happens if a policy that does not include a VoIP profile is accepting SIP sessions. The VoIP profile could have been left out of the policy by mistake or the wrong policy could be accepting SIP sessions.

Consider a configuration with a SIP server on a private network that is contacted by SIP phones on the Internet and on the private network (similar to the configuration in [Figure 264 on page 2524](#)). The FortiGate unit that provides NAT between the private network and the Internet requires a security policy with a firewall virtual IP that allows the SIP phones on the Internet to contact the SIP server. The FortiGate unit also requires outgoing security policies to allow the SIP phones and the SIP server to contact the SIP phones on the Internet.

If a VoIP profile is not added to one of the outgoing security policies the SIP sessions accepted by that policy will be processed by the SIP session helper instead of the SIP ALG. Also, it's possible that some of the SIP sessions could be accepted by a general outgoing policy instead of the policy intended for SIP traffic. You can fix the first problem by adding a VoIP profile to the policy. You can fix the second problem by reviewing the security policy order and source and destination addresses in the security policies and determining if there is a conflict between these and the IP addresses of the SIP server or SIP phones on the Internal network.

You can use `diagnose sys sip` commands to determine if the SIP session helper is processing SIP sessions. For example, the following command displays the overall status of the SIP sessions being processed by the SIP session helper:



The `diagnose sys sip` commands only display current status information. To see activity the SIP session helper has to actually be processing SIP sessions when you enter the command. For example, if the SIP session helper had been used for processing calls that ended 5 minutes ago, the command output would show no SIP session helper activity.

```
diagnose sys sip status
dialogs: max=32768, used=0
mappings: used=0
dialog hash by ID: size=2048, used=0, depth=0
dialog hash by RTP: size=2048, used=0, depth=0
mapping hash: size=2048, used=0, depth=0
count0: 0
count1: 0
count2: 0
count3: 0
count4: 0
```

This command output shows that the session helper is not processing SIP sessions because all of the used and count fields are 0. If any of these fields contains non-zero values then the SIP session helper may be processing SIP sessions.

Also, you can check to see if some ALG-only features are not being applied to all SIP sessions. For example, the VoIP usage widget on the FortiGate dashboard displays statistics for SIP and SCCP calls processed by the ALG but not for calls processed by the session helper. So if you see fewer calls than expected the session helper may be processing some of them.

Other logging and monitoring features such as log messages and DLP archiving are only supported by the ALG.

Finally, you can check the policy usage and session information dashboard widgets to see if SIP sessions are being accepted by the wrong security policies.

Stateful SIP tracking, call termination, and session inactivity timeout

The SIP ALG tracks SIP dialogs over their lifespan between the first INVITE message and the Final 200 OK and ACK messages. For every SIP dialog, stateful SIP tracking reviews every SIP message and makes adjustment to SIP tracking tables as required. These adjustments include source and destination IP addresses, address translation, dialog expiration information, and media stream port changes. Such changes can also result in dynamically opening and closing pinholes. You can use the `diagnose sys sip-proxy stats list` and the `diagnose sys sip-proxy filter` command to view the SIP call data being tracked by the SIP ALG.

The SIP ALG uses the SIP Expires header line to time out a SIP dialog if the dialog is idle and a Re-INVITE or UPDATE message is not received. The SIP ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE message. If the SIP ALG receives an INVITE before the session times out, all timeout values are reset to the settings in the new INVITE message or to default values. As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the FortiGate unit is protected if a call ends prematurely.

When a SIP dialog ends normally, the SIP ALG deletes the SIP call information and closes open pinholes. A SIP call can also end abnormally due to an unexpected signaling or transport event that cuts off the call. When a call ends abnormally the SIP messages to end the call may not be sent or received. A call can end abnormally for the following reasons:

- Phones or servers crash during a call and a BYE message is not received.
- To attack a SIP system, a malicious user never send a BYE message.
- Poor implementations of SIP fail to process Record-Route messages and never send a BYE message.
- Network failures prevent a BYE message from being received.

Any phone or server in a SIP call can cancel the call by sending a CANCEL message. When a CANCEL message is received by the FortiGate unit, the SIP ALG closes open pinholes. Before terminating the call, the ALG waits for the final 200 OK message.

The SIP ALG can be configured to terminate SIP calls if the SIP dialog message flow or the call RTP (media) stream is interrupted and does not recover. You can use the following commands to configure terminating inactive SIP sessions and to set timers or counters to control when the call is terminated by the SIP ALG.

Adding a media stream timeout for SIP calls

Use the following command in a VoIP profile to terminate SIP calls accepted by a security policy containing the VoIP profile when the RTP media stream is idle for 100 seconds.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set call-keepalive 100
    end
  end
```

You can adjust this setting between 1 and 10,080 seconds. The default call keepalive setting of 0 disables terminating a call if the media stream is interrupted. Set call keepalive higher if your network has latency problems that could temporarily interrupt media streams. If you have configured call keepalive and the FortiGate unit terminates calls unexpectedly you can increase the call keepalive time to resolve the problem.



Call keep alive should be used with caution because enabling this feature results in extra FortiGate CPU overhead and can cause delay/jitter for the VoIP call. Also, the FortiGate unit terminates the call without sending SIP messages to end the call. And if the SIP endpoints send SIP messages to terminate the call they will be blocked by the FortiGate unit if they are sent after the FortiGate unit terminates the call.

Adding an idle dialog setting for SIP calls

Use the following command in a VoIP profile to terminate SIP calls when for a single security policy, when the configured number of SIP calls (or dialogs) has stopped receiving SIP messages or has not received legitimate SIP messages. Using this command you can configure how many dialogs that have been accepted by a security policy that the VoIP profile is added to become idle before the SIP ALG deletes the oldest ones. The following command sets the maximum number of idle dialogs to 200:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-idle-dialogs 200
    end
  end
```

Idle dialogs would usually be dialogs that have been interrupted because of errors or problems or as the result of a SIP attack that opens a large number of SIP dialogs without closing them. This command provides a way to remove these dialogs from the dialog table and recover memory and resources being used by these open and idle dialogs.

You can adjust this setting between 1 and a very high number. The default maximum idle dialogs setting of 0 disables this feature. Set maximum dialogs higher if your network has latency problems that could temporarily interrupt SIP messaging. If you have configured max idle dialogs and the FortiGate unit terminates calls unexpectedly you can increase the max idle dialogs number to resolve the problem.

Changing how long to wait for call setup to complete

In some cases and some configurations your SIP system may experience delays during call setup. If this happens, some SIP ALG timers may expire before call setup is complete and drop the call. In some cases you may also want to reduce the amount of time the SIP ALG allows for call setup to complete.

You can use the `provisional-invite-expiry-time` SIP VoIP profile option to control how long the SIP ALG waits for provisional INVITE messages before assuming that the call setup has been interrupted and the SIP call should be dropped. The default value for this timer is 210 seconds. You can change it to between 10 and 3600 seconds.

Use the following command to change the expiry time to 100 seconds.

```
config voip profile
  edit Profile_name
    config sip
      set provisional-invite-expiry-time 100
    end
  end
```

SIP and RTP/RTCP

FortiGate units support the Real Time Protocol (RTP) application layer protocol for the VoIP call audio stream. RTP uses dynamically assigned port numbers that can change during a call. SIP control messages that start a call and that are sent during the call inform callers of the port number to use and of port number changes during the call.

During a call, each RTP session will usually have a corresponding Real Time Control Protocol (RTCP) session. By default, the RTCP session port number is one higher than the RTP port number.

The RTP port number is included in the `m=` part of the SDP profile. In the example above, the SIP INVITE message includes RTP port number is 49170 so the RTCP port number would be 49171. In the SIP response message the RTP port number is 3456 so the RTCP port number would be 3457.

How the SIP ALG creates RTP pinholes

The SIP ALG requires the following information to create a pinhole. The SIP ALG finds this information in SIP messages and some is provided by the SIP ALG:

Protocol	UDP (Extracted from SIP messages by the SIP ALG.)
Source IP	Any
Source port	Any
Destination IP	The SIP ALG extracts the destination IP address from the <code>c=</code> line in the SDP profile. The <code>c=</code> line can appear in either the session or media part of the SDP profile. The SIP ALG uses the IP address in the <code>c=</code> line of the media part of the SDP profile first. If the media part does not contain a <code>c=</code> line, the SIP ALG checks the <code>c=</code> line in the session part of the SDP profile. If the session part of the profile doesn't contain a <code>c=</code> line the packet is dropped. Pinholes for RTP and RTCP sessions share the same destination IP address.
Destination port	The SIP ALG extracts the destination port number for RTP from the <code>m=</code> field and adds 1 to this number to get the RTCP port number.
Lifetime	The length of time during which the pinhole will be open. When the lifetime ends, the SIP ALG removes the pinhole.

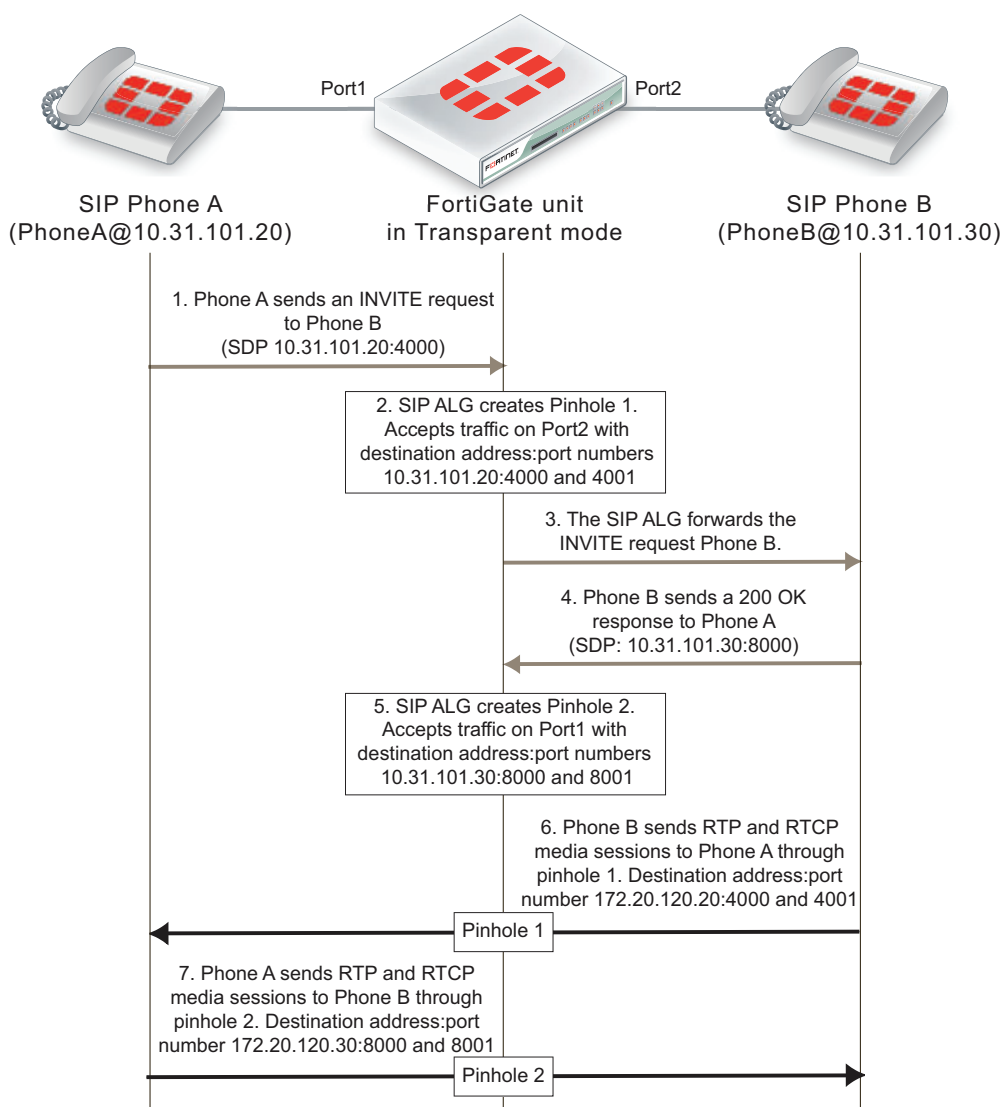
The SIP ALG keeps RTP pinholes open as long as the SIP session is alive. When the associated SIP session is terminated by the SIP ALG or the SIP phones or servers participating in the call, the RTP pinhole is closed.

Figure 269 shows a simplified call setup sequence that shows how the SIP ALG opens pinholes. Phone A and Phone B are installed on either side of a FortiGate unit operating in Transparent mode. Phone A and Phone B are on the same subnet. The FortiGate unit includes a security policy that accepts SIP sessions from port1 to port2 and from port2 to port1. The FortiGate unit does not require an RTP security policy, just the SIP policy.

You can see from this diagram that the SDP profile in the INVITE request from Phone A indicates that Phone A is expecting to receive a media stream sent to its IP address using port 4000 for RTP and port 4001 for RTCP. The SIP ALG creates pinhole 1 to allow this media traffic to pass through the FortiGate unit. Pinhole 1 is opened on the Port2 interface and will accept media traffic sent from Phone B to Phone A.

When Phone B receives the INVITE request from Phone A, Phone B will know to send media streams to Phone A using destination IP address 10.31.101.20 and ports 4000 and 4001. The 200 OK response sent from Phone B indicates that Phone B is expecting to receive a media stream sent to its IP address using ports 8000 and 8001. The SIP ALG creates pinhole 2 to allow this media traffic to pass through the FortiGate unit. Pinhole 2 is opened on the Port1 interface and will accept media traffic sent from Phone A to Phone B.

Figure 269: SIP call setup with a FortiGate unit in Transparent mode




Configuration example: SIP in Transparent Mode

Figure 270 shows an example SIP network consisting of a FortiGate unit operating in Transparent mode between two SIP phones. Since the FortiGate unit is operating in Transparent mode both phones are on the same network and the FortiGate unit and the SIP ALG does not perform NAT. Even though the SIP ALG is not performing NAT you can use this configuration to apply SIP security features to the SIP traffic.

The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

Figure 270: SIP network with FortiGate unit in Transparent mode



General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the security service to ANY to allow traffic other than SIP on UDP port 5060.

- 1 Add firewall addresses for Phone A and Phone B.
- 2 Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
- 3 Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

Configuration steps - web-based manager



Before you begin this procedure you may have to enable VoIP support on the web-based manager by going to *System > Admin > Settings* and selecting the VoIP Support on GUI checkbox.

To add firewall addresses for the SIP phones

- 1 Go to *Firewall Objects > Address*.
- 2 Add the following addresses for Phone A and Phone B:

Address Name	Phone_A
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	port1

Address Name	Phone_B
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.30/255.255.255.255
Interface	port2

To add security policies to apply the SIP ALG to SIP sessions

- 1 Go to *Policy > Policy > Policy*.
- 2 Select Create New to add a security policy.
- 3 Add a security policy to allow Phone A to send SIP request messages to Phone B:

Source Interface/Zone	port1
Source Address	Phone_A
Destination Interface/Zone	port2
Destination Address	Phone_B
Schedule	always
Service	SIP
Action	ACCEPT
UTM	Select
Protocol Options	default
Enable VoIP	Select and select the default VoIP profile.

- 4 Select OK.
- 5 Add a security policy to allow Phone B to send SIP request messages to Phone A:

Source Interface/Zone	port2
Source Address	Phone_B
Destination Interface/Zone	port1
Destination Address	Phone_A
Schedule	always
Service	SIP
Action	ACCEPT
UTM	Select
Protocol Options	default
Enable VoIP	Select (And select the default VoIP profile)

- 6 Select OK.

Configuration steps - CLI**To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions**

- 1 Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
  edit Phone_A
    set associated interface port1
    set type ipmask
    set subnet 10.31.101.20 255.255.255.255
  next
```

```
edit Phone_B
  set associated interface port2
  set type ipmask
  set subnet 10.31.101.30 255.255.255.255
end
```

- 2 Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr Phone_A
    set dstaddr Phone_B
    set action accept
    set schedule always
    set service SIP
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
  next
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr Phone_B
    set dstaddr Phone_A
    set action accept
    set schedule always
    set service SIP
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
  end
```

RTP enable/disable (RTP bypass)

You can configure the SIP ALG to stop from opening RTP pinholes. Called RTP bypass, this configuration can be used when you want to apply SIP ALG features to SIP signalling messages but do not want the RTP media streams to pass through the FortiGate unit. The FortiGate unit only acts as a signalling firewall and RTP media session bypass the FortiGate unit and no pinholes need to be created.

Enter the following command to enable RTP bypass in a VoIP profile by disabling opening RTP pinholes:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set rtp disable
    end
  end
```

Opening and closing SIP register and non-register pinholes

You can use the `open-register-pinhole` and `open-contact-pinhole` VoIP profile CLI options to control whether the FortiGate unit opens register and non-register pinholes. Non-register pinholes are usually opened for SIP INVITE requests.

By default for new VoIP profiles and for both pre-defined VoIP profiles `open-register-pinhole` is enabled and the FortiGate unit opens pinholes for SIP Register request messages. You can disable `open-register-pinhole` so that the FortiGate unit does not open pinholes for SIP Register request messages.

By default for new VoIP profiles and for the default pre-defined VoIP profile `open-contact-pinhole` is enabled and the FortiGate unit opens pinholes for non-Register SIP request messages. You can disable `open-contact-pinhole` so that the FortiGate unit does not open pinholes for non-register requests. This option is not enabled for the strict pre-defined VoIP profile.

Usually you would want to open these pinholes. Keeping them closed may prevent SIP from functioning properly through the FortiGate unit. They can be disabled, however, for interconnect scenarios (where all SIP traffic is between proxies and traveling over a single session). In some cases these settings can also be disabled in access scenarios if it is known that all users will be registering regularly so that their contact information can be learned from the register request.

You might want to prevent pinholes from being opened to avoid creating a pinhole for every register or non-register request. Each pinhole uses additional system memory, which can affect system performance if there are hundreds or thousands of users, and requires refreshing which can take a relatively long amount of time if there are thousands of active calls.

To configure a VoIP profile to prevent opening register and non-register pinholes:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set open-register-pinhole disable
      set open-contact-pinhole disable
    end
  end
```

In some cases you may not want to open pinholes for the port numbers specified in SIP Contact headers. For example, in an interconnect scenario when a FortiGate unit is installed between two SIP servers and the only SIP traffic through the FortiGate unit is between these SIP servers pinholes may not need to be opened for the port numbers specified in the Contact header lines.

If you disable `open-register-pinhole` then pinholes are not opened for ports in Contact header lines in SIP Register messages. If you disable `open-contact-pinhole` then pinholes are not opened for ports in Contact header lines in all SIP messages except SIP Register messages.

Accepting SIP register responses

You can enable the VoIP profile `reg-diff-port` options to accept a SIP Register response message from a SIP server even if the source port of the Register response message is different from the destination port.

Most SIP servers use 5060 as the source port in the SIP register response. Some SIP servers, however, may use a different source port. If your SIP server uses a different source port, you can enable `reg-diff-port` and the SIP ALG will create a temporary pinhole when Register request from a SIP client includes a different source port. The FortiGate unit will accept a SIP Register response with any source port number from the SIP server.

Enter the following command to enable accepting any source port from a SIP server:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set reg-diff-port enable
    end
  end
```

How the SIP ALG performs NAT

In most Network Address Translation (NAT) configurations, multiple hosts in a private network share a single public IP address to access the Internet. For sessions originating on the private network for the Internet, NAT replaces the private IP address of the PC in the private subnet with the public IP address of the NAT device. The NAT device converts the public IP address for responses from the Internet back into the private address before sending the response over the private network to the originator of the session.

Using NAT with SIP is more complex because of the IP addresses and media stream port numbers used in SIP message headers and bodies. When a caller on the private network sends a SIP message to a phone or SIP server on the Internet, the SIP ALG must translate the private network addresses in the SIP message to IP addresses and port numbers that are valid on the Internet. When the response message is sent back to the caller, the SIP ALG must translate these addresses back to valid private network addresses.

In addition, the media streams generated by the SIP session are independent of the SIP message sessions and use varying port numbers that can also change during the media session. The SIP ALG opens pinholes to accept these media sessions, using the information in the SIP messages to determine the pinholes to open. The ALG may also perform port translation on the media sessions.

When an INVITE message is received by the SIP ALG, the FortiGate unit extracts addressing and port number information from the message header and stores it in a SIP dialog table. Similar to an IP session table the data in the dialog table is used to translate addresses in subsequent SIP messages that are part of the same SIP call.

When the SIP ALG receives a response to the INVITE message arrives, (for example, an ACK or 200 OK), the SIP ALG compares the addresses in the message fields against the entries in the SIP dialog table to identify the call context of the message. The SIP ALG then translates addresses in the SIP message before forwarding them to their destination.

The addressing and port number information in SDP fields is used by the ALG to reserve ports for the media session and create a NAT mapping between them and the ports in the SDP fields. Because SDP uses sequential ports for the RTP and RTCP channels, the ALG provides consecutive even-odd ports.

Source address translation

When a SIP call is started by a phone on a private network destined for a phone on the Internet, only source address translation is required. The phone on the private network attempts to contact the actual IP address of the phone on the Internet. However, the source address of the phone on the private network is not routable on the Internet so the SIP ALG must translate all private IP addresses in the SIP message into public IP addresses.

To configure the FortiGate for source address translation you add security policy that accepts sessions from the internal network destined for the Internet. You must enable NAT for the security policy and add a VoIP profile.

When a SIP request is received from the internal to the external network, the SIP ALG replaces the private network IP addresses and port numbers in the SIP message with the IP address of the FortiGate interface connected to the Internet. Depending on the content of the message, the ALG translates addresses in the Via:, Contact:, Route:, and Record-Route: SIP header fields. The message is then forwarded to the destination (either a VoIP phone or a SIP server on the Internet).

The VoIP phone or server in the Internet sends responses to these SIP messages to the external interface of the FortiGate unit. The addresses in the response messages are translated back into private network addresses and the response is forwarded to the originator of the request.

For the RTP communication between the SIP phones, the SIP ALG opens pinholes to allow media through the FortiGate unit on the dynamically assigned ports negotiated based on information in the SDP and the Via:, Contact:, and Record-Route: header fields. The pinholes also allow incoming packets to reach the Contact:, Via:, and Record-Route: IP addresses and ports. When processing return traffic, the SIP ALG inserts the original Contact:, Via:, Route:, and Record-Route: SIP fields back into the packets.

Destination address translation

Incoming calls are directed from a SIP phone on the Internet to the interface of the FortiGate unit connected to the Internet. To receive these calls you must add a security policy to accept SIP sessions from the Internet. The security policy requires a firewall virtual IP. SIP INVITE messages from the Internet connect to the external IP address of the virtual IP. The SIP ALG uses the destination address translation defined in the virtual IP to translated the addresses in the SIP message to addresses on the private network.

When a 200 OK response message arrives from the private network, the SIP ALG translates the addresses in the message to Internet addresses and opens pinholes for media sessions from the private network to the Internet.

When the ACK message is received for the 200 OK, it is also intercepted by the SIP ALG. If the ACK message contains SDP information, the SIP ALG checks to determine if the IP addresses and port numbers are not changed from the previous INVITE. If they are, the SIP ALG deletes pinholes and creates new ones as required. The ALG also monitors the Via:, Contact:, and Record-Route: SIP fields and opens new pinholes as required.

Call Re-invite messages

SIP Re-INVITE messages can dynamically add and remove media sessions during a call. When new media sessions are added to a call the SIP ALG opens new pinholes and update SIP dialog data. When media sessions are ended, the SIP ALG closes pinholes that are no longer needed and removes SIP dialog data.

How the SIP ALG translates IP addresses in SIP headers

The SIP ALG applies NAT to SIP sessions by translating the IP addresses contained in SIP headers. For example, the following SIP message contains most of the SIP fields that contain addresses that need to be translated:

```
INVITE PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 172.20.120.50:5434
From: PhoneA@10.31.101.20
To: PhoneB@172.20.120.30
Call-ID: a12abcde@172.20.120.50
Contact: PhoneA@10.31.101.20:5434
Route: <sip:example@172.20.120.50:5060>
Record-Route: <sip:example@172.20.120.50:5060>
```

How IP address translation is performed depends on whether source NAT or destination NAT is applied to the session containing the message:

Source NAT translation of IP addresses in SIP messages

Source NAT translation occurs for SIP messages sent from a phone or server on a private network to a phone or server on the Internet. The source addresses in the SIP header fields of the message are typically set to IP addresses on the private network. The SIP ALG translates these addresses to the address the FortiGate unit interface connected to the Internet.

Table 137: Source NAT translation of IP addresses in SIP request messages

SIP header	NAT action
To:	None
From:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
Call-ID:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
Via:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
Request-URI:	None
Contact:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
Record-Route:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.
Route:	Replace private network address with IP address of FortiGate unit interface connected to the Internet.

Response messages from phones or servers on the Internet are sent to the FortiGate unit interface connected to the Internet where the destination addresses are translated back to addresses on the private network before forwarding the SIP response message to the private network.

Table 138: Source NAT translation of IP addresses in SIP response messages

SIP header	NAT action
To:	None

Table 138: Source NAT translation of IP addresses in SIP response messages

SIP header	NAT action
From:	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
Call-ID:	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
Via:	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
Request-URI:	N/A
Contact:	None
Record-Route:	Replace IP address of FortiGate unit interface connected to the Internet with private network address.
Route:	Replace IP address of FortiGate unit interface connected to the Internet with private network address.

Destination NAT translation of IP addresses in SIP messages

Destination NAT translation occurs for SIP messages sent from a phone or server on the Internet to a firewall virtual IP address. The destination addresses in the SIP header fields of the message are typically set to the virtual IP address. The SIP ALG translates these addresses to the address of a SIP server or phone on the private network on the other side of the FortiGate unit.

Table 139: Destination NAT translation of IP addresses in SIP request messages

SIP header	NAT action
To:	Replace VIP address with address on the private network as defined in the firewall virtual IP.
From:	None
Call-ID:	None
Via:	None
Request-URI:	Replace VIP address with address on the private network as defined in the firewall virtual IP.
Contact:	None
Record-Route:	None
Route:	None

SIP response messages sent in response to the destination NAT translated messages are sent from a server or a phone on the private network back to the originator of the request messages on the Internet. These reply messages are accepted by the same security policy that accepted the initial request messages. The firewall VIP in the original security policy contains the information that the SIP ALG uses to translate the private network source addresses in the SIP headers into the firewall virtual IP address.

Table 140: Destination NAT translation of IP addresses in SIP response messages

SIP header	NAT action
To:	None

Table 140: Destination NAT translation of IP addresses in SIP response messages

SIP header	NAT action
From:	Replace private network address with firewall VIP address.
Call-ID:	None
Via:	None
Request-URI:	N/A
Contact:	Replace private network address with firewall VIP address.
Record-Route:	Replace private network address with firewall VIP address.
Route:	None

How the SIP ALG translates IP addresses in the SIP body

The SDP session profile attributes in the SIP body include IP addresses and port numbers that the SIP ALG uses to create pinholes for the media stream.

The SIP ALG translates IP addresses and port numbers in the o=, c=, and m= SDP lines. For example, in the following lines the ALG could translate the IP addresses in the o= and c= lines and the port number (49170) in the m= line.

```
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
```

If the SDP session profile includes multiple RTP media streams, the SIP ALG opens pinholes and performs the required address translation for each one.

The two most important SDP attributes for the SIP ALG are c= and m=. The c= attribute is the connection information attribute. This field can appear at the session or media level. The syntax of the connection attribute is:

```
c=IN {IPv4 | IPv6} <destination_ip_address>
```

Where

- `IN` is the network type. FortiGate units support the `IN` or Internet network type.
- `{IPv4 | IPv6}` is the address type. FortiGate units support IPv4 or IPv6 addresses in SDP statements. However, FortiGate units do not support all types of IPv6 address translation. See [“SIP over IPv6” on page 2585](#).
- `<destination_IP_address>` is the unicast numeric destination IP address or domain name of the connection in either IPv4 or IPv6 format.

The syntax of the media attribute is:

```
m=audio <port_number> RTP <format_list>
```

Where

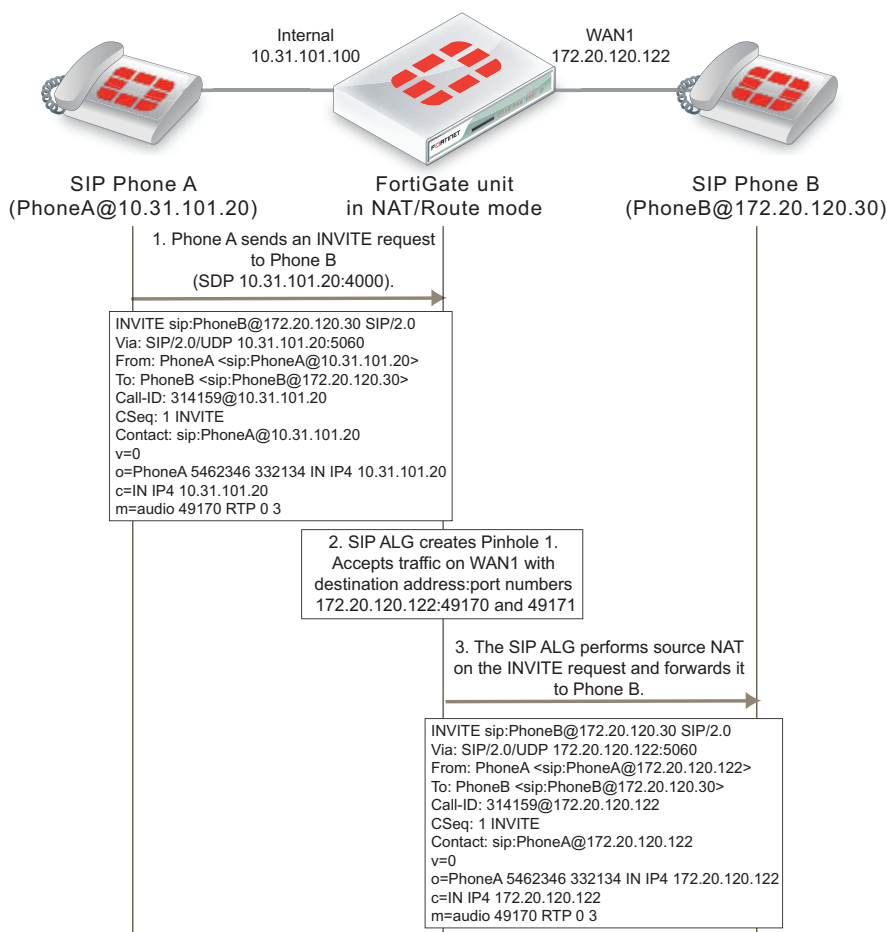
- `audio` is the media type. FortiGate units support the `audio` media type.
- `<port_number>` is the destination port number used by the media stream.
- `RTP` is the application layer transport protocol used for the media stream. FortiGate units support the Real Time Protocol (RTP) transport protocol.
- `<format_list>` is the format list that provides information about the application layer protocol that the media uses.

SIP NAT scenario: source address translation (source NAT)

Figure 271 and Figure 272 show a source address translation scenario involving two SIP phones on different networks, separated by a FortiGate unit. In the scenario, SIP Phone A sends an INVITE request to SIP Phone B and SIP Phone B replies with a 200 OK response and then the two phones start media streams with each other.

To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

Figure 271: SIP source NAT scenario part 1: INVITE request sent from Phone A to Phone B

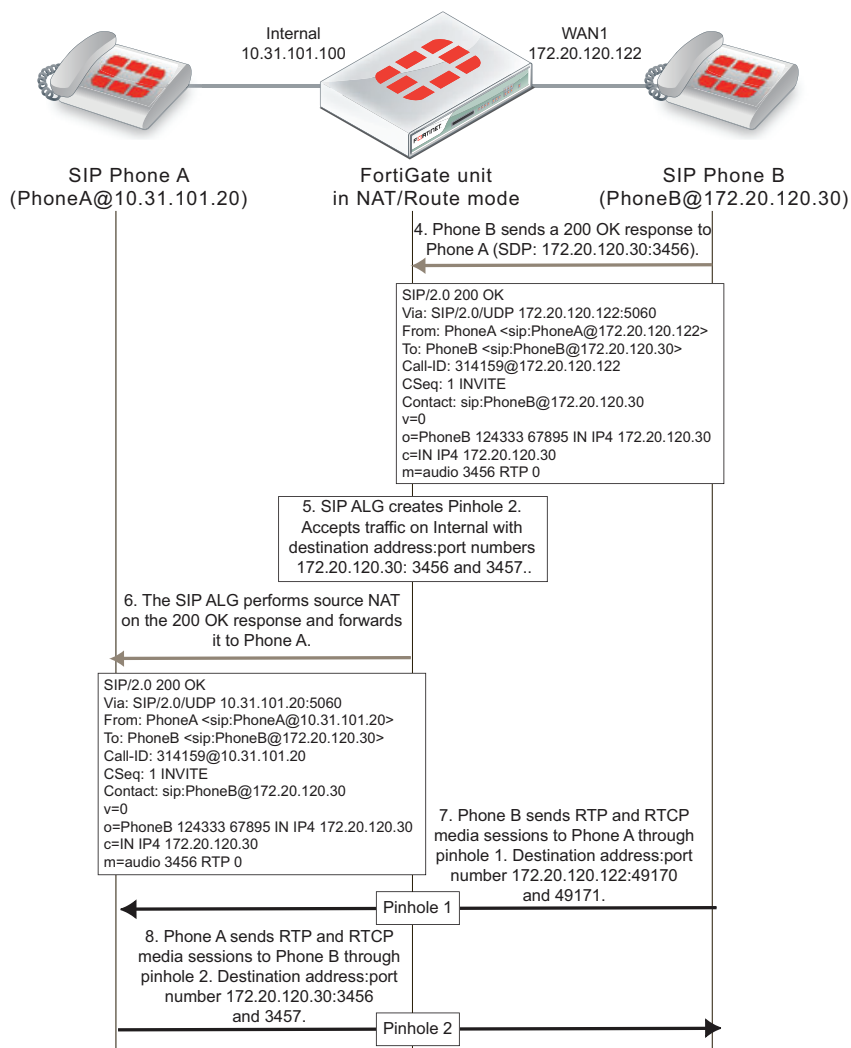


For the replies to SIP packets sent by Phone A to be routable on Phone B's network, the FortiGate unit uses source NAT to change their source address to the address of the WAN1 interface. The SIP ALG makes similar changes to the source addresses in the SIP headers and SDP profile. For example, the original INVITE request from Phone A includes the address of Phone A (10.31.101.20) in the from header line. After the INVITE request passes through the FortiGate unit, the address of Phone A in the From SIP header line is translated to 172.20.120.122, the address of the FortiGate unit WAN1 interface. As a result, Phone B will reply to SIP messages from Phone A using the WAN1 interface IP address.

The FortiGate unit also opens a pinhole so that it can accept media sessions sent to the WAN1 IP address using the port number in the m= line of the INVITE request and forward them to Phone A after translating the destination address to the IP address of Phone A.

Phone B sends the 200 OK response to the INVITE message to the WAN1 interface. The SDP profile includes the port number that Phone B wants to use for its media stream. The FortiGate unit forwards 200 OK response to Phone A after translating the addresses in the SIP and SDP lines back to the IP address of Phone A. The SIP ALG also opens a pinhole on the Internal interface that accepts media stream sessions from Phone A with destination address set to the IP address of Phone B and using the port that Phone B added to the SDP m= line.

Figure 272: SIP source NAT scenario part 2: 200 OK returned and media streams established



SIP NAT scenario: destination address translation (destination NAT)

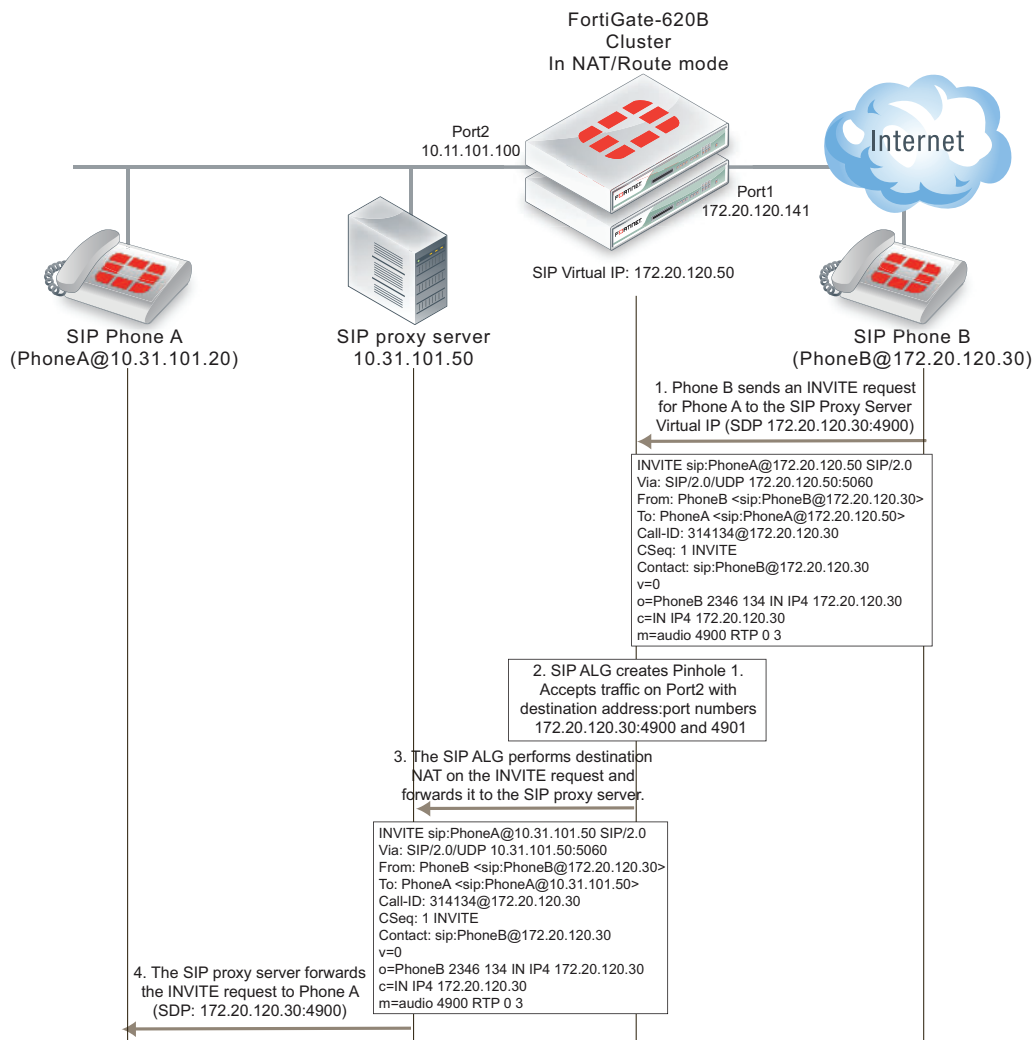
Figure 273 and Figure 274 show how the SIP ALG translates addresses in a SIP INVITE message sent from SIP Phone B on the Internet to SIP Phone A on a private network using the SIP proxy server. Because the addresses on the private network are not visible from the Internet, the security policy on the FortiGate unit that accepts SIP sessions includes a virtual IP. Phone A sends SIP the INVITE message to the virtual IP address. The FortiGate unit accepts the INVITE message packets and using the virtual IP, translates the destination address of the packet to the IP address of the SIP proxy server and forwards the SIP message to it.

To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

The SIP ALG also translates the destination addresses in the SIP message from the virtual IP address (172.20.120.50) to the SIP proxy server address (10.31.101.50). For this configuration to work, the SIP proxy server must be able to change the destination addresses for Phone A in the SIP message from the address of the SIP proxy server to the actual address of Phone A.

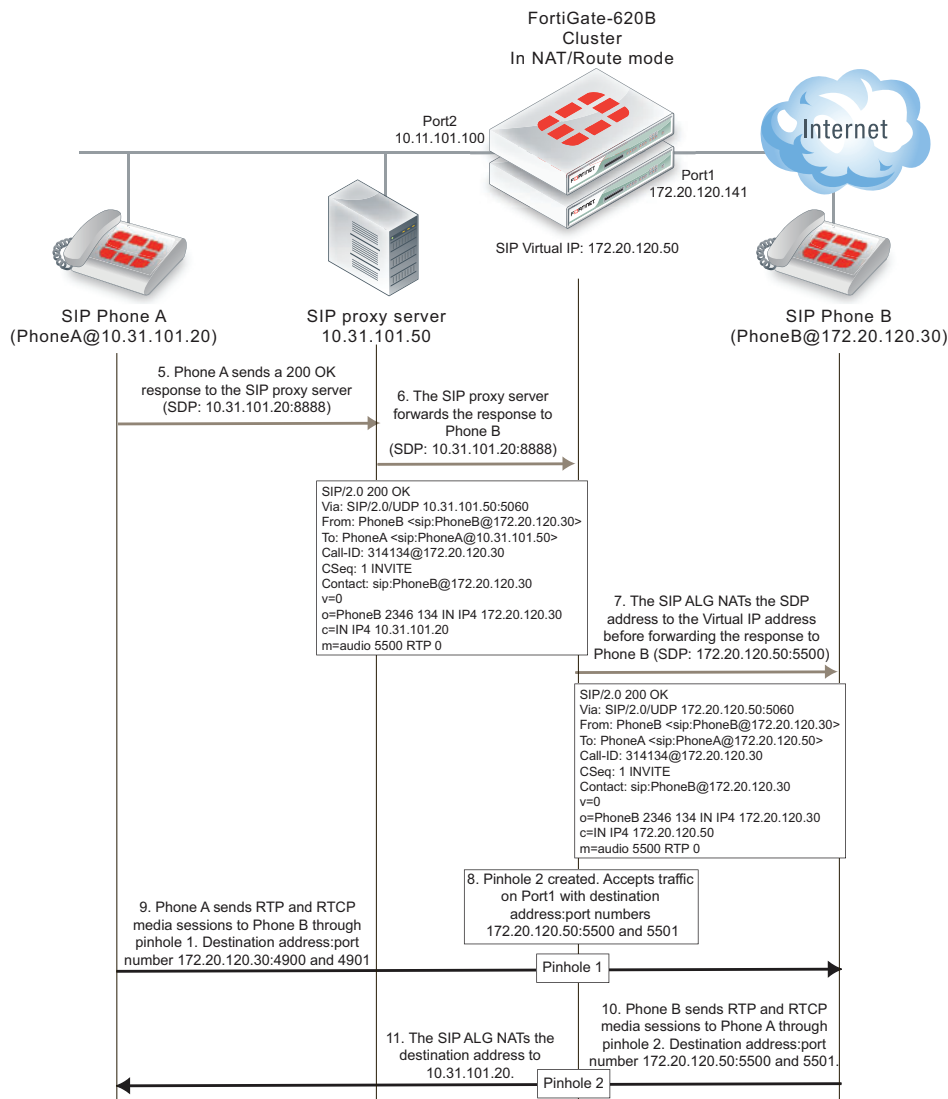
The SIP ALG also opens a pinhole on the Port2 interface that accepts media sessions from the private network to SIP Phone B using ports 4900 and 4901.

Figure 273: SIP destination NAT scenario part 1: INVITE request sent from Phone B to Phone A



Phone A sends a 200 OK response back to the SIP proxy server. The SIP proxy server forwards the response to Phone B. The FortiGate unit accepts the 100 OK response. The SIP ALG translates the Phone A addresses back to the SIP proxy server virtual IP address before forwarding the response back to Phone B. The SIP ALG also opens a pinhole using the SIP proxy server virtual IP which is the address in the o= line of the SDP profile and the port number in the m= line of the SDP code.


Figure 274: SIP destination NAT scenario part 2: 200 OK returned to Phone B and media streams established



The media stream from Phone A is accepted by pinhole one and forwarded to Phone B. The source address of this media stream is changed to the SIP proxy server virtual IP address. The media stream from Phone B is accepted by pinhole 2 and forwarded to Phone B. The destination address of this media stream is changed to the IP address of Phone A.

SIP NAT configuration example: source address translation (source NAT)

This configuration example shows how to configure the FortiGate unit to support the source address translation scenario shown in Figure 275. The FortiGate unit requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A. Both of these policies must include source NAT. In this example the networks are not hidden from each other so destination NAT is not required.

Figure 275: SIP source NAT configuration

General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

- 1 Add firewall addresses for Phone A and Phone B.
- 2 Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
- 3 Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

Configuration steps - web-based manager

To add firewall addresses for the SIP phones

- 1 Go to *Firewall Objects > Address*.
- 2 Add the following addresses for Phone A and Phone B:

Address Name	Phone_A
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	Internal

Address Name	Phone_B
Type	Subnet / IP Range
Subnet / IP Range	172.20.120.30/255.255.255.255
Interface	wan1

To add security policies to apply the SIP ALG to SIP sessions

- 1 Go to *Policy > Policy > Policy*.
- 2 Select Create New to add a security policy.
- 3 Add a security policy to allow Phone A to send SIP request messages to Phone B:

Source Interface/Zone	internal
Source Address	Phone_A

Destination Interface/Zone	wan1
Destination Address	Phone_B
Schedule	always
Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select and select the default VoIP profile.

- 4 Select OK.
- 5 Add a security policy to allow Phone B to send SIP request messages to Phone A:

Source Interface/Zone	wan1
Source Address	Phone_B
Destination Interface/Zone	internal
Destination Address	Phone_A
Schedule	always
Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select (And select the default VoIP profile)

- 6 Select OK.

Configuration steps - CLI

To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions

- 1 Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
  edit Phone_A
    set associated interface internal
    set type ipmask
    set subnet 10.31.101.20 255.255.255.255
  next
  edit Phone_B
    set associated interface wan1
    set type ipmask
    set subnet 172.20.120.30 255.255.255.255
  end
```

- 2 Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

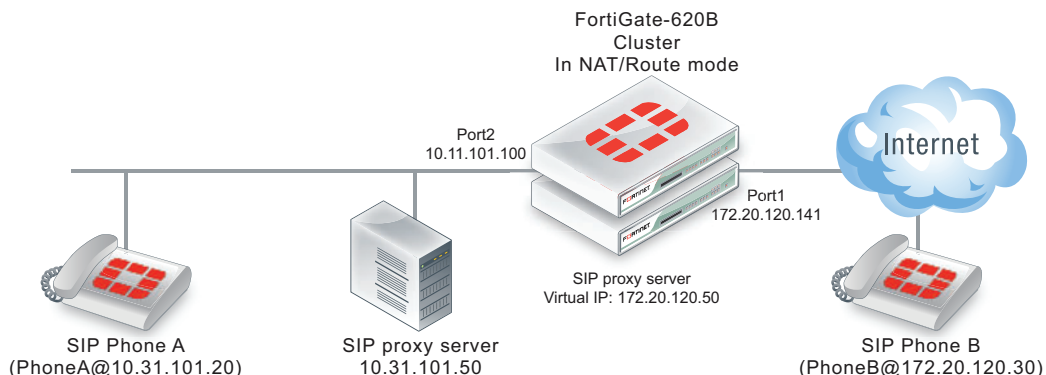
```
config firewall policy
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr Phone_A
    set dstaddr Phone_B
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
  next
  edit 0
    set srcintf wan1
    set dstintf internal
    set srcaddr Phone_B
    set dstaddr Phone_A
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
  end
```

SIP NAT configuration example: destination address translation (destination NAT)

This configuration example shows how to configure the FortiGate unit to support the destination address translation scenario shown in [Figure 276](#). The FortiGate unit requires two SIP security policies:

- A destination NAT security policy that allows SIP messages to be sent from the Internet to the private network. This policy must include destination NAT because the addresses on the private network are not routable on the Internet.
- A source NAT security policy that allows SIP messages to be sent from the private network to the Internet.

Figure 276: SIP destination NAT scenario part two: 200 OK returned to Phone B and media streams established



General configuration steps

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

- 1 Add the SIP proxy server firewall virtual IP.
- 2 Add a firewall address for the SIP proxy server on the private network.
- 3 Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
- 4 Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

Configuration steps - web-based manager

To add the SIP proxy server firewall virtual IP

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Add the SIP proxy server virtual IP.

Name	SIP_Proxy_VIP
External Interface	port1
Type	Static NAT
External IP Address/Range	172.20.120.50
Mapped IP Address/Range	10.31.101.50

To add a firewall address for the SIP proxy server

- 1 Go to *Firewall Objects > Address*.
- 2 Add the following for the SIP proxy server:

Address Name	SIP_Proxy_Server
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.50/255.255.255.255
Interface	port2

To add the security policies

- 1 Go to *Policy > Policy > Policy*.
- 2 Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	SIP_Proxy_VIP
Schedule	always
Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select and select the default VoIP profile.

- 3 Select *OK*.
- 4 Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

Source Interface/Zone	port2
Source Address	SIP_Proxy_Server
Destination Interface/Zone	port1
Destination Address	all
Schedule	always
Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select (And select the default VoIP profile)

- 5 Select *OK*.

Configuration steps - CLI**To add the SIP proxy server firewall virtual IP and firewall address**

- 1 Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
  edit SIP_Proxy_VIP
    set type static-nat
    set extip 172.20.120.50
```

```
set mappedip 10.31.101.50
set extintf port1
end
```

- 2 Enter the following command to add the SIP proxy server firewall address.

```
config firewall address
edit SIP_Proxy_Server
set associated interface port2
set type ipmask
set subnet 10.31.101.50 255.255.255.255
end
```

To add security policies

- 1 Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

```
config firewall policy
edit 0
set srcintf port1
set dstintf port2
set srcaddr all
set dstaddr SIP_Proxy_VIP
set action accept
set schedule always
set service SIP
set nat enable
set utm-status enable
set profile-protocol-options default
set voip-profile default
end
```

- 2 Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

```
config firewall policy
edit 0
set srcintf port2
set dstintf port1
set srcaddr SIP_Proxy_Server
set dstaddr all
set action accept
set schedule always
set service SIP
set nat enable
set utm-status enable
set profile-protocol-options default
set voip-profile default
end
```

Additional SIP NAT scenarios


This section lists some additional SIP NAT scenarios.

Source NAT (SIP and RTP)

In the source NAT scenario shown in Figure 277, a SIP phone connects to the Internet through a FortiGate unit with an IP address configured using PPPoE. The SIP ALG translates all private IPs in the SIP contact header into public IPs.

You need to configure an internal to external SIP security policy with NAT selected, and include a VoIP profile with SIP enabled.


Figure 277: SIP source NAT



Destination NAT (SIP and RTP)

In the following destination NAT scenario, a SIP phone can connect through the FortiGate unit to a private IP address using a firewall virtual IP (VIP). The SIP ALG translates the SIP contact header to the IP of the real SIP proxy server located on the Internet.

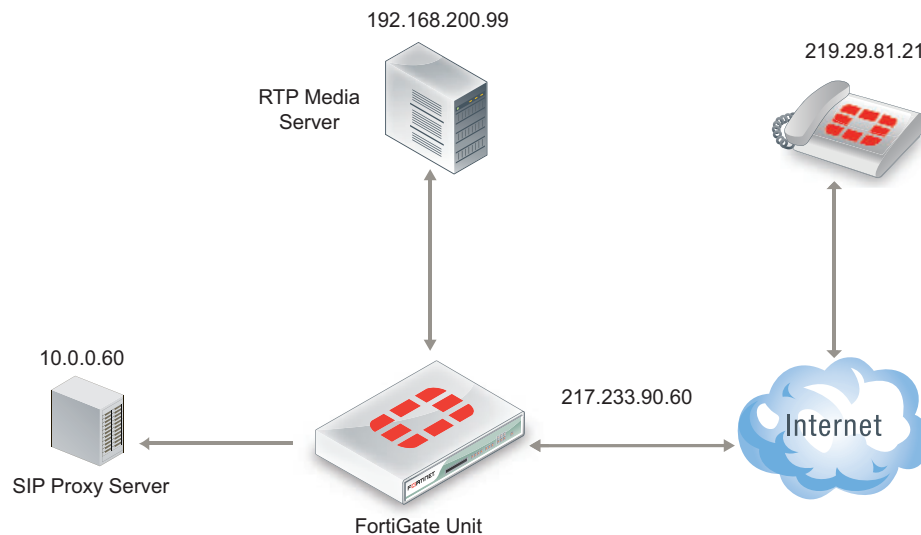
Figure 278: SIP destination NAT



In the scenario, shown in Figure 278, the SIP phone connects to a VIP (10.72.0.60). The SIP ALG translates the SIP contact header to 217.10.79.9, opens RTP pinholes, and manages NAT.

The FortiGate unit also supports a variation of this scenario where the RTP media server's IP address is hidden on a private network or DMZ.

Figure 279: SIP destination NAT-RTP media server hidden



In the scenario shown in [Figure 279](#), a SIP phone connects to the Internet. The VoIP service provider only publishes a single public IP. The FortiGate unit is configured with a firewall VIP. The SIP phone connects to the FortiGate unit (217.233.90.60) and using the VIP the FortiGate unit translates the SIP contact header to the SIP proxy server IP address (10.0.0.60). The SIP proxy server changes the SIP/SDP connection information (which tells the SIP phone which RTP media server IP it should contact) also to 217.233.90.60.

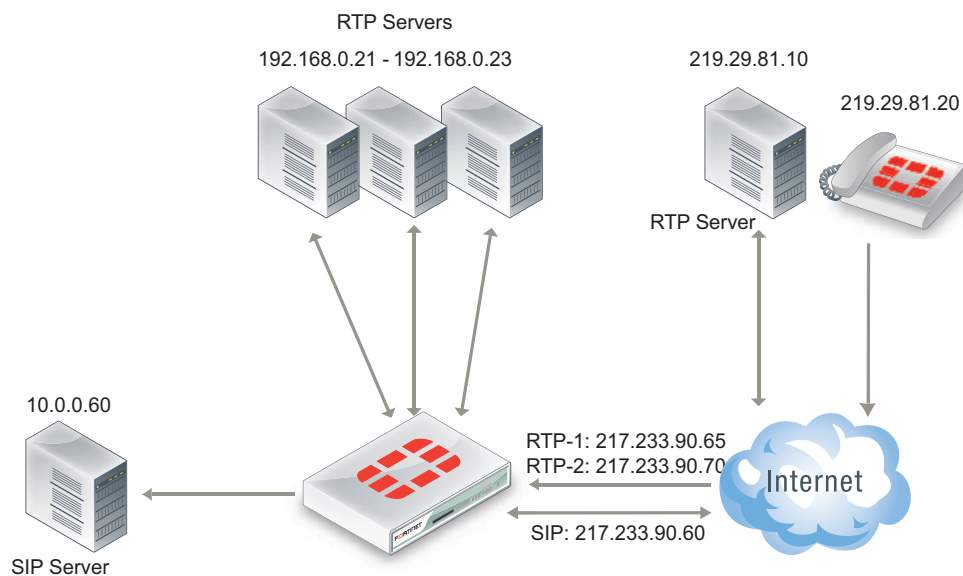
Source NAT with an IP pool

You can choose *NAT* with the *Dynamic IP Pool* option when configuring a security policy if the source IP of the SIP packets is different from the interface IP. The FortiGate ALG interprets this configuration and translates the SIP header accordingly.

This configuration also applies to destination NAT.

Different source and destination NAT for SIP and RTP

This is a more complex scenario that a SIP service provider may use. It can also be deployed in large-scale SIP environments where RTP has to be processed by the FortiGate unit and the RTP server IP has to be translated differently than the SIP server IP.

Figure 280: Different source and destination NAT for SIP and RTP

In this scenario, shown in [Figure 280](#), assume there is a SIP server and a separate media gateway. The SIP server is configured so that the SIP phone (219.29.81.20) will connect to 217.233.90.60. The media gateway (RTP server: 219.29.81.10) will connect to 217.233.90.65.

What happens is as follows:

- 1 The SIP phone connects to the SIP VIP. The FortiGate ALG translates the SIP contact header to the SIP server: 219.29.81.20 > 217.233.90.60 (> 10.0.0.60).
- 2 The SIP server carries out RTP to 217.233.90.65.
- 3 The FortiGate ALG opens pinholes, assuming that it knows the ports to be opened.
- 4 RTP is sent to the RTP-VIP (217.233.90.65.) The FortiGate ALG translates the SIP contact header to 192.168.0.21.

NAT with IP address conservation

In a source or destination NAT security policy that accepts SIP sessions, you can configure the SIP ALG or the SIP session helper to preserve the original source IP address of the SIP message in the `i=` line of the SDP profile. NAT with IP address conservation (also called SIP NAT tracing) changes the contents of SIP messages by adding the source IP address of the originator of the message into the SDP `i=` line of the SIP message. The SDP `i=` line is used for free-form text. However, if your SIP server can retrieve information from the SDP `i=` line, it can be useful for keeping a record of the source IP address of the originator of a SIP message when operating in a NAT environment. You can use this feature for billing purposes by extracting the IP address of the originator of the message.

Configuring SIP IP address conservation for the SIP ALG

You can use the following command to enable or disable SIP IP address conservation in a VoIP profile for the SIP ALG. SIP IP address conservation is enabled by default in a VoIP profile.

```
config voip profile
  edit VoIP_Pro_1
    config sip
```

```

        set nat-trace disable
    end
end

```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate unit would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

You can also use the `preserve-override` option to configure the SIP ALG to either add the original `o=` line to the end of the `i=` line or replace the `i=` line in the original message with a new `i=` line in the same form as above for adding a new `i=` line.

By default, `preserve-override` is disabled and the SIP ALG adds the original `o=` line to the end of the original `i=` line. Use the following command to configure the SIP ALG to replace the original `i=` line:

```

config voip profile
  edit VoIP_Pro_1
    config sip
      set preserve-override enable
    end
  end
end

```

Configuring SIP IP address conservation for the SIP session helper

You can use the following command to enable or disable SIP IP address conservation for the SIP session helper. IP address conservation is enabled by default for the SIP session helper.

```

config system settings
  set sip-nat-trace disable
end

```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate unit would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

Controlling how the SIP ALG NATs SIP contact header line addresses

You can enable `contact-fixup` so that the SIP ALG performs normal SIP NAT translation to SIP contact headers as SIP messages pass through the FortiGate unit.

Disable `contact-fixup` if you do not want the SIP ALG to perform normal NAT translation of the SIP contact header if a Record-Route header is also available. If `contact-fixup` is disabled, the FortiGate ALG does the following with contact headers:

- For Contact in Requests, if a Record-Route header is present and the request comes from the external network, the SIP Contact header is not translated.
- For Contact in Responses, if a Record-Route header is present and the response comes from the external network, the SIP Contact header is not translated.

If `contact-fixup` is disabled, the SIP ALG must be able to identify the external network. To identify the external network, you must use the `config system interface` command to set the `external` keyword to `enable` for the interface that is connected to the external network.

Enter the following command to perform normal NAT translation of the SIP contact header:

```
config voip profile
```

```
edit VoIP_Pro_1
  config sip
    set contact-fixup enable
  end
end
```

Controlling NAT for addresses in SDP lines

You can use the `no-sdp-fixup` option to control whether the FortiGate unit performs NAT on addresses in SDP lines in the SIP message body.

The `no-sdp-fixup` option is disabled by default and the FortiGate unit performs NAT on addresses in SDP lines. Enable this option if you don't want the FortiGate unit to perform NAT on the addresses in SDP lines.

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set no-sdp-fixup enable
    end
  end
```

Translating SIP session destination ports


Using port forwarding virtual IPs you can change the destination port of SIP sessions as they pass through the FortiGate unit.

This section describes:

- [Translating SIP sessions to a different destination port](#)
- [Translating SIP sessions to multiple destination ports](#)

Translating SIP sessions to a different destination port

To configure translating SIP sessions to a different destination port you must add a static NAT virtual IP that translates the SIP destination port to another port destination. In the example the destination port is translated from 5060 to 50601. This configuration can be used if SIP sessions use different destination ports on different networks.

Figure 281: Example translating SIP sessions to a different destination port**To translate SIP sessions to a different destination port****1 Add the static NAT virtual IP.**

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
  edit "sip_port_trans_vip"
    set type static-nat
    set portforward enable
    set protocol tcp
    set extip 172.20.120.20
    set extport 5060
    set extintf "port1"
    set mappedip 192.168.10.20
    set mappedport 50601
    set comment "Translate SIP destination port"
  end
```

2 Add a security policy that includes the virtual IP and the default VoIP profile.


```
config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "all"
    set dstaddr "sip_port_trans_vip"
    set action accept
    set schedule "always"
    set service "ANY"
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
    set comments "Translate SIP destination port"
```

end

Translating SIP sessions to multiple destination ports

You can use a load balance virtual IP to translate SIP session destination ports to a range of destination ports. In this example the destination port is translated from 5060 to the range 50601 to 50603. This configuration can be used if your SIP server is configured to receive SIP traffic on multiple ports.

Figure 282: Example translating SIP traffic to multiple destination ports



To translated SIP sessions to multiple destination ports

1 Add the load balance virtual IP.

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
  edit "sip_port_ldbl_vip"
    set type load-balance
    set portforward enable
    set protocol tcp
    set extip 172.20.120.20
    set extport 5060
    set extintf "port1"
    set mappedip 192.168.10.20
    set mappedport 50601-50603
    set comment "Translate SIP destination port range"
  end
```

2 Add a security policy that includes the virtual IP and VoIP profile.

```
config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
```

```

set srcaddr "all"
set dstaddr "sip_port_ldbl_vip"
set action accept
set schedule "always"
set service "ANY"
set utm-status enable
set profile-protocol-options default
set voip-profile default
set comments "Translate SIP destination port"
end

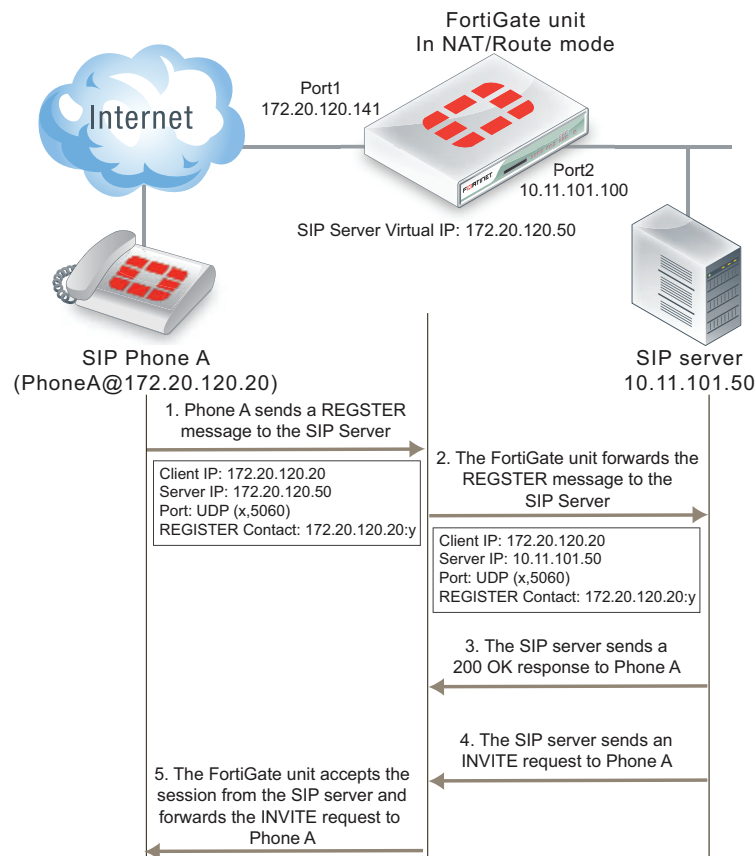
```

Enhancing SIP pinhole security

You can use the `strict-register` option in a SIP VoIP profile to open smaller pinholes.

As shown in Figure 283 when FortiGate unit is protecting a SIP server on a private network, the FortiGate unit does not have to open a pinhole for the SIP server to send INVITE requests to a SIP Phone on the Internet after the SIP Phone has registered with the server.

Figure 283: FortiGate unit protecting a SIP server on a private network



In the example, a client (SIP Phone A) sends a REGISTER request to the SIP server with the following information:

```

Client IP: 10.31.101.20
Server IP: 10.21.101.50

```



```
Port: UDP (x,5060)
REGISTER Contact: 10.31.101.20:y
```

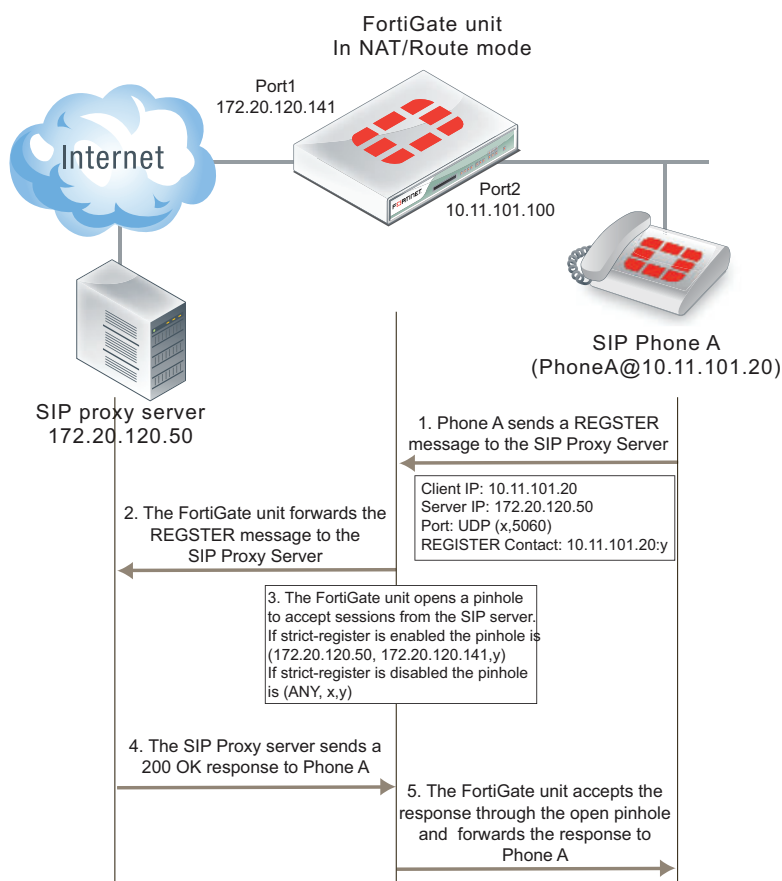
Where x and y are ports chosen by Phone A.

As soon as the server sends the 200 OK reply it can forward INVITE requests from other SIP phones to SIP Phone A. If the SIP proxy server uses the information in the REGISTER message received from SIP Phone A the INVITE messages sent to Phone A will only get through the FortiGate unit if an policy has been added to allow the server to send traffic from the private network to the Internet. Or the SIP ALG must open a pinhole to allow traffic from the server to the Internet. In most cases the FortiGate unit is protecting the SIP server so there is no reason not to add a security policy to all the SIP server to send outbound traffic to the Internet.

In a typical SOHO scenario shown in Figure 284, SIP Phone A is being protected from the Internet by a FortiGate unit. In most cases the FortiGate unit would not allow incoming traffic from the Internet to reach the private network. So the only way that an INVITE request from the SIP server can reach SIP Phone A is if the SIP ALG creates an incoming pinhole. All pinholes have three attributes:

```
(source address, destination address, destination port)
```

Figure 284: SOHO configuration, FortiGate unit protecting a network with SIP phones



The more specific a pinhole is the more secure it is because it will accept less traffic. In this situation, the pinhole would be more secure if it only accepted traffic from the SIP server. This is what happens if `strict-register` is enabled in the VoIP profile that accepts the REGISTER request from Phone A.

(SIP server IP address, client IP address, destination port)

If `strict-register` is disabled (the default configuration) the pinhole is set up with the following attributes

(ANY IP address, client IP address, destination port)

This pinhole allows connections through the FortiGate unit from ANY source address which is a much bigger and less secure pinhole. In most similar network configurations you should enable `strict-register` to improve pinhole security.

Enabling `strict-register` can cause problems when the SIP registrar and SIP proxy server are separate entities with separate IP addresses.

Enter the following command to enable `strict-register` in a VoIP profile.

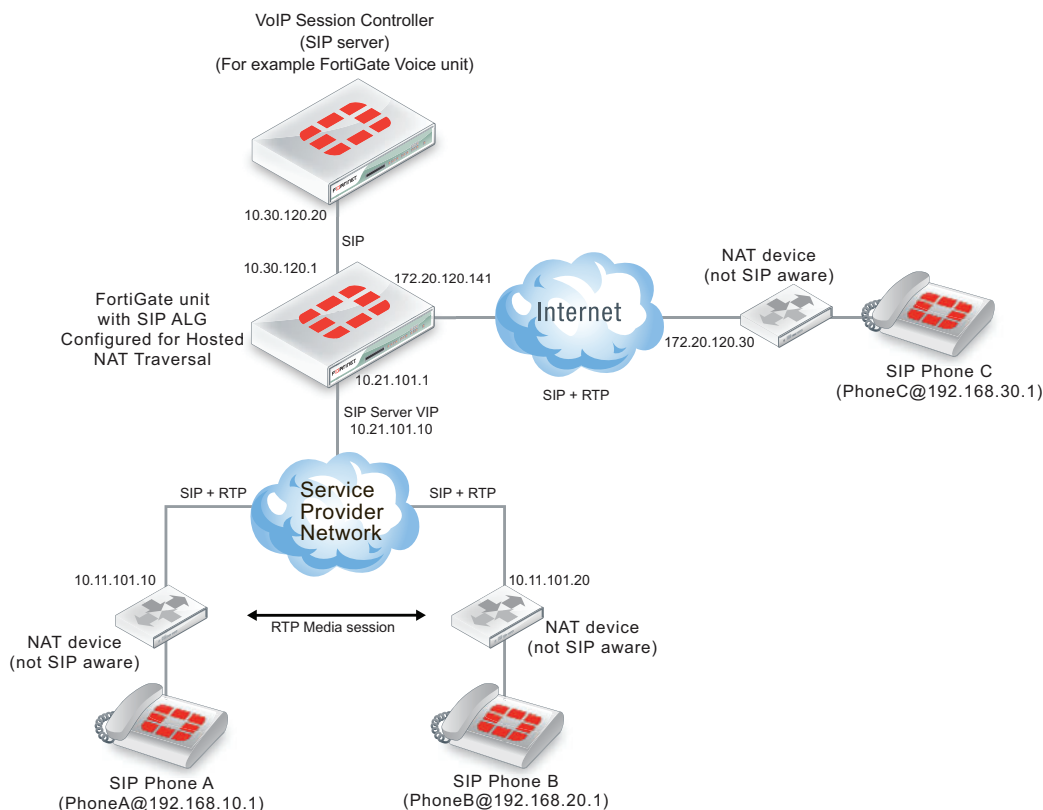
```
config voip profile
  edit Profile_name
    config SIP
      set strict-register enable
    end
```

Hosted NAT traversal

With the increase in the use of VoIP and other media traffic over the Internet, service provider network administrators must defend their networks from threats while allowing voice and multimedia traffic to flow transparently between users and servers and among users. A common scenario could involve providing SIP VoIP services for customers with SIP phones installed behind NAT devices that are not SIP aware. NAT devices that are not SIP aware cannot translate IP addresses in SIP headers and SDP lines in SIP packets but can and do perform source NAT on the source or addresses of the packets. In this scenario the user's SIP phones would communicate with a SIP proxy server to set up calls between SIP phones. Once the calls are set up RTP packets would be communicated directly between the phones through each user's NAT device.

The problem with this configuration is that the SIP headers and SDP lines in the SIP packets sent from the phones and received by the SIP proxy server would contain the private network addresses of the VoIP phones that would not be routable on the service provider network or on the Internet. One solution could be to for each customer to install and configure SIP aware NAT devices. If this is not possible, another solution requires implement hosted NAT traversal.

In a hosted NAT traversal (HNT) configuration (for example, see [Figure 285](#)), a FortiGate unit is installed between the NAT device and the SIP proxy server and configured with a VoIP profile that enables SIP hosted NAT traversal. Security policies that include the VoIP profile also support destination NAT using a firewall virtual IP. When the SIP phones connect to the SIP server IP address the security policy accepts the SIP packets, the virtual IP translates the destination addresses of the packets to the SIP server IP address, and the SIP ALG NAT traversal configuration translates the source IP addresses on the SIP headers and SDP lines to the source address of the SIP packets (which would be the external IP address of the NAT devices). The SIP server then sees the SIP phone IP address as the external IP address of the NAT device. As a result SIP and RTP media sessions are established using the external IP addresses of the NAT devices instead of the actual IP addresses of the SIP phones.

Figure 285: FortiGate SIP Hosted NAT Traversal configuration

Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone B in Figure 285.

- 1 SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1, destination IP address: 10.21.101.10.
- 2 The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.20.
- 3 The SIP packets are received by the FortiGate unit which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
- 4 The SIP server accepts the Invite message and forwards it to SIP Phone B at IP address 10.11.101.20. The SIP server has this address for SIP Phone B because SIP packets from SIP Phone B have also been translated using the hosted NAT traversal configuration of the SIP ALG.
- 5 When the SIP call is established, the RTP session is between 10.11.101.10 and 10.11.101.20 and does not pass through the FortiGate unit. The NAT devices translated the destination address of the RTP packets to the private IP addresses of the SIP phones.

General configuration steps

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

- 1 Add a VoIP profile that enables hosted NAT translation.
- 2 Add a SIP proxy server firewall virtual IP.
- 3 Add a firewall address for the SIP proxy server on the private network.
- 4 Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
- 5 Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

Configuration steps - web-based manager

To add the SIP proxy server firewall virtual IP

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Add the SIP proxy server virtual IP.

Name	SIP_Proxy_VIP
External Interface	port1
Type	Static NAT
External IP Address/Range	172.20.120.50
Mapped IP Address/Range	10.31.101.50

To add a firewall address for the SIP proxy server

- 1 Go to *Firewall Objects > Address*.
- 2 Add the following for the SIP proxy server:

Address Name	SIP_Proxy_Server
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.50/255.255.255.255
Interface	port2

To add the security policies

- 1 Go to *Policy > Policy > Policy*.
- 2 Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	SIP_Proxy_VIP
Schedule	always

Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select and select the default VoIP profile.

- 3 Select OK.
- 4 Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

Source Interface/Zone	port2
Source Address	SIP_Proxy_Server
Destination Interface/Zone	port1
Destination Address	all
Schedule	always
Service	SIP
Action	ACCEPT
Enable NAT	Select
UTM	Select
Protocol Options	default
Enable VoIP	Select (And select the default VoIP profile)

- 5 Select OK.

Configuration steps - CLI

To add a VoIP profile that enables hosted NAT translation.

- 1 Enter the following command to add a VoIP profile named HNT that enables hosted NAT traversal. This command shows how to clone the default VoIP profile and enable hosted NAT traversal.

```
config voip profile
  clone default to HNT
  edit HNT
    config sip
      set hosted-nat-traversal enable
    end
  end
end
```

To add the SIP proxy server firewall virtual IP and firewall address

- 2 Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
  edit SIP_Proxy_VIP
    set type static-nat
    set extip 10.21.101.10
    set mappedip 10.30.120.20
```

```

    set extintf port1
end

```

- 3 Enter the following command to add the SIP proxy server firewall address.

```

config firewall address
  edit SIP_Proxy_Server
    set associated interface port2
    set type ipmask
    set subnet 10.30.120.20 255.255.255.255
  end

```

To add security policies

- 1 Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone A to send SIP request messages to the SIP proxy server.

```

config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr SIP_Proxy_VIP
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set voip-profile HNT
  end

```

- 2 Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B:

```

config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr SIP_Proxy_Server
    set dstaddr all
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set voip-profile default
  end

```

Hosted NAT traversal for calls between SIP Phone A and SIP Phone C

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone C in [Figure 285 on page 2581](#).

- 1 SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1 and destination IP address: 10.21.101.10.

- 2 The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.20.
- 3 The SIP packets are received by the FortiGate unit which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
- 4 The SIP server accepts the Invite message and forwards it to SIP Phone C at IP address 172.20.120.30. The SIP server has this address for SIP Phone C because SIP packets from SIP Phone C have also been translated using the hosted NAT traversal configuration of the SIP ALG.
- 5 When the SIP call is established, the RTP session is between 10.11.101.10 and 172.20.120.30. The packets pass through the FortiGate unit which performs NAT as required.

Restricting the RTP source IP

Use the following command in a VoIP profile to restrict the RTP source IP to be the same as the SIP source IP when hosted NAT traversal is enabled.


```
config voip profile
  edit VoIP_HNT
    config sip
      set hosted-nat-traversal enable
      set hnt-restrict-source-ip enable
    end
  end
```

SIP over IPv6

FortiGate units operating in NAT/Route and in Transparent mode support SIP over IPv6. The SIP ALG can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGate units cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

In the scenario shown in [Figure 286](#), a SIP phone connects to the Internet through a FortiGate unit operating. The phone and the SIP and RTP servers all have IPv6 addresses.

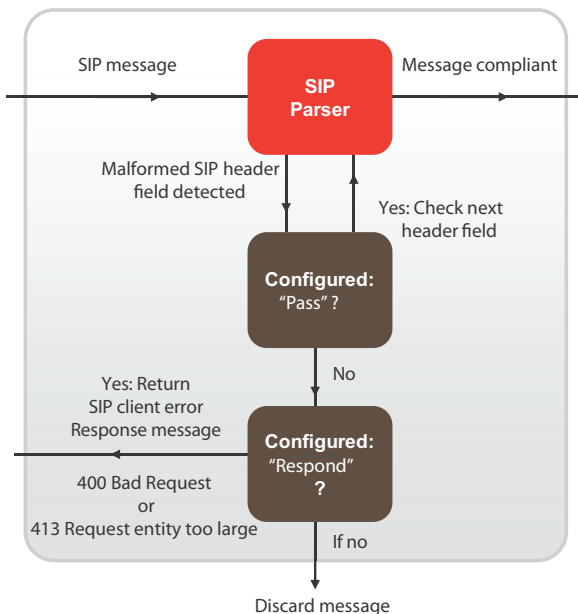
The FortiGate unit has IPv6 security policies that accept SIP sessions. The SIP ALG understands IPv6 addresses and can forward IPv6 sessions to their destinations. Using SIP application control features the SIP ALG can also apply rate limiting and other settings to SIP sessions.

Figure 286: SIP support for IPv6

To enable SIP support for IPv6 add an IPv6 security policy that accepts SIP packets and includes a VoIP profile.

Deep SIP message inspection

Deep SIP message syntax inspection (also called Deep SIP header inspection or SIP fuzzing protection) provides protection against malicious SIP messages by applying SIP header and SDP profile syntax checking. SIP Fuzzing attacks can be used by attackers to discover and exploit vulnerabilities of a SIP entity (for example a SIP proxy server). Most often these attacks could crash or compromise the SIP entity.

Figure 287: Deep SIP message inspection

- Checks the SIP request message Request-line
- Checks the following SIP header fields:
 - Allow, Call-id, Contact, Content-length, Content-type, CSeq, Expires, From, Max-Forwards, P-asserted-identity, Rack, Record-Route, Route, Rseq, To, Via
- Checks all SDP profile lines
- Configurable header and body length checks
- Optional logging of message violations

Deep SIP message inspection checks the syntax of each SIP header and SDP profile line to make sure they conform to the syntax defined in the relevant RFC and IETF standard. You can also configure the SIP ALG to inspect for:

- Unknown SIP message types (message types not defined in a SIP RFC) this option is enabled by default and can be disabled. When enabled unknown message types are discarded. Configured using the `block-unknown` option.
- Unknown line types (message line types that are not defined in any SIP or SDP RFC). Configured using the `unknown-header` option.
- Messages that are longer than a configured maximum size. Configured using the `max-body-length` option.
- Messages that contain one or more lines that are longer than a set maximum line length (default 998 characters). Configured using the `max-line-length` option.

Actions taken when a malformed message line is found

When a malformed message line or other error is found the SIP ALG can be configured to discard the message containing the error, pass the message without any other actions, or responding to the message with a 400 Bad Request or 413 Request entity too large client error SIP response message and then discard the message. (For information about client error SIP response messages, see [“Client error” on page 2529](#).)

If a message line is longer than the configured maximum, the SIP ALG sends the following message:

```
SIP/2.0 413 Request Entity Too Large, <optional_info>
```

If a message line is incorrect or in an unknown message line is found, the SIP ALG sends the following message:

```
SIP/2.0 400 Bad Request, <optional_info>
```

The `<optional_info>` provides more information about why the message was rejected. For example, if the SIP ALG finds a malformed Via header line, the response message may be:

```
SIP/2.0 400 Bad Request, malformed Via header
```

If the SIP ALG finds a malformed message line, and the action for this message line type is discard, the message is discarded with no further checking or responses. If the action is pass, the SIP ALG continues parsing the SIP message for more malformed message lines. If the action is respond, the SIP ALG sends the SIP response message and discards the message containing the malformed line with no further checking or response. If only malformed message line types with action set to pass are found, the SIP ALG extracts as much information as possible from the message (for example for NAT and opening pinholes, and forwards the message to its destination).

If a SIP message containing a malformed line is discarded the SIP ALG will not use the information in the message for call processing. This could result in the call being terminated. If a malformed line in a SIP message includes information required for the SIP call that the SIP ALG cannot interpret (for example, if an IP address required for SIP NAT is corrupted) the SIP ALG may not be able to continue processing the call and it could be terminated. Discarded messages are counted by SIP ALG static message counters.

Logging and statistics

To record a log message each time the SIP ALG finds a malformed header, enable logging SIP violations in a VoIP profile. In all cases, when the SIP ALG finds an error the FortiGate unit records a malformed header log message that contains information about the error. This happens even if the action is set to pass.

If, because of recording log messages for deep message inspection, the CPU performance is affected by a certain amount, the FortiGate unit records a critical log message about this event and stops writing log messages for deep SIP message inspection.

The following information is recorded in malformed header messages:

- The type of message line in which the error was found.
- The content of the message line in which the error was found (it will be truncated if it makes the log message too long)
- The column or character number in which the error was found (to make it easier to determine what caused the error)

Deep SIP message inspection best practices

Because of the risks imposed by SIP header attacks or incorrect data being allowed and because selecting drop or respond does not require more CPU overhead than pass you would want to set all tests to drop or respond. However, in some cases malformed lines may be less of a threat or risk. For example, the SDP `i=` does not usually contain information that is parsed by any SIP device so a malformed `i=` line may not pose a threat.

You can also use the pre-defined VoIP profiles to apply different levels of deep message inspection. The default VoIP profile sets all deep message inspection options to pass and the strict VoIP profile sets all deep message inspection options to discard. From the CLI you can use the `clone` command to copy these pre-defined VoIP profiles and then customize them for your requirements.

Configuring deep SIP message inspection

You configure deep SIP message inspection in a VoIP profile. All deep SIP message inspection options are available only from the CLI.

Enter the following command to configure deep SIP message inspection to discard messages with malformed Request-lines (the first line in a SIP request message):

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set malformed-request-line respond
    end
  end
```



You cannot configure message inspection for the Status-line, which is the first line in a SIP response message.

Table 141 lists the SIP header lines that the SIP ALG can inspect and the CLI command for configuring the action for each line type. The table also lists the RFC that the header line is defined in.

Table 141: SIP header lines that the SIP ALG can inspect for syntax errors

SIP Header line	VoIP profile option	RFC
Allow	malformed-header-allow	RFC 3261
Call-ID	malformed-header-call-id	RFC 3261
Contact	malformed-header-contact	RFC 3261

Table 141: SIP header lines that the SIP ALG can inspect for syntax errors

SIP Header line	VoIP profile option	RFC
Content-Length	malformed-header-content-length	RFC 3261
Content-Type	malformed-header-content-type	RFC 3261
CSeq	malformed-header-cseq	RFC 3261
Expires	malformed-header-expires	RFC 3261
From	malformed-header-from	RFC 3261
Max-forwards	malformed-header-max-forwards	RFC 3261
P-Asserted-Identity	malformed-header-p-asserted-identity	RFC 3325
RAck	malformed-header-rack	RFC 3262
Record-Route	malformed-header-record-route	RFC 3261
Route	malformed-header-route	RFC 3261
RSeq	malformed-header-rseq	RFC 3262
To	malformed-header-to	RFC 3261
Via	malformed-header-via	RFC 3261

Table 142 lists the SDP profile lines that the SIP ALG inspects and the CLI command for configuring the action for each line type. SDP profile lines are defined by RFC 4566 and RFC 2327.

Table 142: SDP profile lines that the SIP ALG can inspect for syntax errors

Attribute	VoIP profile option
a=	malformed-header-sdb-a
b=	malformed-header-sdp-b
c=	malformed-header-sdp-c
i=	malformed-header-sdp-i
k=	malformed-header-sdp-k
m=	malformed-header-sdp-m
o=	malformed-header-sdp-o
r=	malformed-header-sdp-r
s=	malformed-header-sdp-s
t=	malformed-header-sdp-t
v=	malformed-header-sdp-v
z=	malformed-header-sdp-z

Discarding SIP messages with some malformed header and body lines

Enter the following command to configure deep SIP message inspection to discard SIP messages with a malformed Via line, a malformed route line or a malformed m= line but to pass messages with a malformed i= line or a malformed Max-Forwards line

```
config voip profile
  edit VoIP_Pro_Name
    config sip
```

```
        set malformed-header-via discard
        set malformed-header-route discard
        set malformed-header-sdp-m discard
        set malformed-header-sdp-i pass
        set malformed-header-max-forwards pass
    end
end
```

Discarding SIP messages with an unknown SIP message type

Enter the following command to discard SIP messages with an unknown SIP message line type as defined in all current SIP RFCs:

```
config voip profile
    edit VoIP_Pro_Name
        config sip
            set unknown-header discard
        end
    end
```

Discarding SIP messages that exceed a message size

Enter the following command to set the maximum size of a SIP message to 200 bytes. Messages longer than 200 bytes are discarded.

```
config voip profile
    edit VoIP_Pro_Name
        config sip
            set max-body-length 200
        end
    end
```

The `max-body-length` option checks the value in the SIP Content-Length header line to determine body length. The Content-Length can be larger than the actual size of a SIP message if the SIP message content is split over more than one packet. SIP message sizes vary widely. The size of a SIP message can also change with the addition of Via and Record-Route headers as the message is transmitted between users and SIP servers.

Discarding SIP messages with lines longer than 500 characters

Enter the following command to set the length of a SIP message line to 500 characters and to block messages that include lines with 500 or more characters:

```
config voip profile
    edit VoIP_Pro_Name
        config sip
            set max-line-length 500
            set block-long-lines enable
        end
    end
```

Blocking SIP request messages

You may want to block different types of SIP requests:

- to prevent SIP attacks using these messages.

- If your SIP server cannot process some SIP messages because of a temporary issue (for example a bug that crashes or compromises the server when it receives a message of a certain type).
- Your SIP implementation does not use certain message types.

When you enable message blocking for a message type in a VoIP profile, whenever a security policy containing the VoIP profile accepts a SIP message of this type, the SIP ALG silently discards the message and records a log message about the action.

Use the following command to configure a VoIP profile to block SIP CANCEL and Update request messages:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set block-cancel enable
      set block-update enable
    end
  end
```

SIP uses a variety of text-based messages or requests to communicate information about SIP clients and servers to the various components of the SIP network. Since SIP requests are simple text messages and since the requests or their replies can contain information about network components on either side of the FortiGate unit, it may be a security risk to allow these messages to pass through.

[Table 143](#) lists all of the VoIP profile SIP request message blocking options. All of these options are disabled by default.



Blocking SIP OPTIONS messages may prevent a redundant configuration from operating correctly. See [“Supporting geographic redundancy when blocking OPTIONS messages” on page 2595](#) for information about resolving this problem.


Table 143: Options for blocking SIP request messages

SIP request message	SIP message blocking CLI Option
ACK	block-ack
BYE	block-bye
Cancel	block-cancel
INFO	block-info
INVITE	block-invite
Message	block-message
Notify	block-notify
Options	block-options
PRACK	block-prack
Publish	block-publish
Refer	block-refer
Register	block-register
Subscribe	block-subscribe
Update	block-update

SIP rate limiting

Configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.

Figure 288: SIP rate limiting



- **SIP message rate limitation**
- **Individually configurable per SIP method**
- **When threshold is hit additional messages with this method will be discarded**
- **Prevents SIP server from getting overloaded by flash crowds or Denial-of-Service attacks.**
- **May block some methods at all (with extra "block" option)**
- **Can be disabled (unlimited rate)**

FortiGate units support rate limiting for the following types of VoIP traffic:

- Session Initiation Protocol (SIP)
- Skinny Call Control Protocol (SCCP) (most versions)
- Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE).

You can use rate limiting of these VoIP protocols to protect the FortiGate unit and your network from SIP and SCCP Denial of Service (DoS) attacks. Rate limiting protects against SIP DoS attacks by limiting the number of SIP REGISTER and INVITE requests that the FortiGate unit receives per second. Rate limiting protects against SCCP DoS attacks by limiting the number of SCCP call setup messages that the FortiGate unit receives per minute.

You configure rate limiting for a message type by specifying a limit for the number of messages that can be received per second. The rate is limited per security policy. When VoIP rate limiting is enabled for a message type, if the a single security policy accepts more messages per second than the configured rate, the extra messages are dropped and log messages are written when the messages are dropped.

Use the following command to configure a VoIP profile to limit the number of INVITE messages accepted by each security policy that the VoIP profile is added to 100 INVITE messages a second:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set invite-rate 100
    end
```

end

If you are experiencing denial of service attacks from traffic using these VoIP protocols, you can enable VoIP rate limiting and limit the rates for your network. Limit the rates depending on the amount of SIP and SCCP traffic that you expect the FortiGate unit to be handling. You can adjust the settings if some calls are lost or if the amount of SIP or SCCP traffic is affecting FortiGate unit performance.

Table 144 lists all of the VoIP profile SIP rate limiting options. All of these options are set to 0 so are disabled by default.

Table 144: Options for SIP rate limiting

SIP request message	Rate Limiting CLI Option
ACK	ack-rate
BYE	bye-rate
Cancel	cancel-rate
INFO	info-rate
INVITE	invite-rate
Message	message-rate
Notify	notify-rate
Options	options-rate
PRACK	prack-rate
Publish	publish-rate
Refer	refer-rate
Register	register-rate
Subscribe	subscribe-rate
Update	update-rate

Limiting the number of SIP dialogs accepted by a security policy

In addition to limiting the rates for receiving SIP messages, you can use the following command to limit the number of SIP dialogs (or SIP calls) that the FortiGate unit accepts.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-dialogs 2000
    end
  end
end
```

This command sets the maximum number of SIP dialogs that can be open for SIP sessions accepted by any security policy that you add the VoIP profile to. The default setting of 0 does not limit the number of dialogs. You can add a limit to control the number of open dialogs and raise and lower it as required. You might want to limit the number of open dialogs for protection against SIP-based attackers opening large numbers of SIP dialogs. Every dialog takes memory and FortiGate CPU resources to process. Limiting the number of dialogs may improve the overall performance of the FortiGate unit. Limiting the number of dialogs will not drop calls in progress but may prevent new calls from connecting.

SIP logging and DLP archiving

You can enable SIP logging and logging of SIP violations, and SIP DLP archiving a VoIP profile. To record SIP log messages you must also enable VoIP event logging in the FortiGate unit event logging configuration.

To view SIP log messages go to *Log&Report > Log Access > Event*.

To view SIP DLP archive messages to go *Log&Report > Archive Access > VoIP*.

Use the following command enable SIP logging, SIP archiving, and logging of SIP violations in a VoIP profile:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set log-call-summary enable
      set log-violations enable
    end
  end
```

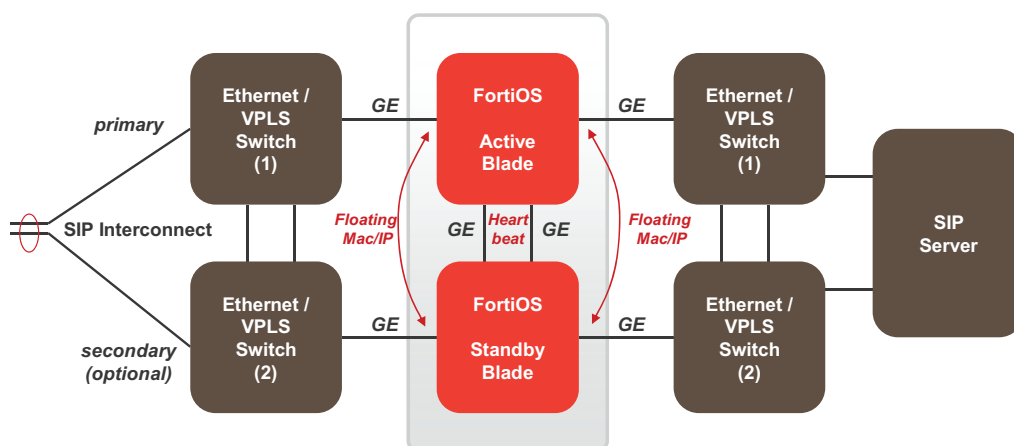
SIP and HA: session failover and geographic redundancy

FortiGate high availability supports SIP session failover (also called stateful failover) for active-passive HA. To support SIP session failover, create a standard HA configuration and select the Enable Session Pick-up option.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

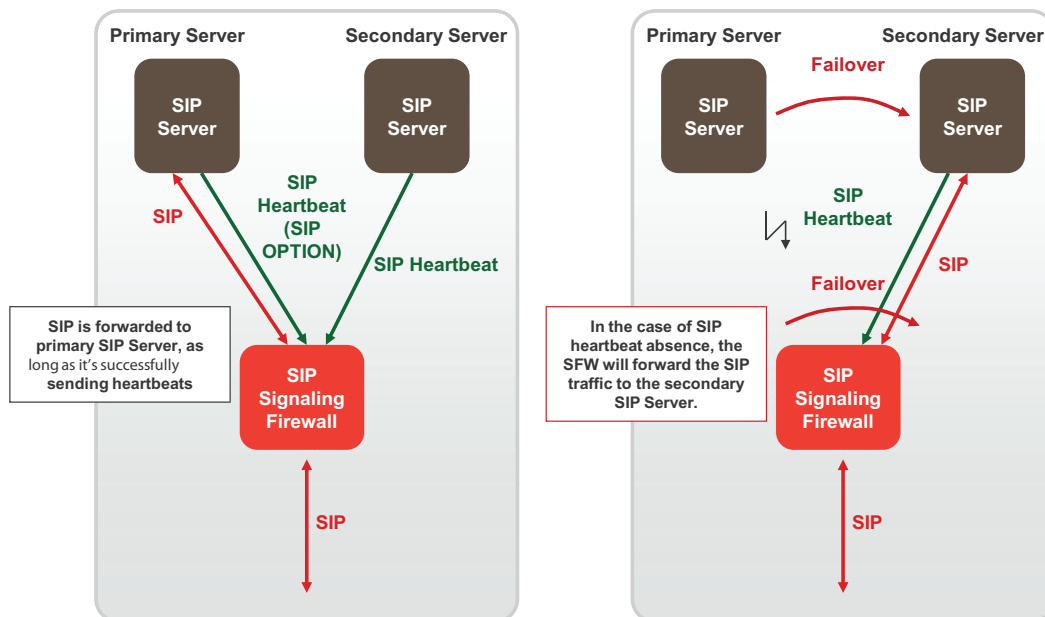
Figure 289: SIP HA session failover



SIP geographic redundancy

Maintains an active-standby SIP server configuration, which even supports geographical distribution. If the active SIP server fails (missing SIP heartbeat messages or SIP traffic) FortiOS will redirect the SIP traffic to a secondary SIP server.

Figure 290: SIP geographic redundancy



Supporting geographic redundancy when blocking OPTIONS messages

For some geographic redundant SIP configurations, the SIP servers may use SIP OPTIONS messages as heartbeats to notify the FortiGate unit that they are still operating (or alive). This is a kind of passive SIP monitoring mechanism where the FortiGate unit isn't actively monitoring the SIP servers and instead the FortiGate unit passively receives and analyzes OPTIONS messages from the SIP servers.

If FortiGate units block SIP OPTIONS messages because `block-options` is enabled, the configuration may fail to operate correctly because the OPTIONS messages are blocked by one or more FortiGate units.

However, you can work around this problem by enabling the `block-geo-red-options` application control list option. This option causes the FortiGate unit to refresh the local SIP server status when it receives an OPTIONS message before dropping the message. The end result is the heartbeat signals between geographically redundant SIP servers are maintained but OPTIONS messages do not pass through the FortiGate unit.

Use the following command to block OPTIONS messages while still supporting geographic redundancy:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set block-options disable
      set block-geo-red-options enable
    end
```

end



The `block-options` option setting overrides the `block-geo-red-options` option. If `block-options` is enabled the FortiGate unit only blocks SIP OPTIONS messages and does not refresh local SIP server status.

Support for RFC 2543-compliant branch parameters

RFC 3261 is the most recent SIP RFC, it obsoletes RFC 2543. However, some SIP implementations may use RFC 2543-compliant SIP calls.

The `rfc2543-branch` VoIP profile option allows the FortiGate unit to support SIP calls that include an RFC 2543-compliant branch parameter in the SIP Via header. This option also allows FortiGate units to support SIP calls that include Via headers that are missing the branch parameter.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set rfc2543-branch enable
    end
  end
```

SIP and IPS

You can enable IPS in security policies that also accept SIP sessions to protect the SIP traffic from SIP-based attacks. If you enable IPS in this way then by default the pinholes that the SIP ALG creates to allow RTP and RTCP to flow through the firewall will also have IPS enabled.

This inheritance of the IPS setting can cause performance problems if the RTP traffic volume is high since IPS checking may reduce performance in some cases. Also if you are using network processor (NP) interfaces to accelerate VoIP performance, when IPS is enabled for the pinhole traffic is diverted to the IPS and as a result is not accelerated by the network processors.

You can use the following CLI command to disable IPS for the RTP pinhole traffic.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set ips-rtp disable
    end
  end
```

SIP debugging

SIP debug log format

Assuming that `diagnose debug console timestamp` is enabled then the following shows the debug that is generated for an INVITE if `diag debug appl sip -1` is enabled:

```
2010-01-04 21:39:59 sip port 26 locate session for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip sess 0x979df38 found for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 192.168.2.134:5061 -> 172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 read [(0,515)]
```

```
(494e56495445207369703a73657276696365403139322e3136382e322e3130303a35303630205349502f322e300d0
a5669613a205349502f322e302f554450203132372e302e312e313a353036313b6272616e63683d7a39684734624b2
d363832372d3632302d300d0a46726f6d3a2073697070203c7369703a73697070403132372e302e312e313a3530363
13e3b7461673d363832375349507054616730303632300d0a546f3a20737574203c7369703a7365727669636540313
9322e3136382e322e3130303a353036303e0d0a43616c6c2d49443a203632302d36383237403132372e302e312e310
d0a435365713a203120494e564954450d0a436f6e746163743a207369703a73697070403132372e302e312e313a353
036310d0a4d61782d466f7277617264733a2037300d0a5375626a6563743a20506572666f726d616e6365205465737
40d0a436f6e74656e742d547970653a206170706c69636174696f6e2f7364700d0a436f6e74656e742d4c656e67746
83a20203132390d0a0d0a763d300d0a6f3d7573657231203533363535373635203233353336383736333720494e204
95034203132372e302e312e310d0a733d2d0d0a633d494e20495034203132372e302e312e310d0a743d3020300d0a6
d3d617564696f2036303031205254502f41565020300d0a613d7274706d61703a302050434d552f383030300d0a) (I
NVITE
sip:service@192.168.2.100:5060 SIP/2.0..Via: SIP/2.0/UDP
127.0.1.1:5061;branch=z9hG4bK-6827-620-0..From: sipp
<tag=sipp@127.0.1.1:5061>;tag=6827SIPpTag00620..To: sut
<tag=sipp:service@192.168.2.100:5060>..Call-ID: 620-6827@127.0.1.1..CSeq: 1
INVITE..Contact: sip:sipp@127.0.1.1:5061..Max-Forwards: 70..Subject: Performance
Test..Content-Type: application/sdp..Content-Length: 129....v=0..o=user1 53655765
2353687637 IN IP4 127.0.1.1..s=-..c=IN IP4 127.0.1.1..t=0 0..m=audio 6001 RTP/AVP
0..a=rtpmap:0 PCMU/8000..)]
2010-01-04 21:39:59 sip port 26 len 515
2010-01-04 21:39:59 sip port 26 INVITE '192.168.2.100:5060' addr 192.168.2.100:5060
2010-01-04 21:39:59 sip port 26 CSeq: 1 INVITE
2010-01-04 21:39:59 sip port 26 Via: UDP 127.0.1.1:5061 len 14 received 0 rport 0 0 branch
'z9hG4bK-6827-620-0'
2010-01-04 21:39:59 sip port 26 From: 'sipp ;tag=6827SIPpTag00620' URI
'sip:sipp@127.0.1.1:5061' tag '6827SIPpTag00620'
2010-01-04 21:39:59 sip port 26 To: 'sut' URI 'sip:service@192.168.2.100:5060' tag ''
2010-01-04 21:39:59 sip port 26 Call-ID: '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 Contact: '127.0.1.1:5061' addr 127.0.1.1:5061 expires 0
2010-01-04 21:39:59 sip port 26 Content-Length: 129 len 3
2010-01-04 21:39:59 sip port 26 sdp o=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp c=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp m=6001 len=4
2010-01-04 21:39:59 sip port 26 find call 0 '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 not found
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open (collision (nil))
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open txn 0x979f7f8 INVITE dir 0
2010-01-04 21:39:59 sip port 26 sdp i: 127.0.1.1:6001
2010-01-04 21:39:59 sip port 26 policy id 1 is_client_vs_policy 1 policy_dir_rev 0
2010-01-04 21:39:59 sip port 26 policy 1 not RTP policy
2010-01-04 21:39:59 sip port 26 learn sdp from stream address
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 sdp 172.16.67.198:43722
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address
and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address
and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address
and port
2010-01-04 21:39:59 sip port 30 write 192.168.2.134:5061 -> 172.16.67.192:5060 (13,539)
2010-01-04 21:39:59 sip port 30 write [(13,539)
(494e56495445207369703a73657276696365403137322e31362e36372e3139323a35303630205349502f322e300d0
a5669613a205349502f322e302f554450203137322e31362e36372e3139383a35323036353b6272616e63683d7a396
84734624b2d363832372d3632302d300d0a46726f6d3a2073697070203c7369703a73697070403137322e31362e363
72e3139383a34333732343e3b7461673d363832375349507054616730303632300d0a546f3a20737574203c7369703
a73657276696365403137322e31362e36372e3139323a353036303e0d0a43616c6c2d49443a203632302d363832374
03132372e302e312e310d0a435365713a203120494e564954450d0a436f6e746163743a207369703a7369707040313
7322e31362e36372e3139383a34333732350d0a4d61782d466f7277617264733a2037300d0a5375626a6563743a205
06572666f726d616e636520546573740d0a436f6e74656e742d547970653a206170706c69636174696f6e2f7364700
d0a436f6e74656e742d4c656e6774683a20203133380d0a0d0a763d300d0a6f3d75736572312035333635353736352
03233353336383736333720494e20495034203137322e31362e36372e3139380d0a733d2d0d0a633d494e204950342
03137322e31362e36372e3139380d0a743d3020300d0a6d3d617564696f203433373232205254502f41565020300d0
a613d7274706d61703a302050434d552f383030300d0a) (INVITE sip:service@172.16.67.192:5060
SIP/2.0..Via: SIP/2.0/UDP 172.16.67.198:52065;branch=z9hG4bK-6827-620-0..From: sipp
;tag=6827SIPpTag00620..To: sut ..Call-ID: 620-6827@127.0.1.1..CSeq: 1 INVITE..Contact:
sip:sipp@172.16.67.198:43725..Max-Forwards: 70..Subject: Performance Test..Content-Type:
application/sdp..Content-Length: 138....v=0..o=user1 53655765 2353687637 IN IP4
172.16.67.198..s=-..c=IN IP4 172.16.67.198..t=0 0..m=audio 43722 RTP/AVP 0..a=rtpmap:0
PCMU/8000..)]
```

SIP-proxy filter per VDOM

You can use the `diagnose sys sip-proxy xxx` command in a VDOM to get info about how SIP is operating in each VDOM.

SIP-proxy filter command

Use the `diagnose system sip-proxy filter` to filter diagnose information for the SIP ALG. The following filters are available:

```
diag sys sip-proxy filter vd
diag sys sip-proxy filter dst-addr4
diag sys sip-proxy filter dst-addr6
diag sys sip-proxy filter dst-port
diag sys sip-proxy filter identity-policy
diag sys sip-proxy filter negate
diag sys sip-proxy filter policy
diag sys sip-proxy filter policy-type
diag sys sip-proxy filter profile-group
diag sys sip-proxy filter src-addr4
diag sys sip-proxy filter src-addr6
diag sys sip-proxy filter src-port
diag sys sip-proxy filter vd
diag sys sip-proxy filter voip-profile
```

You can clear, view and negate/invert the sense of a filter using these commands:

```
diag sys sip-proxy filter clear
diag sys sip-proxy filter list
diag sys sip-proxy filter negate
```

SIP debug log filtering

You can filter by VDOM/IP/PORT and by policy and VoIP profile. The filtering can be controlled by:

```
diagnose system sip-proxy log-filter
```

The list of filters is:

```
diag sys sip-proxy log-filter vd
diag sys sip-proxy log-filter dst-addr4
diag sys sip-proxy log-filter dst-addr6
diag sys sip-proxy log-filter dst-port
diag sys sip-proxy log-filter identity-policy
diag sys sip-proxy log-filter policy
diag sys sip-proxy log-filter policy-type
diag sys sip-proxy log-filter profile-group
diag sys sip-proxy log-filter src-addr4
diag sys sip-proxy log-filter src-addr6
diag sys sip-proxy log-filter src-port
diag sys sip-proxy log-filter vd
diag sys sip-proxy log-filter voip-profile
```

You can clear, view and negate/invert the sense of a filter using these commands:

```
diag sys sip-proxy log-filter clear
diag sys sip-proxy log-filter list
diag sys sip-proxy log-filter negate
```

SIP debug setting

Control of the SIP debug output is governed by the following command

```
diagnose debug application sip <debug_level_int>
```

Where the `<debug_level_int>` is a bitmask and the individual values determine whether the listed items are logged or not. The `<debug_level_int>` can be

- 1 - configuration changes. Mainly addition/deletion/modification of virtual domains.
- 2 - (TCP) connection accepts or connects, redirect creation
- 4 - create or delete a session
- 16 - any IO read or write
- 32 - an ASCII dump of all data read or written
- 64 - Include HEX dump in the above output
- 128 - any activity related to the use of the FortiCarrier dynamic profile feature to determine the correct profile-group to use
- 256 - log summary of interesting fields in a SIP call
- 1024 - any activity related to SIP geo-redundancy.
- 2048 - any activity related to HA syncing of SIP calls.

SIP test commands

Use the following command to control or inspect the behavior of the SIP ALG.

```
diagnose test application sip <test_level_int>
```

Where <test_level_int> can be

- 1 - Display memory statistics summary
- 2 - Display all memory statistics
- 3 - Display debug consoles
- 4 - Display all SIP redirects
- 20 - Display SIP per-policy configurations
- 21 - Display SIP VoIP profiles
- 22 - Display SIP meters
- 23 - Display SIP VIPs
- 24 - Display SIP RTP policies
- 30 - Display SIP stats summary
- 31 - Display per VDOM SIP stats
- 50 - Display all SIP idle calls
- 51 - Display all SIP sessions
- 70 - Start measuring scheduler times
- 71 - Stop measuring scheduler times
- 72 - Display scheduler times
- 99 - Restart SIP -- this will drop all SIP calls as well as all IM and SCCP

Display SIP rate-limit data

You can use the `diagnose sys sip-proxy meters` command to display SIP rate limiting data.

For the following command output `rate 1` shows that the current (over last second) measured rate for INVITE/ACK and BYTE was 1 per second, the `peak 1` shows that the peak rate recorded is 1 per second, the `max 0` shows that there is no maximum limit set, the `count 18` indicates that 18 messages were received and `drop 0` indicates that none were dropped due to being over the limit.

```
diag sys sip-proxy meters
```

```

sip
sip vd: 0
sip policy: 1
sip identity-policy: 0
sip policy-type: IPv4
sip profile-group:
sip dialogs: 18
sip dialog-limit: 0
sip UNKNOWN: rate 0 peak 0 max 0 count 0 drop 0
sip ACK: rate 1 peak 1 max 0 count 18 drop 0
sip BYE: rate 1 peak 1 max 0 count 18 drop 0
sip CANCEL: rate 0 peak 0 max 0 count 0 drop 0
sip INFO: rate 0 peak 0 max 0 count 0 drop 0
sip INVITE: rate 1 peak 1 max 0 count 18 drop 0
sip MESSAGE: rate 0 peak 0 max 0 count 0 drop 0
sip NOTIFY: rate 0 peak 0 max 0 count 0 drop 0
sip OPTIONS: rate 0 peak 0 max 0 count 0 drop 0
sip PRACK: rate 0 peak 0 max 0 count 0 drop 0
sip PUBLISH: rate 0 peak 0 max 0 count 0 drop 0
sip REFER: rate 0 peak 0 max 0 count 0 drop 0
sip REGISTER: rate 0 peak 0 max 0 count 0 drop 0
sip SUBSCRIBE: rate 0 peak 0 max 0 count 0 drop 0
sip UPDATE: rate 0 peak 0 max 0 count 0 drop 0
sip PING: rate 0 peak 0 max 0 count 0 drop 0
sip YAHOOREF: rate 0 peak 0 max 0 count 0 drop 0

```

VoIP Profile options

The following are VoIP profile configuration settings in *UTM Profiles > VoIP > Profile*.

Profile page

Lists the profiles that you created for SIP and SCCP protocols. On this page, you can edit, delete or create a new profile for VoIP protocols.

You are redirected to this page when you select *View List* on the Edit VoIP Profile page.

Create New

Creates a new VoIP profile. When you select *Create New*, you are automatically redirected to the New VoIP Profile page.

Edit

Modifies settings within a VoIP profile. When you select *Edit*, you are automatically redirected to the Edit VoIP Profile page.

Removes a VoIP profile from the list on the Profile page.

Delete

To remove multiple VoIP profiles from within the list, on the Profile page, in each of the rows of the profiles you want removed, select the check box and then select *Delete*.

To remove all VoIP profiles from the list, on the Profile page, select the check box in the check box column and then select *Delete*.

Name

The name of the profile.

Comments

A description about the profile. This is an optional setting.

Ref.

Displays the number of times the VoIP is referenced to other objects.

New VoIP Profile page

Provides settings for configuring SIP and SCCP options within the profile.

This page appears when you select *Create New* on the Edit VoIP Profile page. If you are on the Profile page, and you select *Create New*, you will be redirected to the New VoIP Profile page.

Name	Enter a name for the profile.
Comments	Enter a description about the profile. This is optional.
SIP	Configuration settings for SIP protocols.
Limit REGISTER requests	Enter a number for limiting the time it takes to register requests.
Limit INVITE requests	Enter a number to limit invitation requests.
SCCP	Configuration settings for SCCP protocols.
Limit Call Setup	Enter a number to limit call setup time.



Example FortiGate Voice branch office configuration

This section describes how to configure a FortiGate Voice-80C unit to operate in NAT/Route mode and provide basic UTM and SIP services for the example branch office network shown in [Figure 291 on page 2604](#). The non-PSTN parts of this example configuration also apply to FortiGate Voice models that do not include PSTN interfaces.

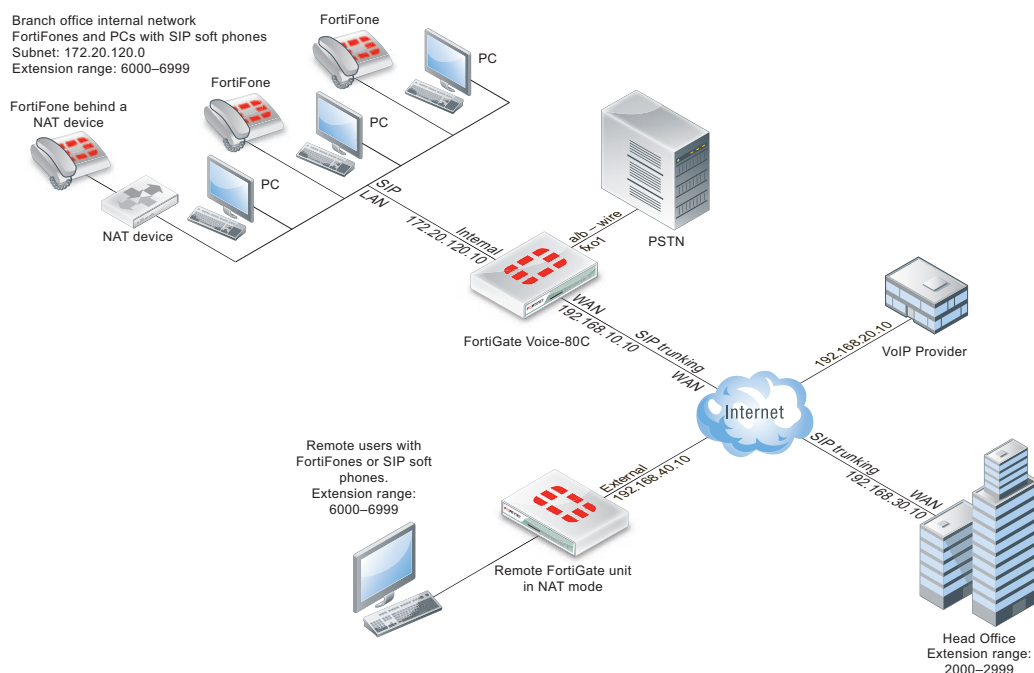
In this example the FortiGate Voice-80C unit provides:

- Internet connectivity, networking, and UTM features for the PCs on the branch office internal network.
- A single line a/b wire connection between the FortiGate Voice-80C fxo1 interface and a public switched telephone network (PSTN) line so that branch office phones can call the PSTN or receive calls from the PSTN.
- VoIP PBX services for FortiFones and SIP soft phones connected to the branch office internal network. PBX features include:
 - Extensions to the FortiFones and SIP soft phones in the internal network. The branch office phones use four digit numeric extensions that begin with the number 6. Example valid extensions are 6123, 6456, and 6899.
 - Extensions for phones behind NAT devices on the internal network.
 - Extensions for phones behind NAT devices on a remote network.
 - To collect voicemail the branch office phones dial *97.
 - Configure the Voice Menu to program the IP phone keys for various things such as recording a custom welcome message, providing access to company directory and adding a shortcut for checking voicemail.
 - SIP trunking to a VoIP provider for calling the head office.
 - To call a phone number on the PSTN from a branch office phone, dial 9 followed by the phone number. PSTN support includes:
 - Dialing 911 for emergencies
 - Support for dialing international calls
 - Support for dialing toll-free calls
 - Support for long distance calls
 - The FortiGate Voice unit sends email notifications to users when they receive voicemail.
 - To call the head office, the branch office phones dial a head office extension directly. The head office extension range is 2000-2999.

This configuration example describes configuring the FortiGate Voice-80C unit to support these services and where required also provides configuration steps for other devices such as the FortiFones and the remote FortiGate unit operating in NAT mode.

Details about the PSTN connection requirements, SIP trunking for the VoIP provider and the Head Office SIP configuration are not described.

Figure 291: Example Branch Office network configuration



This section describes:

- General configuration steps
- Connecting the FortiGate Voice unit
- Configuring basic FortiGate Voice network and UTM settings
- Configuring network settings for the devices on the Internal network
- Configuring the FortiGate Voice PSTN and PBX settings
- Configuring the FortiFones on the internal network
- Adding extensions and configuring FortiFones for users behind a NAT device
- FortiGate Voice IVR configuration
- Providing access to the company directory
- Adding a shortcut for checking voicemail

General configuration steps

- 1 Connect the FortiGate Voice unit to the Internet, the internal network and the PSTN.
- 2 Configure FortiGate Voice unit network and UTM settings.

The network configuration includes enabling the *SIP Traffic* option on the internal and wan1 interfaces. You must enable SIP traffic on these interfaces to accept and process SIP calls. No other special network configuration, firewall policies, or routing is required for the FortiGate Voice to accept and process SIP calls.



You do not have to add SIP firewall policies to enable SIP traffic for the FortiGate Voice unit to function as a PBX. Also, with PBX functionality enabled, you cannot apply FortiGate SIP application control features to SIP traffic received by FortiGate Voice interfaces for which you have enabled the *SIP Traffic* option.

This example also describes how to configure the FortiGate Voice as a DHCP server and DNS server for the branch office internal network. As a DHCP server, the FortiGate Voice unit can supply network configuration settings for the PCs and FortiFones on the internal network.

- 3 Configure network settings for the PCs on the internal network.
- 4 Configure the FortiGate Voice PSTN and PBX settings.
- 5 Configure the FortiFones on the internal network.
- 6 Configure the FortiGate Voice unit to SIP phone users behind a remote NAT device.

Connecting the FortiGate Voice unit

The following procedure describes how to connect the FortiGate Voice unit to the Internet, the branch office internal network, and the PSTN (supported by some FortiGate Voice models).

To connect the FortiGate Voice unit

- 1 Use an Ethernet cable to connect the FortiGate Voice wan1 interface to the device that connects the branch office to the Internet.

The device could be a cable or DSL modem or other device depending on how the Branch Office connects to the Internet.
- 2 Use Ethernet cables to connect the PCs and FortiFones on the internal network to the FortiGate Voice internal interface switch connectors.

You can connect up to 8 PCs and FortiFones directly to the FortiGate Voice Internal interface switch connectors. To connect more devices, add Ethernet switches to your network as required.
- 3 Use an RJ-45 telephone cable to connect the FortiGate Voice fxo1 port to the branch office PSTN phone line supplied by your local telephone service provider.

Configuring basic FortiGate Voice network and UTM settings

The following procedures describe how to configure a FortiGate Voice to provide basic Internet connectivity, network services, and UTM services for the branch office internal network. Network services include configuring the FortiGate Voice to be the DHCP server and DNS server for the internal network.

As part of the FortiGate Voice network interface configuration you must enable *SIP Traffic* on the internal and wan1 interfaces so that the FortiGate Voice unit accepts SIP sessions received by these interfaces. No other special network configuration, firewall policies, or routing is required for the FortiGate Voice to accept SIP sessions from configured extensions.

To configure basic network settings

- 1 Connect to the FortiGate Voice web-based manager.
- 2 Go to *System > Network > Interface*.

- 3 Edit the *internal* interface and configure the following settings:

Addressing Mode	Manual
IP/Netmask	172.20.120.10/255.255.255.0
Enable DNS Query	Select <i>Enable DNS Query</i> then select <i>recursive</i> from the drop-down menu.
SIP Traffic	Select Enable

Configure other network interface settings as required and select OK.

The procedure “[To configure the FortiGate Voice to be a DHCP server for the internal network](#)” on page 2606 describes how to configure the FortiGate DHCP server to configure PCs on the internal network to use the FortiGate Voice internal interface as a DNS server.

- 4 Edit the *wan1* interface and configure the following settings:

Addressing Mode	Manual
IP/Netmask	192.168.10.10/255.255.255.0
SIP Traffic	Select Enable

Configure other network interface settings as required and select OK.



You can also set the *Addressing mode* to *DHCP* or *PPPoE* for the *wan1* interface depending on the requirements of your ISP. In the example the *wan1* interface has a static IP address.

- 5 Go to *System > Network > Options*.
- 6 Add the IP addresses of the primary and secondary DNS servers used by the branch office provided by your ISP.
- 7 Select Apply.
- 8 Go to *Router > Static > Static Route*.
- 9 Edit the default static route and configure the following settings:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	wan1
Gateway	Enter the IP address of the default gateway provided by your ISP.
Distance	10

- 10 Select OK.

To configure the FortiGate Voice to be a DHCP server for the internal network

Use this procedure to add a new DHCP server for the internal network or to change the configuration of the default FortiGateVoice DHCP server. The DHCP server will give PCs on the Internal network IP addresses in the range 172.20.120.110 to 172.20.120.210 and set their default gateway and DNS server to the IP address of the FortiGate Voice internal interface.

- 1 Go to *System > DHCP > Service*.

2 Select the *Create New*.

If a DHCP server has already been added for the internal interface, select the Edit icon to change its configuration.

3 Configure the following settings.

Interface Name	internal
Mode	Server
Enable	Select
Type	Regular
IP Range	172.20.120.110 - 172.20.120.210
Network Mask	255.255.255.0
Default Gateway	172.20.120.10
DNS Service	Specify
DNS Server 0	172.20.120.10

4 Change other settings if required and select OK.**To configure FortiGuard services for the FortiGate Voice unit**

Use the following procedure to configure the FortiGate Voice unit to connect to the FortiGuard Distribution Network (FDN) to update the antivirus, antispam and IPS attack definitions. Before you can begin receiving updates, you must register the FortiGate Voice unit from the Fortinet Support web site. For more information, see [“Registering your Fortinet product” on page 3040](#).

1 Go to *System > Maintenance > FortiGuard*.**2** Select the expand arrow for *AntiVirus and IPS Options* to expand the options.**3** Select *Update Now* to update the FortiGuard services and definitions.

If the connection to the FDN is successful, the web-based manager displays a message similar to the following:

```
Your update request has been sent. Your database will be updated
in a few minutes. Please check your update page for the
status of the update.
```

After a few minutes, if an update is available, the FortiGuard page lists new version information for the FortiGate services and definitions. The system dashboard license information widget also displays new dates and version numbers for the FortiGuard definitions. Messages are recorded to the event log indicating whether the update was successful or not.

To configure basic Internet access and UTM features

This procedure describes how to add a firewall policy that allows users on the internal network to connect to the Internet with antivirus protection. This configuration is not required for VoIP support. It just provides users on the internal network with UTM-protected access to the Internet.

1 Go to *Firewall > Policy* and select *Create New* to add a new firewall policy.**2** Configure the following settings.

Source Interface/Zone	internal
Source Address	all

Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 3 Select UTM and Select Enable Antivirus.
- 4 Select Create New from the drop-down menu.
- 5 Customize the antivirus profile and select OK to apply UTM virus scanning to the traffic accepted by the firewall policy.
- 6 Select OK to save the firewall policy.

Configuring network settings for the devices on the Internal network

You can configure the PCs and other devices on the internal network to get their network configuration automatically using DHCP. If required you can also configure devices on the internal network with static IP addresses on the 172.20.120.0 subnet but outside the range awarded by the FortiGate Voice DHCP server. Example static TCP/IP configuration:

IP Address	172.20.120.20
Subnet Mask	255.255.255.0
Default Gateway	172.20.120.10
DNS Server	172.20.120.10

You can also use the same network configuration for the SIP phones on the internal network.

Configuring the FortiGate Voice PSTN and PBX settings

The procedures in this section describe how to configure the FortiGate Voice unit as the PBX for SIP phones on the branch office internal network. These procedures describe how to configure many of the FortiGate Voice PSTN and PBX features. PSTN features are supported on some FortiGate Voice models. The following procedures are included:

- [To configure the fxo1 PSTN interface](#)
- [To configure basic PBX system and voicemail notification settings](#)
- [To add a VoIP provider](#)
- [To add a dial plan for dialing the PSTN and the main office](#)
- [To add the extensions that are on the branch office internal network](#)

To configure the fxo1 PSTN interface

This procedure describes how to configure the FortiGate Voice fxo1 PSTN interface to connect the FortiGate Voice unit to one PSTN phone line. If you have more PSTN phone lines you can connect and configure more fxo interfaces. Skip this procedure if your FortiGate Voice unit does not include PSTN interfaces.

- 1 Go to *PBX > Service Providers > PSTN Interface* and edit the fxo1 interface.
- 2 Configure the following settings.

Phone Number	Enter the phone number of the PSTN phone line as provided by your phone service provider. The phone number is used for caller ID for calls from the FortiGate Voice unit to the PSTN. It can be any number, but is usually the actual phone number of the PSTN line connected to the fxo1 interface. Area code and country codes are optional.
Display Name	This name is used for caller ID for calls from the FortiGate Voice unit to the PSTN. It can be any name, such as a company name, that identifies the branch office.
Caller ID Options	Configure the following options to support caller ID functions for calls from the internal network to the PSTN.
Catch Caller ID	Select to enable the FortiGate Voice unit to receive caller ID information from calls originating on the PSTN and send the caller ID information to the extension that answers the call.
Caller ID Protocol	Select the caller ID protocol required by PSTN line that the fxo interface is connected to. Contact your service provider for the name of the protocol to use.
Caller ID Indicator	Select the caller ID indicator required by the PSTN line. Contact your service provider for details.
Ring #	Set the number of rings to wait before receiving caller ID information. In most cases, enter 1 to send caller ID information between the first and second ring. Contact your service provider for details.
Hang-up Options	Configure the following options to configure how the FortiGate Voice unit hangs up calls from the PSTN.
Hang up on Polarity Reversal	Select if the PSTN line uses polarity reversal to indicate a call has been hung up. Contact your service provider for details.
Hang up on Busy Tone	Select if you want the FortiGate Voice unit to hang up automatically when it receives a busy tone when attempting to dial a number on the PSTN.
Busy Tone Detection #	The number of busy tones that the FortiGate Voice receives before hanging up if <i>Hang up on Busy Tone</i> is selected.

Busy Tone Duration	Tune the FortiGate Voice unit to accurately detect busy tones on this PSTN line. You can change the default settings if busy tones are not accurately detected.
Busy Tone Interval	
Administrative Status	Set to <i>Up</i> if the fxo interface is connected to the PSTN and you want to be able to receive and send calls on this PSTN interface.

- 3 Select OK.

To configure basic PBX system and voicemail notification settings

Use the following procedure to configure PBX system settings and voicemail notification email settings that affect the overall performance of the PBX service and all of the users of it. Usually you would configure these settings once and rarely thereafter.

- 1 Go to *PBX > Calling Rules > Setting*.
- 2 Configure the following settings.

Extension Pattern	Select <i>Other</i> and enter 6XXX The example extension range means that every extension added to the FortiGate Voice unit must be a four digit number starting with a 6.
Country/Area	Select the country from the drop-down menu.
Country Code	Enter the international country calling code for the country or region in which you are installing the FortiGate Voice unit.
Local Area Code	Enter the local area code for the country or region in which you are installing the FortiGate Voice unit.
Voicemail Access	*97 Phone users on the internal network can dial *97 to get their voicemail.
Outgoing Prefix	9 Phone users must dial 9 to get an outside line. The outgoing prefix should not be the same as the first number of the extension range.
Max Voicemail Duration	60 seconds Limits a single voicemail message to 60 seconds.

- 3 Select Apply to save the changes.
- 4 Go to *Log&Report > Log Config > Alert E-mail*.
- 5 Configure the following settings.

SMTP Server	The name or IP address of an email server that the FortiGate Voice unit can send email notifications to when PBX users receive a voicemail. For example: <i>mail.example.com</i> . You can optionally create an email account on the email server for the FortiGate Voice unit.
Email from	Enter the email address that the alert email messages will come from.

Email to	Enter up to three email address recipients for alert email message.
Authentication	Select if the email server requires authentication.
SMTP User	Enter a valid username for an account on the email server.
Password	Enter the password for the account on the email sever.

- 6 Select Apply to save the changes.

To add a VoIP provider

Use the following procedure to add the information required by the FortiGate Voice unit to use a VoIP provider for routing SIP calls on the main office. In the example, the organization uses a third-party VoIP provider to handle VoIP calls between the head office and the branch office.

- 1 Go to *PBX > Service Providers > SIP Trunk*.
- 2 Select *Create New*.
- 3 Configure the following settings.

Name	VoIP_Provider_1 A name for the VoIP provider. This can be any name.
Domain	192.168.20.10 The VoIP provider's IP address. This could also be the VoIP providers domain name (for example, voip.example.com).
User Name	Enter a valid user name for an account on the VoIP provider's server. This could also be a phone number including area code, depending on the requirements of the VoIP provider.
Password	Enter the password for the account on the VoIP provider's SIP sever.
Authorization User Name	Enter a valid authorization user name for an account on the VoIP provider's server if required by the VoIP provider.
Display User Name	Enter a valid display user name for an account on the VoIP provider's server if required by the VoIP provider.
Account Type	Select Static or Dynamic depending on the account with the VoIP provider.
Registration Interval	If this is a dynamic account with the VoIP provider, enter the registration interval as required by the VoIP provider. After each registration interval the FortiGate Voice renews the registration of the account with the VoIP provider.
DTMF Method	Auto Auto means the VoIP provider's server and the FortiGate Voice unit will negotiate to select a DTMF method. You could also select a specific DTMF method if required.

- 4 Select OK to add the VoIP provider.

To add a dial plan for dialing the PSTN and the main office

Dial plans are used to route calls made from an extension to an external phone system. The external phone system can be the PSTN or a VoIP provider. To route calls to an external phone system you add dial plan rules that include a dial pattern and list of outgoing destinations. When the FortiGate Voice unit receives a call from an extension and the number dialed matches a pattern in a dial plan rule, the FortiGate Voice unit routes the call to the outgoing destination added to the dial plan.

In addition to PSTN and head office support, the dial plan must also support emergency, international, toll free and long distance dialing.

Use the following steps to add a dial plan with the following dial plan rules:

- Allows the branch office to call the PSTN
 - Dialing 911 for emergencies
 - Dialing 9 followed by a country code for international calls
 - Dialing 9 followed by 18 for toll free calls
 - Dialing 9 followed by 1 for long distance calls
 - Dialing 9 for all other PSTN calls
 - Allows the branch office to dial head office extensions directly. The dial plan rule sends calls starting with 2 to the VoIP provider where they are routed to the head office. This dial plan does not include any other settings because users dial the head office extension number directly without a prefix.
- 1 Go to *PBX > Calling Rules > Dial Plan* and select *Create New*.
 - 2 Add a name for the new dial plan, for example, *Dial_Plan_1*.
 - 3 Select OK.
 - 4 Select *Create New* to add the dial plan rule for dialing 911 for emergencies.

Name	Emergency
Use Default Outgoing Prefix ("9")	Not selected
Phone number Begin with	911
Action	Allow
Outgoing Selected	PSTN - fxo1

- 5 Select *Create New* to add the dial plan rule for dialing 9 followed by a country code for international calls.

Name	International
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	011
Action	Block

- 6 Select *Create New* to add the dial plan rule for dialing 9 followed by 18 for toll free calls.

Name	Toll_Free
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	18
Action	Allow
Outgoing Selected	PSTN - fxo1

- 7 Select *Create New* to add the dial plan rule for dialing 9 followed by 1 for long distance calls.

Name	Long_Distance
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	1
Action	Allow
Outgoing Selected	PSTN - fxo2

- 8 Select *Create New* to add the dial plan rule for dialing 9 for all other PSTN calls.

Name	Other_PSTN_Numbers
Use Default Outgoing Prefix ("9")	Selected
Action	Allow
Outgoing Selected	Move <i>PSTN - fxo1</i> to the <i>Selected</i> list to send calls to the PSTN out the fxo1 interface.

- 9 Select *Create New* to add the dial plan rule for dialing the Head Office.

Name	Head_Office_Dial_Rule
Use Default Outgoing Prefix ("9")	Deselect.
Phone number Begin with	2 Indicates that outgoing calls to the Head Office must start with a 2.
Action	Allow
Outgoing	Move <i>VoIP - VoIP_Provider_1</i> to the <i>Selected</i> list to send calls to the PSTN out the fxo1 interface.

- 10 Select OK.

To add the extensions that are on the branch office internal network

Use the following steps to add extensions to the FortiGate Voice unit for the IP phones that are to be connected to the internal network. You add identifying information to each extension entry. The IP phone must be configured with identifying information that matches an entry in the extension list in order to get an extension from the FortiGate Voice unit. Extension numbers are independent of the IP address of the IP phone.

- 1 Go to *PBX > Extension > Extension* and select *Create New*.
- 2 Configure the following settings to add extension 6001.

Extension	6001
Password	<p>The SIP phone user password for the phone assigned to this extension.</p> <p>For a FortiFone on the internal network to be able to register with the FortiGate Voice unit to get this extension, the FortiFone <i>Register Name</i> must consist of the extension <i>First Name</i> followed by the <i>Last Name</i> separated by one space. The FortiFone must also be configured with this <i>Password</i> and the IP address of the FortiGate Voice internal interface.</p> <p>The password must be 8 or more characters, must contain at least an uppercase character, a lowercase character and a number, non-alphanumeric characters, like (- \$, are not supported in the password field.</p>
Type	SIP Phone
First Name	The first name assigned to this extension. Usually a person's first name.
Last Name	<p>The last name assigned to this extension. Usually a person's last name.</p> <p>When this extension calls another phone the caller ID displayed on the called phone consists of the extension <i>First Name</i> followed by the <i>Last Name</i>.</p>
Email	The email address of the person assigned to this extension. The FortiGate Voice unit sends voicemail notifications for the extension to this email address.
MAC Address	Enter the MAC address of the SIP phone.
Dial Plan	Dial_Plan_1
Voicemail	Select
Voicemail Password	Enter the numeric password that the SIP user must enter to get voicemail. The password can contain numbers only.
Email Notification	Select
Voicemail to Email Attachment	Select to attach a recording of the user's voicemail message to the voicemail notification email.
Maximum Message #	<p>50</p> <p>The FortiGate Voice unit keeps up to 50 voicemail messages for this extension.</p>

- 3 Select OK to add the extension.
- 4 Repeat to add more extensions.

Configuring the FortiFones on the internal network

This section contains high-level instructions for installing and configuring FortiFones for the example configuration. For more detailed information see the FortiFone documentation.

To configure FortiFones on the internal network

The following steps describe how to configure a FortiFone on the internal network with extension number 6001. This procedure would also apply to configuring a FortiFone for most networks. See the documentation supplied with the FortiFone for details.

- 1 Connect and power on the FortiFone handset.
- 2 Connect to the handset web configuration interface.
The default web configuration interface address is `http://192.168.0.1`. To connect to this address from a PC, your PC should have an IP address on the 192.168.0.0 subnet, for example: 192.168.0.10/255.255.255.0.
The default Username is root. No password is required.
- 3 Go to *Network > LAN Settings* and set the *IP Type* to DHCP Client and select Submit.
- 4 Select Save & Reboot to save the IP addressing change.
- 5 Log into the FortiFone using the IP address it acquired from the DHCP server.
- 6 Go to *SIP Settings > Service Domain* and add the following configuration information:

Active	On
Display Name	The name to be displayed on the phone. This name is only displayed on this phone. When this phone calls another phone the name displayed is the <i>First Name</i> and <i>Last Name</i> added to the FortiGate Voice <i>Extension</i> configuration.
User Name	6001 This is actually the Line Number or Extension Number and must match the Extension Number added to the FortiGate Voice Extension configuration for this phone.
Register Name	6001 The Register Name is used to authenticate the FortiFone and must match the Extension Number added to the FortiGate Voice Extension configuration for this phone. Both the User Name and Register Name are required.
Register Password	The <i>Password</i> added to the FortiGate Voice Extension configuration for this phone. The Register Name and Register Password are used to authenticate the phone with the FortiGate Voice unit.
Domain Server	Leave this field blank. Not required since the configuration uses the FortiGate Voice unit as a SIP proxy. This field is only used to add the phone to a SIP service domain.

Proxy Server	172.20.120.10 The IP address of the FortiGate Voice internal interface.
Outbound Proxy	Leave this field blank.

- 7 Select Submit.
- 8 Select Save & Reboot to save the service domain information.
- 9 If the FortiFone can successfully connect to and register with the FortiGate Voice unit the *Status* of the FortiFone changes to *Registered*.
If *Status* does not change to *Registered* you should verify the *Register Name* or re-enter the *Password*. You should also confirm that the *Domain Server* and *Proxy Server* IP addresses are correct



If you manually configure FortiFone registration settings, set register interval to 60 seconds where it applies. Refer to FortiFone documentation for more information about FortiFone configuration.

Adding extensions and configuring FortiPhones for users behind a NAT device

When adding an extension for any SIP phone with a NAT device between the phone and the FortiGate Voice unit you must enable NAT in the FortiGate Voice extension configuration for the phone. You can enable NAT only from the CLI. This applies whether the phone is on a remote network behind a NAT device or behind a NAT device on the internal network.

To add an extension for a SIP phone behind a NAT device

The following procedure describes adding the extension from the FortiGate Voice CLI because you must use the CLI to enable NAT. You could add the extension from the web-based manager and then edit the extension from the CLI to enable NAT.

The following configuration is the same whether the phone is behind a NAT device on the internal network or on a remote network,

- 1 Connect to the FortiGate CLI.
- 2 Enter the following command to add extension 6010.

The command includes setting `nat` to `yes` to enable NAT.

```
config pbx extension
  edit 6010
    set first-name <first_name_str>
    set last-name <last_name_str>
    set email <email_str>
    set secret <password_str>
    set dialplan Dial_Plan_1
    set vm-secret <voicemail_password_str>
    set email-notify enable
    set attach enable
    set nat yes
    set macaddress <mac_address>
  end
```

To configure FortiFones behind a NAT device on the internal network

The configuration for FortiFones behind a NAT device on the internal network is the same as for FortiFones directly on the Internal network. See [“To configure FortiFones on the internal network” on page 2615](#).

You may have to configure the NAT device to allow SIP sessions between the FortiFone and the FortiGate Voice unit.

To configure FortiFones behind a NAT device on a remote network

The following steps describe how to configure a FortiFone on the remote network with extension number 6010.

- 1 Connect and power on the FortiFone handset.
- 2 Connect to the handset web configuration interface.
The default web configuration interface address is <http://192.168.0.1>. To connect to this address from a PC, your PC should have an IP address on the 192.168.0.0 subnet, for example: 192.168.0.10/255.255.255.0.
The default Username is root. No password is required.
- 3 Go to *Network > LAN Settings* and set the *IP Type* to DHCP Client and select Submit.
- 4 Select Save & Reboot to save the IP addressing change.
- 5 Log into the FortiFone using the IP address it acquired from the DHCP server.
- 6 Go to *SIP Settings > Service Domain* and add the following configuration information:

Active	On
Display Name	The name to be displayed on the phone. This name is only displayed on this phone. When this phone calls another phone the name displayed is the <i>First Name</i> and <i>Last Name</i> added to the FortiGate Voice <i>Extension</i> configuration.
User Name	6010 This is actually the Line Number or Extension Number and must match the Extension Number added to the FortiGate Voice Extension configuration for this phone.
Register Name	6010 The Register Name is used to authenticate the FortiFone and must match the Extension Number added to the FortiGate Voice Extension configuration for this phone. Both the User Name and Register Name are required.
Register Password	The <i>Password</i> added to the FortiGate Voice Extension configuration for this phone. The Register Name and Register Password are used to authenticate the phone with the FortiGate Voice unit.
Domain Server	Leave this field blank. Not required since the configuration uses the FortiGate Voice unit as a SIP proxy. This field is only used to add the phone to a SIP service domain.

Proxy Server	172.20.120.10 The IP address of the FortiGate Voice internal interface.
Outbound Proxy	Leave this field blank.

- 7 If the FortiFone can successfully connect to and register with the FortiGate Voice unit the *Status* of the FortiFone changes to *Registered*.

If *Status* does not change to *Registered* you should verify the *Register Name* or re-enter the *Password*. You should also confirm that the *Domain Server* and *Proxy Server* IP addresses are correct.

To configure the remote FortiGate unit in NAT mode

The remote FortiGate unit in NAT mode must be configured to allow SIP sessions between the remote users on the remote network and the FortiGate Voice external interface. To do this you need to:

- Add an internal to external firewall policy that allows SIP sessions so that the remote users can start SIP sessions with the FortiGate Voice unit
- Add a virtual IP and an external to internal firewall policy that allows SIP sessions from the FortiGate Voice wan1 interface to connect to the phones in the remote network

For higher security, you could configure IPSec tunneling between the branch office network and the remote network and send SIP traffic over the IPSec tunnel.

FortiGate Voice IVR configuration

By default, when callers call into the FortiGate Voice PBX from a remote system such as the PSTN, the call is picked up by the PBX system which plays a default message asking the caller to dial the extension number that they want to reach or to dial 0 for assistance. If the caller dials 0 they can use the number keys on their phone to spell out the *First Name* or *Last Name* of an extension to connect with that extension.

You can use the following procedure to add a custom welcome message.

To add a custom welcome message

- 1 Log into the FortiGate Voice web-based manager.
- 2 Go to *PBX > Calling Rules > Voice Menu*.
- 3 Enter a *Recorder Extension*.
- 4 Enter a *Password for Recording*.
The password should include numbers only.
- 5 Select OK.
- 6 From a SIP phone that is registered with the FortiGate Voice unit, dial the Extension added in step 3.
- 7 Follow the prompts to record a new welcome message.

Providing access to the company directory

Use the following procedure to allow phone users to dial 3 to access the FortiGate Voice PBX directory. Phone users can use the directory to call an extension by using the number keys on their phone to spell out the *First Name* or *Last Name* of an extension to connect with that extension.

To provide access to the company directory from any extension

- 1 Log into the FortiGate Voice web-based manager.
- 2 Go to *PBX > Calling Rules > Voice Menu*.
- 3 Select the Edit icon for Key 3.
You can select any available key, but this example uses 3.
- 4 Set Action to *Go to Company Directory* and select OK.

Adding a shortcut for checking voicemail


Use the following procedure to allow phone users to dial 7 to access their voicemail.

To provide access to the company directory form any extension

- 1 Log into the FortiGate Voice web-based manager.
- 2 Go to *PBX > Calling Rules > Voice Menu*.
- 3 Select the Edit icon for Key 7.
You can select any available key, but this example uses 7.
- 4 Set Action to *Check Voicemail* and select OK.

Checking voicemail

Once users connect to their voicemail using the *Voicemail Access* number configured from *PBX > Calling Rules > Setting* or by pressing the configured voicemail key they can follow the prompts to listen to, store, and delete messages. Users can also change their voicemail password.



FortiGate Voice web-based manager configuration reference

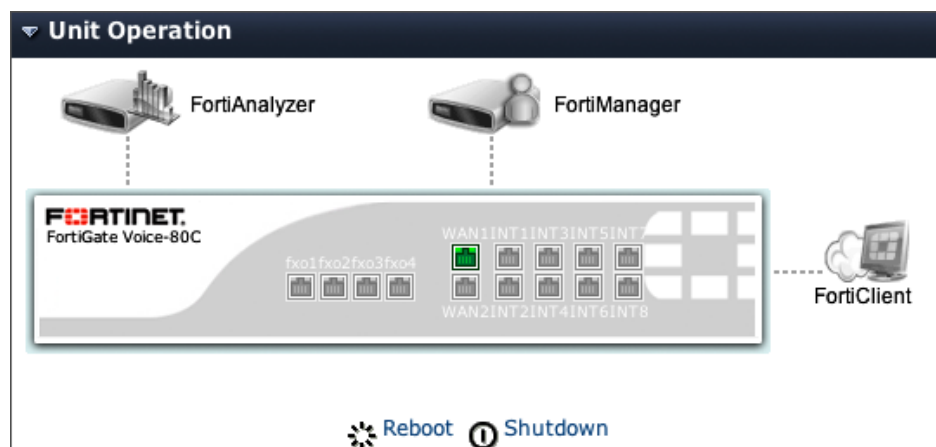
This section describes the following FortiGate Voice web-based manager configuration settings.

- [Unit operation dashboard widget](#)
- [Configuring interface settings to support VoIP PBX features](#)
- [PBX configuration](#)
- [Logging of PBX activities](#)

Unit operation dashboard widget

Go to *System > Dashboard > Status* and view the *Unit Operation* widget to see the status of the FortiGate Voice unit and its Ethernet and fxo interfaces. The fxo interfaces appear if your FortiGate Voice unit includes PSTN interfaces.

Figure 292: FortiGate Voice-80C Unit operation widget



Configuring interface settings to support VoIP PBX features

You can configure the following options on one or more FortiGate Voice ethernet interfaces to support PBX functionality for FortiGate Voice PBX users:

- [Configuring an interface to accept SIP traffic](#)
- [Enabling access to the PBX user web portal](#)

- [SIP phone auto-provisioning](#)



These options may compromise the security of the FortiGate Voice unit interface on which they are enabled (for example by opening a TCP or UDP port). The best practice is to only enable them on an interface connected to a secure network. Use caution when enabling them on an interface that is connected to an unsecure network such as the Internet.

Configuring an interface to accept SIP traffic

For PBX users to access the PBX features of a FortiGate Voice unit, you must configure the FortiGate Voice interface that PBX users connect to accept SIP traffic.



This option allows the FortiGate Voice interface to accept SIP traffic on UDP port 5060. The best practice is to enable this option on an interface connected to a secure network. Use caution when enabling this option on an interface that is connected to an unsecure network such as the Internet.

To enable SIP traffic on a FortiGate Voice interface

- 1 Go to *System > Network > Interface*.
- 2 Select the interface that you want to configure VoIP settings for.
The interface connected to the network containing SIP phones.
- 3 Select the *SIP* check box to enable SIP traffic.
- 4 Select *OK*.

Enabling access to the PBX user web portal

PBX users can log into the FortiGate Voice PBX user web portal using their PBX extension and password. The connection to the web portal is secured using HTTPS. From the portal, users can view their PBX configuration, change some configuration settings such as their password, listen to their voicemail, configure call forwarding, access conference calls, and listen to recorded conference calls.



This option allows users to log into the FortiGate Voice unit using an HTTPS connection to the configured port (8443 by default). The best practice is to enable this option only on an interface connected to a secure network. If you enable this option on an interface that is connected to an unsecure network such as the Internet make sure the extensions have secure passwords.

For information about how to log into and use the PBX user web portal, see [“Using the PBX user web portal” on page 2649](#).

To allow access to the FortiGate Voice PBX user web portal

You can configure one or more FortiGate Voice interfaces to accept connections to the PBX user web portal.

- 1 Go to *System > Network > Interface* and edit an interface.
- 2 Enable the *PBX User Portal* for this interface.

To log into the PBX user web portal

- 1 Browse to the following address:

`https://<interface_address>:8443`

where:

`<interface_address>` is the IP address or domain name of a FortiGate Voice interface on which you have enabled the PBX user portal.

8443 is the default port that users must browse to connect to the portal.

SIP phone auto-provisioning

You can configure the FortiGate Voice unit to auto-provision SIP phones on your network. The SIP phones must support auto-provisioning using TFTP. See the FortiGate Voice release notes for a list of supported SIP phones.

With auto-provisioning configured, when a supported SIP phone is connected to the network and powered on it automatically receives all of its PBX setup information from the FortiGate Voice unit. In most cases the administrator does not have to make configuration changes to the SIP phone itself.

To configure FortiGate Voice auto-provisioning you need to enable *Phone Auto-Provision* for the FortiGate Voice interface connected to the same network as the SIP phones to be auto-provisioned.



Enabling Phone Auto-Provision on a FortiGate Voice interface opens UDP port 69 (the TFTP port) on that interface. You should only enable phone auto-provisioning on a secure network.

You must also configure a DHCP server on this interface. In the DHCP server configuration, select DHCP option 66 (an advanced option on the web-based manager) and include the IP address of this FortiGate Voice interface.



DHCP server option 66 identifies a TFTP server and includes the IP address of the TFTP server and downloads the TFTP server identity to the device that gets an IP address from the DHCP server. DHCP option 66 is defined in [RFC 2132](#).

The address must be input in hexadecimal ASCII format. In this format, each ASCII character in the dotted decimal IP address (including periods) is represented by a 2-digit hexadecimal number that maps to the corresponding ASCII character. For example, for the FortiGate Voice default configuration, the IP address of the FortiGate Voice internal interface (192.168.1.99) is input in hexadecimal format as follows:

Hexadecimal	31	39	32	2E	31	36	38	2E	31	2E	39	39
IPv4	1	9	2	.	1	6	8	.	1	.	9	9

For a hexadecimal representation of all ASCII characters, see the [ASCII printable characters](#) list on Wikipedia.

With this configuration in place, when a SIP phone that supports auto-provisioning with TFTP receives its IP address via DHCP from the FortiGate Voice unit, it also receives the TFTP server address. The SIP phone then downloads its PBX configuration from TFTP server running on the FortiGate Voice unit.

The FortiGate Voice unit identifies the SIP phone by its MAC address. When you add an extension for the phone to the PBX configuration you include its MAC address. The FortiGate Voice unit uses this MAC address to identify the SIP phone and match it with its extension in the PBX configuration. In most cases the MAC address is available from a sticker attached to the phone or its packaging.

Default FortiGate Voice auto-provisioning configuration

By default the FortiGate Voice unit enables auto-provisioning on the Internal interface for SIP phones attached to the internal network. The default IP address of the internal interface is 192.168.1.99. The FortiGate Voice default configuration also includes a DHCP server for the Internal interface that provides IP addresses in the range of 192.168.1.110 to 192.168.1.210. You can configure all of the devices on your internal network (including PCs and SIP phones and so on) to use DHCP to automatically get their IP addresses from the FortiGate Voice unit DHCP server. SIP phones that support auto-provisioning can also get their PBX configuration if an extension for the SIP phone's MAC address has been added to the FortiGate Voice PBX configuration.

If you want to configure the network connected to the FortiGate Voice internal interface to use a different subnet but still support auto-provisioning, you must:

- Change the IP address of the FortiGate Voice internal interface.
- Change the IP addresses that the DHCP server provides.
- Change the TFTP server IP address added to the DHCP server option 66.

Configuring SIP phones for auto-provisioning

Before you begin, make sure your FortiGate Voice unit is set up for auto-provisioning and the configuration meets the following:

- A DHCP server set up as described.
- A PBX extension set up with SIP phone's MAC address.

This section describes how to configure auto-provisioning for the following phones:

- FortiFone-110
- FortiFone-210
- Polycom
- Aastra

To configure a FortiFone-110 phone for auto-provisioning

- 1 Connect the FortiFone-110 phone to the VoIP network.
Refer to the phone's documentation for instructions for connecting the phone.
- 2 Reset the phone to its factory default settings by selecting *Menu > Phone Settings > Reset to Factory*.
- 3 Set DHCP IP client as WAN port option by selecting *Menu > Network > WAN Port Option > DHCP IP Client*.

The FortiFone-110 phone reboots and starts acquiring an IP address and PBX extension settings from the FortiGate Voice unit. During this time **invalid account** followed by **Updating configuration, please wait** appears on the screen. After a few minutes, the phone reboots one last time. The phone extension settings appear on the screen, the IP address and the other connection details can be found by navigating through the phone's menu.

To configure a FortiFone-210 phone for auto-provisioning

- 1 Connect the FortiFone-210 phone to the VoIP network.
Refer to the phone's documentation for instructions for connecting the phone.

- 2 Reset the phone to its factory default settings.

Refer to the phone's documentation.

The FortiFone-210 phone reboots and starts acquiring an IP address and PBX extension settings from the FortiGate Voice unit.

To configure Polycom phones for auto-provisioning

- 1 Connect the Polycom phone to the VoIP network.

Refer to the phone's documentation for instructions for connecting the phone.

- 2 Select *Setup* on the screen.

- 3 Enter the administrator password.

The default password is 456.

- 4 Set the server type to Trivial FTP by selecting *V > Server Menu > Server Type > Trivial FTP*.

- 5 Select *Exit*.

The Polycom phone reboots and starts acquiring an IP address and PBX extension settings from the FortiGate Voice unit.

To configure Aastra phones for auto-provisioning

- 1 Connect the Aastra phone to the VoIP network.

Refer to phone's documentation for instructions for connecting the phone.

- 2 Reset the phone to its factory default settings.

Refer to the phone's documentation.

The Aastra phone reboots and starts acquiring an IP address and PBX extension settings from the FortiGate Voice unit.



If you manually configure any FortiFone's registration settings, set register interval to 60 seconds where it applies. Refer to FortiFone documentation for more information about FortiFone configuration.

PBX configuration

The following explains how to configure PBX settings for your network environment. These settings include, adding extensions to your PBX, voicemail notification settings, configuring the FortiGuard Voice service, configuring other VoIP providers, as well as system settings such as a voicemail access code and a maximum voicemail duration time limit.

Configuring extensions

Add PBX extensions to configure the settings for the SIP phones that the FortiGate Voice unit supplies PBX services for. Each extension configuration includes the extension number, the name and email address of the person assigned to the extension, the MAC address of the device to be used for the extension (to support auto-provisioning), whether the device includes video calling, and voicemail settings. You can also use the extension settings to add conference bridges.

To add a new extension, go to *PBX > Extension > Extension*, select *Create New*, enter the information and then select *OK*.



FortiGate Voice unit uses the alertmail settings to access an SMTP server and send email notifications. Alertmail is configured in *Log&Report > Log Config > Alert E-mail*.

Figure 293: Configuring extensions

Extension	Type	Name	Dial Plan	MAC Address	Status
6555	SIP Phone	first last	company-default	00:00:00:00:00:00	Unregistered
6777	Conference	N/A	N/A	00:00:00:00:00:00	N/A

New Extension	
Extension (Pattern:6XXX)	6111
Password	*****
Type	SIP Phone
First Name	First
Last Name	Last
Email	flast@example.com
MAC Address	00:30:4F:74:5F:63
Dial Plan	company-default
Video Capability	<input type="checkbox"/>
Voicemail	<input checked="" type="checkbox"/>
Voicemail Password	*****
Email Notification	<input checked="" type="checkbox"/>
Voicemail to Email Attachment	<input type="checkbox"/>
Maximum Message#	50 (1-9999)

General extension settings

Create New	Select to create an extension.
Extension	The extension number.
Type	The type of extension the number is. Type can be: <ul style="list-style-type: none"> SIP Phone to configure a SIP phone extension Conference to configure a conference bridge. For the Conference extension you can add an extension number and a password. PBX users can call this extension number and enter the password to join a conference call.
Name	The name of the extension.
Dial Plan	The dial plan that will be used for that extension.

Extension configuration settings

Extension	Enter the extension number. The web-based manager shows the pattern that the extension number must follow. For example, (Pattern: XXX) means the extension must consist of three numbers.
------------------	---

Password	<p>Enter a password for the extension. The password is used to log into the PBX user web portal. The password cannot be blank, must be 8 or more characters, must contain at least one uppercase character, one lowercase character and one number. Non-alphanumeric characters, like (- \$, are not supported in the password field.</p> <p>If you are entering a password for a conference bridge, the password cannot be blank and must contain only numbers. This becomes the conference PIN number.</p>
Type	Select the type of extension. You can choose from SIP Phone or Conference. If you select SIP phone you can add a SIP extension to the PBX configuration. If you select Conference you can add a conference bridge to the PBX configuration. To configure a conference bridge you specify the conference bridge extension, password, conference host, host pin, video capability, and whether conferences can be recorded.
Conference Host	The name of the PBX extension that can host conference calls using this conference bridge. The user of the conference host extension can manage the conference bridge from the PBX user web portal. The conference host user can change the conference password and host pin and change the video and recordable settings for the conference bridge.
Host PIN	Enter the number to be entered by the conference host to be able to host a conference call.
Conference Recordable	Select to enable recording of conference calls that use this conference bridge. You cannot select this option if the conference bridge supports video calls.
First Name	Enter the first name of the person that will be using this extension.
Last Name	Enter the surname of the person that will be using this extension.
Email	Enter the email address of the person that will be using this extension.
MAC Address	The MAC address of the SIP phone to be used for this extension. The MAC address is required if the SIP is to be configured using auto-provisioning. See “SIP phone auto-provisioning” on page 2623 .
Video Capability	Select if the SIP phone can display video and handle video calls.
Dial Plan	Select the dial plan that will be used with this extension from the drop-down list.
Voicemail	Select if you want to have voicemail available for this extension.
Voicemail Password	Enter a voicemail password for accessing the voicemail.
Email Notification	Select to have an email sent to the email address given in the Email field so that the person is notified when a voicemail message is in their voicemail message inbox.

Voicemail to Email Attachment	Select to attach the actual voicemail message to the notification email.
Maximum Message #	Enter a number for the maximum number of messages that can be stored in the extension's voicemail inbox before automatically deleting the oldest messages.

Configuring extension groups (ring groups)

Extension groups (also called ring groups) are a group of extensions that can be called using one number. The extension group can be used to call all the extensions in the group at the same time or to call the extensions one at a time until someone answers.



The order in which the members are added to the ring group does not match the order in which the FortiGate Voice unit calls them.

To add an extension group, go to *PBX > Extension > Group*, select *Create New*, enter the information, and then select *OK*.

Figure 294: Configuring extension groups

The screenshot shows the 'New Group' configuration form in the FortiGate Voice web-based manager. The form is open for extension 6001, with a description of 'Support'. The 'Ring Strategy' is set to 'Sequential'. The 'No Answer Action' is set to 'Voicemail'. The 'Voicemail Extension' is set to '6111 - User Name'. The 'Member' list is empty, and the 'Available' list shows '6111 - User Name' and '6555 - first last'.

Extension Number	The number to call to reach extension group. This number must be a valid extension number for the FortiGate Voice configuration.
Description	A description of the extension group.
# of Members	The number of extensions in the extension group.
Ring Strategy	Select a type from the drop-down list. You can choose either Sequential or Ring All.
No Answer Action	Select the action to take when there is no answer for the incoming caller. You can select Voicemail, which routes the caller to voicemail, IVR, or Hangup. If you select Voicemail, the Voicemail Extension list appears and you need to select the voicemail extension number.

Voicemail Extension	Select the voicemail extension number from the drop-down list. This option appears only when Voicemail is selected in No Answer Action.
Member	Select an extension in the Available column and then use the -> arrow to move it to the Selected column. To remove an extension from the Selected column, select the extension and use the <- arrow to move it back to the Available column.

Configuring service providers (the FortiGuard Voice service)

If your FortiGate Voice unit is installed in North America and the Country Code is set to 1 then you can use the FortiGuard Voice service as your SIP service provider. (The default Country Code is 1 and is set from *PBX > Calling Rules > Setting*.) The FortiGuard Voice service is supported only in North America. If you install the FortiGate Voice unit elsewhere in the world and change the Country Code, the FortiGuard Voice Service configuration is not available.

Configuring PSTN interfaces

Some FortiGate Voice models include public switched telephone network (PSTN) interfaces that you can use to connect the FortiGate Voice PBX to your local public telephone network. Using these interfaces you can route calls from your FortiGate Voice network to the public telephone network. The PSTN interfaces are named fxo1, fxo2, and so on.

To configure the PSTN interfaces, go to *PBX > Service Providers > PSTN Interface*, configure settings for the fxo interface and then select *OK*.

Figure 295: Configuring PSTN interfaces

	Name	Phone Number	Display Name	Catch Caller ID	Administrative Status
<input type="checkbox"/>	fxo1	613-222-5555	Example	✓	↑
<input type="checkbox"/>	fxo2			✓	↑
<input type="checkbox"/>	fxo3			✓	↑
<input type="checkbox"/>	fxo4			✓	↑

Edit PSTN Interface

Basic Options

Name: fxo1

Phone Number: 613-222-5555

Display Name: Example

Caller ID Options

Catch Caller ID: ☒

Caller ID Protocol: Bell

Caller ID Indicator: ☒ Ring ☐ Polarity

Ring #: 1 (1-4)

Hang-up Options

Hang up on Polarity Reversal: ☒

Hang up on Busy Tone: ☒

Busy Tone Detection #: 4 (1-8)

Busy Tone Duration: 500 (milliseconds)

Busy Tone Interval: 500 (milliseconds)

Administrative Status: ☒ Up ☐ Down

OK Cancel

General PSTN interface settings

Name	The name of the PSTN interface.
Phone Number	The phone number that is associated with that PSTN interface.
Display Name	The name that displays on the phone's LCD.
Catch Caller ID	If enabled, a green checkmark appears. If Catch Caller ID is disabled, a gray X appears.
Administrative Status	Status of the PSTN interface. A red down arrow indicates that the interface is down; a green up arrow indicates that the interface is up.

PSTN interface configuration settings

Basic Options	The basic options for the interface.
Name	The name of the PSTN interface.
Phone Number	<p>Enter the phone number of the PSTN phone line as provided by your phone service provider.</p> <p>The phone number is used for caller ID for calls from the FortiGate Voice unit to the PSTN. It can be any number, but is usually the actual phone number of the PSTN line connected to the fxo1 interface. Area code and country codes are optional.</p>
Display Name	This name is used for caller ID for calls from the FortiGate Voice unit to the PSTN. It can be any name, such as a company name, that identifies the branch office.

Caller ID Options	Configure the following options to support caller ID functions for calls from the internal network to the PSTN.
Catch Caller ID	Select to enable the FortiGate Voice unit to receive caller ID information from calls originating on the PSTN and send the caller ID information to the extension that answers the call.
Caller ID Protocol	Select the caller ID protocol required by PSTN line that the fxo interface is connected to. Contact your service provider for the name of the protocol to use.
Caller ID Indicator	Select the caller ID indicator required by the PSTN line. Contact your service provider for details.
Ring #	Set the number of rings to wait before receiving caller ID information. In most cases, enter 1 to send caller ID information between the first and second ring. Contact your service provider for details.
Hang-up Options	Configure the following options to configure how the FortiGate Voice unit hangs up calls from the PSTN.
Hang up on Polarity Reversal	Select if the PSTN line uses polarity reversal to indicate a call has been hung up. Contact your service provider for details.
Hang up on Busy Tone	Select if you want the FortiGate Voice unit to hang up automatically when it receives a busy tone when attempting to dial a number on the PSTN.
Busy Tone Detection #	The number of busy tones that the FortiGate Voice receives before hanging up if <i>Hang up on Busy Tone</i> is selected.
Busy Tone Duration	Tune the FortiGate Voice unit to accurately detect busy tones on this PSTN line. You can change the default settings if busy tones are not accurately detected.
Busy Tone Interval	
Administrative Status	Set to <i>Up</i> if the fxo interface is connected to the PSTN and you want to be able to receive and send calls on this PSTN interface.

Configuring the FortiGuard Voice service

The FortiGuard Voice service is a SIP trunking or SIP VoIP service provided by Fortinet. The FortiGuard Voice service includes the following features in addition to the being able to make calls from the FortiGate Voice unit to other VoIP providers and to the PSTN:

- Direct inward dial numbers
- FortiFAX eFAX service (see [“FortiFAX service” on page 2644](#))
- Toll free numbers
- Different calling card packages and credit levels

The FortiGuard Voice service is currently available only in North America. To use the FortiGuard Voice service the FortiGate Voice unit Country Code must be set to 1 and the FortiGate Voice unit must have a valid static public Internet IP address.

You can connect your FortiGate Voice unit to the FortiGuard Voice service by purchasing a subscription for your FortiGate Voice unit. Once you have purchased a subscription, if your FortiGate Voice unit meets the above criteria and is connected to the Internet it is also automatically connected to the FortiGuard Voice service network.

You can also go to the License Information dashboard widget or to *PBX > Service Providers > FortiGuard Voice Service* to subscribe to the FortiGuard Voice service. If you have already subscribed to but have not activated the service you can select *Activate Now* from either of these locations to activate the service.

To activate the service you must fill out the following fields:

- Area Code: The local area code used by the FortiGate Voice unit.
- Company Name: Name of the company.
- Directory List Address: Local number to call for directory listings, for example 411.
- E911 Emergency Address: Local emergency number, for example 911.

When you activate the service, FortiGuard sends the direct inward dial, EFax number, and toll free number settings assigned to the FortiGuard Voice service license assigned to the FortiGate Voice unit.

When the FortiGuard Voice service is active it is integrated into your FortiGate Voice default configuration, the company-default dial plan sends all outgoing calls to the FortiGuard Voice service.

To view the status of your FortiGuard Voice service subscription find the *Voice Service* entry in the License Information dashboard widget. You can also go to *System > Maintenance > FortiGuard* and view the status of the Voice Service entry in the FortiGuard Subscription Services list. The status of the FortiGuard Voice service should indicate that the FortiGate Voice unit is licensed for the service.

Viewing FortiGuard Voice service status

You can also view more detailed status information about the FortiGuard Voice service by going to *PBX > Service Providers > FortiGuard Voice Service*. You can view the overall status of your FortiGuard Voice service account as well as details about the service that you have purchased.

Figure 296: FortiGuard Voice service status



Subscription	The status of the FortiGuard Voice service subscription for the FortiGate Voice unit. Status should be Yes if everything is properly configured.
Account Status	The status of the FortiGuard Voice service account used by the FortiGate Voice unit. If account status is active the FortiGate Voice unit can receive and send calls from and to the FortiGuard Voice service.

DIDs	The direct inward dial numbers available from FortiGuard Voice service. DID allows the FortiGate Voice unit to direct calls from external callers directly to PBX extensions. For more information about DID, see “Configuring direct inward dialing” on page 2640 .
FortiFAX	The FortiGate Voice unit’s FortiFAX eFax number if this is part of the FortiGuard Voice service license for the FortiGate Voice unit. Third parties can send faxes to the FortiGate Voice unit using this number.
Toll Frees	The FortiGate Voice unit’s toll free number if this is part of the FortiGuard Voice service license for the FortiGate Voice unit. Third parties can call this number toll free to reach extensions connected to the FortiGate Voice unit.
Packages	Information about additional packages that are part of the FortiGuard Voice services purchased for the FortiGate Voice unit.
Package Type	The package type depends on the FortiGuard Voice services that have been purchased. Multiple packages are available with different features.
Calling Card Credit Left	The amount of money available for making phone calls to the FortiGuard Voice service from the FortiGate Voice unit for each package.
Calling Card Credit Used	The amount of money spent for making phone calls to the FortiGuard Voice service from the FortiGate Voice unit for each package.
Expiration Date	The date on which the package expires. When the package expires all unused calling card credit is lost.
SIP Status	An indicator of the SIP status of the FortiGate Voice service. If it is operating normally the status should show a green check mark icon.

Adding SIP trunks

You can configure multiple VoIP providers for your PBX configuration.

To configure VoIP providers, go to *PBX > Service Providers > SIP Trunk*, select *Create New*, configure the settings and then select *OK*.

Figure 297: VoIP Provider

	Name	Domain	User Name	Account Type	DTMF Method	Status
<input type="checkbox"/>	_FtgdVoice_1	208.91.115.145	10009	Dynamic	RFC2833	Registered

New VoIP Provider

Name

voip-123

Domain

172.16.26.155

User Name

user1

Password

Authorization User Name

user1

Display User Name

example name

Account Type

☒ Static
 ☐ Dynamic

Registration Interval

0 (seconds)

DTMF Method

Auto

Video Capability

☐

OK

Cancel

Name	Enter the name for the VoIP provider configuration. This can be any name.
Domain	The VoIP provider's domain name or IP address. For example, 172.20.120.11 or voip.example.com.
User Name	Enter a valid user name for an account on the VoIP provider's server. This could also be a phone number including area code, depending on the requirements of the VoIP provider.
Password	Enter the password for the account on the VoIP provider's SIP sever.
Authorization User Name	Enter a valid authorization user name for an account on the VoIP provider's server if required by the VoIP provider.
Display User Name	Enter a valid display user name for an account on the VoIP provider's server if required by the VoIP provider.
Account Type	Select Static or Dynamic depending on the account with the VoIP provider.
Registration Interval	If this is a dynamic account with the VoIP provider, enter the registration interval as required by the VoIP provider. After each registration interval the FortiGate Voice renews the registration of the account with the VoIP provider.
DTMF Method	Select the DTMF method used by the VoIP provider. Options are RFC2833, Inband, Info, and Auto. Auto means the VoIP provider's server and the FortiGate Voice unit will negotiate to select a DTMF method. You could also select a specific DTMF method if required.
Video Capability	Select if the SIP service provider supports video calling over SIP.

Branch Office

The branch office feature allows you to easily configure communication between two FortiGate Voice phone systems. A common application is the easy linking of branch offices in separate locations.

To configure Branch Office, go to *PBX > Service Providers > Branch Office*, select *Create New*, configure the settings and then select *OK*.

Figure 298: Branch office

	Name	Prefix	Pattern	IP Address	Dialplan	Status
<input type="checkbox"/>	Branch 15		7XXX	172.20.120.124		N/A

Create New

Edit

Delete

New Branch Office

Name

Prefix

Pattern

IP Address

Registration ☒ No ☐ Yes

Dialplan [Please Select]

OK

Cancel

Name	The name for the branch office configuration. This can be any name.
Prefix	The prefix required to dial an extension to the branch office. If there is no overlap between the extensions in your office and the branch office, no prefix is required. If no prefix is specified, extensions in the branch office are dialed exactly as local extensions.
Pattern	The extension number pattern of the branch office. For example, XXXX is any four digit number while 7XXX is a four digit number that always starts with 7.
IP Address	The IP address of the branch office FortiGate Voice unit.
Registration	<p>The registration determines the means by which a caller from another branch is authenticated.</p> <p>The No setting uses authentication by IP address. This is the preferred method because it is simple and convenient.</p> <p>The Yes setting requires extension and extension password for authentication. This option is provided for compatibility with equipment that does not support IP address authentication.</p>
Dialplan	<p>Specify a dialplan to allow a user to call out of a branch office. The user calls will be routed as specified in the dial plan.</p> <p>If no dial plan is selected, users will be able to dial only branch extensions.</p>

Configuring dial plans

Dial plans route calls made from a FortiGate Voice extension to an external phone system. The external phone system can be one or more PSTN lines if your FortiGate Voice unit includes PSTN interfaces, or a VoIP service provider. To route calls to an external phone system you add dial plan rules that define the extra digits that extension users must dial to call out of the PBX. The rules also control how the FortiGate Voice unit handles these calls including whether to block or allow the call, the destinations the calls are routed to and whether to add digits to the beginning of the dialed number (called prepending).

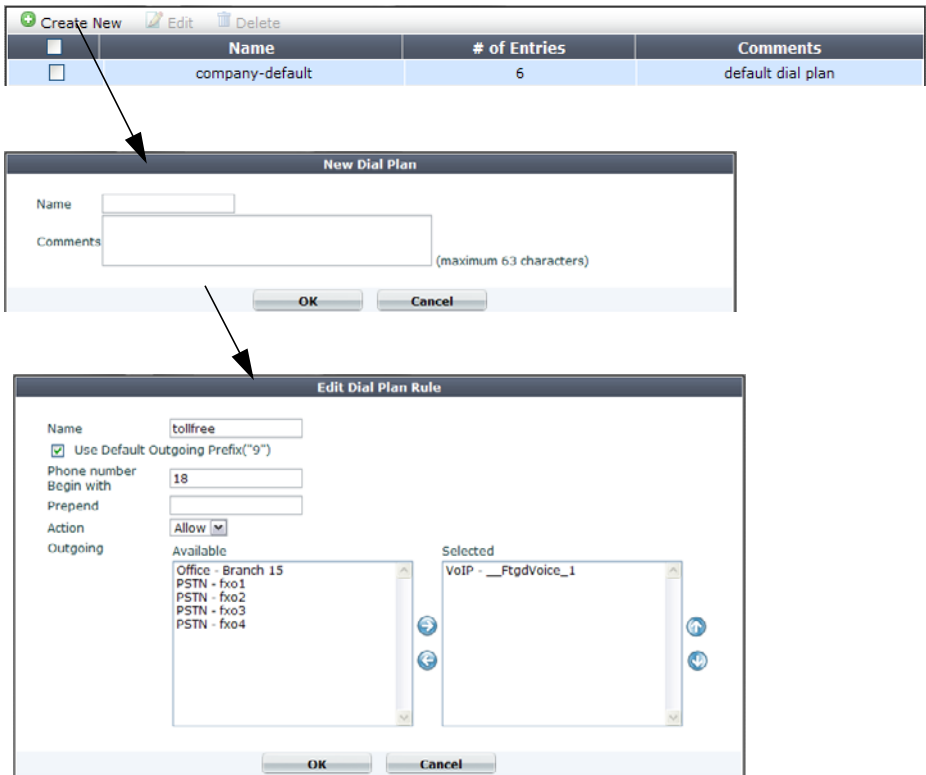
For example, if PBX users should be able to dial 911 for emergencies you should include a dial plan rule that sends all calls that begin with 911 to an external phone system. This rule should also override the default outgoing prefix so that users can dial 911 without having to dial 9 first.

You can also use dial plan rules to block some calls. For example, if you want to block extensions from making international calls you can add dial plan rule that blocks calls that start with the default outgoing prefix followed by 011.

When the FortiGate Voice unit receives a call from an extension that does not match the FortiGate Voice unit's extension range, the call is processed according to the dial plan added to the extension. (If the extension does not have a dial plan the call is blocked). To process the call, the FortiGate Voice unit selects the dial plan rule that best matches the dialed numbers and processes the call using the settings in the dial plan rule. For example, the emergency dial plan rule could route calls out a local PSTN line (if your FortiGate Voice unit includes them) or to a remote VoIP provider.

To configure dial plan, go to *PBX > Calling Rules > Dial Plan*, select *Create New*, configure the settings and then select *OK*.

Figure 299: Configuring a dial plan



General dial plan list settings	
Create New	Select to configure a dial plan. You can add multiple dial plans and assign them to different extensions. For example, you might want to have a dial plan that allows long distance calls and a dial plan that does not.
Name	The name of the dial plan.
# of Entries	The number of entries in each dial plan.
Comments	An optional description of the dial plan.
Dial plan rule configuration settings	
Name	Enter a descriptive name for the dial plan rule.
Use Default Outgoing Prefix("9")	Select this checkbox if the dial plan rule should use the default outgoing prefix (usually 9).
Outgoing Prefix	If you clear the Use Default Outgoing Prefix checkbox you can enter a different outgoing prefix for this dial plan.

Phone number Begin with	Enter the leading digits of the phone number that this dial plan rule should match with. For example, a dial plan rule for toll free numbers in North America should begin with 18. The FortiGate Voice uses a best match to match a dialed number with a dial plan. So each dial plan should have a different Phone number Begin with setting. But you should plan your dial plan to make sure that unexpected matches do not occur.
Prepend	Add digits that should be prepended or added to the beginning of the dialed number before the call is forwarded to its destination. You can prepend digits at the beginning of a call of special dialing is required to reach and external phone system.
Action	Set the action to Allow if this dial plan rule should allow a call. Set the action to Block if the dial plan should block a call. For example, if you want to block international calls you could set the Phone Number begin with to 011 and set the action to block.
Outgoing	<p>In the <i>Available</i> column, select one or more PSTN interfaces (if your FortiGate Voice unit includes them) and/or VoIP service providers that the calls matching this dial plan should be routed to and use the arrow to move to them to the <i>Selected</i> column.</p> <p>If you need to remove a PSTN interface or VoIP provider from the <i>Selected</i> list, select the item and use the arrow to move it back to the <i>Available</i> column list.</p> <p>The FortiGate Voice unit uses the PSTN interfaces and VoIP providers in the <i>Selected</i> list in the order in which they are arranged in the list. You can arrange the PSTN interfaces and VoIP providers in the <i>Selected</i> column using the up and down arrows beside the <i>Selected</i> column. Select a PSTN interface or VoIP provider and then use the arrows to arrange them in the list.</p>

Example dial plan

This simplified example dial plan is similar the default FortiGate Voice dial plan. The default dial plan that routes all external calls to the FortiGuard Voice service. The following example includes 5 dial plan rules that:

- Routes emergency calls (dialing 911) to the fxo1 PSTN interface
- Blocks international calls (the phone number begins with 011)
- Routes Toll Free calls (beginning with 18) to the FortiGuard Voice service
- Routes non-international long distance calls (beginning with 1) to the FortiGuard Voice service
- Routes all other external calls to the fxo2 and fxo3 PSTN interfaces

In this example, all outgoing calls are routed to the PSTN and not to other VoIP service providers. On a FortiGate Voice unit without PSTN interfaces, the dial plan would route all calls to the FortiGuard Voice service or to one or more VoIP service providers.

Table 145: Rule 1: emergency calls using 911

Name	Emergency
Use Default Outgoing Prefix ("9")	Not selected
Phone number Begin with	911
Action	Allow
Outgoing Selected	PSTN - fxo1

Table 146: Rule 2: international calls beginning with 011

Name	International
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	011
Action	Block

Table 147: Rule 3: Toll free calls starting with 18

Name	Toll_Free
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	18
Action	Allow
Outgoing Selected	VoIP - __FtgdVoice_1

Table 148: Rule 4: Long Distance calls starting with 1

Name	Long_Distance
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	1
Action	Allow
Outgoing Selected	VoIP - __FtgdVoice_1

Table 149: Rule 5: Other outgoing calls

Name	Other_PSTN_Numbers
Use Default Outgoing Prefix ("9")	Selected
Phone number Begin with	
Action	Allow
Outgoing Selected	PSTN - fxo2, PSTN - fxo3

Configuring voice menu options

The operator voice mail message can be accessed and reprogrammed by configuring recorder extension. To configure a recorder extension, go to *PBX > Calling rules > Voice Menu* and enter an extension number and a password. The recorder extension can be dialed from any PBX extension and used for recording a new operator voice mail message.

Configure voice menu options to provide PBX users with shortcuts to PBX functions such as accessing their voice mail, finding numbers in the company directory, or dialing a ring group.

To access voice menu functions PBX users dial a single number on their phones and wait a few seconds for the PBX to respond. For example, you can use voice menu options to allow PBX users to simply dial 3 to access their voicemail.

To configure voice menu options

- 1 Go to *PBX > Calling Rules > Voice Menu*.
- 2 Edit the row of the key that you want to configure voice menu options for.
- 3 In the *Action* drop-down list, select one of the following:

None	No action will be taken when a caller dial this number.
Ring Group	The PBX user calls a ring group. Select the ring group to call. A ring group is also called an extension group. To add ring groups, see “Configuring extension groups (ring groups)” on page 2628.
Check Voicemail	Provides direct access to the PBX user's voice mail inbox.
Go to Company Directory	Provides direct access to the PBX company phone directory.

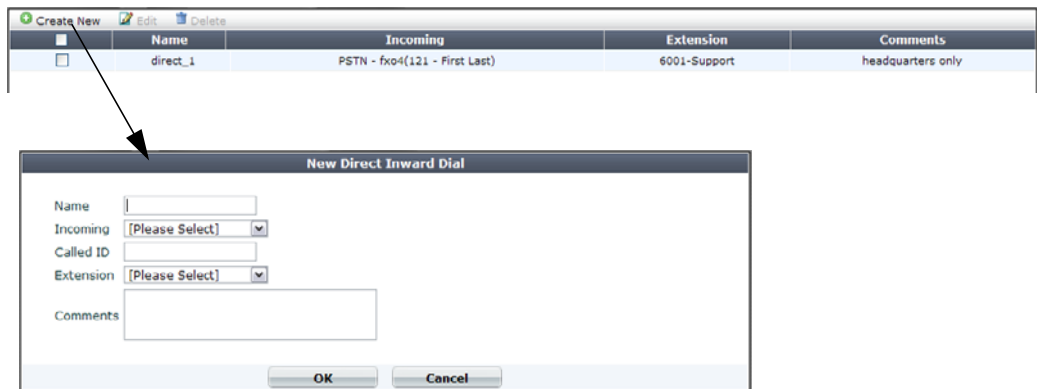
- 4 Select *OK*.

Configuring direct inward dialing

You can configure direct inward dialing (DID) for calls. DID allows the FortiGate Voice unit to direct calls from external callers directly to PBX extensions. For example, you could set up DID so that external users call 555-1234 and DID directs the call to extension 1234. Using the FortiGate Voice unit direct inward dial settings you associate an incoming PSTN interface (if supported by your FortiGate Voice unit), the FortiGuard Voice service, or VoIP service provider with a PBX extension. When an incoming call is received from one of these sources, if the last digits of the dialed number match the selected extension number the FortiGate Voice unit directs the call directly to the extension. For this to work you must obtain an external phone number with the last digits matching the selected extension.

To configure direct inward dialing, go to *PBX > Calling Rules > Direct Inward Dial*, enter the information, and then select *OK*.

Figure 300: Direct inward dialing




Name	A name for the DID configuration.
Incoming	Select a PSTN interface, the FortiGuard Voice service, or VoIP service provider from which to accept calls to this DID number.
Called ID	Enter the ID displayed on the callers phone when they dial this DID number.
Extension	Select an extension from the drop-down list.
Comments	Enter a description, if applicable, for the direct inward dialing configuration.

From the CLI you can use the `cid-number` option of the `config pbx did` command to specify the number called from an external line that is re-directed to the selected extension. Use this option if the extension number cannot be matched with the external number. In the following example, DID sends calls received on the `fxo1` PSTN interface that end with 5555 to extension 1234.

```
config pbx did
  edit did_example
    set external-line fxo1
    set cid-number 5555
    set extension 1234
  end
end
```

Configuring PBX global settings

Configure PBX global settings that affect the overall performance of the PBX service and all of the users of it. Settings include the extension pattern for the PBX, information about the country or area in which the FortiGate Voice unit is installed, and the outgoing dial prefix. Usually you would configure these settings once and rarely thereafter.



The Country Code must be set to 1 to use the FortiGuard Voice service. This service is available only in North America.

To configure PBX settings group, go to *PBX > Calling Rules > Setting*, make configuration changes as required and then select Apply.

Figure 301: Configuring PBX settings

PBX Settings

PBX Global Setting

Extension Pattern

☐ Two digits

☐ Three digits

☐ Four digits

☒ Others

6XXX

Country/Area

United States

Country Code

1

Local Area Code

408

VoiceMail Access

*97

Outgoing Prefix

9

FAX Admin Email

Max VoiceMail Duration

No Limit

60

(seconds)

Voice Prompt Package

[\[Import\]](#)

Call Parking Code

700

Parking Slots

20

Parking Time

45

(seconds)

Call Parking DTMF Command

#72

Apply

PBX Global Settings	
Extension Pattern	<p>Select two, three, or four digit extensions, or choose <i>Others</i> to specify an extension pattern.</p> <p>An extension pattern of two, three, or four digits allows you to choose any extension number of the selected length.</p> <p>Other allows you to enter a pattern that defines the valid extensions that can be added to the FortiGate Voice configuration. The pattern can include numbers that must be in every extension and upper case Xs to indicate the number of digits. The extension range can only contain numbers and the letter X.</p> <ul style="list-style-type: none">If you add numbers to the extension range, all extensions added to this FortiGate Voice unit must include the same numbers in the same location in the extension number. For example, if you include a 6 as the first digit, all extensions added this FortiGate Voice unit must begin with the number 6.The Xs indicate the number of digits in addition to the required number that each extension must have. For example, 6XXX indicates the extensions must start with the number 6 and be followed by any three numbers. <p>Usually you would add one or two numbers to the start of the extension range to identify the extensions for this PBX and follow this with enough Xs to be able to add the required number of extensions.</p> <p>The extension range should not begin with the same number as the outgoing prefix.</p>
Country/Area	<p>Select the country or region in which the FortiGate Voice unit is installed.</p>

Country Code	Enter the international country calling code for the country or region in which you are installing the FortiGate Voice unit.
Local Area Code	Enter the local area code for the location in which you are installing the FortiGate Voice unit.
Voicemail Access	Enter the exact pattern that PBX users dial to get their voicemail. For example, for users to dial *99 to get their voice mail, enter *99.
Outgoing Prefix	The number that PBX users must dial to get an outside line. For example, if users should dial 9 to get an outside line, add 9 to this field. The outgoing prefix should not be the same as the first number of the extension range.
FAX Admin Email	Enter the email address of the fax administrator.
Max Voicemail Duration	Select <i>No Limit</i> if you don't want to limit the voice mail duration. Otherwise enter a maximum time in seconds for voice mail recordings.
Voice Prompt Package	Select <i>Import</i> to import a new voice prompt file. The new file replaces the current voice prompt. See “Importing a new voice prompt file” on page 2643 .
Call Parking Code	Enter the extension number used as the transfer destination for parking calls.
Parking Slots	Enter the number of parking slots.
Parking Time	Enter the maximum parking time, in seconds. If a call is left parked for the parking time duration, the call will ring to the originally dialed extension.
Call Parking DTMF Command	Enter the command used to park a call.

Importing a new voice prompt file

You can replace the default voice prompt by importing a new voice prompt file. You can create your own voice prompt file or you can obtain one in a different language from Fortinet. The new file overwrites the current voice prompt when imported.

To import a new pbx voice prompt file go to *PBX > Calling Rules > Setting* then select *Import*. The voice prompt file should be added to a tar file and zipped. This file would usually have the extension tgz.

Parking calls

Parking a call is similar to putting the caller on hold, except that the call can then be picked up from any extension.

For example, if an urgent call came in to an office for the manager, the receptionist may call the manager's office to make sure he is present to take the call rather than blindly transferring the call and possibly sending the caller to voicemail. If the manager is not in his office, the receptionist can park the call and use the PA system to inform the manager of the call and the number to dial to receive it. He can then use the nearest extension to dial the number and take the call. Further, if the call is not taken within 45 seconds, it will ring back to the extension from which it was parked.

Without call parking, the receptionist must put the caller on hold and then determine where the manager is, direct him to a nearby extension, and finally forward the call to the extension.

To configure call parking, go to *PBX > Calling Rules > Setting*.

FortiFAX service

The FortiFAX service is available as part of the FortiGuard Voice service for sending and receiving faxes. If your FortiGuard Voice service includes FortiFAX, the FortiGate Voice unit stores and then forwards faxes sent to and received from the FortiGuard Voice service. You can go to *PBX > FortiFAX > Received Fax* to view faxes received from the FortiGuard Voice service. You can go to *PBX FortiFAX > Sent Fax* to view faxes sent from a FAX device connected to the FortiGate Voice unit to the FortiGuard Voice service. For all faxes sent and received you can view information about the sender and receiver of the fax, the date and time the fax was received and the status of the fax. You can also download or delete any listed fax.

Incoming faxes are received and stored by the FortiGate Voice unit. PBX users can view their faxes from the FortiGate Voice user portal. See [“Using the PBX user web portal” on page 2649](#). The FortiGate Voice unit sends email notifications to the PBX user’s email address when a fax is received for them.

Users upload outgoing faxes to the FortiGate Voice unit using the FortiGate Voice user portal. The FortiGate Voice unit sends the faxes to their destination and records the result of sending the fax. If the fax cannot be sent right away the FortiGate Voice unit continues polling and will send the fax when possible. When the fax is sent the FortiGate Voice unit sends an email to the PBX user’s email address.

Monitoring calls

You can monitor incoming and outgoing calls from *PBX > Monitor > Active Call*.

You can view information for all active calls including the originator of the call (From) the destination of the call (To), how long the call has been active (Duration), the codec used for transmitting voice packets, and the status of the call.

Monitoring recorded conference calls

You can view a list of recorded conference calls from *PBX > Monitor > Recorded Conference*.

You can view information for all recorded conference calls including the host of the call, the recording time of the conference and the size of the recorded file. You can also download recordings and delete them from the FortiGate Voice unit.

All recorded conference calls are saved on the FortiGate Voice hard disk. Recordings are also available on the user portal.

Monitoring voice mail storage

You can view information about each PBX user’s stored voicemail messages from *PBX > Monitor > Voice Mail Storage*.

The information displayed includes each extension’s voicemail status and amount of disk space used for recorded voice messages. You can delete voice mail for any or all extensions to recover disk space.

All recorded voice messages are saved on the FortiGate Voice hard disk. Recordings are also available on the user portal.

Monitoring active phones

You can view information about each active phone from *PBX > Monitor > Phone*.

The information displayed includes each phone's status, IP address, MAC address, and VCI.

Logging of PBX activities

After configuring PBX settings, you can configure logging of PBX activities and events. In addition to configuring required FortiGate logging settings you can also configure logging of PBX events.

To configure logging of PBX settings

- 1 Go to *Log&Report > Log Config > Log Setting*.
- 2 Select the check box beside *Enable* to make the other event log options available.
- 3 Select the check box beside *PBX event*.

Viewing log messages

You can view the PBX activities and events from *Log&Report > Log & Archive Access*. The log messages can be filtered so that you are viewing specific information, or you can display them in Raw format. Raw format is the format of what a log message actually appears in the log file.

Example PBX log messages

The following log message indicates that the phone with FortiGate Voice extension number 6005 called 914036085000 and the call was routed to the skype-088adb08 service provider. The call was answered and lasted for 1869 seconds.

```
2010-03-12 12:53:27 log_id=0162043782 type=event subtype=pbx pri=information
fwver=040000 vd=root action=PBX-call clid="6005", src="6005"
dst="914036085000" channel="SIP/6005-088a7c08" dstchannel="SIP/skype-088adb08"
duration=1869 start="Fri Mar 12 12:22:18 2010 " end="Fri Mar 12 12:53:27 2010 "
disposition="ANSWERED" msg="call from 6005=>914036085000, ANSWERED, for 1869
seconds"
```

The following log message indicates that the phone with FortiGate Voice extension number 6012 with caller ID Example Caller called extension 6036. And that the call was answered and lasted for 23 seconds.

```
2010-03-12 01:12:42 log_id=0162043782 type=event subtype=pbx pri=information
fwver=040000 vd=root action=PBX-call clid="Example Caller" <6012>", src="6012"
dst="6036" channel="SIP/6012-084a9aa0" dstchannel="SIP/6036-08464150"
duration=23 start="Fri Mar 12 01:12:19 2010 " end="Fri Mar 12 01:12:42 2010 "
disposition="ANSWERED" msg="call from 6012=>6036, ANSWERED, for 23 seconds"
```

VoIP interface reference

The unit can effectively secure VoIP solutions since it supports VoIP protocols and associates state at the signaling layer with packet flows at the media layer. By using SIP ALG controls, the unit can interpret the VoIP signaling protocols used in the network and dynamically open and close ports (pinholes) for each specific VoIP call to maintain security.

In *UTM Profiles > VoIP > Profile*, you can configure multiple profiles for applying to firewall policies that concern only VoIP protocols.

Profile

The Profile menu allows you to configure VoIP profiles for applying to firewall policies. A profile is specific information that defines how the traffic within a policy is examined and what action may be taken based on the examination.

VoIP profile configuration settings

The following are VoIP profile configuration settings in *UTM Profiles > VoIP > Profile*. If the VoIP option does not appear, use this CLI command to enable it.

```
config system global
    set gui-voip-profile enable
end
```


Profile page Lists the profiles that you created for SIP and SCCP protocols. On this page, you can edit, delete or create a new profile for VoIP protocols. You are redirected to this page when you select <i>View List</i> on the Edit VoIP Profile page.	
Create New	Creates a new VoIP profile. When you select <i>Create New</i> , you are automatically redirected to the New VoIP Profile page.
Edit	Modifies settings within a VoIP profile. When you select <i>Edit</i> , you are automatically redirected to the Edit VoIP Profile page.
Delete	Removes a VoIP profile from the list on the Profile page. To remove multiple VoIP profiles from within the list, on the Profile page, in each of the rows of the profiles you want removed, select the check box and then select <i>Delete</i> . To remove all VoIP profiles from the list, on the Profile page, select the check box in the check box column and then select <i>Delete</i> .
Name	The name of the profile.
Comments	A description about the profile. This is an optional setting.
Ref.	Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a firewall policy; on the Profile page (<i>UTM Profiles > Antivirus > Profile</i>), 1 appears in <i>Ref</i> . To view the location of the referenced object, select the number in <i>Ref.</i> , and the Object Usage window appears displaying the various locations of the referenced object. To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window: <ul style="list-style-type: none"> View the list page for these objects – automatically redirects you to the list page where the object is referenced at. Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy and so, when this icon is selected, the user is redirected to the Edit Policy page. View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a firewall policy, and that firewall policy's settings appear within the table.

New VoIP Profile page

Provides settings for configuring SIP and SCCP options within the profile.

This page appears when you select *Create New* on the Edit VoIP Profile page. If you are on the Profile page, and you select *Create New*, you will be redirected to the New VoIP Profile page.

Name	Enter a name for the profile.
Comments	Enter a description about the profile. This is optional.
SIP	Configuration settings for SIP protocols.
Limit REGISTER requests	Enter a number for limiting the time it takes to register requests.
Limit INVITE requests	Enter a number to limit invitation requests.
SCCP	Configuration settings for SCCP protocols.
Limit Call Setup	Enter a number to limit call setup time.



Using the PBX user web portal

This section describes how to log into and use the FortiGate Voice PBX portal. FortiGate Voice PBX users can use the PBX web user portal to configure some of their extension's PBX settings, retrieve their voicemail, configure call forwarding for their extension, and review conference calls and conference call recordings.

Logging into and out of the FortiGate Voice PBX user web portal

FortiGate Voice PBX users use their extension number and password to log in to the PBX user web portal.

To log into the PBX user web portal

- 1 Open any web browser and browse to the following address:

`https://<interface_address>:8443`

where:

`<interface_address>` is the IP address or domain name of a FortiGate Voice interface from which users can access the FortiGate Voice PBX user web portal.

8443 is the default port that users must browse to connect to the portal. FortiGate Voice system administrators can change this port number.

- 2 On the login form enter your extension number and extension password as provided by your system administrator and select Login.

If you successfully log in the user portal web-based manager is displayed.

To log out of the PBX user web portal

- 1 Select the Logout icon that appears at the top of every PBX user web portal page.

Configuring PBX extension settings

Go to *Configuration > Setting* to view your extension settings and to change your PBX user web portal password, and voicemail password and other PBX settings.

Extension	The extension number used to log into the portal (view only).
First Name Last Name	The first and last names assigned to this extension number (view only).
Email	The email address assigned to this extension number (view only).

Password	<p>The password used to log into the portal. You can change the password by entering a new password into this field. When you start typing the new password a <i>Repeat Password</i> field appears. Re-type the new password and select OK to change the password. The next time you log into the portal you must use this password.</p> <p>The password cannot be blank, must be 8 or more characters, must contain at least one uppercase character, one lowercase character and one number. Non-alphanumeric characters, like (- \$, are not supported in the password field.</p>
Voicemail	Indicates whether or not this extension includes voicemail (ready only).
Voicemail password	<p>The password used to retrieve your voicemail using your SIP phone. This password can only contain numbers. You can change the password by entering a new password into this field. When you start typing the new password a <i>Repeat Password</i> field appears. Re-type the new password and select OK to change the password. The next time you retrieve voicemail into the portal you must use this password.</p> <p>The password cannot be blank. Using a secure password is a best practice.</p>
Email Notification	Select to receive email notifications from the FortiGate Voice unit when a caller leaves you a voicemail.
Voicemail to Email Attachment	Select to have the FortiGate Voice unit attach the voicemail messages to the notification email as an audio file. When you receive an email notification of a voicemail you can listen to the voicemail by playing the audio file.
Maximum Message#	Shows the number of saved voicemail messages that can be saved for this extension.

Voicemail

You can view lists of your extension's voicemail messages. Any message can be downloaded to be listened to or deleted. You can also save voicemail messages. Saved messages appear on the saved message list. If you have listened to voicemail messages from your SIP phone but not saved them they are added to the Old Message list.

To download or delete a message, select the message and select Download or Delete.

Configuring call forwarding

Use call forwarding options to configure forward your calls. The following options are available. Select an option and select OK to save the change.

None	Disable or turn off call forwarding.
Forward Calls to VoiceMail Box	Send all calls directly to your voicemail. The caller hears your welcome message and can leave a message for your voicemail. The extension does not ring.

Forward All Calls	Forward all calls to a different phone number. This can be a different extension or an external phone number (for example, your cell phone number). In the Forward Number field, enter the number that the call is forwarded to. Include the area code, country code etc. if required.
Forward Unanswer Calls	Forward calls to a different phone number if you don't answer the phone. The call is forwarded to the external number instead of to voice mail after the same number of rings. This can be a different extension or an external phone number (for example, your cell phone number). In the Forward Number field, enter the number that the call is forwarded to. Include the area code, country code etc. if required.

Sending a Fax using FortiFAX

To send a fax, create a pdf file containing the pages that you want to send. Use the following instructions to upload the pdf to the FortiGate Voice PBX and send it as a fax.



FortiFAX is only available via active FortiGuard Voice Service subscriptions that include FortiFAX service.

To send a fax, go to *FortiFAX > Fax Document Upload*. Enter the fax number that you want to send a fax message to. Select *Browse* to upload the pdf file that you want to send. You can also type in the file location manually into the File to Fax field. Select *OK* to send the fax. A message appears to confirm that the fax has been sent.

The FortiGate Voice PBX stores a copy of the pdf file, which is only viewable by a system administrator.

Conference calls

If the PBX administrator has made your extension a conference bridge host you can view the conference bridges that you can host, change their configuration, and see lists of the participants in an active conference call using your conference bridges. You can also manage conference call recordings for your conference bridges.

To change your conference bridge settings, log into the PBX user web portal and go to *Conference > Conference List* and edit the conference bridge.

Conference ID	The number that participants dial to call into a conference using this conference bridge (view only).
Password	Change the password for this conference bridge. The password cannot be blank and must contain only numbers. This becomes the conference PIN number.
Host	The extension number of the conference bridge host, which is your extension number (view only).
Host Pin	Change the host pin for this conference bridge. The pin only contain numbers.

Video	Select if the conference bridge can display video and handle video calls.
Recordable	Select to enable recording of conference calls that use this conference bridge. You cannot select this option if the conference bridge supports video calls.

Conference call recordings are saved on the FortiGate Voice hard disk. You can go to *Conference > Recorded Conference* to download and delete conference call recordings.

Managing conference calls

The conference bridge host must dial into the conference bridge and enter the host pin for the conference call to start. During the call the host can also do the following from the phone used to start the conference call:

- Press * to hear a conference bridge menu
- Press 8 to start recording the conference call
- Press 9 to stop recording the conference call

Conference call recordings are saved on the FortiGate Voice hard disk. After the recording has been stopped and the call ended the host can log into the FortiGate Voice PBX user web portal and download or delete conference call recordings.



FortiGate Voice VoIP, PBX, and PSTN CLI Reference

This section describes FortiGate Voice VoIP, PBX, and PSTN configuration settings. PSTN interfaces are not available on all FortiGate Voice models. For information about other FortiGate Voice CLI commands see the [FortiGate CLI Reference](#).

This section describes:

- [config pbx dialplan](#)
- [config pbx did](#)
- [config pbx extension](#)
- [config pbx global](#)
- [config pbx ringgrp](#)
- [config pbx voice-menu](#)
- [config pbx sip-trunk](#)
- [config system pstn](#)
- [config system interface](#)
- [execute pbx](#)
- [get pbx ftgd-voice-pkg](#)
- [get pbx global](#)
- [get pbx voice-menu](#)
- [diagnose pbx restart](#)

config pbx dialplan

Use this command to add a dial plan and add rules to the dial plan. A dial plan rule indicates an outgoing destination to send calls to. You can add multiple rules to a dial plan. You add dial plans to extensions to control how to handle outgoing calls from the extension.

Syntax

```
config pbx dialplan
  edit <pbx_dialplan_name>
    set comments <comment_string>
    config rule
      edit <rule_name_str>
        set action {allow | block}
        set callthrough {fx01 | fx02 | fx03 | fx04 |
          <voip_providers>}
        set outgoing-prefix <pattern_str>
        set phone-no-beginwith <patern_str>
        set prepend <pattern_str>
```

```

        set use-global-outgoing-prefix {no | yes}
    end
end

```

Variables	Description	Default
edit <pbx_dialplan_name>	Enter the name for the dial plan. If you entering an existing dial plan, select Tab to get to the dial plan that you want to edit.	No default
comments <comment_string>	Optionally enter a description of the dial plan.	No default
config rule	Configure a new dial plan rule.	No default
edit <rule_name_str>	Enter the name of the dial plan rule to configure.	No default
action {allow block}	Set the action to <code>allow</code> if this dial plan rule should allow a call. Set the action to <code>block</code> if the dial plan should block a call. For example, if you want to block international calls you could set the Phone Number begin with to 011 and set the action to block.	No default
callthrough {fxo1 fxo2 fxo3 fxo4 <voip_providers>}	Select one or more destinations that the dial plan rule sends outgoing calls to. <code>fxo1</code> , <code>fxo2</code> , <code>fxo3</code> , and <code>fxo4</code> are the 4 PSTN interfaces. <code><voip_providers></code> are the VoIP providers added to the FortiGate Voice. A dial plan rule can send calls to one or more destinations.	No default
outgoing-prefix <pattern_str>	If you set <code>use-global-outgoing-prefix</code> to <code>no</code> you can enter a different outgoing prefix for this dial plan.	null
phone-no-beginwith <patern_str>	Enter the leading digits of the phone number that this dial plan rule should match with. For example, a dial plan rule for toll free numbers in North America should begin with 18. The FortiGate Voice uses a best match to match a dialed number with a dial plan. So each dial plan should have a different Phone number Begin with setting. But you should plan your dial plan to make sure that unexpected matches do not occur.	null
prepend <pattern_str>	Add digits that should be prepended or added to the beginning of the dialed number before the call is forwarded to its destination. You can prepend digits at the beginning of a call of special dialing is required to reach and external phone system.	null
use-global-outgoing-prefix {no yes}	Select <code>yes</code> if the dial plan rule should use the default outgoing prefix (usually 9). Select <code>no</code> to add a different <code>outgoing-prefix</code> .	yes

config pbx did

Use this command to configure Direct Inward Dialing (DID). DID allows calls from external phone systems to dial directly to extensions added to the FortiGate Voice unit.

Syntax

```
config pbx did
  edit <pbx_did_name>
    set external-line {fxo1 | fxo2 | fxo3 | fx04 |
      <voip_providers>}
    set cid-number <phone_number>
    set extension <extension_number>
    set comment <comment_string>
  end
```

Variables	Description	Default
edit <pbx_did_name>	Enter the name for the Direct Inward Dial.	No default
external-line {fxo1 fxo2 fxo3 fx04 <voip_providers>}	Select one external system that can dial directly to an extension. fxo1, fxo2, fxo3, and fx04 are the 4 PSTN interfaces. <voip_providers> are the VoIP providers added to the FortiGate Voice.	null
cid-number <phone_number>	Enter the phone number dialed by a caller on the external system.	null
extension <extension_number>	Enter the FortiGate Voice extension number the call is directed to.	null
comment <comment_string>	Enter a description, if applicable, about the direct inward dial configuration.	null

config pbx extension

Use this command to add SIP phone extensions to the FortiGate Voice unit. You can add new extensions or reconfigure the existing ones. For example, you can label an extension by user name, or you can add an extension and set it as a host for conference calls, or you can get FortiGate Voice unit to send email notifications to the users when they receive new voicemail messages.



FortiGate Voice uses the alertmail settings to access an SMTP server and send email notifications. Alertmail can be configured through `config system alertmail` command. For more information about alertmail CLI command configuration refer to FortiGate CLI Reference.

Syntax

```
config pbx extension
  edit <extension_number>
    set attach {enable | disable}
    set auto-delete {enable | disable}
    set conference-host <extension_number>
    set dialplan <dialplan_name>
```

```

set email <user_email>
set email-notify {enable | disable}
set first-name <first_name>
set host-pin <host_password>
set last-name <surname_name>
set macaddress <mac_address>
set max-msg <max_messages_allowed>
set nat {no | yes}
set recordable-flag {enable | disable}
set secret <user_password>
set type {conference | sip-phone}
set video {enable | disable}
set vm-secret <user_password>
set voicemail {enable | disable}
end

```

Variables	Description	Default
edit <extension_number>	Enter the extension number. The extension number has to match the config pbx global extension pattern.	No default
attach {enable disable}	Enable the voicemail message as an attachment in an email.	disable
auto-delete {enable disable}	Enable to automatically delete voice mail.	disable
conference-host <extension_number>	Enter the extension number that will host the conference.	null
dialplan <dialplan_name>	Enter the dial plan that you want to use for the extension.	null
email <user_email>	Enter the user's email address. This email address can be used to notify the user when they have a new voicemail message.	null
email-notify {enable disable}	Enable email notification. When email notification is enabled the user gets notified of each new voicemail messages.	disable
first-name <first_name>	Enter the person's first name.	null
host-pin <host_password>	Enter the password for the conference call. The password must contain only numbers. The users need to enter this password to join the conference call.	
last-name <surname_name>	Enter the surname of the person.	null
macaddress <mac_address>	Enter the MAC address of the SIP phone for the current extension. A typical MAC address consists of six double digit alpha-numeric characters separated by colons. Colons must be used when entering the MAC address.	00:00:00:00:00:00

Variables	Description	Default
max-msg <max_messages_allowed>	Enter the maximum number of voicemail messages that are allowed in a user's voicemail inbox.	50
nat {no yes}	Enter to indicate that the phone is behind a NAT device.	no
recordable-flag {enable disable}	Enable conference recording. When enabled the conference call are recorded on FortiGate Voice unit's hard drive.	disable
secret <user_password>	Enter the user's password for voicemail.	No default
type {conference sip-phone}	Enter the type of extension to configure. <ul style="list-style-type: none"> sip-phone to configure a SIP phone extension conference to add a conference bridge. Multiple users can call the conference bridge extension number enter the <code>secret</code> and have a conference call. A <code>conference</code> bridge only requires an extension number and a <code>secret</code>. 	sip-phone
video {enable disable}	Enable video conferencing.	disable
vm-secret <user_password>	Enter the user's password for accessing their voicemail inbox.	No default
voicemail {enable disable}	Enable the extension to have voicemail.	enable

config pbx global

Use this command to configure voicemail settings such as limiting the length of voicemail messages, as well as the country and the extension pattern of the user.

Syntax

```
config pbx global
  set atxfer-dtmf <str>
  set blindxfer-dtmf <str>
  set block-blacklist {enable | disable}
  set code-callpark <str>
  set country-area <country_name>
  set country-code <country_code>
  set dtmf-callpark <str>
  set efax-check-interval <integer>
  set extension-pattern <extension_pattern>
  set fax-admin-email <email_address>
  set ftgd-voice-server <server_address>
  set local-area-code <code_string>
  set max-voicemail <max_length_seconds>
  set outgoing-prefix <pattern_str>
  set parking-slots <int>
  set parking-time <int>
  set ring-timeout <time_int>
  set rtp-hold-timeout <time_int>
```

```

set rtp-timeout <time_int>
set voicemail-extension <access_number>
end

```

Variables	Description	Default
atxfer-dtmf <str>	The DTMF command to trigger an attended transfer.	*2
blindxfer-dtmf <str>	The DTMF command to trigger a blind transfer.	#1
block-blacklist {enable disable}	Enable to block blacklist IP addresses.	enable
code-callpark <str>	Enter this numeric code to park the current call.	700
country-area <country_name>	Enter the name of the country in which the FortiGate Voice unit is installed.	USA
country-code <country_code>	Enter the country code in which the FortiGate Voice unit is installed.	1
dtmf-callpark <str>	The DTMF command to trigger a call park.	#72
efax-check-interval <integer>	Enter the efax polling interval from FortiGuard fax server. The value range is 5 to 120 in minutes.	5
extension-pattern <extension_pattern>	<p>Enter a pattern that defines the valid extensions that can be added to the FortiGate Voice configuration. The pattern can include numbers that must be in every extension and upper case xs to indicate the number of digits. The extension range can only contain numbers and the letter x.</p> <ul style="list-style-type: none"> If you add numbers to the extension range, all extensions added to this FortiGate Voice unit must include the same numbers in the same location in the extension number. For example, if you include a 6 as the first digit, all extensions added this FortiGate Voice unit must begin with the number 6. The xs indicate the number of digits in addition to the required number that each extension must have. For example, 6xxx indicates the extensions must start with the number 6 and be followed by any three numbers. <p>Usually you would add one or two numbers to the start of the extension range to identify the extensions for this PBX and follow this with enough Xs to be able to add the required number of extensions.</p> <p>The extension range should not begin with the same number as the outgoing prefix.</p>	null

Variables	Description	Default
fax-admin-email <email_address>	Enter the email address of the fax administrator.	null
ftgd-voice-server <server_address>	Enter the FortiGuard voice server address. Default: service.fortivoice.com	
local-area-code <code_string>	Enter the local area code for the country or region in which you are installing the FortiGate Voice unit.	408
max-voicemail <max_length_seconds>	Limit the length of voicemail messages in seconds. Set to 0 for no limit.	60
outgoing-prefix <pattern_str>	The number that PBX users must dial to get an outside line. For example, if users should dial 9 to get an outside line, add 9 to this field. The outgoing prefix should not be the same as the first number of the extension range.	9
parking-slots <int>	The maximum number of calls that can be parked at the same time.	20
parking-time <int>	The length of time, in seconds, a call can be parked. If this time expires without the call being answered, the parked call will ring back to the extension from which it was parked.	45
ring-timeout <time_int>	The number of seconds that an extension should be allowed to ring before going to voicemail.	20
rtp-hold-timeout <time_int>	The amount of time in seconds that the extension will wait on hold for RTP packets before hanging up the call. 0 means no time limit.	0
rtp-timeout <time_int>	The amount of time in seconds during an active call that the extension will wait for RTP packets before hanging up the call. 0 means no time limit.	60
voicemail-extension <access_number>	Enter the voicemail extension number that a user will use to access their voicemail inbox.	*97

config pbx ringgrp

Use this command to add and configure the extension groups. An extension group here is referred to a ring group and is a group of extensions that can be called using one number. You can configure the ring group to call all of the extensions in the group at the same time or to call the extensions one at a time until someone answers.



The order in which the members are added to the ring group does not match the order in which the FortiGate Voice unit calls them.

Syntax

```
config pbx ringgrp
  edit <ring_group_name>
    set description <description_str>
    set member <acd_group_member>
    set no-answer-action {hangup | ivr | voicemail}
    set strategy {ring-all | sequential}
    set voicemail-of-extension <extension_number>
  end
```

Variables	Description	Default
edit <ring_group_name>	Enter the name for the group.	No default.
description <description_str>	A description of the extension group.	null
member <acd_group_member>	Enter the ACD member for the group.	No default
no-answer-action {hangup ivr voicemail}	Enter the action that will be taken when none of the extensions in the ring group answers. <ul style="list-style-type: none"> hangup hand up and end the call. ivr return the caller to the attendant where they can try another extension. voicemail the caller is directed to the voicemail system where they can leave a message. 	voicemail
strategy {ring-all sequential}	Control how the extensions in the group are called by the ring group. <ul style="list-style-type: none"> ring-all calls all of the extensions in the group at the same time. sequential calls the extensions in the group one at a time in the order in which they have been added to the group. 	sequential
voicemail-of-extension <extension_number>	Enter the extension number to use for voicemail if no one answers the call and no-answer-action is set to voicemail.	null

config pbx voice-menu

Use this command to configure the menu that callers will access when they call. The variable `config press-<number>` configures the settings for the type of ring group and the type of group associated with that number.

Syntax

```
config pbx voice-menu
  set comment <comment_string>
  set password <ext_password>
  set recorder-exten <extension_str>
  config [press-0 | press-1 | press-2 | press-3 | press-4 |
    press-5 | press-6 | press-7 | press-8 | press-9]
    set type {directory | none | ring-group | voicemail}
```

```

        set ring-group <group_string>
    end
end

```

Variables	Description	Default
comment <comment_string>	Enter a description of the voice-menu settings, if applicable.	No default
password <ext_password>	Enter the password to access recording a new IVR message.	null
recorder-exten <extension_str>	Enter the extension number for recording a new IVR message.	*30
config [press-0 press-1 press-2 press-3 press-4 press-5 press-6 press-7 press-8 press-9]	Use this command when configuring what action each number on the phone's keypad will take. For example, you want the personnel directory to come up every time someone presses 1; config press-1 variable would have the type directory selected in type.	No default
type {directory none ring-group voicemail}	Enter the type of action that is associated with the specific number on the phone's keypad. For example, the office phone directory is heard when a caller presses 0 because config press-0 has directory as its type.	No default
ring-group <group_string>	Enter to include a specific ring-group if you have select ring-group in type. This variable appears only when ring-group is selected in type.	null

config pbx sip-trunk

Use this command to configure SIP server providers for the PBX. If your FortiGate Voice unit is installed in North America and the Country Code is set to 1 then you can use the FortiGuard Voice service as your SIP service provider. (The default Country Code is 1, see [“config pbx global” on page 2657](#) for changing county code.) The FortiGuard Voice service is supported only in North America. If you install the FortiGate Voice unit elsewhere in the world and change the Country Code, the FortiGuard Voice Service configuration is replaced by the SIP trunk configuration. You can use the SIP trunk configuration to add one or more SIP service providers to the FortiGate Voice configuration.

Syntax

```

config pbx voip-provider
    edit <provider_name>
        set user <user_name>
        set domain {<VoIP_provider_address_ipv4> |
                    <VoIP_provider_domain> }
        set secret <password>
        set authuser <authuser>
        set display-name <display_name>
        set reigstration-interval <refresh_interval>
    end
end

```

```

set account-type {static | dynamic}
set dtmf-metod {auto | inband | info | rfc2833}
set codec {alaw | g729 | none | ulaw}
set codec1 {alaw | g729 | none | ulaw}
set codec2 {alaw | g729 | none | ulaw}
set video {enable | disable}
end

```

Variables	Description	Default
edit <provider_name>	Enter the VoIP provider's name.	No default
user <user_name>	Enter the user name for the provider. You can enter the phone number registered with this provider instead.	No default
secret <password>	Enter the password associated with the provider.	No default
domain {<VoIP_provider_addresses_ipv4> <VoIP_provider_domain> }	The VoIP provider's domain name or IP address. For example, 172.20.120.11 or voip.example.com.	No default
authuser <authuser>	Enter the authentication user for the account.	No default
display-name <display_name>	Enter the name that will be used as the caller ID name if the provider supports this feature.	No default
reigstration-interval <refresh_interval>	Enter a number for the refresh interval.	No default
account-type {static dynamic}	Enter to define the type of account.	No default.
dtmf-metod {auto inband info rfc2833}	Enter the DTMF method that will be used.	No default
codec {alaw g729 none ulaw}	Enter the most preferred Codec for the VoIP provider.	ulaw
codec1 {alaw g729 none ulaw}	Enter the second most preferred Codec for the VoIP provider.	none
codec2 {alaw g729 none ulaw}	Enter the third most preferred Codec for the VoIP provider.	none
video {enable disable}	Enable video capability if the provider supports this feature.	disable

config system pstn

Use this command to configure the PSTN interfaces. PSTN interfaces are available on some FortiGate Voice models.

Syntax

```
config system pstn
  edit <fxo_name>
    set cid-name <caller_name>
    set cid-number <caller_name>
    set status {enable | disable}
    set use-callerid {enable | disable}
    set cid-signalling {bell | dtmf | v23 | v23-jp}
    set cid-start {polarity | ring}
    set send-callerid-after <integer>
    set hangup-on-polarity-reversal {enable | disable}
    set hangup-on-zero-voltage {enable | disable}
    set hangup-on-busy-tone {enable | disable}
    set busycount <integer>
    set busy-tone-length <integer>
    set busy-quiet-length <integer>
    set codec {alaw | ulaw}
  end
```

Variables	Description	Default
edit <fxo_name>	Enter the name of the FXO.	No default
cid-name <caller_name>	This name is used for caller ID for calls from the FortiGate Voice unit to the PSTN. It can be any name, such as a company name, that identifies the branch office.	No default
cid-number <caller_name>	Enter the phone number of the PSTN phone line as provided by your phone service provider.	No default
status {enable disable}	Enable the status of the port.	enable
use-callerid {enable disable}	Enable to catch the caller ID.	enable
cid-signalling {bell dtmf v23 v23-jp}	Enter the caller ID protocol. The protocol v23-jp is the v23 protocol for Japan.	bell
cid-start {polarity ring}	Enter to start transmitting the caller ID.	ring
send-callerid-after <integer>	Enter a number for the number of rings after that the caller ID began to transmit.	1
hangup-on-polarity- reversal {enable disable}	Enable to have the phone hang up when there is polarity reversal.	enable

Variables	Description	Default
hangup-on-zero-voltage {enable disable}	Enable to have the phone hang up when there is zero voltage.	disable
hangup-on-busy-tone {enable disable}	Enable to have the phone hang up when a busy tone is detected.	enable
busycount <integer>	Enter a number for the accurate number of busy tones that are detected.	4
busy-tone-length <integer>	Enter a number that determines how long the busy tone is on in milliseconds.	500
busy-quiet-length <integer>	Enter a number that determines how long the busy tone is off in milliseconds.	500
codec {alaw ulaw}	Enter the Codec preference type based on the country.	ulaw
ring detect {ring-cross-threshold ring-full-wave ring-half-wave ring-validate}	Enter the appropriate ring detection method for your phone system.	ring-validate
ring-timeout {128ms 256ms 384ms 512ms 640ms 768ms 896ms 1024ms 1152ms 1280ms 1408ms 1536ms 1664ms 1792ms 1920ms}	Enter the appropriate ring time-out for your phone system.	640ms
ring-threshold {level-1 level-2 level-3}	Enter the appropriate ring threshold for your phone system. The ring-threshold is based on voltage: <ul style="list-style-type: none"> level-1: 13.5V to 16.5V level-2: 19.35V to 23.65V level-3: 40.5V to 49.5V 	level-1
ring-delay-time {256ms 512ms 768ms 1024ms 1280ms 1536ms 1792ms}	Enter the appropriate ring delay time for your phone system.	512ms
ring-confirm-time {100ms 150ms 200ms 256ms 384ms 512ms 640ms 1024ms}	Enter the appropriate ring confirmation time for your phone system.	512ms
ring-max-assertion-count <int>	Enter the appropriate ring maximum assertion count for your phone system.	22
ring-assertion-time <int>	Enter the appropriate ring assertion time for your phone system.	25

Variables	Description	Default
tx-gain <int>	Enter the gain for the transmitted signal, in dB, from -15 to 12.	0
rx-gain <int>	Enter the gain for the received signal, in dB, from -15 to 12.	0

config system interface

Use this command to allow traffic for the VoIP protocol, SIP, to flow on a specific interface. You can also allow users to access PBX user portal and enable auto-provisioning for SIP phone configuration on the same interface.

Syntax

```
config system interface
edit <interface_name>
set pbx-user-portal {enable | disable}
set phone-auto-provision {enable | disable}
set voip {enable | disable}
end
```

Variables	Description	Default
edit <interface_name>	Enter the interface that you want to allow SIP traffic on.	No default
pbx-user-portal {enable disable}	Enable PBX user portal on the interface.	disable
phone-auto-provision {enable disable}	Enable SIP phone auto-provisioning on the interface.	disable
voip {enable disable}	Enable the VoIP SIP protocol for allowing SIP traffic on the interface.	disable

execute pbx

Use this command to view active channels and to delete, list or upload music files for when music is playing while a caller is on hold.

Syntax

```
execute pbx active-call <list>
execute pbx extension <list>
execute pbx ftgd-voice-pkg {sip-trunk}
execute pbx music-on-hold {delete | list | upload}
execute pbx prompt upload ftp <file.tgz>
<ftp_server_address>[:port] [<username>] [password]
execute pbx prompt upload tftp <file.tgz>
<ftp_server_address>[:port] [<username>] [password]
execute pbx prompt upload usb <file.tgz>
<ftp_server_address>[:port] [<username>] [password]
execute pbx restore-default-prompts
```

```
execute pbx sip-trunk list
```

Variables	Description
active-call <list>	Enter to display a list of the active calls being processed by the FortiGate Voice unit.
extension <list>	Enter to display the status of all extensions with SIP phones that have connected to the FortiGate Voice unit.
ftgd-voice-pkg {sip-trunk}	Enter to retrieve FortiGuard voice package sip trunk information.
music-on-hold {delete list upload}	Enter to either delete, list or upload music on hold files. You can upload music on hold files using FTP, TFTP, or from a USB drive plugged into the FortiGate Voice unit.
prompt upload ftp <file.tgz> <ftp_server_address> [:port] [<username>] [password]	Upload new pbx voice prompt files using FTP. The voice prompt files should be added to a tar file and zipped. This file would usually have the extension tgz. You must include the filename, FTP server address (domain name of IPv4 address) and if required the username and password for the server.
prompt upload tftp <file.tgz> <ftp_server_address> [:port] [<username>] [password]	Upload new pbx voice prompt files using TFTP. The voice prompt files should be added to a tar file and zipped. This file would usually have the extension tgz. You must include the filename and TFTP server IP address.
prompt upload usb <file.tgz> <ftp_server_address> [:port] [<username>] [password]	Upload new pbx voice prompt files from a USB drive plugged into the FortiGate Voice unit. The voice prompt files should be added to a tar file and zipped. This file would usually have the extension tgz. You must include the filename.
restore-default-prompts	Restore default English voicemail and other PBX system prompts. Use this command if you have changed the default prompts and want to restore the default settings.
sip-trunk list	Enter to display the status of all SIP trunks that have been added to the FortiGate Voice configuration.

Example command output

Enter the following command to view active calls:

```
execute pbx active-call
```

```
Call-From    Call-To    Durationed
6016         6006      00:00:46
```

Enter the following command to display the status of all extensions

```
execute pbx extension list
```

```
Extension    Host          Dialplan
6052         Unregister    company-default
6051         Unregister    company-default
6050         Unregister    company-default
6022         Unregister    company-default
6021/6021    172.30.63.34 company-default
```



```
6020      Unregister      company-default
Enter the following command to display the status of all SIP trunks
execute pbx sip-trunk list
Name      Host      Username      Account-Type      State
Provider_1 192.169.20.1 +5555555      Static           N/A
```

get pbx branch-office

Use this command to list the configured branch offices.

Syntax

```
get pbx branch-office
```

Example output

```
== [ Branch 15 ]
name: Branch 15
== [ Branch 12 ]
name: Branch 12
```

get pbx dialplan

Use this command to list the configured dial plans.

Syntax

```
get pbx dialplan
```

Example output

```
== [ company-default ]
name: company-default
== [ inbound ]
name: inbound
```

get pbx did

Use this command to list the configured direct inward dial (DID) numbers.

Syntax

```
get pbx did
```

Example output

```
== [ Operator ]
name: Operator
== [ Emergency ]
name: Emergency
```

get pbx extension

Use this command to list the configured extensions.

Syntax

```
get pbx extension
```

Example output

```
== [ 6555 ]  
extension: 6555  
== [ 6777 ]  
extension: 6777  
== [ 6111 ]  
extension: 6111
```

get pbx ftgd-voice-pkg

Use this command to display the current FortiGate Voice service package status.

Syntax

```
get pbx ftgd-voice-pkg status
```

Example output

```
Status: Activated  
Total 1 Packages:  
Package Type: B, Credit Left: 50.00, Credit Used: 0.00,  
Expiration Date: 2011-01-01 12:00:00  
  
Total 1 Dids:  
12345678901  
Total 1 Efxas:  
12345678902  
Total 0 Tollfrees:
```

get pbx global

Use this command to display the current global pbx settings.

Syntax

```
get pbx global
```

Example output

```
block-blacklist      : enable  
country-area         : USA  
country-code         : 1  
efax-check-interval  : 5  
extension-pattern    : 6XXX  
fax-admin-email      : faxad@example.com  
ftgd-voice-server    : service.fortivoice.com  
local-area-code      : 408  
max-voicemail        : 60  
outgoing-prefix      : 9  
ring-timeout         : 20  
rtp-hold-timeout     : 0  
rtp-timeout          : 60
```

```
voicemail-extension : *97
```

get pbx ringgrp

Use this command to display the currently configured ring groups.

Syntax

```
get pbx ringgrp
```

Example output

```
== [ 6001 ]  
name: 6001  
== [ 6002 ]  
name: 6002
```

get pbx sip-trunk

Use this command to display the currently configured SIP trunks.

Syntax

```
get pbx sip-trunk
```

Example output

```
== [ __FtgdVoice_1 ]  
name: __FtgdVoice_1
```

get pbx voice-menu

Use this command to display the current voice menu and recorder extension configuration.

Syntax

```
get pbx voice-menu
```

Example output

```
comment           : general  
password          : *  
press-0:  
    ring-group     : 6001  
    type           : ring-group  
press-1:  
    type           : voicemail  
press-2:  
    type           : directory  
press-3:  
    type           : none  
press-4:  
    type           : none  
press-5:  
    type           : none  
press-6:  
    type           : none
```

```
press-7:
  type           : none
press-8:
  type           : none
press-9:
  type           : none
recorder-exten   : *30
```

diagnose pbx restart

Use this diagnose command to restart the FortiGate Voice PBX daemon.

```
diagnose pbx restart
```



Chapter 17 WAN Optimization, Web Cache, Explicit Proxy, and WCCP

The FortiOS Handbook chapter contains the following sections:

WAN optimization, web cache, explicit proxy, and WCCP concepts: Provides an overview of FortiGate WAN optimization best practices and technologies and some of the concepts and rules for using them. We recommend that you begin with this chapter before attempting to configure your FortiGate unit to use WAN optimization.

WAN optimization and Web cache storage: Describes how to configure WAN optimization storage settings to control how data is stored for web caching and byte caching.

WAN optimization peers and authentication groups: Describes how to use WAN optimization peers and authentication groups to control access to WAN optimization tunnels.

Configuring WAN optimization rules: Provides basic configuration for WAN optimization rules, including adding rules, organizing rules in the rule list and using WAN optimization addresses. This chapter also explains how WAN optimization accepts sessions, as well as how and when you can apply UTM features to WAN optimization traffic.

WAN optimization configuration examples: Describes basic active-passive and peer-to-peer WAN optimization configuration examples. This chapter is a good place to start learning how to put an actual WAN optimization network together.

Web caching: Describes how WAN optimization web caching works to cache different session types, including HTTPS, and includes web caching configuration examples.

Advanced configuration example: Provides a configuration example that combines WAN optimization, web caching, out-of-path WAN optimization, and the use of multiple VDOMs to apply UTM features to sessions being optimized.

SSL offloading for WAN optimization and web caching: Describes how to offload SSL processing from web sites to FortiGate units to improve WAN performance for SSL-protected web sites on a WAN.

FortiClient WAN optimization: Describes how FortiGate and FortiClient WAN optimization work together and includes an example configuration.

The FortiGate explicit web proxy: Describes how to configure the FortiGate explicit web proxy, how users connect to the explicit web proxy, and how to add web caching to the explicit web proxy.

The FortiGate explicit FTP proxy: Describes how to configure the FortiGate explicit FTP proxy and how users connect to the explicit FTP proxy.

FortiGate WCCP: Describes FortiGate WCCP and how to configure WCCP and the WCCP client.

WAN optimization, web cache, explicit proxy and WCCP get and diagnose commands: describes get and diagnose commands available for troubleshooting WAN optimization, web cache, and WCCP.



WAN optimization, web cache, explicit proxy, and WCCP concepts

FortiGate WAN optimization consists of a number of techniques that you can apply to improve the efficiency of communication across your WAN. These techniques include protocol optimization, byte caching, web caching, SSL offloading, and secure tunnelling. Protocol optimization can improve the efficiency of traffic that uses the CIFS, FTP, HTTP, or MAPI protocol, as well as general TCP traffic. Byte caching caches files and other data on FortiGate units to reduce the amount of data transmitted across the WAN. Web caching stores web pages on FortiGate units to reduce latency and delays between the WAN and web servers. SSL offloading offloads SSL decryption and encryption from web servers onto FortiGate SSL acceleration hardware. Secure tunnelling secures traffic as it crosses the WAN.

You can apply different combinations of these WAN optimization techniques to a single traffic stream depending on the traffic type. For example, you can apply byte caching and secure tunneling to any TCP traffic. For HTTP and HTTPS traffic, you can also apply protocol optimization and web caching.

You can configure a FortiGate unit to be an explicit web proxy server and an explicit FTP proxy server. Users on your internal network can browse the Internet through the explicit web proxy server or connect to FTP servers through the explicit FTP proxy server. You can also configure these proxies to protect access to web or FTP servers behind the FortiGate unit using a reverse proxy configuration.

FortiGate units that support WAN optimization can also be configured to support web caching. Both WAN optimization and web caching require that the FortiGate unit include a hard disk. Either an internal hard disk or AMC or other hard disk module. Web caching can be applied to any HTTP, this includes HTTP traffic accepted by a security policy, explicit web proxy traffic, and HTTP and HTTPS WAN optimization traffic.

You can also configure a FortiGate unit to operate as a Web Cache Communication Protocol (WCCP) client or server. WCCP provides the ability to offload web caching to one or more redundant web caching servers.

This chapter describes:

- [WAN optimization topologies](#)
- [Explicit Web proxy topologies](#)
- [Explicit FTP proxy topologies](#)
- [Web caching topologies](#)
- [WCCP topologies](#)
- [WAN optimization client/server architecture](#)
- [WAN optimization tunnels](#)
- [Protocol optimization](#)
- [Byte caching](#)
- [WAN optimization and HA](#)

- [WAN optimization, web caching and memory usage](#)
- [Monitoring WAN optimization performance](#)
- [Configuring WAN optimization traffic usage logs](#)
- [WAN optimization best practices](#)

WAN optimization topologies


This section describes some common WAN optimization topologies:

- [“Basic WAN optimization topologies” on page 2674](#)
- [“Out-of-path topology” on page 2675](#)
- [“Topology for multiple networks” on page 2677](#)
- [“WAN optimization with web caching” on page 2678](#)
- [“WAN optimization and web caching with FortiClient peers” on page 2679](#)

Basic WAN optimization topologies

The basic FortiGate WAN optimization topology consists of two FortiGate units operating as WAN optimization peers intercepting and optimizing traffic crossing the WAN between the private networks.

Figure 302: Security device and WAN optimization topology



As shown in [Figure 302](#), the FortiGate units can be deployed as security devices that protect private networks connected to the WAN and also perform WAN optimization. In this configuration, the FortiGate units are configured as typical security devices for the private networks and are also configured for WAN optimization. The WAN optimization configuration intercepts traffic to be optimized as it passes through the FortiGate unit and uses a WAN optimization tunnel with another FortiGate unit to optimize the traffic that crosses the WAN.

As shown in Figure 303, you can also deploy WAN optimization on single-purpose FortiGate units that only perform WAN optimization. In Figure 303, the WAN optimization FortiGate units are located on the WAN outside of the private networks. You can also install the WAN optimization FortiGate units behind the security devices on the private networks.

The WAN optimization configuration is the same for FortiGate units deployed as security devices and for single-purpose WAN optimization FortiGate units. The only differences would result from the different network topologies.

Out-of-path topology

In an out-of-path topology, one or both of the FortiGate units configured for WAN optimization are not directly in the main data path. Instead, the out-of-path FortiGate unit is connected to a device on the data path, and the device is configured to redirect sessions to be optimized to the out-of-path FortiGate unit.

Figure 303: Single-purpose WAN optimization topology





Figure 304 shows out-of-path FortiGate units configured for WAN optimization and connected directly to FortiGate units in the data path. The FortiGate units in the data path use a method such as policy routing to redirect traffic to be optimized to the out-of-path FortiGate units. The out-of-path FortiGate units establish a WAN optimization tunnel between each other and optimize the redirected traffic.

Figure 304: Out-of-path WAN optimization

One of the benefits of out-of-path WAN optimization is that out-of-path FortiGate units only perform WAN optimization and do not have to process other traffic. An in-path FortiGate unit configured for WAN optimization also has to process other non-optimized traffic on the data path.


The out-of-path FortiGate units can operate in NAT/Route or Transparent mode.

Other out-of-path topologies are also possible. For example, you can install the out-of-path FortiGate units on the private networks instead of on the WAN. Also, the out-of-path FortiGate units can have one connection to the network instead of two. In a one-arm configuration such as this, security policies and routing have to be configured to send the WAN optimization tunnel out the same interface as the one that received the traffic.

Topology for multiple networks

As shown in Figure 305, you can create multiple WAN optimization configurations between many private networks. Whenever WAN optimization occurs, it is always between two FortiGate units, but you can configure any FortiGate unit to perform WAN optimization with any of the other FortiGate units that are part of your WAN.

Figure 305: WAN optimization among multiple networks




You can also configure WAN optimization between FortiGate units with different roles on the WAN. FortiGate units configured as security devices and for WAN optimization can perform WAN optimization as if they are single-purpose FortiGate units just configured for WAN optimization.

WAN optimization with web caching

You can add web caching to a WAN optimization topology when users on a private network communicate with web servers located across the WAN on another private network.

Figure 306: WAN optimization with web caching topology




The topology in [Figure 306](#) is the same as that of [Figure 302 on page 2674](#) with the addition of web caching to the FortiGate unit in front of the private network that includes the web servers. In a similar way, you can add web caching to all of the topologies shown in “[WAN optimization topologies](#)” on [page 2674](#).

WAN optimization and web caching with FortiClient peers

FortiClient WAN optimization works with FortiGate WAN optimization to accelerate remote user access to the private networks behind FortiGate units. The FortiClient application requires a simple WAN optimization configuration to automatically detect if WAN optimization is enabled on the FortiGate unit. Once WAN optimization is enabled, the FortiClient application transparently makes use of the WAN optimization and web caching features available.


Figure 307: FortiClient WAN optimization topology



Explicit Web proxy topologies


You can configure a FortiGate unit to be an explicit web proxy server for Internet web browsing. To use the explicit web proxy, users must add the IP address of the FortiGate interface configured for the explicit web proxy to their web browser proxy configuration.

Figure 308: Explicit web proxy topology



If the FortiGate unit supports web caching, you can also add web caching to the security policy that accepts explicit web proxy sessions. The FortiGate unit then caches Internet web pages on a hard disk to improve web browsing performance.


Figure 309: Explicit web proxy with web caching topology



Explicit FTP proxy topologies


You can configure a FortiGate unit to be an explicit FTP proxy server for FTP users. To use the explicit web proxy, FTP users must connect to and authenticate with the explicit FTP proxy before connecting to an FTP server.

Figure 310: Explicit FTP proxy topology



You can also configure reverse explicit FTP proxy (Figure 311). In this configuration, users on the Internet connect to the explicit web proxy before connecting to an FTP server installed behind a FortiGate unit.

Figure 311: Reverse explicit FTP proxy topology




Web caching topologies


FortiGate web caching can be added to any security policy and any HTTP or HTTPS traffic accepted by that security policy can be cached on the FortiGate unit hard disk. You can also add web caching explicit web proxy security policies to cache explicit web proxy traffic. You can also create web caching only WAN optimization rules. The network topologies for all of these scenarios are very similar. They involve a FortiGate unit installed between users and web servers with web caching enabled.

A typical web-caching topology includes one FortiGate unit that acts as a web cache server (Figure 312). Web caching is enabled in a security policy and the FortiGate unit intercepts web page requests accepted by the security policy, requests web pages from the web servers, caches the web page contents, and returns the web page contents to the users. When the FortiGate unit intercepts subsequent requests for cached web pages, the FortiGate unit contacts the destination web server just to check for changes.

Figure 312: Web caching topology




You can also configure reverse proxy web-caching (Figure 313). In this configuration, users on the Internet browse to a web server installed behind a FortiGate unit. The FortiGate unit intercepts the web traffic (HTTP and HTTPS) and caches pages from the web server. Reverse proxy web caching on the FortiGate unit reduces the number of requests that the web server must handle, leaving it free to process new requests that it has not serviced before.

Figure 313: Reverse proxy web caching topology

WCCP topologies

You can operate a FortiGate unit as a Web Cache Communication Protocol (WCCP) router or cache engine. As a router, the FortiGate unit intercepts web browsing requests from client web browsers and forwards them to a WCCP cache engine. The cache engine returns the required cached content to the client web browser. If the cache server does not have the required content it accesses the content, caches it and returns the content to the client web browser.

Figure 314: WCCP topology




FortiGate units can also operate as WCCP cache servers, communicating with WCCP routers, caching web content and providing it to client web browsers as required.

WCCP is transparent to client web browsers. The web browsers do not have to be configured to use a web proxy.

WAN optimization client/server architecture

Traffic across a WAN typically consists of clients on a client network communicating across a WAN with a remote server network. The clients do this by starting communication sessions from the client network to the server network. To optimize these sessions, you add security policies to the client-side FortiGate unit (which is located between the client network and the WAN, see [Figure 315](#)) to accept sessions from the client network that are destined for the server network. To apply WAN optimization to these sessions, you must also add WAN optimization rules to the client-side FortiGate unit. The WAN optimization rules intercept sessions accepted by security policies and apply WAN optimization to them.

Figure 315: Client/server architecture



When a client-side FortiGate unit matches a session with a WAN optimization rule, it uses the information in the rule to attempt to start a WAN optimization tunnel with a server-side FortiGate unit installed in front of the server network. The client-side and server side FortiGate units must be able to identify each other. To do this the client-side FortiGate unit configuration must include the IP address and peer host ID of the server-side FortiGate unit and the configuration of the server-side FortiGate unit must include the IP address and peer host ID of the client-side FortiGate unit. With this information available, when the client-side FortiGate unit attempts to contact the server-side FortiGate unit, the two units share their IP addresses and peer host IDs and confirm that they can create a WAN optimization tunnel between each other.


Security policies are not required for WAN optimization on the server-side FortiGate unit. Sessions from the client-side to the server-side FortiGate unit are WAN optimization tunnel requests. As long as the client-side and server-side FortiGate units can identify each other according to peer host ID and IP address the server-side FortiGate unit will accept WAN optimization tunnel requests from the client-side FortiGate unit.

In addition to basic identification by peer host ID and IP address you can configure authentication options to impose authentication using certificates and pre-shared keys. In addition to you can configure FortiGate units involved in WAN optimization to accept connections from any identified peer or restrict connections to specific peers.

WAN optimization peers

The client-side and server-side FortiGate units are called WAN optimization peers (see [Figure 316](#)) because all of the FortiGate units in a WAN optimization network have the same peer relationship with each other. The client and server roles just relate to how a session is started. Any FortiGate unit configured for WAN optimization can be a client-side and a server-side FortiGate unit at the same time, depending on the direction of the traffic. Client-side FortiGate units initiate WAN optimization sessions and server-side FortiGate units respond to the session requests. Any FortiGate unit can simultaneously be a client-side FortiGate unit for some sessions and a server-side FortiGate unit for others.

Figure 316: WAN optimization peer and tunnel architecture



To identify all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with, you add host IDs and IP addresses of all of the peers to the FortiGate unit configuration. The peer IP address is actually the IP address of the peer unit interface that communicates with the FortiGate unit.

Peer-to-peer and active-passive WAN optimization

You can create peer-to-peer and active-passive WAN optimization configurations. Peer-to-peer configurations are less complex because they only require the creation of a WAN optimization rule in the client side FortiGate unit. Active-passive WAN optimization configurations require an active rule on the client side FortiGate unit and a passive rule on the server-side FortiGate unit. For more details about peer to peer and active-passive WAN optimization, see [“Configuring WAN optimization rules”](#) on page 2705.

WAN optimization and the FortiClient application

PCs running the FortiClient application are client-side peers that initiate WAN optimization tunnels with server-side peer FortiGate units. However, you can have an ever-changing number of FortiClient peers with IP addresses that also change regularly. To avoid maintaining a list of such peers, you can instead configure WAN optimization to accept any peer and use authentication to identify FortiClient peers.

Together, the WAN optimization peers apply the WAN optimization features to optimize the traffic flow over the WAN between the clients and servers. WAN optimization reduces bandwidth requirements, increases throughput, reduces latency, offloads SSL encryption/decryption and improves privacy for traffic on the WAN.

Operating modes and VDOMs

To use WAN optimization, the FortiGate units can operate in either NAT/Route or Transparent mode. The client-side and server-side FortiGate units do not have to be operating in the same mode.

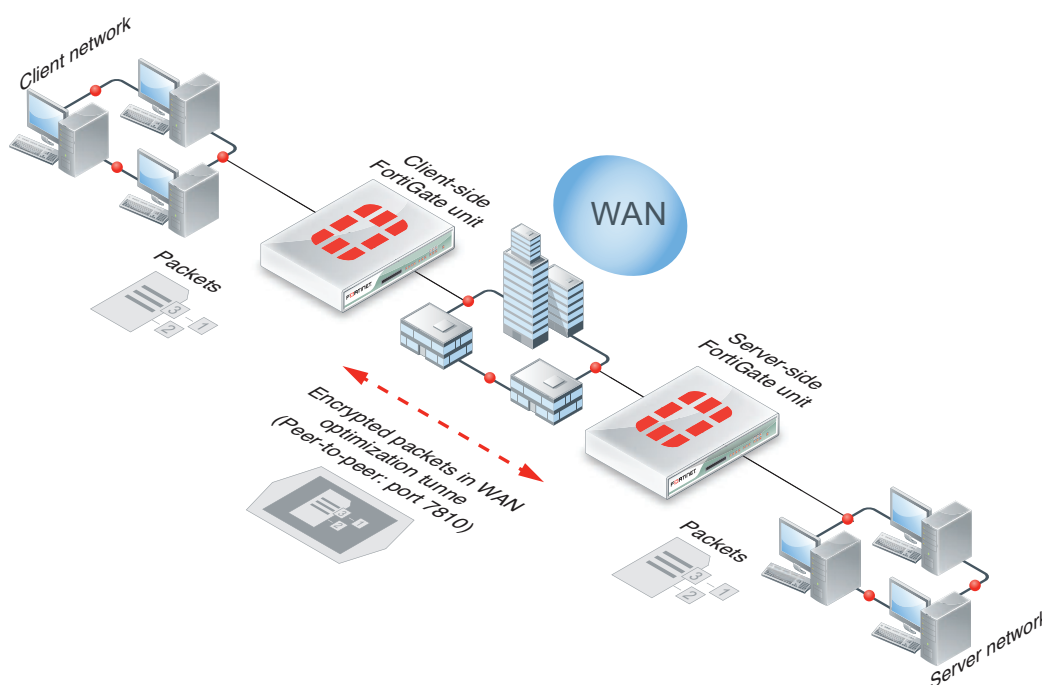
As well, the FortiGate units can be configured for multiple virtual domain (VDOM) operation. You configure WAN optimization for each VDOM and configure one or both of the units to operate with multiple VDOMs enabled.

If a FortiGate unit or VDOM is operating in Transparent mode with WAN optimization enabled, WAN optimization uses the management IP address as the peer IP address of the FortiGate unit instead of the address of an interface.

WAN optimization tunnels

All optimized traffic passes between the FortiGate units or between a FortiClient peer and a FortiGate unit over a WAN optimization tunnel. Traffic in the tunnel can be sent in plain text or encrypted using AES-128bit-CBC SSL.

Figure 317: WAN optimization tunnels



Both plain text and the encrypted peer-to-peer tunnels use TCP destination port 7810.

Before a tunnel can be started, the peers must be configured to authenticate with each other and to agree on the tunnel configuration. Then, the client-side peer attempts to start a WAN optimization tunnel with the server-side peer. Once the peers authenticate with each other, they bring up the tunnel and WAN optimization communication over the tunnel starts. After a tunnel has been established, multiple WAN optimization sessions can start and stop between peers without restarting the tunnel.

Tunnel sharing

You can use the `tunnel-sharing` WAN optimization rule CLI keyword to configure tunnel sharing for WAN optimization rules with `auto-detect` set to `off`. Tunnel sharing means multiple WAN optimization sessions share the same WAN optimization tunnel. Tunnel sharing can improve WAN performance by reducing the number of WAN optimization tunnels between FortiGate units. Having fewer tunnels means less data to manage. Also, tunnel setup requires more than one exchange of information between the ends of the tunnel. Once the tunnel is set up, each new session that shares the tunnel avoids tunnel setup delays.

Tunnel sharing also uses bandwidth more efficiently by reducing the chances that small packets will be sent down the tunnel. Processing small packets reduces network throughput, so reducing the number of small packets improves performance. A shared tunnel can combine all the data from the sessions being processed by the tunnel and send the data together. For example, suppose a FortiGate unit is processing five WAN optimization sessions and each session has 100 bytes to send. If these sessions use a shared tunnel, WAN optimization combines the packets from all five sessions into one 500-byte packet. If each session uses its own private tunnel, five 100-byte packets will be sent instead. Each packet also requires a TCP ACK reply. The combined packet in the shared tunnel requires one TCP ACK packet. The separate packets in the private tunnels require five.

Tunnel sharing is not always recommended and may not always be the best practice. Aggressive and non-aggressive protocols should not share the same tunnel. An aggressive protocol can be defined as a protocol that is able to get more bandwidth than a non-aggressive protocol. (The aggressive protocols can “starve” the non-aggressive protocols.) HTTP and FTP are considered aggressive protocols. If aggressive and non-aggressive protocols share the same tunnel, the aggressive protocols may take all of the available bandwidth. As a result, the performance of less aggressive protocols could be reduced. To avoid this problem, rules for HTTP and FTP traffic should have their own tunnel. To do this, set `tunnel-sharing` to `private` for WAN optimization rules that accept HTTP or FTP traffic.

It is also useful to set `tunnel-sharing` to `express-sharing` for applications, such as Telnet, that are very interactive but not aggressive. Express sharing optimizes tunnel sharing for Telnet and other interactive applications where latency or delays would seriously affect the user’s experience with the protocol.

Set `tunnel-sharing` to `sharing` for applications that are not aggressive and are not sensitive to latency or delays. WAN optimization rules set to `sharing` and `express-sharing` can share the same tunnel.

Protocol optimization

Protocol optimization techniques optimize bandwidth use across the WAN. These techniques can improve the efficiency of communication across the WAN optimization tunnel by reducing the amount of traffic required by communication protocols. You can apply protocol optimization to Common Internet File System (CIFS), FTP, HTTP, MAPI, and general TCP sessions. You can apply general TCP optimization to MAPI sessions.

For example, CIFS provides file access, record locking, read/write privileges, change notification, server name resolution, request batching, and server authentication. CIFS is a fairly “chatty” protocol, requiring many background transactions to successfully transfer a single file. This is usually not a problem across a LAN. However, across a WAN, latency and bandwidth reduction can slow down CIFS performance.

When you set *Protocol* to *CIFS* in a WAN optimization rule, the FortiGate units at both ends of the WAN optimization tunnel use a number of techniques to reduce the number of background transactions that occur over the WAN for CIFS traffic.

You can select only one protocol in a WAN optimization rule. For best performance, you should separate the traffic by protocol by creating different WAN optimization rules for each protocol. For example, to optimize HTTP traffic, you should set *Port* to 80 so that only HTTP traffic is accepted by this WAN optimization rule. For an example configuration that uses multiple rules for different protocols, see [“Example: Active-passive WAN optimization” on page 2721](#).

If the WAN optimization accepts a range of different types of traffic, you can set *Protocol* to *TCP* to apply general optimization techniques to TCP traffic. However, applying this TCP optimization to a range of different types of traffic is not as effective as applying more protocol-specific optimization to specific types of traffic. TCP protocol optimization uses techniques such as TCP SACK support, TCP window scaling and window size adjustment, and TCP connection pooling to remove TCP bottlenecks.

Protocol optimization and MAPI

By default the MAPI service uses port number 135 for RPC port mapping and may use random ports for MAPI messages. The random ports are negotiated through sessions using port 135. The FortiOS DCE-RPC session helper learns these ports and opens pinholes for the messages. WAN optimization is also aware of these ports and attempts to apply protocol optimization to MAPI messages that use them. However, to configure protocol optimization for MAPI you should set the WAN optimization rule to a single port number (usually port 135). Specifying a range of ports may reduce performance.

Byte caching

Byte caching breaks large units of application data (for example, a file being downloaded from a web page) into small chunks of data, labelling each chunk of data with a hash of the chunk and storing those chunks and their hashes in a database. The database is stored on a WAN optimization storage device. Then, instead of sending the actual data over the WAN tunnel, the FortiGate unit sends the hashes. The FortiGate unit at the other end of the tunnel receives the hashes and compares them with the hashes in its local byte caching database. If any hashes match, that data does not have to be transmitted over the WAN optimization tunnel. The data for any hashes that does not match is transferred over the tunnel and added to that byte caching database. Then the unit of application data (the file being downloaded) is reassembled and sent to its destination.

The stored byte caches are not application specific. Byte caches from a file in an email can be used to optimize downloading that same file or a similar file from a web page.

The result is less data transmitted over the WAN. Initially, byte caching may reduce performance until a large enough byte caching database is built up.

To enable byte caching, you select *Enable Byte Cache* in a WAN optimization rule. The *Protocol* setting does not affect byte caching. Data is byte cached when it is processed by a WAN optimization rule that includes byte caching.

Byte caching cannot determine whether or not a file is compressed (for example a zip file), and caches compressed and non-compressed versions of the same file separately.

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended best practice HA configuration for WAN optimization is active-passive mode. When the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

You can also form a WAN optimization tunnel between a cluster and a standalone FortiGate unit or between two clusters.

In a cluster, the primary unit stores only web cache and byte cache databases. These databases are not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its web and byte caches.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate units that it is participating with in WAN optimization tunnels.

WAN optimization, web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency FortiOS WAN optimization and web caching uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, WAN optimization and web cache require a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

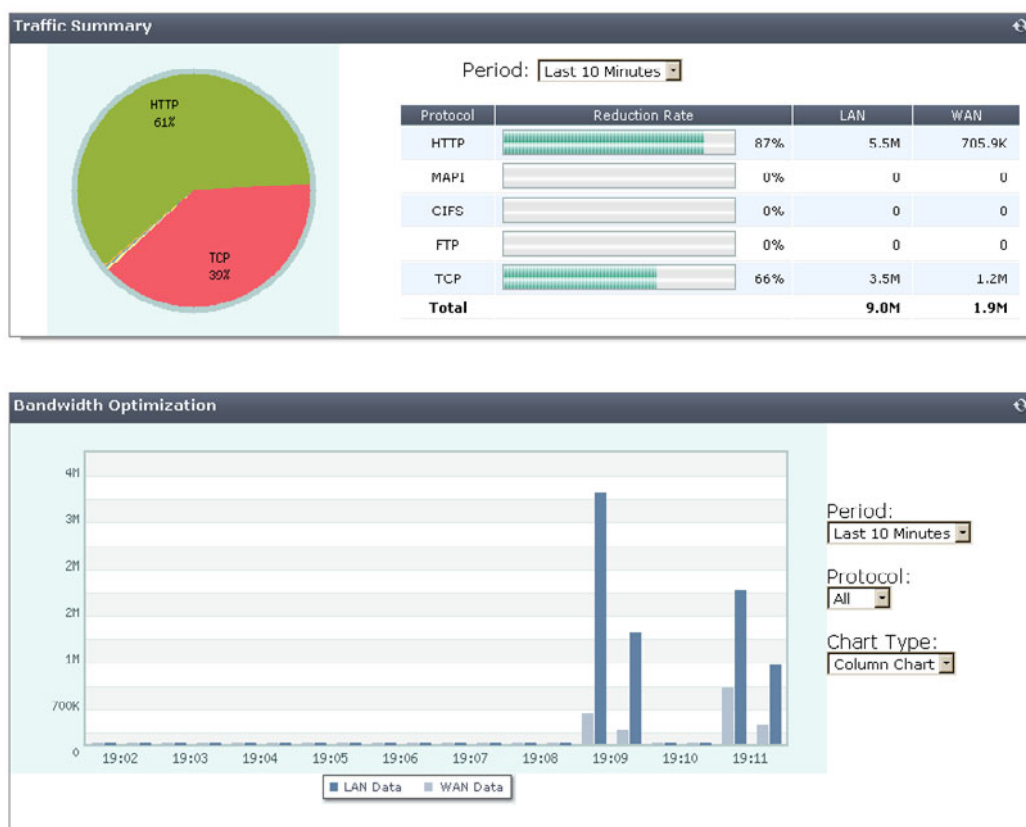
When WAN optimization and web caching are enabled you will see a reduction in available memory. The amount of reduction will increase when more WAN optimization and web cache sessions are being processed. If you are thinking of enabling WAN optimization or web caching on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods before you enable these features.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 1` command to see how much memory is currently being used by WAN optimization and web caching. See [“get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>” on page 2859](#) for more information.

Monitoring WAN optimization performance

Using WAN optimization monitoring, you can confirm that a FortiGate unit is optimizing traffic and view estimates of the amount of bandwidth saved. The WAN optimization monitor presents collected log information in a graphical format to show network traffic summary and bandwidth optimization information.

To view the WAN optimization monitor, go to *WAN Opt. & Cache > Monitor > WAN Opt Monitor*.

Figure 318: WAN optimization monitor

Traffic Summary

The traffic summary shows how WAN optimization is reducing the amount of traffic on the WAN for each WAN optimization protocol by showing the traffic reduction rate as a percentage of the total traffic. The traffic summary also shows the amount of WAN and LAN traffic. If WAN optimization is being effective the amount of WAN traffic should be lower than the amount of LAN traffic.

You can use the refresh icon to update the traffic summary display at any time. You can also set the amount of time for which the traffic summary shows data. The time period can vary from the last 10 minutes to the last month.

Bandwidth Optimization

This section shows network bandwidth optimization per time period. A line or column chart compares an application's pre-optimized size (LAN data) with its optimized size (WAN data). You can select the chart type, the monitoring time period, and the protocol for which to display data. If WAN optimization is being effective the WAN bandwidth should be lower than the LAN bandwidth.

Configuring WAN optimization traffic usage logs

Use the following command to generate WAN optimization traffic log messages for each WAN optimization protocol. WAN optimization traffic logs are required to generate WAN optimization usage reports. By default WAN optimization traffic log messages are not generated:

```
config wanopt settings
  set log-traffic {cifs | ftp | http | mapi | tcp}
end
```

For example, to enable WAN optimization traffic logging for CIFS and HTTP traffic enter:

```
config wanopt settings
  set log-traffic cifs http
end
```

You must also enable traffic logging in security policies that accept the traffic to be optimized.

As a result of this configuration, traffic log messages with a sub type of `wanopt-traffic` are generated by the FortiGate unit.

You can review, filter, and analyze these messages in the same way as other traffic log messages. For example, to view WAN optimization traffic log messages, go to *Log&Report > Log & Archive Access > Traffic Log* and configure a filter to view all log messages with a sub type of `wanopt-traffic`.

You can also use the following commands to enable or disable sending WAN optimization traffic logs to memory, FortiAnalyzer, or to a remote syslog server. These are all enabled by default.

```
config log memory filter
  set wanopt-traffic enable
end
config log fortianalyzer filter
  set wanopt-traffic enable
end
config log syslogd filter
  set wanopt-traffic enable
end
```

WAN optimization best practices

This is a short list of WAN optimization and explicit proxy best practices.

- WAN optimization tunnel sharing is recommended for similar types of WAN optimization traffic. However, tunnel sharing for different types of traffic is not recommended. For example, aggressive and non-aggressive protocols should not share the same tunnel. See [“Tunnel sharing” on page 2688](#).
- Active-passive HA is the recommended HA configuration for WAN optimization. See [“Tunnel sharing” on page 2688](#).
- Configure WAN optimization authentication with specific peers. Accepting any peer is not recommended as this can be less secure. See [“Accepting any peers” on page 2697](#).
- Set the explicit HTTP proxy *Default Policy Action* to *Deny*. This means that a security policy is required to use the explicit web proxy. See [“Explicit web proxy configuration overview” on page 2809](#).
- Set the explicit FTP proxy *Default Policy Action* to *Deny*. This means that a security policy is required to use the explicit FTP proxy. See [“Explicit FTP proxy configuration overview” on page 2834](#).



WAN optimization and Web cache storage

WAN optimization storage is used for storing the byte cache and web cache databases. In most cases, you can accept the default WAN optimization storage configuration because all of the disk space available on the FortiGate unit is in one partition. By default WAN optimization and logging and archiving are configured to use this partition.

You only have to configure WAN optimization storage if you have more than one possible storage location. This can happen if you have multiple partitions that you can use for storage locations. If you have more than one storage location you can move WAN optimization storage to it. You can also configure WAN optimization to use multiple storage locations.

You can also optionally configure WAN optimization storage if you want to adjust the relative amounts of disk space available for byte caching and web caching.

This chapter contains the following topics:

- [Formatting the hard disk](#)
- [Configuring WAN optimization and Web cache storage](#)

Formatting the hard disk

In most cases the hard disks on your FortiGate unit should be formatted with one partition that is used for WAN optimization and Logging and Archiving. If for some reason the hard disk is not formatted you can use the following information to format it. In some cases you might also want to use the following options to erase all data from the hard disk by reformatting it.

From the web-based manager go to *System > Config > Advanced > Disk Management* to display information about the hard disk or disks available to the FortiGate unit. To format a hard disk, select the format icon. The hard disk format takes a few minutes and the FortiGate unit restarts after formatting is complete.

From this web-based manager page you can also view and change the WAN optimization and Web Cache Storage size and view how much of the WAN optimization and web cache storage has been used.

From the CLI you can use the following command to view the current disk format and partition status. See the following example for a FortiGate-51B unit.

```
execute disk list
```

```
Device I1          29.9 GB      ref: 256 SUPER TALENT (IDE)
  partition 1      29.9 GB      ref: 257 label: 2B6375792136C707
```

You can use the following command to reformat the hard disk. Use this command if for some reason the disk is not formatted correctly. The command includes the device partition reference number (256) so formats the entire disk and not just the partition.

```
execute disk format 256
```

You can use the following command to reformat the partition. The command includes the partition reference number so formats the partition, removing all data from it. You can use this command to delete all data from the partition and to fix partition errors.

```
execute disk format 257
```

Configuring WAN optimization and Web cache storage

You can use the following command to add multiple WAN optimization storage locations if your FortiGate unit has multiple disk partitions and you want to use more than one for WAN optimization storage:

```
config system storage
```

Enter `get` to see the name of the default storage location. You cannot edit this storage location, but you can add new ones:

```
config system storage
edit new_storage
set partition <partition_number>
end
```

Where `<partition_number>` is the number of the partition to create a storage location in. This cannot be the same as the partition added to the default storage location. This command automatically adds a WAN optimization storage location with the name `new_storage`.

Changing the amount of space allocated for WAN optimization and Web cache storage

From the web-based manager you can go to *System > Config > Advanced > Disk Management* to edit the WAN optimization & Web Cache storage and change the allocation size to limit the amount of storage available for WAN optimization byte caching and web caching. The size is in Mbytes.

You can use the following command to change the size of any WAN optimization storage location. For example, in the FortiGate-51B the default WAN optimization storage is `Internal`. Use the following command to limit the amount of space allocated for WAN optimization to 20 Gbytes

```
config wanopt storage
edit Internal
set size 20000
end
```

Adjusting the relative amount of disk space available for byte caching and web caching

By default the `config wanopt storage` command allocates the same amount disk for byte caching and for web caching. In some cases you may want to adjust the relative amounts of disk space available for these two uses. For example, if you have not implemented web caching you may want to reduce the amount of disk space used for web caching and increase the amount of space used for byte caching.

You can adjust the relative amount of disk space used for byte caching using the `webcache-storage-percentage` option of the `config wanopt storage` command. This option adjusts the percentage in the range of 0 to 100. The default percentage is 50.

To reduce the percentage of space allocated on the Internal disk for web caching to 10% (resulting in the amount of space for byte caching increasing to 90%) enter:

```
config wanopt storage
  edit Internal
    set webcache-storage-percentage 10
  end
```

You can enter this command at any time without disrupting web caching or byte caching performance. Data may be lost from the cache that is reduced in size.



WAN optimization peers and authentication groups

All communication between WAN optimization peers begins with one WAN optimization peer (or client-side FortiGate unit) sending a WAN optimization tunnel request to another peer (or server-side FortiGate unit). During this process, the WAN optimization peers identify and optionally authenticate each other.

This chapter describes:

- [Basic WAN optimization peer requirements](#)
- [How FortiGate units process tunnel requests for peer authentication](#)
- [Configuring peers](#)
- [Configuring authentication groups](#)
- [Secure tunneling](#)
- [Monitoring WAN optimization peer performance](#)

Basic WAN optimization peer requirements

WAN optimization requires the following configuration on each peer. For information about configuring local and peer host IDs, see [“Configuring peers” on page 2699](#).

- The peer must have a unique host ID.
Unless authentication groups are used, peers authenticate each other using host ID values. Do not leave the local host ID at its default value.
- The peer must know the host IDs and IP addresses of all of the other peers that it can start WAN optimization tunnels with. This does not apply if you use authentication groups that accept all peers.

All peers must have the same local certificate installed on their FortiGate units if the units authenticate by local certificate (see [“Certificate-based authentication” on page 1263](#)). Similarly, if the units authenticate by pre-shared key (password), administrators must know the password. The type of authentication is selected in the authentication group. This applies only if you use authentication groups.

Accepting any peers

Strictly speaking, you do not need to add peers. Instead you can configure authentication groups that accept any peer. However, for this to work, both peers must have the same authentication group (with the same name) and both peers must have the same certificate or pre-shared key.

Accepting any peer is useful if you have many peers or if peer IP addresses change. For example, you could have many travelling FortiClient peers with IP addresses that are always changing as the users travel to different customer sites. This configuration is also useful if you have FortiGate units with dynamic external IP addresses (using DHCP or PPPoE). For most other situations, this method is not recommended and is not a best practice as it is less secure than accepting defined peers or a single peer. For more information, see [“Configuring authentication groups” on page 2700](#).

How FortiGate units process tunnel requests for peer authentication

When a client-side FortiGate unit attempts to start a WAN optimization tunnel with a peer server-side FortiGate unit, the tunnel request includes the following information:

- the client-side local host ID
- the name of an authentication group, if included in the rule that initiates the tunnel
- if an authentication group is used, the authentication method it specifies: pre-shared key or certificate
- the type of tunnel (secure or not).

For information about configuring the local host ID, peers and authentication groups, see [“Configuring peers” on page 2699](#) and [“Configuring authentication groups” on page 2700](#).

The authentication group is optional unless the tunnel is a secure tunnel. For more information, see [“Secure tunneling” on page 2702](#).

If the tunnel request includes an authentication group, the authentication will be based on the settings of this group as follows:

- The server-side FortiGate unit searches its own configuration for the name of the authentication group in the tunnel request. If no match is found, the authentication fails.
- If a match is found, the server-side FortiGate unit compares the authentication method in the client and server authentication groups. If the methods do not match, the authentication fails.
- If the authentication methods match, the server-side FortiGate unit tests the peer acceptance settings in its copy of the authentication group.
 - If the setting is *Accept Any Peer*, the authentication is successful.
 - If the setting is *Specify Peer*, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the peer name in the server-side authentication group. If the names match, authentication is successful. If a match is not found, authentication fails.
 - If the setting is *Accept Defined Peers*, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the server-side peer list. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the tunnel request does not include an authentication group, authentication will be based on the client-side local host ID in the tunnel request. The server-side FortiGate unit searches its peer list to match the client-side local host ID in the tunnel request. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the server-side FortiGate unit successfully authenticates the tunnel request, the server-side FortiGate unit sends back a tunnel setup response message. This message includes the server-side local host ID and the authentication group that matches the one in the tunnel request.

The client-side FortiGate unit then performs the same authentication procedure as the server-side FortiGate unit did. If both sides succeed, tunnel setup continues.

Configuring peers

When you configure peers, you first need to add the local host ID that identifies the FortiGate unit for WAN optimization and then add the peer host ID and IP address of each FortiGate unit with which a FortiGate unit can create WAN optimization tunnels.

To configure WAN optimization peers - web-based manager

- 1 Go to *Wan Opt. & Cache > WAN Opt. Peer > Peer*.
- 2 For *Local Host ID*, enter the local host ID of **this** FortiGate unit and select *Apply*. If you add this FortiGate unit as a peer to another FortiGate unit, use this ID as its **peer** host ID.

The local or host ID can contain up to 25 characters and can include spaces.

- 3 Select *Create New* to add a new peer.
- 4 For *Peer Host ID*, enter the peer host ID of the peer FortiGate unit. This is the local host ID added to the peer FortiGate unit.
- 5 For *IP Address*, add the IP address of the peer FortiGate unit. This is the source IP address of tunnel requests sent by the peer, usually the IP address of the FortiGate interface connected to the WAN.
- 6 Select *OK*.

To configure WAN optimization peers - CLI

In this example, the local host ID is named `HQ_Peer` and has an IP address of `172.20.120.100`. Three peers are added, but you can add any number of peers that are on the WAN.

- 1 Enter the following command to set the local host ID to `HQ_Peer`.

```
config wanopt settings
  set host-id HQ_peer
end
```

- 2 Enter the following commands to add three peers.

```
config wanopt peer
  edit Wan_opt_peer_1
    set ip 172.20.120.100
  next
  edit Wan_opt_peer_2
    set ip 172.30.120.100
  next
  edit Wan_opt_peer_3
    set ip 172.40.120.100
end
```

Configuring authentication groups

You need to add authentication groups to support authentication and secure tunneling between WAN optimization peers.

To perform authentication, WAN optimization peers use a certificate or a pre-shared key added to an authentication group so they can identify each other before forming a WAN optimization tunnel. Both peers must have an authentication group with the same name and settings. You add the authentication group to a peer-to-peer or active rule on the client-side FortiGate unit. When the server-side FortiGate unit receives a tunnel start request from the client-side FortiGate unit that includes an authentication group, the server-side FortiGate unit finds an authentication group in its configuration with the same name. If both authentication groups have the same certificate or pre-shared key, the peers can authenticate and set up the tunnel.

Authentication groups are also required for secure tunneling. See [“Secure tunneling” on page 2702](#).

To add authentication groups, go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group*.

To add an authentication group - web-based manager

Use the following steps to add any kind of authentication group. It is assumed that if you are using a local certificate to authenticate, it is already added to the FortiGate unit. For more information about FortiGate units and certificates, see the [FortiGate Certificate Management Guide](#).

- 1 Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.

- 2 Select *Create New*.

- 3 Add a *Name* for the authentication group.

You will select this name when you add the authentication group to a WAN optimization rule.

- 4 Select the *Authentication Method*.

Select *Certificate* if you want to use a certificate to authenticate and encrypt WAN optimization tunnels. You must select a local certificate that has been added to this FortiGate unit. (To add a local certificate, go to *System > Certificates > Local Certificates*.) Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and certificate.

Select *Pre-shared key* if you want to use a pre-shared key or password to authenticate and encrypt WAN optimization tunnels. You must add the *Password* (or pre-shared key) used by the authentication group. Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and password. The password must contain at least 6 printable characters and should be known only by network administrators. For optimum protection against currently known attacks, the key should consist of a minimum of 16 randomly chosen alphanumeric characters.

5 Configure *Peer Acceptance* for the authentication group.

Select *Accept Any Peer* if you do not know the peer host IDs or IP addresses of the peers that will use this authentication group. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP.

Select *Accept Defined Peers* if you want to authenticate with peers added to the peer list only.

Select *Specify Peer* and select one of the peers added to the peer list to authenticate with the selected peer only.

For more information, see [“Configuring peers” on page 2699](#).

6 Select OK.

7 Add the authentication group to a WAN optimization rule to apply the authentication settings in the authentication group to the rule.

For more information, see [“Configuring WAN optimization rules” on page 2710](#).

To add an authentication group that uses a certificate- CLI

Enter the following command to add an authentication group that uses a certificate and can authenticate all peers added to the FortiGate unit configuration.

In this example, the authentication group is named `auth_grp_1` and uses a certificate named `Example_Cert`.

```
config wanopt auth-group
  edit auth_grp_1
    set auth-method cert
    set cert Example_Cert
    set peer-accept defined
  end
```

To add an authentication group that uses a pre-shared key - CLI

Enter the following command to add an authentication group that uses a pre-shared key and can authenticate only the peer added to the authentication group.

In this example, the authentication group is named `auth_peer`, the peer that the group can authenticate is named `Server_net`, and the authentication group uses `123456` as the pre-shared key. In practice you should use a more secure pre-shared key.

```
config wanopt auth-group
  edit auth_peer
    set auth-method psk
    set psk 123456
    set peer-accept one
    set peer Server_net
  end
```

To add an authentication group that accepts WAN optimization connections from any peer - web-based manager

Add an authentication group that accepts any peer for situations where you do not have the *Peer Host IDs* or *IP Addresses* of the peers that you want to perform WAN optimization with. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP. An authentication group that accepts any peer is less secure than an authentication group that accepts defined peers or a single peer.

The example below sets the authentication method to *Pre-shared key*. You must add the same password to all FortiGate units using this authentication group.

- 1 Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.
- 2 Select *Create New* to add a new authentication group.
- 3 Configure the authentication group:

Name	Specify any name.
Authentication Method	Pre-shared key
Password	Enter a pre-shared key.
Peer Acceptance	Accept Any Peer

To add an authentication group that accepts WAN optimization connections from any peer - CLI

In this example, the authentication group is named `auth_grp_1`. It uses a certificate named `WAN_Cert` and accepts any peer.

```
config wanopt auth-group
edit auth_grp_1
set auth-method cert
set cert WAN_Cert
set peer-accept any
end
```

Secure tunneling

You can configure WAN optimization rules to use AES-128bit-CBC SSL to encrypt the traffic in the WAN optimization tunnel. WAN optimization uses FortiASIC acceleration to accelerate SSL decryption and encryption of the secure tunnel. Peer-to-peer secure tunnels use the same TCP port as non-secure peer-to-peer tunnels (TCP port 7810).

To use secure tunneling, you must select *Enable Secure Tunnel* in a WAN optimization rule and add an authentication group. The authentication group specifies the certificate or pre-shared key used to set up the secure tunnel. The *Peer Acceptance* setting of the authentication group does not affect secure tunneling.

The FortiGate units at each end of the secure tunnel must have the same authentication group with the same name and the same configuration, including the same pre-shared key or certificate. To use certificates you must install the same certificate on both FortiGate units.

For active-passive WAN optimization you can select *Enable Secure Tunnel* only in the active rule. In peer-to-peer WAN optimization you select *Enable Secure Tunnel* in the WAN optimization rule on both FortiGate units. For information about active-passive and peer-to-peer WAN optimization, see [“Configuring WAN optimization rules” on page 2705](#).

For a secure tunneling configuration example, see [“Example: Adding secure tunneling to an active-passive WAN optimization configuration” on page 2728](#). Secure tunneling is also used in the configuration example: [“Example: SSL offloading for a WAN optimization tunnel” on page 2786](#).

Monitoring WAN optimization peer performance

The WAN optimization peer monitor lists all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with. These include peers manually added to the configuration as well as discovered peers.

The monitor lists each peer's name, IP address, and peer type. The peer type indicates whether the peer was manually added or discovered. To show WAN optimization performance, for each peer the monitor lists the percent of traffic reduced by the peer in client-side WAN optimization configurations and in server-side configurations (also called gateway configurations).

To view the peer monitor, go to *WAN Opt. & Cache > Monitor > Peer Monitor*.



Configuring WAN optimization rules

To configure WAN optimization, you add WAN optimization rules. Similar to security policies, when a FortiGate unit receives a connection packet, it analyzes the packet's source address, destination address, and service (by destination port number), and attempts to locate a matching WAN optimization rule that decides how to optimize the traffic over the WAN. WAN optimization rules also apply features such as byte-caching and protocol optimization to optimized traffic.

You can add one of two types of WAN optimization rules: peer-to-peer and active-passive.

A **peer-to-peer WAN optimization** rule includes a peer host ID. WAN optimization sessions matched by a client-side peer-to-peer rule can only connect to the named server-side peer. When the client-side peer unit initiates a tunnel with the server-side peer, the packets that initiate the tunnel include extra information so that the server-side peer can determine that it is a peer-to-peer tunnel request. This extra information is required because the server-side peer does not require a WAN optimization rule; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list. Peer to peer WAN optimization tunnels use port 7810.

For **active-passive WAN optimization**, you add active rules to client-side FortiGate units and passive rules to server-side FortiGate units. A single passive rule can accept tunnel requests from multiple active rules. The configuration of the active rule enables WAN optimization features. The passive rule uses the configuration of the active rules. The one exception is web caching, which is enabled in passive rules.

This chapter describes:

- [WAN optimization rules, security policies, and UTM protection](#)
- [WAN optimization transparent mode](#)
- [WAN optimization rule list](#)
- [WAN optimization address formats](#)
- [Configuring WAN optimization rules](#)

WAN optimization rules, security policies, and UTM protection

The FortiGate unit applies security policies to communication sessions before WAN optimization rules. A WAN optimization rule can be applied to a packet only after the packet is accepted by a security policy. WAN optimization processes all sessions accepted by a security policy that also match a WAN optimization rule. However, if the security policy includes any UTM features, communication sessions accepted by the policy are processed by the UTM engine and not by WAN optimization. Before you add WAN optimization rules, you must add security policies to accept the traffic that you want to optimize.

To apply WAN optimization to traffic that is accepted by a security policy containing UTM features, you can use multiple FortiGate units or multiple VDOMs. You apply the UTM features in the first FortiGate unit or VDOM and then apply WAN optimization in the second FortiGate unit or VDOM. You also add inter-VDOM links between the VDOMs. See the configuration example “[Out-of-path WAN optimization with inter-VDOM routing](#)” on page 2757.

WAN optimization does not apply source and destination NAT settings included in security policies. This means that selecting NAT or adding virtual IPs in a security policy does not affect WAN optimized traffic. WAN optimization is also not compatible with firewall load balancing. However, traffic accepted by these policies that is not WAN optimized is processed as expected.

WAN optimization is compatible with identity-based security policies. If a session is allowed after authentication and if the identity-based policy that allows the session does not include UTM features, the session can be processed by matching WAN optimization rules.

traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.

WAN optimization transparent mode

WAN optimization is transparent to users. This means that with WAN optimization in place, clients connect to servers in the same way as they would without WAN optimization. However, servers receiving packets after WAN optimization “see” different source addresses depending on whether or not transparent mode is selected for WAN optimization. If transparent mode is selected, WAN optimization keeps the original source address of the packets, so servers appear to receive traffic directly from clients. Routing on the server network should be configured to route traffic with client source IP addresses from the server-side FortiGate unit to the server and back to the server-side FortiGate unit.



Some protocols, for example CIFS, may not function as expected if transparent mode is **not** selected. In most cases, for CIFS WAN optimization you should select transparent mode and make sure the server network can route traffic as described to support transparent mode.

If transparent mode is not selected, the source address of the packets received by servers is changed to the address of the server-side FortiGate unit interface that sends the packets to the servers. So servers appear to receive packets from the server FortiGate unit. Routing on the server network is simpler in this case because client addresses are not involved. All traffic appears to come from the server FortiGate unit and not from individual clients.



Do not confuse WAN optimization transparent mode with FortiGate transparent mode. WAN optimization transparent mode is configured in individual WAN optimization rules. FortiGate Transparent mode is a system setting that controls how the FortiGate unit (or a VDOM) processes traffic.

WAN optimization rule list

The WAN optimization rule list displays WAN optimization rules in their order of matching precedence. WAN optimization rule order affects rule matching. For details about arranging rules in the rule list, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

For information about WAN optimization rules and security policies, see [“WAN optimization rules, security policies, and UTM protection” on page 2705](#).

Then you add WAN optimization rules that:

- match WAN traffic to be optimized that is accepted by a security policy according to source and destination addresses and destination port of the traffic
- add the WAN optimization techniques to be applied to the traffic.

To view the WAN optimization rule list, go to *WAN Opt. & Cache > WAN Opt. Rule > Rule*.

Create New	Add a new WAN optimization rule. New rules are added to the bottom of the list.
Status	Select to enable a rule or clear to disable a rule. A disabled rule is out of service.
ID	The rule identifier. Rules are numbered in the order they are added to the rule list.
Source	The source address or address range that the rule matches. For more information, see “WAN optimization address formats” on page 2709 .
Destination	The destination address or address range that the rule matches. For more information, see “WAN optimization address formats” on page 2709 .
Port	The destination port number or port number range that the rule matches.
Method	Indicates whether you have selected byte caching in the WAN optimization rule.
Auto-Detect	Indicates whether the rule is an active (client) rule, a passive (server) rule or if auto-detect is off. If auto-detect is off, the rule can be peer-to-peer or Web Cache Only.
Protocol	The protocol optimization WAN optimization technique applied by the rule. For more information, see “Protocol optimization” on page 2689 .
Peer	For a peer-to-peer rule, the name of the peer WAN optimizer at the other end of the link.
Mode	Indicates whether the rule applies Full Optimization or Web Cache Only.
SSL	Indicates whether the rule is configured for SSL offloading.
Secure Tunnel	Indicates whether the rule is configured to use a WAN optimization tunnel.
Delete icon	Delete a rule from the list.
Edit icon	Edit a rule.

Insert WAN Optimization Rule Before icon	Add a new rule above the corresponding rule.
Move To icon	Move the corresponding rule before or after another rule in the list. For more information, see “How list order affects rule matching” on page 2708 and “Moving a rule to a different position in the rule list” on page 2709 .

How list order affects rule matching

Similar to security policies, you add WAN optimization rules to the WAN optimization rule list. The FortiGate unit uses the first-matching technique to select the WAN optimization rule to apply to a communication session.

When WAN optimization rules have been added, each time the FortiGate security accepts a communication session, it then searches the WAN optimization rule list for a matching rule. Matching rules are determined by comparing the rule with the session source and destination addresses and destination port. The search begins at the top of the rule list and progresses in order towards the bottom. Each rule in the rule list is compared with the communication session until a match is found. When the FortiGate unit finds the first matching rule, it applies that rule's specified WAN optimization features to the session and disregards subsequent rules.

If no WAN optimization rule matches, the session is processed according to the security policy that originally accepted the session.

As a general rule, you should order the WAN optimization rule list from most specific to most general because of the order in which rules are evaluated for a match, and because only the **first** matching rule is applied to a session. Subsequent possible matches are not considered or applied. Ordering rules from most specific to most general prevents rules that match a wide range of traffic from superseding and effectively masking rules that match exceptions.

For example, you might have a general WAN optimization rule that applies WAN optimization features but does not apply secure tunneling to most WAN traffic. However, you want to apply secure tunneling to FTP traffic (FTP traffic uses port 21). In this case, you would add a rule that creates a secure tunnel for FTP sessions above the general rule.

Figure 319: Example: secure tunneling for FTP — correct rule order

Status	ID	Source	Destination	Port	Method	Auto-Detect	Protocol	Peer	Mode	SSL Secure Tunnel	
<input checked="" type="checkbox"/>	2	192.168.20.*	172.20.120.*	21 - 21	Byte Caching	Active	FTP		Full Optimization		
<input checked="" type="checkbox"/>	3	192.168.20.*	172.20.120.*	1 - 65535	Byte Caching	Active	TCP		Full Optimization		

Exception
General

FTP sessions (using port 21) would immediately match the secure tunnel rule. Other kinds of services would not match the FTP rule, so rule evaluation would continue until the search reaches the matching general rule. This rule order has the intended effect. But if you reversed the order of the two rules, positioning the general rule before the FTP rule, all session, including FTP, would immediately match the general rule, and the rule to secure FTP would never be applied. This rule order would not have the intended effect.

Figure 320: Example: secure tunneling for FTP — incorrect rule order

Status	ID	Source	Destination	Port	Method	Auto-Detect	Protocol	Peer	Mode	SSL	Secure Tunnel	
<input checked="" type="checkbox"/>	3	192.168.20.*	172.20.120.*	1 - 65535	Byte Caching	Active	TCP		Full Optimization			
<input checked="" type="checkbox"/>	2	192.168.20.*	172.20.120.*	21 - 21	Byte Caching	Active	FTP		Full Optimization			

General
Exception

Similarly, if specific traffic requires exceptional WAN optimization rule settings, you would position those rules above other potential matches in the rule list. Otherwise, the other matching rules would take precedence, and the required exceptional settings might never be used.

Moving a rule to a different position in the rule list

When more than one rule has been defined, the first matching rule is applied to the traffic session. You can arrange the WAN optimization rule list to influence the order in which rules are evaluated for matches with incoming traffic. For more information, see [“How list order affects rule matching” on page 2708](#).

Moving a rule in the rule list does not change its ID, which only indicates the order in which the rule was created.

To move a rule in the WAN optimization rule list - web-based manager

- 1 Go to *WAN Opt & Cache > WAN Opt. Rule > Rule*.
- 2 In the rule list, note the ID of a rule that is before or after your intended destination.
- 3 In the row corresponding to the rule that you want to move, select the *Move To* icon.
- 4 Select *Before* or *After*, and enter the ID of the rule that is before or after your intended destination. This specifies the rule's new position in the WAN optimization rule list.
- 5 Select *OK*.

To move a rule in the WAN optimization rule list - CLI

- 1 Use the following command to move a WAN optimization rule with ID 34 above the rule in the rule list with ID 10.

```
config wanopt rule
  move 34 before 10
end
```
- 2 Use the following command to move a WAN optimization rule with ID 5 after the rule in the rule list with ID 1.

```
config wanopt rule
  move 5 after 1
end
```

WAN optimization address formats

A WAN optimization source or destination address can contain one or more network addresses. Network addresses can be represented by an IP address with a netmask or an IP address range.

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a source or destination address can be:

- a single computer, for example, 192.45.46.45
- a subnetwork, for example, 192.168.1.* for a class C subnet

- 0.0.0.0, matches any IP address.

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR-formatted netmasks to dotted decimal format. Example formats:

- netmask for a single computer: 255.255.255.255, or /32
- netmask for a class A subnet: 255.0.0.0, or /8
- netmask for a class B subnet: 255.255.0.0, or /16
- netmask for a class C subnet: 255.255.255.0, or /24
- netmask including all IP addresses: 0.0.0.0

Valid IP address and netmask formats include:

- x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0
- x.x.x.x/x, such as 192.168.1.0/24



An IP address 0.0.0.0 with netmask 255.255.255.255 is not a valid source or destination address.

When representing hosts by an IP range, the range indicates hosts with continuous IP addresses in a subnet, such as 192.168.1.[2-10], or 192.168.1.* to indicate the complete range of hosts on that subnet. You can also indicate the complete range of hosts on a subnet by entering 192.168.1.[0-255] or 192.168.1.0-192.168.1.255. Valid IP range formats include:

- x.x.x.x-x.x.x.x, for example, 192.168.110.100-192.168.110.120
- x.x.x.[x-x], for example, 192.168.110.[100-120]
- x.x.x.*, for a complete subnet, for example: 192.168.110.*
- x.x.x.[0-255] for a complete subnet, such as 192.168.110.[0-255]
- x.x.x.0 -x.x.x.255 for a complete subnet, such as 192.168.110.0 - 192.168.110.255



You cannot use square brackets [] or asterisks * when adding addresses to the CLI. Instead you must enter the start and end addresses of the subnet range separated by a dash -. For example, 192.168.20.0-192.168.20.255 for a complete subnet and 192.168.10.10-192.168.10.100 for a range of addresses.

Configuring WAN optimization rules

This section describes all the details that you can configure for the WAN optimization rules. The options available depend on how you configure a specific rule. The conditions are noted.

To add a WAN optimization rule - web-based manager

- 1 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.

- 2 Configure the WAN optimization rule, using the guidance in the following table, and select *OK*.

Mode	<p>Select <i>Full Optimization</i> to add a rule that can apply all WAN optimization features.</p> <p>Select <i>Web Cache Only</i> to add a rule that just applies web caching. If you select <i>Web Cache Only</i>, you can configure the source and destination address and port for the rule. You can also select <i>Transparent Mode</i> and <i>Enable SSL</i>.</p>
Source	<p>Enter an IP address, followed by a forward slash (/), then subnet mask, or enter an IP address range separated by a hyphen. For more information, see “WAN optimization address formats” on page 2709.</p> <p>Only packets whose source address header contains an IP address matching this IP address or address range will be accepted by and subject to this rule.</p> <p>For a passive rule, the server (passive) source address range should be compatible with the source addresses of the matching client (active) rule. To match one passive rule with many active rules, the passive rule source address range should include the source addresses of all of the active rules.</p>
Destination	<p>Enter an IP address, followed by a forward slash (/), then subnet mask, or enter an IP address range separated by a hyphen. For more information, see “WAN optimization address formats” on page 2709.</p> <p>Only a packet whose destination address header contains an IP address matching this IP address or address range will be accepted by and subject to this rule.</p> <p>For a Web Cache Only rule, if you set <i>Destination</i> to 0.0.0.0, the rule caches web pages on the Internet or any network.</p> <p>For a passive rule, the server (passive) destination address range should be compatible with the destination addresses of the matching client (active) rule. To match one passive rule with many active rules, the passive rule destination address range should include the destination addresses of all of the active rules.</p>
Port	<p>Enter a single port number or port number range. Only packets whose destination port number matches this port number or port number range will be accepted by and subject to this rule.</p> <p>For a passive rule, the server (passive) port range should be compatible with the port range of the matching client (active) rule. To match one passive rule with many active rules, the passive rule port range should include the port ranges of all of the active rules.</p>

Auto-Detect	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>.</p> <p>Specify whether the rule is <i>Active</i> (client), <i>Passive</i> (server) or if <i>Auto-Detect</i> is <i>Off</i>. If <i>Auto-Detect</i> is <i>Off</i>, the rule is a peer-to-peer rule.</p> <p>For an <i>Active</i> (client) rule, you must select all of the WAN optimization features to be applied by the rule. You can select the protocol to optimize, transparent mode, byte caching, SSL offloading, secure tunneling, and an authentication group.</p> <p>A <i>Passive</i> (server) rule uses the settings in the active rule on the client FortiGate unit to apply WAN optimization settings. You can also select web caching for a passive rule.</p> <p>If <i>Auto-Detect</i> is <i>Off</i>, the rule must include all required WAN optimization features and you must select a <i>Peer</i> for the rule. Select this option to configure peer-to-peer WAN optimization where this rule can start a WAN optimization tunnel with this peer only.</p>
Protocol	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Off</i> or <i>Active</i>.</p> <p>Select CIFS, FTP, HTTP or MAPI to apply protocol optimization for one of these protocols. For information about protocol optimization, see “Protocol optimization” on page 2689.</p> <p>Select TCP if the WAN optimization tunnel accepts sessions that use more than one protocol or that do not use the CIFS, FTP, HTTP, or MAPI protocol.</p>
Peer	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Off</i>.</p> <p>Select the peer host ID of the peer that this peer-to-peer WAN optimization rule will start a WAN optimization tunnel with. You can also select <i>[Create New...]</i> from the list to add a new peer.</p>
Enable Web Cache	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Off</i> or <i>Passive</i>. If <i>Auto-Detect</i> is set to <i>Off</i>, then <i>Protocol</i> must be set to <i>HTTP</i>.</p> <p>Select to apply WAN optimization web caching to the sessions accepted by this rule. For more information, see “Web caching” on page 2735.</p>
Transparent Mode	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i> and <i>Auto-Detect</i> is set to <i>Active</i> or <i>Off</i>, or if <i>Mode</i> is set to <i>Web Cache Only</i>.</p> <p>Servers receiving packets after WAN optimization “see” different source addresses depending on whether or not you select <i>Transparent Mode</i>.</p> <p>For more information, see “WAN optimization transparent mode” on page 2706.</p>
Enable Byte Caching	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Off</i> or <i>Active</i>.</p> <p>Select to apply WAN optimization byte caching to the sessions accepted by this rule. For more information, see “Byte caching” on page 2689.</p>

Enable SSL	<p>Available only if <i>Auto-Detect</i> is set to <i>Active</i> or <i>Off</i>.</p> <p>Select to apply SSL offloading for HTTPS traffic. You can use SSL offloading to offload SSL encryption and decryption from one or more HTTP servers to the FortiGate unit. If you enable this option, you must configure the rule to accept SSL-encrypted traffic. For example, you can configure the rule to accept HTTPS traffic by setting <i>Port</i> to 443.</p> <p>If you enable SSL offloading, you must also use the CLI command <code>config wanopt ssl-server</code> to add an SSL server for each HTTP server that you want to offload SSL encryption/decryption for. For more information, see “SSL offloading for WAN optimization and web caching” on page 2781.</p>
Enable Secure Tunnel	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Active</i> or <i>Off</i>.</p> <p>If you select <i>Enable Secure Tunnel</i>, the WAN optimization tunnel is encrypted using SSL encryption. You must also add an authentication group to the rule. For more information, see “Secure tunneling” on page 2702.</p>
Authentication Group	<p>Available only if <i>Mode</i> is set to <i>Full Optimization</i>, and <i>Auto-Detect</i> is set to <i>Active</i> or <i>Off</i>.</p> <p>Select this option and select an authentication group from the list if you want groups of FortiGate units to authenticate with each other before starting the WAN optimization tunnel. You must also select an authentication group if you select <i>Enable Secure Tunnel</i>.</p> <p>You must add identical authentication groups to both of the FortiGate units that will participate in the WAN optimization tunnel started by the rule. For more information, see “Configuring authentication groups” on page 2700.</p>

To add a WAN optimization rule - CLI

Using the guidance in the previous table, enter the following commands. For more information, see the `wanopt` and `rules` listings in the [FortiGate CLI Reference](#).

```
config wanopt rule
edit <index_int>
    set auth-group <auth_group_name>
    set auto-detect {active | off | passive}
    set byte-caching {disable | enable}
    set dst-ip <address_ipv4>[-<address-ipv4>]
    set mode {full | webcache-only}
    set peer <peer_name>
    set port <port_int>[-<port-int>]
    set proto {cifs | ftp | http | mapi | tcp}
    set secure-tunnel {disable | enable}
    set src-ip <address_ipv4>[-<address-ipv4>]
    set ssl {disable | enable}
    set status {disable | enable}
    set transparent {disable | enable}
    set tunnel-non-http {disable | enable}
    set tunnel-sharing {express-shared | private | shared}
    set unknown-http-version {best-effort | reject | tunnel}
    set webcache {disable | enable}
end
```

Processing non-HTTP sessions accepted by an HTTP rule

From the CLI, use the `tunnel-non-http` keyword of the `config wanopt rule` command to configure how to process non-HTTP sessions when a rule configured to accept and optimize HTTP traffic accepts a non-HTTP session. This can occur if an application sends non-HTTP sessions using an HTTP destination port.

To drop non-HTTP sessions accepted by the rule set `tunnel-non-http` to `disable`, or set it to `enable` to pass non-HTTP sessions through the tunnel without applying protocol optimization, byte-caching, or web caching. In this case, the FortiGate unit applies TCP protocol optimization to non-HTTP sessions.

Processing unknown HTTP sessions

Unknown HTTP sessions are HTTP sessions that do not comply with HTTP 0.9, 1.0, or 1.1. From the CLI, use the `unknown-http-version` keyword of the `config wanopt rule` command to specify how a rule handles such HTTP sessions.

To assume that all HTTP sessions accepted by the rule comply with HTTP 0.9, 1.0, or 1.1, select `best-effort`. If a session uses a different HTTP version, WAN optimization may not parse it correctly. As a result, the FortiGate unit may stop forwarding the session and the connection may be lost. To reject HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, select `reject`.

To pass HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, but without applying HTTP protocol optimization, byte-caching, or web caching, you can also select `tunnel`. TCP protocol optimization is applied to these HTTP sessions.



WAN optimization configuration examples

This chapter provides the following basic examples to illustrate WAN optimization configurations introduced in the previous chapters:

- [Example: Basic peer-to-peer WAN optimization configuration](#)
- [Example: Active-passive WAN optimization](#)
- [Example: Adding secure tunneling to an active-passive WAN optimization configuration](#)

Example: Basic peer-to-peer WAN optimization configuration

Peer-to-peer WAN optimization is the simplest WAN optimization configuration. In a peer to peer configuration the WAN optimization tunnel can be set up only between one client-side FortiGate unit and one server-side FortiGate unit named in the WAN optimization rule added to the client-side FortiGate unit. When the client-side FortiGate unit initiates a tunnel with the server-side FortiGate unit, the packets that initiate the tunnel include extra information so that this server-side FortiGate unit can determine that it is a peer-to-peer tunnel request. This extra information is required because the server-side FortiGate unit does not require a WAN optimization rule; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list.

The extra information in the communication session plus the peer list entry allow the server-side FortiGate unit to set up the WAN optimization tunnel with the client-side FortiGate unit by using only the settings on the client-side WAN optimization rule.




Traffic shaping is ignored for peer-to-peer WAN optimization.

In a peer-to-peer WAN optimization configuration you create a peer-to-peer WAN optimization rule on the client-side FortiGate unit with *Auto-Detect* to *Off* and include the peer host ID of the server-side FortiGate unit. Using this rule, the client-side FortiGate unit can create a WAN optimization tunnel only with the peer that is added to the rule.

You do not have to add a rule to the server-side FortiGate unit. But the server-side FortiGate unit peer list must include the Peer Host ID and IP address of the client FortiGate unit. The server-side FortiGate unit uses the WAN optimization settings in the client-side rule.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Peer_Fgt_1 with a WAN IP address of 172.20.34.12. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Peer_Fgt_2 with a WAN IP address of 192.168.30.12. This unit is in front of a web server network with IP address 192.168.10.0.

Figure 321: Example peer-to-peer topology

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the client-side FortiGate unit by adding peers and a security policy that accepts traffic to be optimized.
- 2 Configure the server-side FortiGate unit.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring basic peer-to-peer WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure the client-side FortiGate unit and security policy

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

Local Host ID	Peer_Fgt_1
----------------------	------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

Peer Host ID	Peer_Fgt_2
IP Address	192.168.30.12

- 4 Select *OK*.

- 5 Go to *Policy > Policy > Policy* and add a security policy to the client-side FortiGate unit that accepts traffic to be optimized:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 6 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.

- 7 Configure the rule:

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	135
Auto-Detect	Off
Protocol	MAPI
Peer	Peer_Fgt_2
Transparent Mode	Select
Enable Byte Caching	Select

- 8 Select *OK*.

The rule is added to the bottom of the WAN optimization list.

- 9 If required, move the rule to a different position in the list so that the rule accepts the required MAPI sessions that use port 135. Depending on your rule list configuration, this may involve moving the rule above more general rules that would also match MAPI traffic.

For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	Peer_Fgt_2
----------------------	------------

- 2 Select *Apply* to save your setting.

- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the peer side FortiGate unit:

Peer Host ID	Peer_Fgt_1
IP Address	172.20.34.12

- 4 Select *OK*.

Configuring basic peer-to-peer WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit and security policy

- 1 Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
  set host-id Peer_Fgt_1
end
```

- 2 Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Peer_Fgt_2
  set ip 192.168.30.12
end
```

- 3 Add a security policy to the client-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 23
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
end
```

- 4 Add the following peer-to-peer rule:

```
config wanopt rule
  edit 2
    set src-ip 172.20.120.0-172.20.120.255
    set dst-ip 192.168.10.0-192.168.10.255
    set port 135
    set proto mapi
    set peer Peer_Fgt_2
  end
```

Accept default settings for auto-detect (off), transparent (enable), status (enable), mode (full), byte-caching (enable), ssl (disable), secure-tunnel (disable), auth-group (null), unknown-http-version (tunnel), and tunnel-non-http (disable).

- 5 If required, move the rule to a different position in the list.

For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

- 6 If required, use the `move` command to change the order of the rules in the list so that the rule accepts the required MAPI sessions that use port 135. Depending on your rule list configuration, this may involve moving the rule above more general rules that would also match MAPI traffic.

For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
    set host-id Peer_Fgt_2
end
```

- 2 Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
    edit Peer_Fgt_1
        set ip 192.168.30.12
    end
```

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the user network and the web server network. For example, from a PC on the user network browse to the IP address of a web server on the web server network, for example <http://192.168.10.100>. Even though this address is not on the user network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring (go to *WAN Opt. & Cache > Monitor > Monitor*). If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include UTM options. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy
- Check routing on the FortiGate units and on the user and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the user network must allow packets destined for the web server network to be received by the client side FortiGate unit, and packets from the server side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command on the client-side FortiGate unit to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
  http tunnel
    bytes_in=1751767 bytes_out=325468
  ftp tunnel
    bytes_in=0 bytes_out=0
  cifs tunnel
    bytes_in=0 bytes_out=0
  mapi tunnel
    bytes_in=0 bytes_out=0
```

```
tcp tunnel
  bytes_in=3182253 bytes_out=200702
maintenance tunnel
  bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output for the client side FortiGate unit shows one peer with IP address 192.168.20.1, peer name Web_servers, and with 10 active tunnels.

```
get test wad 26
peer name=Web_servers ip=192.168.20.1 vd=0 version=1
  tunnels(active/connecting/failover)=10/0/0
  sessions=0 n_retries=0 version_valid=true
```

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output for the client-side FortiGate unit shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list
```

```
Tunnel: id=100 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=100 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384
```

```
Tunnel: id=99 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=99 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384
```

```
Tunnel: id=98 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=98 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384
```

```
Tunnel: id=39 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=39 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1068 bytes_out=1104
```

```
Tunnel: id=7 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=7 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=8 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=8 ip=192.168.30.12
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=5 type=manual
      vd=0 shared=no uses=0 state=3
      peer name=Web_servers id=5 ip=192.168.30.12
      SSL-secured-tunnel=no auth-grp=
      bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
      vd=0 shared=no uses=0 state=3
      peer name=Web_servers id=4 ip=192.168.30.12
      SSL-secured-tunnel=no auth-grp=
      bytes_in=1228 bytes_out=1264

Tunnel: id=1 type=manual
      vd=0 shared=no uses=0 state=3
      peer name=Web_servers id=1 ip=192.168.30.12
      SSL-secured-tunnel=no auth-grp=
      bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
      vd=0 shared=no uses=0 state=3
      peer name=Web_servers id=2 ip=192.168.30.12
      SSL-secured-tunnel=no auth-grp=
      bytes_in=1228 bytes_out=1264

Tunnels total=10 manual=10 auto=0
```

Example: Active-passive WAN optimization

In active-passive WAN optimization you add active WAN optimization rules on the client-side FortiGate unit by setting WAN optimization *Auto-Detect* to *Active*. You configure passive WAN optimization rules on the server-side FortiGate unit by setting WAN optimization *Auto-Detect* to *Passive*.


You can add multiple active rules for one passive rule to optimize different protocols. Since you do not configure the protocol in the passive rule, one passive rule can be used for each of the active rules. Adding fewer passive rules simplifies the WAN optimization configuration.

Network topology and assumptions

This example configuration includes three active rules on the client-side FortiGate unit and one passive rule in the server-side FortiGate unit. The active rules do the following:

- optimize CIFS traffic from IP addresses 172.20.120.100 to 172.20.120.200
- optimize HTTP traffic from IP addresses 172.20.120.100 to 172.20.120.150
- optimize FTP traffic from IP addresses 172.20.120.151 172.20.120.200.

You can do this by adding three active WAN optimization rules to the client-side FortiGate unit, one for each protocol—with port set to 80 for the HTTP rule, 21 for the FTP rule and 1-65535 for the CIFS rule. Then you arrange the rules in the WAN optimization rule list with the CIFS rule last because the HTTP and FTP rules include single port numbers.

Figure 322: Example active-passive WAN optimization topology

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the client-side FortiGate unit by adding peers and a security policy that accepts traffic to be optimized.
- 2 Add WAN optimization rules to the FortiGate unit.
- 3 Configure the server-side FortiGate unit.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring basic active-passive WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure peers on the client-side FortiGate unit and add a security policy

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

Local Host ID	User_net
----------------------	----------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

Peer Host ID	Web_servers
IP Address	192.168.20.1

- 4 Select *OK*.
- 5 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy to the client-side FortiGate unit to accept the traffic to be optimized:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

To add the active rules to the client-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule*.
- 2 Select *Create New* to add the active rule to optimize CIFS traffic from IP addresses 172.20.120.100 to 172.20.120.200:

Mode	Full Optimization
Source	172.20.120.[100-200]
Destination	192.168.10.*
Port	1 - 65535
Auto-Detect	Active
Protocol	CIFS
Transparent Mode	Select
Enable Byte Caching	Select

- 3 Select *OK*.
- 4 Select *Create New* to add the active rule to optimize HTTP traffic for IP addresses 172.20.120.100 to 172.20.120.150:

Mode	Full Optimization
Source	172.20.120.[100-150]
Destination	192.168.10.*
Port	80
Auto-Detect	Active
Protocol	HTTP
Transparent Mode	Select
Enable Byte Caching	Select

- 5 Select *OK*.

- 6 Select *Create New* to add the active rule to optimize FTP traffic from IP addresses 172.20.120.151 172.20.120.200:

Mode	Full Optimization
Source	172.20.120.[151-200]
Destination	192.168.10.*
Port	21
Auto-Detect	Active
Protocol	FTP
Transparent Mode	Select
Enable Byte Caching	Select

- 7 Select *OK*.
- 8 If required, use the *Move To* icon to change the order of the rules in the list so that the HTTP and FTP rules are above the CIFS rule in the list. You may need to do this if you have other WAN optimization rules in the list.

For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	Web_servers
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

Peer Host ID	User_net
IP Address	172.30.120.1

- 4 Select *OK*.
- 5 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.
- 6 Add the passive rule. The source address matches the 172.20.120.100 to 172.20.120.200 IP address range and the 1-65535 port range. You can also enable web caching for the HTTP traffic:

Mode	Full Optimization
Source	172.20.120.[100-200]
Destination	192.168.10.*
Port	1-65535
Auto-Detect	Passive
Enable Web Cache	Select

- 7 Select *OK*.
- The rule is added to the bottom of the rule list.

- 8 If required, move the rule to a different position in the list so that the tunnel request from the client-side FortiGate unit matches with this rule.

For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

Configuring basic active-passive WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure peers on the client-side FortiGate unit and add a security policy

- 1 Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
  set host-id User_net
end
```

- 2 Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Web_servers
  set ip 192.168.20.1
end
```

- 3 Add a security policy to the client-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 20
  set srcintf port1
  set dstintf port2
  set srcaddr all
  set dstaddr all
  set action accept
  set service ANY
  set schedule always
end
```

To add the active rules to the client-side FortiGate unit

- 1 Add the following active rule to optimize CIFS traffic for IP addresses 172.20.120.100 to 172.20.120.200:

```
config wanopt rule
  edit 2
  set auto-detect active
  set src-ip 172.20.120.100-172.20.120.200
  set dst-ip 192.168.10.0-192.168.10.255
  set port 1-65535
  set proto cifs
end
```

Accept default settings for transparent (enable), status (enable), mode (full), byte-caching (enable), ssl (disable), secure-tunnel (disable), auth-group (null), unknown-http-version (tunnel), and tunnel-non-http (disable).

- 2 Add the following active rule to optimize HTTP traffic for IP addresses 172.20.120.100 to 172.20.120.150:

```

config wanopt rule
  edit 3
    set auto-detect active
    set src-ip 172.20.120.100-172.20.120.150
    set dst-ip 192.168.10.0-192.168.10.255
    set port 80
  end

```

Accept default settings for transparent (enable), proto (http), status (enable), mode (full), byte-caching (enable), ssl (disable), secure-tunnel (disable), auth-group (null), unknown-http-version (tunnel), and tunnel-non-http (disable).

- 3 Add the following active rule to optimize FTP traffic from IP addresses 172.20.120.151-172.20.120.200:

```

config wanopt rule
  edit 4
    set auto-detect active
    set src-ip 172.20.120.151-172.20.120.200
    set dst-ip 192.168.10.0-192.168.10.255
    set port 21
    set proto ftp
  end

```

Accept default settings for transparent (enable), status (enable), mode (full), byte-caching (enable), ssl (disable), secure-tunnel (disable), auth-group (null), unknown-http-version (tunnel), and tunnel-non-http (disable).

- 4 If required, use the `move` command to change the order of the rules in the list so that the HTTP and FTP rules are above the CIFS rule in the list. You may need to do this if you have other WAN optimization rules in the list.

For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Add the Local Host ID to the server-side FortiGate configuration:

```

config wanopt settings
  set host-id Web_servers
end

```

- 2 Add the client-side Local Host ID to the server-side peer list:

```

config wanopt peer
  edit User_net
    set ip 172.20.120.1
  end

```

- 3 Add the following passive rule to the server-side FortiGate unit:

```

config wanopt rule
  edit 5
    set auto-detect passive
    set src-ip 172.20.120.[100-200]
    set dst-ip 192.168.10.0-192.168.10.255
    set port 1-65535
    set webcache enable
  end

```

```
end
```

Accept default settings for `status` (`enable`) and `mode` (`full`).

- 4 If required, use the `move` command to move the rule to a different position in the list so that the tunnel request from the client-side FortiGate unit matches with this rule.

For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the user network and the web server network. For example, from a PC on the user network browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the user network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring (go to *WAN Opt. & Cache > Monitor > Monitor*). If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include UTM options. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy
- Check routing on the FortiGate units and on the user and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the user network must allow packets destined for the web server network to be received by the client side FortiGate unit, and packets from the server side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command to display WAN optimization tunnel protocol statistics. The `http` tunnel and `tcp` tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
  http tunnel
    bytes_in=1751767 bytes_out=325468
  ftp tunnel
    bytes_in=0 bytes_out=0
  cifs tunnel
    bytes_in=0 bytes_out=0
  mapi tunnel
    bytes_in=0 bytes_out=0
  tcp tunnel
    bytes_in=3182253 bytes_out=200702
  maintenance tunnel
    bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output for the client side FortiGate unit shows one peer with IP address 192.168.20.1, peer name Web_servers, and with 10 active tunnels.

```
get test wad 26
peer name=Web_servers ip=192.168.20.1 vd=0 version=1
tunnels(active/connecting/failover)=10/0/0
sessions=0 n_retries=0 version_valid=true
```

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 3 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to on).

```
diagnose wad tunnel list

Tunnel: id=139 type=auto
vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=744 bytes_out=76

Tunnel: id=141 type=auto
vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=727 bytes_out=76

Tunnel: id=142 type=auto
vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=727 bytes_out=76

Tunnels total=3 manual=0 auto=3
```

Example: Adding secure tunneling to an active-passive WAN optimization configuration


This example shows how to configure two FortiGate units for active-passive WAN optimization with secure tunneling. The same authentication group is added to both FortiGate units. The authentication group includes a password (or pre-shared key) and has *Peer Acceptance* set to *Accept any Peer*. An active rule is added to the client-side FortiGate unit and a passive rule to the server-side FortiGate unit. The active rule uses secure tunneling, optimizes HTTP traffic, and uses Transparent Mode and byte caching.

The authentication group is named *Auth_Secure_Tunnel* and the password for the pre-shared key is 2345678. The topology for this example is shown in [Figure 323](#). This example includes web-based manager configuration steps followed by equivalent CLI configuration steps. For information about secure tunneling, see [“Secure tunneling” on page 2702](#).

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called User_net with a WAN IP address of 172.30.120.1. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Web_servers and has a WAN IP address of 192.168.20.1. This unit is in front of a web server network with IP address 192.168.10.0.

Figure 323: Example active-passive WAN optimization and secure tunneling topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the client-side FortiGate unit by adding peers and a security policy that accepts traffic to be optimized.
- 2 Add an authentication group and WAN optimization rule to the client-side FortiGate unit.
- 3 Configure peers on the server-side FortiGate unit.
- 4 Add the same authentication group and add a WAN optimization rule to the server-side FortiGate unit.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring WAN optimization with secure tunneling - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure peers on the client-side FortiGate unit and add a security policy

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate unit:

Local Host ID	User_net
----------------------	----------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

Peer Host ID	Web_servers
IP Address	192.168.20.1

- 4 Select *OK*.
- 5 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy to the client-side FortiGate unit to accept the traffic to be optimized:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

To add the authentication group and WAN optimization rule to the client-side FortiGate unit

- 1 Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.
- 2 Select *Create New* to add a new authentication group to be used for secure tunneling:

Name	Auth_Secure_Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

- 3 Select *OK*.
- 4 Go to *Wan Opt. & Cache > WAN Opt. Rule > Rule*.
- 5 Select *Create New* to add an active rule that enables secure tunneling and includes the authentication group:

Mode	Full Optimization
Source	172.20.120.[100-200]
Destination	192.168.10.*
Port	80
Auto-Detect	Active
Protocol	HTTP

Transparent Mode	Select
Enable Byte Caching	Select
Enable Secure Tunnel	Select
Authentication Group	Auth_Secure_Tunnel

- 6 Select *OK*.

To configure peers on the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	Web_servers
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

Peer Host ID	User_net
IP Address	172.30.120.1

- 4 Select *OK*.

To add the authentication group and WAN optimization rule to the server-side FortiGate unit

- 1 Go to *Wan Opt. & Cache > WAN Opt. Peer > Authentication Group*.
- 2 Select *Create New* and add a new authentication group to be used for secure tunneling:

Name	Auth_Secure_Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

- 3 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.
- 4 Add the passive rule. The source address matches the 172.20.120.100 to 172.20.120.200 IP address range and the 1-65535 port range. You can also enable web caching for HTTP traffic:

Mode	Full Optimization
Source	172.20.120.[100-200]
Destination	192.168.10.*
Port	1-65535
Auto-Detect	Passive
Enable Web Cache	Select

- 5 Select *OK*.

Configuring WAN optimization with secure tunneling - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure peers on the client-side FortiGate unit and add a security policy

- 1 Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
  set host-id User_net
end
```

- 2 Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Web_servers
  set ip 192.168.20.1
end
```

- 3 Add a security policy to the server-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 20
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
end
```

To add the authentication group and WAN optimization rule to the client-side FortiGate unit

- 1 Add a new authentication group to be used for secure tunneling:

```
config wanopt auth-group
  edit Auth_Secure_Tunnel
    set auth-method psk
    set psk 2345678
  end
```

Leave peer-accept at its default value.

- 2 Add the following active rule to optimize HTTP traffic for IP addresses 172.20.120.100 to 172.20.120.200:

```
config wanopt rule
  edit 1
    set auto-detect active
    set src-ip 172.20.120.100-172.20.120.200
    set dst-ip 192.168.10.0-192.168.10.255
    set port 80
    set proto http
    set secure-tunnel enable
    set auth-group Auth_Secure_Tunnel
  end
```

Leave the rest of the settings at their default values.

To configure peers on the server-side FortiGate unit

- 1 Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
    set host-id Web_servers
end
```

- 2 Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
    edit User_net
        set ip 172.20.120.1
    end
```

To add the authentication group and WAN optimization rule to the server-side FortiGate unit

- 1 Add a new authentication group to be used for secure tunneling:

```
config wanopt auth-group
    edit Auth_Secure_Tunnel
        set auth-method psk
        set psk 2345678
    end
```

Leave `peer-accept` at its default value.

- 2 Add the following passive rule to the server-side FortiGate unit:

```
config wanopt rule
    edit 5
        set auto-detect passive
        set src-ip 172.20.120.[100-200]
        set dst-ip 192.168.10.0-192.168.10.255
        set port 1-65535
        set webcache enable
    end
```

Leave `status` (`enable`) and `mode` (`full`) at their default values.



Web caching

FortiGate web caching is a form of object caching that accelerates web applications and web servers by reducing bandwidth usage, server load, and perceived latency. Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites. Web caching can also support caching HTTPS sessions provided that you import the correct certificate. See [RFC 2616](#) for information about web caching for HTTP 1.1. Web caching does not cache audio and video streams including Flash videos and streaming content.

Web caching involves storing HTML pages, images, servlet responses and other web-based objects for later retrieval. These objects are stored in the web cache storage location defined by the `config wanopt storage` command. You can also go to *System > Config > Advanced > Disk Management* to view the storage locations on the FortiGate unit hard disks.

There are three significant advantages to using web caching to improve HTTP and WAN performance:

- reduced bandwidth consumption because fewer requests and responses go over the WAN or Internet.
- reduced web server load because there are fewer requests for web servers to handle.
- reduced latency because responses for cached requests are available from a local FortiGate unit instead of from across the WAN or Internet.

You can use web caching to cache any web traffic that passes through the FortiGate unit, including web pages from web servers on a LAN, WAN or on the Internet. You apply web caching by enabling the web caching option in any security policy and WAN optimization rule. When enabled in a security policy, web caching is applied to all HTTP sessions accepted by the security policy. If the security policy is an explicit web proxy security policy, the FortiGate unit caches explicit web proxy sessions.

When enabled in a WAN optimization rule, the FortiGate unit caches HTTP traffic processed by that WAN optimization rule. You can add WAN optimization rules that only apply web caching. You can also add web caching to WAN optimization rules for HTTP traffic that also include byte caching, protocol optimization, and other WAN optimization features.

Web caching caches compressed and non-compressed versions of the same file separately. If the HTTP protocol considers the compressed and uncompressed versions of a file the same object, only the compressed or uncompressed file will be cached.

This section contains the following topics:

- [Web caching in security policies](#)
- [Web Caching only WAN optimization](#)
- [Web caching for active-passive WAN optimization](#)
- [Web caching for peer-to-peer WAN optimization](#)
- [Exempting web sites from web caching](#)
- [Changing web cache settings](#)
- [Monitoring Web caching performance](#)

Web caching in security policies

Web caching can be applied to any HTTP traffic, including explicit web proxy traffic, accepted by any security policy by enabling web caching in that security policy.

Enabling web caching in a security policy cannot apply web caching to HTTPS traffic. To apply web caching to HTTPS traffic you need to create a WAN optimization rule.

Web Caching in a security policy takes place before web caching in a WAN Optimization rule. So traffic accepted by a security policy that includes web caching will not be cached by the WAN optimization rule.

By default FortiOS assumes HTTP traffic uses TCP port 80. So web caching in a security policy caches all HTTP traffic accepted by the policy on TCP port 80. If you want to cache HTTP traffic on other ports, you can enable UTM for the security policy and configure a protocol options profile to that looks for HTTP traffic on other TCP ports or on multiple ports.

If you set the HTTP port to 0 (for auto detection) in a protocol options profile, FortiOS looks for HTTP traffic on any port. In this case, HTTP traffic is detected by the IPS. Setting the HTTP port to 0 in a protocol options profile is not compatible with web caching. If you set the HTTP port to 0, web caching only caches HTTP traffic on port 80.

You can add web caching to a security policy to:

- Cache Internet HTTP traffic for users on an internal network to reduce Internet bandwidth use. Do this by selecting the web cache option for security policies that allow users on the internal network to browse web sites on the Internet.
- Reduce the requirements of a public facing web server by caching objects on the FortiGate unit to offload processing from the web server. A reverse proxy with web caching configuration. Do this by selecting the web cache option for security policy that allows users on the Internet to connect to the web server.
- Cache outgoing explicit web proxy traffic when the explicit proxy is used to proxy users in an internal network who are connecting to the web servers on the Internet. Do this by selecting the web cache option for explicit web proxy security policies that allow users on the internal network to browse web sites on the Internet.


Example: Web caching of Internet content for users on an internal network

This example describes how to configure web caching for users on a private network connecting to the Internet. This example uses web caching in a security policy to cache HTTP traffic.

Network topology and assumptions

This example includes a client network with subnet address 10.31.101.0 connecting to web servers on the Internet (Figure 324). In this basic example, all of the users on the private network access the Internet through a single general security policy on the FortiGate unit that accepts all sessions connecting to the Internet. Web caching is just added to this security policy.

Initially, UTM is not selected so the example caches all HTTP traffic on TCP port 80. The example also describes how to configure the security policy to cache HTTP traffic on port 80 and 8080 by adding a protocol options profile that looks for HTTP traffic on TCP ports 80 and 8080.

Figure 324: Example web caching topology

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Add web caching to the security policy that all users on the private network use to connect to the Internet.
- 2 Add a protocol options profile to look for HTTP traffic on ports 80 and 8080 and add this protocol options profile to the security policy to enable web caching HTTP traffic on ports 80 and 8080.

If you perform any additional actions between procedures, your configuration may have different results.

Configuration Steps - web-based manager

Use the following steps to configure the example configuration from the FortiGate web-based manager.

To add web caching to a security policy

- 1 Go to *Policy > Policy > Policy* and add a security policy that allows all users on the internal network to access the Internet.

Source Interface/Zone	Internal
Source Address	all
Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
Enable web cache	Select
NAT	Enable NAT

- 2 Select OK to save the security policy.
- 3 If required, adjust the position of the security policy in the internal > wan1 policy list to make sure all outgoing connections to the Internet use this policy

To cache HTTP traffic on port 80 and 8080

- 1 Go to *Policy > Policy > Protocol Options* and either add a new protocol options profile or edit the default profile.
- 2 Change the HTTP settings of the protocol options profile to look for HTTP traffic on ports 80 and 8080:

Port	80, 8080
-------------	----------

Adjust other settings as required or leave them as the defaults.

- 3 Edit the security policy created in the previous procedure. Select UTM and set Protocol Options to the protocol options profile you added or changed.

Configuration Steps - CLI

Use the following steps to configure the example configuration from the FortiGate CLI.

To add web caching to a security policy

- 1 Enter the following command to add a security policy that allows all users on the internal network to access the Internet.

```
config firewall policy
  edit 0
    set srcintf internal
    set srcaddr all
    set dstintf wan1
    set dstintf all
    set schedule always
    set service ANY
    set action accept
    set webcache enable
    set nat enable
  end
```

- 2 If required, adjust the position of the security policy in the internal > wan1 policy list to make sure all outgoing connections to the Internet use this policy

To cache HTTP traffic on port 80 and 8080

- 1 Go to *Policy > Policy > Protocol Options* and either add a new protocol options profile or edit the default profile.

- 1 Enter the following command to add a protocol option profile that looks for HTTP traffic on ports 80 and 8080.

```
config firewall profile-protocol-options
  edit custom-HTTP-proto
    config http
      set port 80 8080
    end
  end
```

- 2 Enter the following command to add the protocol options profile to the security policy:

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options custom-HTTP-proto
  end
```


Web Caching only WAN optimization

You can use Web Cache Only WAN optimization to cache web pages from any web server. In a Web Cache Only configuration, only one FortiGate unit is involved. All traffic between a client network and one or more web servers is intercepted by a Web Cache Only WAN optimization rule. This rule causes the FortiGate unit to cache pages from the web servers on the FortiGate unit and makes the cached pages available to users on the client network. Web cache only WAN optimization can be configured for standard and reverse web caching.

In a standard web caching configuration, the FortiGate unit caches pages for users on a client network. The FortiGate unit is installed between the client network and the WAN or Internet, and the web server or servers are located elsewhere on the WAN or Internet. See [“Example: Web Cache Only WAN optimization” on page 2739](#) for an example of this configuration.

You can also create a reverse proxy web caching configuration where the FortiGate unit is dedicated to providing web caching for a single web server or server farm. In this second configuration, the FortiGate unit is installed between the server network and the WAN or Internet, and users are located elsewhere on the WAN or Internet. See [“Example: SSL offloading and reverse proxy web caching for an Internet web server using static one-to-one virtual IPs” on page 2789](#) for an example of this configuration.

WAN optimization rule order affects Web Cache Only rules in the same way as other WAN optimization rules. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).




Since only one FortiGate unit is involved in a Web Cache Only configuration, you do not need to change the WAN optimization peer configuration.

Example: Web Cache Only WAN optimization

This example describes how to configure web caching for users in a client network connecting to a web server network across a WAN.

Network topology and assumptions

This example includes a client network with subnet address 172.20.120.0 connecting to web servers on a network with subnet address 192.168.10.0. Only the communication between the client network and the web server network using Port 80 is to be cached, so the Web Cache Only WAN optimization rule includes the IP addresses of the networks and the *Port* is set to 80. As well, the security policy used in this example includes the addresses of the client and sever subnets instead of more general security addresses.

Figure 325: Example Web Cache Only topology

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Add firewall addresses and a security policy that accepts traffic to be optimized to the FortiGate unit.
- 2 Add a Web Cache Only WAN optimization rule to the FortiGate unit.

If you perform any additional actions between procedures, your configuration may have different results.

Configuring Web Cache Only WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the FortiGate unit web-based manager.

To add the firewall addresses and security policy

- 1 Go to *Firewall Objects > Address > Address* and select *Create New* to add the firewall address for the client network:

Address Name	Client_Net
Type	Subnet/IP Range
Subnet / IP Range	172.20.120.*
Interface	Any

- 2 Add the firewall address for the web server network:

Address Name	Web_Server_Net
Type	Subnet/IP Range

Subnet / IP Range	192.168.10.*
Interface	Any

- 3 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts traffic to be web cached:

Source Interface/Zone	port1
Source Address	Client_Net
Destination Interface/Zone	port2
Destination Address	Web_Server_Net
Schedule	always
Service	HTTP
Action	ACCEPT

To add a Web Cache Only WAN optimization rule

- 1 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.
- 2 Select *Web Cache Only*.
- 3 Configure the Web Cache Only rule:

Mode	Web Cache Only
Source	172.20.120.*
Destination	192.168.10.*
Port	80 Usually you would set the port to 80 to cache normal HTTP traffic. But you can change the Port to a different number (for example 8080) or to a port number range so that the FortiGate unit provides web caching for HTTP traffic using other ports.
Transparent Mode	Select Transparent Mode
Enable SSL	Do not select Enable SSL. In this example SSL offloading is disabled. For an example of a reverse proxy Web Cache Only configuration that also includes SSL offloading, see “Example: SSL offloading for a WAN optimization tunnel” on page 2786.

- 4 Select *OK*.
The rule is added to the bottom of the WAN optimization list.
- 5 If required, use the *Move To* icon to move the rule to a different position in the list.
The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching”](#) on page 2708 and [“Moving a rule to a different position in the rule list”](#) on page 2709.

Configuring Web Cache Only WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the FortiGate unit CLI.

To add the firewall addresses and security policy

- 1 Add the firewall address for the client network:

```
config firewall address
  edit Client_Net
    set type iprange
    set start-ip 172.20.120.0
    set end-ip 172.20.120.255
  end
```

2 Add the firewall address for the web server network:

```
config firewall address
  edit Web_Server_Net
    set type iprange
    set start-ip 192.168.10.0
    set end-ip 192.168.10.255
  end
```

3 Add a security policy that accepts traffic to be web cached:

```
config firewall policy
  edit 2
    set srcintf port1
    set dstintf port2
    set srcaddr Client_Net
    set dstaddr Web_Server_Net
    set action accept
    set service HTTP
    set schedule always
  end
end
```

To add a Web Cache Only WAN optimization rule

1 Add the following Web Cache Only rule:

```
config wanopt rule
  edit 2
    set mode webcache-only
    set src-ip 172.20.120.0-172.20.120.255
    set dst-ip 192.168.10.0-192.168.10.255
    set port 80
    set peer Peer_Fgt_2
  end
```

Accept default settings for transparent (enable), status (enable), ssl (disable), unknown-http-version (tunnel), and tunnel-non-http (disable).



In this example, SSL offloading is disabled. For an example of a reverse proxy Web Cache Only configuration that also includes SSL offloading, see [“Example: SSL offloading for a WAN optimization tunnel”](#) on page 2786.

2 If required, use the `move` command to move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching”](#) on page 2708 and [“Moving a rule to a different position in the rule list”](#) on page 2709.

Testing and troubleshooting the configuration

To test the configuration, attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network, browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the user network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring in *WAN Opt. & Cache > Monitor > Monitor*. If WAN optimization has been forwarding the traffic, the WAN optimization monitor should show the HTTP protocol that has been optimized and the reduction rate in WAN bandwidth usage.

If you cannot connect, try the following to diagnose the problem:

- Review your configuration and make sure all details, such as address ranges, peer names and IP addresses, are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include UTM options. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy
- Check routing on the FortiGate units and on the user and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the user network must allow packets destined for the web server network to be received by the client side FortiGate unit, and packets from the server side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command on the client-side FortiGate unit to display WAN optimization tunnel protocol statistics. The `http tunnel` and `tcp tunnel` parts of the command output below shows that WAN optimization has been processing HTTP packets. If the `http bytes in` and `bytes out` fields are zero, then WAN optimization is not accepting HTTP packets.

```
get test wad 11
wad tunnel protocol stats:
  http tunnel
    bytes_in=1749865 bytes_out=25926
  ftp tunnel
    bytes_in=0 bytes_out=0
  cifs tunnel
    bytes_in=0 bytes_out=0
  mapi tunnel
    bytes_in=0 bytes_out=0
  tcp tunnel
    bytes_in=0 bytes_out=0
  maintenance tunnel
    bytes_in=0 bytes_out=0
```

You can use the following command to display information about the WAN optimization web cache daemon. The command will only display information if the web cache daemon is running and the statistics displayed show the number of open connections and other indications of activity:

```
diagnose wacs stats
Disk 0 /Internal-2B6375792136C707/wa_cs
```

```
Current number of open connections: 2
Number of terminated connections: 7
Number of requests -- Adds: 206 (0 repetitive keys),
Lookups: 860, Conflict incidents: 0
Percentage of missed lookups: 88.49
Communication is blocked for 0 client(s)
Disk usage: 5196 KB (11%)
```

Web caching for active-passive WAN optimization

You add web caching support to the passive or server side of an active-passive WAN optimization configuration. Web pages are cached on the server-side FortiGate unit so you should also select *Enable Byte Caching* for optimum WAN optimization performance.


For web caching to work, the WAN optimization tunnel must accept HTTP (and optionally HTTPS) traffic. To do this, the active rule on the client side must include the ports used for HTTP (and HTTPS) traffic. Set *Protocol* to *HTTP* to perform protocol optimization of the HTTP traffic. You can also enable SSL offloading and secure tunneling, as well as add an authentication group.

Example: Active-passive Web Caching

This example describes how to configure active-passive web caching for users in a client network connecting to a web server network across a WAN.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called *Client_Side* with a WAN IP address of 172.10.10.1 in front of a user network with IP address 172.20.120.0. The server-side FortiGate unit is called *Server_Side* and has a WAN IP address of 172.20.20.1. This server-side unit is in front of a web server network with IP address 192.168.10.0. Web caching is enabled on the server-side FortiGate unit.

Figure 326: Example active-passive web cache topology

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the client-side FortiGate unit by adding peers, a security policy that accepts traffic to be optimized, and an active WAN optimization rule.
- 2 Configure the server-side FortiGate unit by adding peers and a passive WAN optimization rule that includes web caching.

If you perform any additional actions between procedures, your configuration may have different results.

Configuring active-passive web caching - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure the client-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client FortiGate unit:

Local Host ID	Client_Side
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

Peer Host ID	Server_Side
IP Address	172.20.20.1

- 4 Select *OK*.

- 5 Go to *Policy > Policy > Policy* and add a security policy that accepts traffic to be web cached:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 6 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.

- 7 Configure the rule:

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	1-65535
Auto-Detect	Active
Protocol	HTTP
Transparent Mode	Select Transparent Mode
Enable Byte Caching	Select Enable Byte Caching

- 8 Select *OK*.

The rule is added to the bottom of the WAN optimization list.

- 9 If required, use the *Move To* icon to move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	Server_Side
----------------------	-------------

- 2 Select *Apply* to save your setting.

- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

Peer Host ID	Client_Side
IP Address	172.10.10.1

- 4 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.

- 5 Configure the passive web cache rule:

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	1-65535
Auto-Detect	Passive
Enable Web Cache	Select

- 6 Select *OK*.

The rule is added to the bottom of the WAN optimization rule list.

- 7 If required, use the *Move To* icon to move the rule to a different position in the list.
For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

Configuring active-passive web caching - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

- 1 Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
  set host-id Client_Side
end
```

- 2 Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Server_Side
  set ip 172.20.20.1
end
```

- 3 Add a security policy to the server-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 23
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
end
```

- 4 Configure the following active rule:

```
config wanopt rule
  edit 2
    set auto-detect active
    set src-ip 172.20.120.0-172.20.120.255
    set dst-ip 192.168.10.0-192.168.10.255
    set port 1-65535
```

```

    set proto http
end

```

Accept default settings for `transparent (enable)`, `status (enable)`, `mode (full)`, `byte-caching (enable)`, `ssl (disable)`, `secure-tunnel (disable)`, `auth-group (null)`, `unknown-http-version (tunnel)`, and `tunnel-non-http (disable)`.

- 5 If required, use the `move` command to move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Add the Local Host ID to the server-side FortiGate configuration:

```

config wanopt settings
    set host-id Server_Side
end

```

- 2 Add the client-side Local Host ID to the server-side peer list:

```

config wanopt peer
    edit Client_Side
        set ip 172.10.10.1
    end

```

- 3 Add the following passive web cache rule:

```

config wanopt rule
    edit 5
        set auto-detect passive
        set src-ip 172.20.120.0-172.20.120.255
        set dst-ip 192.168.10.0-192.168.10.255
        set port 1-65535
        set webcache enable
    end

```

Accept default settings for `status (enable)` and `mode (full)`.

- 4 If required, use the `move` command to move the rule to a different position in the list so that the tunnel request from the client-side FortiGate unit matches with this rule.

For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

Web caching for peer-to-peer WAN optimization

In a peer-to-peer web caching configuration, you create a peer-to-peer WAN optimization rule on the client-side FortiGate unit and include the peer host ID of the server-side FortiGate unit. In the rule, you set *Auto-Detect* to *Off* and select *Enable Web Cache*. By using this rule, the client-side FortiGate unit can create a WAN optimization tunnel only with the peer that is added to the rule.

In a peer-to-peer configuration, you do not have to add a rule to the server-side FortiGate unit. If the server-side FortiGate unit peer list contains the client FortiGate unit, the server FortiGate unit accepts WAN optimization tunnel connections from the client FortiGate unit and the two units can form a WAN optimization tunnel. The server-side FortiGate unit uses the settings in the rule added to the client-side FortiGate unit.

For web caching to work, the WAN optimization tunnel must allow HTTP (and optionally HTTPS) traffic. To do this, the WAN optimization rule must include the ports used for HTTP (and HTTPS) traffic. Set *Protocol* to *HTTP* to perform protocol optimization of the HTTP traffic. You can also enable WAN optimization transparent mode, byte caching, SSL offloading, and secure tunneling, as well as add an authentication group.


Example: Peer-to-peer web caching

This example describes how to configure peer-to-peer web caching for users in a client network connecting to a web server network across a WAN.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called *Client_Side* with a WAN IP address of 172.10.10.1 in front of a user network with IP address 172.20.120.0. The server-side FortiGate unit is called *Server_Side* and has a WAN IP address of 172.20.20.1. This server-side unit is in front of a web server network with IP address 192.168.10.0. Web caching is enabled on the server-side FortiGate unit.

Figure 327: Example peer-to-peer web cache topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the client-side FortiGate unit by adding peers, a security policy that accepts traffic to be optimized, and a peer-to-peer WAN optimization rule that includes web caching.
- 2 Configure the server-side FortiGate unit.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring peer-to-peer web caching - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure the client-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client FortiGate unit:

Local Host ID	Client_Side
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the server-side FortiGate unit:

Peer Host ID	Server_Side
IP Address	192.168.30.12

- 4 Select *OK*.
- 5 Go to *Policy > Policy > Policy* and add a security policy that accepts traffic to be web cached:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT

- 6 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.
- 7 Configure the rule:

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	80
Auto-Detect	Off
Protocol	HTTP
Peer	Server_Side
Enable Web Cache	Select
Transparent Mode	Select
Enable Byte Caching	Select

- 8 Select *OK*.
The rule is added to the bottom of the WAN optimization list.
- 9 If required, use the *Move To* icon to move the rule to a different position in the list.
The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server FortiGate unit:

Local Host ID	Server_Side
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the client-side FortiGate unit:

Peer Host ID	Client_Side
IP Address	172.20.34.12

- 4 Select *OK*.

Configuring peer-to-peer web caching - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

- 1 Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
  set host-id Client_Side
end
```

- 2 Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Server_Side
  set ip 192.168.30.12
end
```

- 3 Add a security policy to the server-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 23
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
end
```

- 4 Configure the following active rule:

```
config wanopt rule
  edit 5
    set auto-detect off
    set src-ip 172.20.120.*
    set dst-ip 192.168.10.*
    set port 80
    set proto http
    set peer Server_Side
    set web cache enable
```

```
end
```

Accept default settings for `transparent (enable)`, `status (enable)`, `mode (full)`, `byte-caching (enable)`, `ssl (disable)`, `secure-tunnel (disable)`, `auth-group (null)`, `unknown-http-version (tunnel)`, and `tunnel-non-http (disable)`.

- 5 If required, use the `move` command to the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the server-side FortiGate unit

- 1 Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
  set host-id Server_Side
end
```

- 2 Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
  edit Client_Side
  set ip 172.20.34.12
end
```

Exempting web sites from web caching

You may want to exempt some URLs from web caching for a number of reasons. For example, if your users access websites that are not compatible with FortiGate web caching you can add the URLs of these web sites to the web caching exempt list. All traffic accepted by WAN optimization and the explicit web proxy for these websites will not be cached. You can add URLs and numeric IP addresses to the web cache exempt list. When enabled the web cache exempt list applies to web caching in security policies and in WAN optimization rules.

Enter the following command to add `www.example.com` to the web cache exempt list.

```
config wanopt webcache
  set cache-exemption enable
  config cache-exemption-list
    edit 1
      set url-pattern www.example.com
      set status enable
    end
  end
```

Enter the following command to enable the web cache exempt list and add two IP address URLs and a web page URL to the list.

```
config wanopt webcache
  set explicit enable
  set cache-exemption enable
  config cache-exemption-list
    edit 1
      set url-pattern "192.168.1.121"
    next
    edit 2
      set url-pattern "google.com/test123/321"
```

```
next
edit 3
    set url-pattern "1.1.1.1"
next
end
end
```

You can also add URLs to the web cache exempt list by going to *WAN Opt. & Cache > Cache > Exempt List*. However, the web cache exempt feature must be enabled from the CLI.

Changing web cache settings

In most cases, the default settings for the WAN optimization web cache are acceptable. However, you may want to change them to improve performance or optimize the cache for your configuration. To change these settings, go to *WAN Opt. & Cache > Cache > Settings*.

From the FortiGate CLI, you can use the `config wanopt webcache` command to change these WAN optimization web cache settings.



For more information about many of these web cache settings, see [RFC 2616](#).

Always revalidate

Select to always revalidate requested cached objects with content on the server before serving them to the client.

Max cache object size

Set the maximum size of objects (files) that are cached. The default size is 512000 KB and the range is 1 to 4294967 KB. This setting determines the maximum object size to store in the web cache. Objects that are larger than this size are still delivered to the client but are not stored in the FortiGate web cache.

Negative response duration

Set how long in minutes that the FortiGate unit caches error responses from web servers. If error responses are cached, then subsequent requests to the web cache from users will receive the error responses regardless of the actual object status.

The default is 0, meaning error responses are not cached. The content server might send a client error code (4xx HTTP response) or a server error code (5xx HTTP response) as a response to some requests. If the web cache is configured to cache these negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes.

Fresh factor

Set the fresh factor as a percentage. The default is 100, and the range is 1 to 100%. For cached objects that do not have an expiry time, the web cache periodically checks the server to see if the objects have expired. The higher the *Fresh Factor* the less often the checks occur.

For example, if you set the *Max TTL* value and *Default TTL* to 7200 minutes (5 days) and set the *Fresh Factor* to 20, the web cache checks the cached objects 5 times before they expire, but if you set the *Fresh Factor* to 100, the web cache will check once.

Max TTL

The maximum amount of time (Time to Live) an object can stay in the web cache without the cache checking to see if it has expired on the server. The default is 7200 minutes (120 hours or 5 days) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Min TTL

The minimum amount of time an object can stay in the web cache before the web cache checks to see if it has expired on the server. The default is 5 minutes and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Default TTL

The default expiry time for objects that do not have an expiry time set by the web server. The default expiry time is 1440 minutes (24 hours) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Proxy FQDN

The fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server. This field is for information only and can be changed from the explicit web proxy configuration.

Max HTTP request length

The maximum length of an HTTP request that can be cached. Larger requests will be rejected. This field is for information only and can be changed from the explicit web proxy configuration.

Max HTTP message length

The maximum length of an HTTP message that can be cached. Larger messages will be rejected. This field is for information only and can be changed from the explicit web proxy configuration.

Ignore

Select the following options to ignore some web caching features.

- If-modified-since

By default, if the time specified by the if-modified-since (IMS) header in the client's conditional request is greater than the last modified time of the object in the cache, it is a strong indication that the copy in the cache is stale. If so, HTTP does a conditional GET to the Overlay Caching Scheme (OCS), based on the last modified time of the cached object.

Enable ignoring if-modified-since to override this behavior.

- HTTP 1.1 conditionals

HTTP 1.1 provides additional controls to the client over the behavior of caches toward stale objects. Depending on various cache-control headers, the FortiGate unit can be forced to consult the OCS before serving the object from the cache. For more information about the behavior of cache-control header values, see [RFC 2616](http://tools.ietf.org/html/rfc2616).

Enable ignoring HTTP 1.1 Conditionals to override this behavior.

- **Pragma-no-cache**

Typically, if a client sends an HTTP GET request with a pragma no-cache (PNC) or cache-control no-cache header, a cache must consult the OCS before serving the content. This means that the FortiGate unit always re-fetches the entire object from the OCS, even if the cached copy of the object is fresh.

Because of this behavior, PNC requests can degrade performance and increase server-side bandwidth utilization. However, if you enable ignoring Pragma-no-cache, then the PNC header from the client request is ignored. The FortiGate unit treats the request as if the PNC header is not present.

- **IE Reload**

Some versions of Internet Explorer issue Accept / header instead of Pragma no-cache header when you select *Refresh*. When an Accept header has only the / value, the FortiGate unit treats it as a PNC header if it is a type-N object.

Enable ignoring IE reload to cause the FortiGate unit to ignore the PNC interpretation of the Accept / header.

Cache Expired Objects

Applies only to type-1 objects. When this option is selected, expired type-1 objects are cached (if all other conditions make the object cacheable).

Revalidated Pragma-no-cache

The pragma-no-cache (PNC) header in a client's request can affect how efficiently the FortiGate unit uses bandwidth. If you do not want to completely ignore PNC in client requests (which you can do by selecting to ignore Pragma-no-cache, above), you can nonetheless lower the impact on bandwidth usage by selecting *Revalidate Pragma-no-cache*.

When you select *Revalidate Pragma-no-cache*, a client's non-conditional PNC-GET request results in a conditional GET request sent to the OCS if the object is already in the cache. This gives the OCS a chance to return the 304 Not Modified response, which consumes less server-side bandwidth, because the OCS has not been forced to otherwise return full content.

By default, *Revalidate Pragma-no-cache* is disabled and is not affected by changes in the top-level profile.

Most download managers make byte-range requests with a PNC header. To serve such requests from the cache, you should also configure byte-range support when you configure the *Revalidate pragma-no-cache* option.

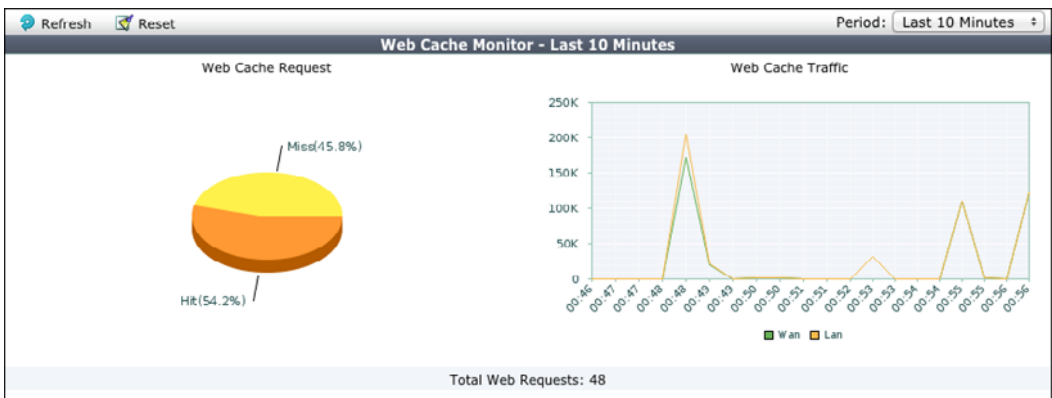
Monitoring Web caching performance

The web cache monitor shows the percentage of web cache requests that retrieved content from the cache (hits) and the percentage that did not receive content from the cache (misses). A higher the number of hits usually indicates that the web cache is being more effective at reducing WAN traffic.

The web cache monitor also shows a graph of web traffic on the WAN and LAN. A lower WAN line on the graph indicates the web cache is reducing traffic on the WAN. The web cache monitor also displays the total number of web requests processed by the web cache.

To view the web cache monitor, go to *WAN Opt. & Cache > Monitor > Cache Monitor*.

Figure 328: Web cache monitor





Advanced configuration example

This chapter contains an advanced WAN optimization configuration example that combines many of the concepts described in the previous chapters of this document. The configuration example described here includes active-passive rules, web caching, policy routes for out-of-path WAN optimization, and multiple VDOMs with inter-VDOM routing to apply virus scanning (and optionally other UTM features) to traffic before it is optimized.

Out-of-path WAN optimization with inter-VDOM routing

This example describes how to configure out-of-path WAN optimization to optimize web browsing and FTP file transfers between a client network and a server network.

Network topology and assumptions

The client network connects to the Internet through a FortiGate-300A unit, and the server network connects to the Internet through a cluster of two FortiGate-1000A units.

Adding in-path WAN optimization requires replacing these FortiGate units with models that support WAN optimization or adding new FortiGate units in the data path. In either of these in-path configurations, the optimizing FortiGate units would also be required to support all traffic on the data path plus provide WAN optimization.

The out-of-path topology shown in [Figure](#) offloads WAN optimization to out-of-path FortiGate units that only process sessions to be optimized. The topology includes a FortiGate-311B unit installed at the client network and a single FortiGate-620B unit installed at the server network.




The FortiGate-620B unit is installed at the server network because other client networks also use it for WAN optimization. The configuration for those other client networks is not described in this example.

The client-side FortiGate-300A unit uses policy routing to offload WAN optimization of HTTP and FTP sessions by re-directing all HTTP and FTP sessions to the FortiGate-311B unit. The FortiGate-311B and 620B units work together to apply web caching, byte caching, and HTTP and FTP protocol optimization to HTTP and FTP sessions. The WAN optimization tunnel between the 311B and the 620B operates in Transparent mode. The FortiGate-311B unit also web caches all Internet HTTP traffic from the client network.

The client-side FortiGate-311B unit also applies virus scanning (and optionally other UTM features) to the HTTP and FTP traffic. To do this, the FortiGate-311B unit is configured for multiple VDOM operation. A new VDOM named Wanopt is added to the FortiGate-311B. HTTP and FTP sessions are received by the “root” VDOM. Security policies in the root VDOM accept HTTP and FTP sessions and apply virus scanning (and optionally other UTM features) to them. To preserve the source addresses of the HTTP and FTP sessions, NAT is not enabled for these policies.

The sessions are then routed through an inter-VDOM link to the Wanopt VDOM. The Wanopt VDOM includes security policies that accept the HTTP and FTP sessions and WAN optimization rules that apply WAN optimization and web caching to the sessions.

Figure 329: Out-of-path WAN optimization

The server-side FortiGate-620B unit includes a passive WAN optimization rule that accepts WAN optimization tunnel requests from the FortiGate-311B unit. Only one passive rule is required on the FortiGate-620B unit. The FortiGate-620B unit also forwards sessions to the server-side FortiGate-1000A cluster which forwards them to the server network.

WAN optimization is operating in Transparent mode, so the packets from the client network include their client network source IP addresses. To preserve these source IP addresses, the security policies on the FortiGate-1000A cluster that accept the sessions from the FortiGate-620B unit should not apply NAT. If the security policies were to apply NAT, the client network addresses would be replaced with the port1 IP address of the FortiGate-1000A cluster and the client network source IP addresses would be lost.

The optimizing FortiGate units operate in NAT/Route mode and are directly connected to the Internet. This configuration requires two Internet connections and two Internet IP addresses for each network. (Reminder: All of the example IP addresses shown in [Figure](#) are private IP addresses because all Fortinet documentation examples use only private IP addresses.) If these extra Internet IP addresses are not available, you can install a router between the WAN and the FortiGate units or install the optimizing FortiGate units out of path on the private networks and configure routing on the private networks to route HTTP and FTP sessions to the optimizing FortiGate units.

Configuration steps

This example is divided into client-side and the server-side steps, as configured through the web-based manager and the CLI. Use either method, but for best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

This example includes the following sections:

- [“Client-side configuration steps - web-based manager”](#) on page 2759
- [“Server-side configuration steps - web-based manager”](#) on page 2766

- “Client-side configuration steps - CLI” on page 2769
- “Server-side configuration steps - CLI” on page 2776

Client-side configuration steps - web-based manager

This section describes the configuration steps required to redirect HTTP and FTP sessions from the client-side FortiGate-300A unit and to configure the client-side FortiGate-311B unit to optimize HTTP and FTP sessions to the server network and to apply web caching to all other HTTP sessions from the client network.

The section breaks down the client-side configuration into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the FortiGate-300A unit to redirect all HTTP and FTP sessions to the FortiGate-311B unit.
- 2 Configure the FortiGate-311B unit for multiple VDOM operation and add an inter-VDOM link.
- 3 Configure routing for the FortiGate-311B root VDOM.
- 4 Add security policies to the FortiGate-311B root VDOM to accept HTTP and FTP sessions received at port1 and destined for Vlink0, and apply virus scanning (and optionally other UTM features).
- 5 Configure routing for the FortiGate-311B Wanopt VDOM.
- 6 Add security policies to the FortiGate-311B Wanopt VDOM to accept HTTP and FTP sessions received at the Vlink1 interface of the inter-VDOM link and destined for port10.
- 7 Configure peers for the FortiGate-311B Wanopt VDOM.
- 8 Add WAN optimization rules for HTTP and FTP to the FortiGate-311B Wanopt VDOM.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

To configure the FortiGate-300A unit to redirect all HTTP and FTP sessions to the FortiGate-311B unit

- 1 Go to *System > Network > Interface*, edit *port4*, and set the port4 IP address to 172.10.10.1/24.
- 2 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that allows all port5 to port4 HTTP sessions:

Source Interface/Zone	port5
Source Address	all
Destination Interface/Zone	port4
Destination Address	all
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select

Configure other policy settings that you may require.

- 3 Select *Create New* to add a security policy that allows all port5 to port4 FTP sessions:

Source Interface/Zone	port5
Source Address	all
Destination Interface/Zone	port4
Destination Address	all
Schedule	always
Service	FTP
Action	ACCEPT
NAT	Select

Configure other policy settings that you may require.

- 4 Select *OK*.
- 5 If required, use the *Move To* icon to change the order of the security policies.
Follow the normal rules for ordering security policies in the policy list. For example, move specific rules above general rules.
- 6 Go to *Router > Static > Policy Route* and select *Create New* to add a policy route to redirect HTTP traffic received at port5 to exit the FortiGate unit using port4. Set the gateway address of the route to 172.10.10.2 so that the HTTP sessions are directed to the FortiGate-311B port1 interface. For HTTP traffic, the protocol is 6 (TCP) and the destination port is 80:

Protocol	6
Incoming interface	port5
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0
Destination Ports	From 80 to 80
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)
Outgoing interface	port4
Gateway Address	172.10.10.2

- 7 Select *OK*.
- 8 Select *Create New* to add a policy route to redirect FTP traffic received at port5 to exit the FortiGate unit using port4. Set the gateway address of the route to 172.10.10.2 so that the HTTP sessions are directed to the FortiGate-311B port1 interface. For FTP traffic, the protocol is 6 (TCP) and the destination port is 21:

Protocol	6
Incoming interface	port5
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0
Destination Ports	From 21 to 21
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)

Outgoing interface	port4
Gateway Address	172.10.10.2

- 9 Select OK.

To configure the FortiGate-311B unit for multiple VDOM operation and add an inter-VDOM link

- 1 Go to *System > Status > Dashboard*.
- 2 In the *System Information* widget, select *Enable* beside *Virtual Domain* to enable multiple VDOM operation and log back in to the web-based manager.
- 3 Go to *System > VDOM* and select *Create New* to add a new virtual domain named *Wanopt*.
- 4 Select OK twice to add the *Wanopt* VDOM with default resource limits.
- 5 Go to *System > Network*, edit the *port10* interface, and configure the following settings to add the port10 interface to the *Wanopt* VDOM:

Virtual Domain	Wanopt
Addressing Mode	Manual
IP/Netmask	10.10.10.2/24

Configure other settings that you may require.

- 6 Select OK.
- 7 Select *Create New > VDOM Link* and add an inter-VDOM link with the following settings:

Name	Vlink
Interface #0	
Virtual Domain	root
IP/Netmask	172.1.1.1/24
Interface #1	
Virtual Domain	Wanopt
IP/Netmask	172.1.1.2/24

- 8 Select OK.

To configure routing for the FortiGate-311B root VDOM

- 1 Log in to the root VDOM.
- 2 Go to *Router > Static* and select *Create New* to add a default route. The destination of the default route is the inter-VDOM link interface in the root VDOM. The gateway of the default route is the IP address of the inter-VDOM link interface in the *Wanopt* VDOM. The result is the default route sends all traffic out the inter-VDOM link and into the *Wanopt* VDOM:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	Vlink0
Gateway	172.1.1.2
Distance	10

- 3 Select OK.

- 4 Select *Create New* to add a route to send return traffic from the server network destined for the client network out the port1 interface to the port4 interface of the FortiGate-300A which has IP address 172.10.10.1:

Destination IP/Mask	172.20.120.0/24
Device	port1
Gateway	172.10.10.1
Distance	10

- 5 Select *OK*.

To add security policies to the FortiGate-311B root VDOM to accept HTTP and FTP sessions received at port1 destined for Vlink0 and apply virus scanning (and optionally other UTM features)

- 1 Log in to the root VDOM.
- 2 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts HTTP sessions received at port1 destined for Vlink0 and applies virus scanning and other UTM features:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	Vlink0
Destination Address	all
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Do not select. To preserve the source addresses of the HTTP sessions, NAT should not be enabled for this policy.
UTM	Select UTM, select a protocol options profile and select an antivirus profile. Optionally select other UTM profiles.

Configure other policy settings that you may require. You can also use more specific security addresses or add one security policy that accepts both FTP and HTTP traffic.

- 3 Select *OK*.

- 4 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts FTP sessions received at port1 and destined for Vlink0 and applies virus scanning and other UTM features to them:

Source Interface/Zone	port1
Source Address	all
Destination Interface/Zone	Vlink0
Destination Address	all
Schedule	always
Service	FTP
Action	ACCEPT
NAT	Do not select. To preserve the source addresses of the FTP sessions, NAT should not be enabled for this policy.
UTM	Select UTM, select a protocol options profile and select an antivirus profile. Optionally select other UTM profiles.

Configure other policy settings that you may require. You can also use more specific security addresses or add one security policy that accepts both FTP and HTTP traffic.

- 5 Select *OK*.

To configure routing for the FortiGate-311B Wanopt VDOM

- 1 Log in to the Wanopt VDOM.
- 2 Go to *Router > Static* and select *Create New* to add a default route. The destination of the default route is the port10 interface. The gateway of the default route is the next hop router that the port10 interface connects with:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port10
Gateway	(next hop router IP address)
Distance	10

- 3 Select *OK*.
- 4 Select *Create New* to add a route to send return traffic from the server network destined for the client network out the Vlink1 interface to the Vlink0 interface in the root VDOM, which has the IP address 172.1.1.2:

Destination IP/Mask	172.20.120.0/24
Device	Vlink1
Gateway	172.1.1.2
Distance	10

- 5 Select *OK*.

To add security policies to the FortiGate-311B Wanopt VDOM to accept HTTP and FTP sessions received at the Vlink1 interface of the inter-VDOM link and destined for port10

- 1 Log in to the Wanopt VDOM.

- 2 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts HTTP sessions received at Vlink1 and destined for port10:

Source Interface/Zone	Vlink1
Source Address	all
Destination Interface/Zone	port10
Destination Address	all
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select NAT is ignored for all HTTP sessions for the server network because these sessions are intercepted by a full optimization WAN optimization rule. However, HTTP sessions for the Internet are intercepted by the Web Cache Only rule, so source NAT is required for replies.
UTM	Do not select. Do not select UTM because you cannot apply UTM and WAN optimization to the same session in the same VDOM. UTM was applied to the session in the root VDOM.

Configure other settings that you may require.

- 3 Select *OK*.
- 4 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts FTP sessions received at Vlink1 and destined for port10:

Source Interface/Zone	Vlink1
Source Address	all
Destination Interface/Zone	port10
Destination Address	all
Schedule	always
Service	FTP
Action	ACCEPT
NAT	Select NAT is ignored for all FTP sessions for the server network because these sessions are intercepted by a full optimization WAN optimization rule. However, FTP sessions for the Internet are allowed to reach their destination, so source NAT is required for replies.
UTM	Do not select. Do not select UTM because you cannot apply UTM and WAN optimization to the same session in the same VDOM. UTM was applied to the session in the root VDOM.

Configure other settings that you may require.

- 5 Select *OK*.

To configure peers for the FortiGate-311B Wanopt VDOM

- 1 Log in to the Wanopt VDOM.
- 2 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the client-side FortiGate-311B unit:

Local Host ID	Client_Fgt
----------------------	------------

- 3 Select *Apply* to save your setting.
- 4 Select *Create New* and add a Peer Host ID and the *IP Address* for the server-side FortiGate-620B unit:

Peer Host ID	Server_Fgt
IP Address	10.20.20.2

- 5 Select *OK*.

To add WAN optimization rules for HTTP and FTP to the FortiGate-311B Wanopt VDOM

- 1 Log in to the Wanopt VDOM.
- 2 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule*.
- 3 Select *Create New* to add an active rule to optimize HTTP traffic from IP addresses on the Client network (172.20.120.0) with a destination address on the server network (192.168.10.0):

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	80
Auto-Detect	Active
Protocol	HTTP
Transparent Mode	Select
Enable Byte Caching	Select
Enable SSL	Do not select.
Enable Secure Tunnel	Do not select. For improved privacy you can select this option and add an authentication group to both optimizing FortiGate units.
Authentication Group	Do not select.

- 4 Select *OK*.
- 5 Select *Create New* to add an active rule to optimize FTP traffic from IP addresses on the Client network (172.20.120.0) with a destination address on the server network (192.168.10.0):

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	21

Auto-Detect	Active
Protocol	FTP
Transparent Mode	Select
Enable Byte Caching	Select
Enable SSL	Do not select.
Enable Secure Tunnel	Do not select. For improved privacy you can select this option and add an authentication group to both optimizing FortiGate units.
Authentication Group	Do not select.

- 6 Select *OK*.
- 7 Select *Create New* to add a rule to web cache HTTP traffic from IP addresses on the Client network (172.20.120.0) with any destination address:

Mode	Web Cache Only
Source	172.20.120.*
Destination	0.0.0.0
Port	80
Transparent Mode	Select
Enable SSL	Do not select.

- 8 Select *OK*.
- 9 If required, use the *Move To* icon to move the Web Cache Only rule below the full optimization HTTP and FTP rules in the list. The Web Cache Only rule should be below the full optimization rules because it will match all HTTP traffic and you need HTTP sessions with destination address 192.168.10.0 to match the full optimization HTTP rule.

Server-side configuration steps - web-based manager

This section describes the configuration steps required for the server-side FortiGate-620B unit to perform WAN optimization with the client-side FortiGate-311B unit and to send HTTP and FTP sessions to the server-side FortiGate-1000A cluster. This section also describes how to configure the FortiGate-1000A cluster to forward HTTP and FTP sessions from the client network to the server network.

The section breaks down the client-side configuration into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure routing for the FortiGate-620B unit.
- 2 Configure peers for the server-side FortiGate-620B unit.
- 3 Add a passive WAN optimization rule to the server-side FortiGate-620B unit.
- 4 Configure the FortiGate-1000A cluster to accept HTTP and FTP connections at port5 and forward them out port1 to the server network.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

To configure routing for the FortiGate-620B unit

- 1 Go to *Router > Static* and select *Create New* to add a default route. The destination of the default route is the port16 interface. The gateway of the default route is the next hop router that the port16 interface connects with:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port16
Gateway	(next hop router IP address)
Distance	10

- 2 Select *OK*.
- 3 Select *Create New* to add a route to send traffic for the server network out port1 to the port5 interface of the FortiGate-1000A cluster, which has the IP address 192.20.20.1:

Destination IP/Mask	192.168.10.0/24
Device	port1
Gateway	192.20.20.1
Distance	10

- 4 Select *OK*.

To configure peers for the server-side FortiGate-620B unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate-620B unit:

Local Host ID	Server_Fgt
----------------------	------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a Peer Host ID and the *IP Address* for the client-side FortiGate-311B unit:

Peer Host ID	Client_Fgt
IP Address	10.10.10.2

- 4 Select *OK*.

To add a passive WAN optimization rule to the server-side FortiGate-620B unit

You can add one passive WAN optimization rule to the server-side FortiGate-620B unit for both active rules on the FortiGate-311B unit. This rule can also allow the FortiGate-620B to perform WAN optimization with other client-side devices as long as the required Peer Host IDs are added to the FortiGate-620B configuration and to the client-side configurations.

- 1 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New* to add a passive rule that accepts any WAN optimization tunnel request:

Mode	Full Optimization
Source	0.0.0.0
Destination	192.168.10.*
Port	1-65535 You can also use a narrower port range such as 21-80 or add two rules, one with port set to 80 and one with port set to 21.

Auto-Detect	Passive
Enable Web Cache	Select

- 2 Select *OK*.
- 3 If required, use the *Move To* icon to move the rule to a different position in the list so that the tunnel request from the client-side FortiGate unit matches with this rule.
For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate-1000A cluster to accept HTTP and FTP connections at port5 and forward them out port1 to the server network

- 1 Go to *Firewall Objects > Address* and select *Create New* to add an address for the server network:

Address Name	Server_Net
Type	Subnet / IP Range
Subnet / IP Range	192.168.10.*
Interface	Any

- 2 Select *OK*.
- 3 Go to *Firewall Objects > Address* and select *Create New* to add an address for the client network:

Address Name	Client_Net
Type	Subnet / IP Range
Subnet / IP Range	172.20.120.*
Interface	Any

- 4 Select *OK*.
- 5 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts HTTP sessions at port5 destined for port1 and the server network:

Source Interface/Zone	port5
Source Address	Client_Net
Destination Interface/Zone	port1
Destination Address	Server_Net
Schedule	always
Service	HTTP

Action	ACCEPT
NAT	Do not select. WAN optimization is operating in Transparent mode so the packets from the client network include their client network source IP addresses. To preserve these source IP addresses the security policies on the FortiGate-1000A cluster that accept the sessions from the FortiGate- 620B unit should not apply NAT. If the policies were to apply NAT, the client network addresses would be replaced with the port1 IP address of the FortiGate-1000A cluster and the client network source IP addresses would be lost.

6 Select *OK*.

7 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts FTP sessions at port5 destined for port1 and the server network:

Source Interface/Zone	port5
Source Address	Client_Net
Destination Interface/Zone	port1
Destination Address	Server_Net
Schedule	always
Service	FTP
Action	ACCEPT
NAT	Do not select As described above, selecting NAT would cause the loss of client network source IP addresses.

8 Select *OK*.

Client-side configuration steps - CLI

This section describes the configuration steps required to redirect HTTP and FTP sessions from the client-side FortiGate-300A unit and to configure the client-side FortiGate-311B unit to optimize HTTP and FTP sessions to the server network and to apply web caching to all other HTTP sessions from the client network.

The section breaks down the client-side configuration into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the FortiGate-300A unit to redirect all HTTP and FTP sessions to the FortiGate-311B unit.
- 2 Configure the FortiGate-311B unit for multiple VDOM operation and add an inter-VDOM link.
- 3 Configure routing for the FortiGate-311B root VDOM.
- 4 Add security policies to the FortiGate-311B root VDOM to accept HTTP and FTP sessions received at port1 and destined for Vink0, and apply virus scanning (and optionally other UTM features).
- 5 Configure routing for the FortiGate-311B Wanopt VDOM.

- 6 Add security policies to the FortiGate-311B Wanopt VDOM to accept HTTP and FTP sessions received at the Vlink1 interface of the inter-VDOM link and destined for port10.
- 7 Configure peers for the FortiGate-311B Wanopt VDOM.
- 8 Add WAN optimization rules for HTTP and FTP to the FortiGate-311B Wanopt VDOM.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

To configure the FortiGate-300A unit to redirect all HTTP and FTP sessions to the FortiGate-311B unit

- 1 Set the FortiGate-300A port4 IP address to 172.10.10.1:

```
config system interface
  edit port4
    set ip 172.10.10.1/24
  end
end
```

- 2 Add a security policy that allows all port5 to port4 HTTP sessions:

```
config firewall policy
  edit 1
    set srcintf port5
    set dstintf port4
    set srcaddr all
    set dstaddr all
    set action accept
    set service HTTP
    set schedule always
    set nat enable
  end
end
```

Configure other policy settings that you may require. For example, you could add virus scanning (and optionally other UTM features).

- 3 Add a security policy that allows all port5 to port4 FTP sessions:

```
config firewall policy
  edit 2
    set srcintf port5
    set dstintf port4
    set srcaddr all
    set dstaddr all
    set action accept
    set service FTP
    set schedule always
    set nat enable
  end
end
```

Configure other policy settings that you may require.

- 4 If required, use the `move` command to change the order of the policies in the policy list.

Follow the normal rules for ordering security policies in the policy list. For example, move specific rules above general rules.

- 5 Add a policy route to redirect HTTP traffic received at port5 to exit the FortiGate unit using port4. Set the gateway address of the route to 172.10.10.2 so that the HTTP sessions are directed to the FortiGate-311B port1 interface. For HTTP traffic, the protocol is 6 (TCP) and the destination port is 80:

```
config router policy
  edit 1
    set protocol 6
    set input-device port5
    set output-device port4
    set src 0.0.0.0/0.0.0.0
    set dst 0.0.0.0/0.0.0.0
    set start-port 80
    set end port 80
    set gateway 172.10.10.2
  end
end
```

Accept default settings for `tos (0x00)` and `tos-mask (0x00)`.

- 6 Add a policy route to redirect FTP traffic received at port5 to exit the FortiGate unit using port4. Set the gateway address of the route to 172.10.10.2 so that the FTP sessions are directed to the FortiGate-311B port1 interface. For FTP traffic, the protocol is 6 (TCP) and the destination port is 21:

```
config router policy
  edit 1
    set protocol 6
    set input-device port5
    set output-device port4
    set src 0.0.0.0/0.0.0.0
    set dst 0.0.0.0/0.0.0.0
    set start-port 21
    set end port 21
    set gateway 172.10.10.2
  end
end
```

Accept default settings for `tos (0x00)` and `tos-mask (0x00)`.

To configure the FortiGate-311B unit for multiple VDOM operation and add an inter-VDOM link

- 1 Enable multiple VDOM operation and log back in to the web-based manager:

```
config system global
  set vdom-admin enable
end
```

- 2 Log back in to the CLI.

- 3 Add a new virtual domain named Wanopt.

```
config vdom
  edit Wanopt
end
```

- 4 Add the port10 interface to the Wanopt VDOM:

```
config global
  config system interface
    edit port10
      set vdom Wanopt
    end
  end
end
```

```

        set IP 10.10.10.2/24
    end
end

```

5 Add an inter-VDOM named Vlink and configure the Vlink0 and Vlink1 interfaces:

```

config global
    config system vdom-link
        edit Vlink
        end
    config system interface
        edit Vlink0
            set vdom root
            set ip 172.1.1.1/24
        next
        edit Vlink1
            set vdom Wanopt
            set ip 172.1.1.2/24
        end
    end
end

```

To configure routing for the FortiGate-311B root VDOM

- 1** Log in to the root VDOM from the CLI.
- 2** Add a default route. The destination of the default route is the inter-VDOM link interface in the root VDOM. The gateway of the default route is the IP address of the inter-VDOM link interface in the Wanopt VDOM. The result is the default route sends all traffic out the inter-VDOM link and into the Wanopt VDOM:

```

config router static
    edit 1
        set dst 0.0.0.0/0.0.0.0
        set device Vlink0
        set gateway 172.1.1.2
        set distance 10
    end

```

- 3** Add a route to send return traffic from the server network destined for the client network out the port1 interface to the port4 interface of the FortiGate-300A which has IP address 172.10.10.1:

```

config router static
    edit 2
        set dst 172.20.120.0/24
        set device port1
        set gateway 172.10.10.1
        set distance 10
    end

```

To add security policies to the FortiGate-311B root VDOM to accept HTTP and FTP sessions received at port1 and destined for Vlink0 and apply virus scanning and optionally other UTM features)

- 1** Log in to the root VDOM from the CLI.
- 2** Add a security policy that accepts HTTP sessions received at port1 and applies virus scanning to them:

```

config firewall policy
    edit 20

```

```

set srcintf port1
set dstintf Vlink0
set srcaddr all
set dstaddr all
set action accept
set service HTTP
set schedule always
set utm-status enable
set profile-protocol-options default
set av-profile scan
end

```



To preserve the source addresses of the HTTP sessions, NAT should not be enabled for this policy.

Configure other policy settings that you may require. You can also use more specific firewall addresses or add one security policy that accepts both FTP and HTTP traffic.

- 3** Add a security policy that accepts FTP sessions received at port1 and applies virus scanning to them:

```

config firewall policy
edit 20
set srcintf port1
set dstintf Vlink0
set srcaddr all
set dstaddr all
set action accept
set service FTP
set schedule always
set utm-status enable
set profile-protocol-options default
set av-profile scan
end

```



To preserve the source addresses of the HTTP sessions, NAT should not be enabled for this policy.

Configure other policy settings that you may require. You can also use more specific firewall addresses or add one security policy that accepts both FTP and HTTP traffic.

To configure routing for the FortiGate-311B Wanopt VDOM

- 1** Log in to the Wanopt VDOM from the CLI.
- 2** Add a default route. The destination of the default route is the port10 interface. The gateway of the default route is the next hop router that the port10 interface connects with:

```

config router static
edit 1
set dst 0.0.0.0/0.0.0.0
set device port10
set gateway (next hop router IP address)
set distance 10
end

```

- 3 Add a route to send return traffic from the server network destined for the client network out the Vlink1 interface to the Vlink0 interface in the root VDOM, which has the IP address 172.1.1.2:

```
config router static
edit 2
set dst 172.20.120.0/24
set device Vlink1
set gateway 172.1.1.2
set distance 10
end
```

To add security policies to the FortiGate-311B Wanopt VDOM to accept HTTP and FTP sessions received at the Vlink1 interface of the inter-VDOM link destined for port10

- 1 Log in to the Wanopt VDOM from the CLI.
- 2 Add a security policy that accepts HTTP sessions received at Vlink1 and destined for port10:

```
config firewall policy
edit 20
set srcintf Vlink1
set dstintf port10
set srcaddr all
set dstaddr all
set action accept
set service HTTP
set schedule always
set nat enable
end
```



NAT is ignored for all HTTP sessions for the server network because these sessions are intercepted by a full optimization WAN optimization rule. However, HTTP sessions for the Internet are intercepted by the Web Cache Only rule, so source NAT is required for replies.



Do not enable UTM because you cannot apply UTM features and WAN optimization to the same session in the same VDOM. Virus scanning was applied to the session in the root VDOM.

Configure other settings that you may require.

- 3 Go to *Policy > Policy > Policy* and select *Create New* to add a security policy that accepts FTP sessions received at Vlink1 and destined for port10:

```
config firewall policy
edit 20
set srcintf Vlink1
set dstintf port10
set srcaddr all
set dstaddr all
set action accept
set service FTP
set schedule always
set nat enable
```

end



NAT is ignored for all HTTP sessions for the server network because these sessions are intercepted by a full optimization WAN optimization rule. However, HTTP sessions for the Internet are intercepted by the Web Cache Only rule, so source NAT is required for replies.



Do not enable UTM because you cannot apply UTM features and WAN optimization to the same session in the same VDOM. Virus scanning was applied to the session in the root VDOM.

Configure other settings that you may require.

To configure peers for the FortiGate-311B Wanopt VDOM

- 1 Log in to the Wanopt VDOM from the CLI.
- 2 Add the Local Host ID for the client-side FortiGate-311B unit:

```
config wanopt settings
  set host-id Client_Fgt
end
```

- 3 Add a Peer Host ID and the IP Address for the server-side FortiGate-620B unit.

```
config wanopt peer
  edit Server_Fgt
  set ip 10.20.20.2
end
```

To add WAN optimization rules for HTTP and FTP to the FortiGate-311B Wanopt VDOM

- 1 Log in to the Wanopt VDOM from the CLI.
- 2 Add an active rule to optimize HTTP traffic from IP addresses on the Client network (172.20.120.0) with a destination address on the server network (192.168.10.0):

```
config wanopt rule
  edit 4
    set auto-detect active
    set src-ip 172.20.120.0-172.20.120.255
    set dst-ip 192.168.10.0-192.168.10.255
    set port 80
    set proto http
  end
```

Accept default settings for transparent (enable), status (enable), mode (full), byte-caching (enable), ssl (disable), secure-tunnel (disable), auth-group (null), unknown-http-version (tunnel), and tunnel-non-http (disable).



For improved privacy you can enable `secure-tunnel` and add an authentication group to both optimizing FortiGate units.

- 3 Add an active rule to optimize FTP traffic from IP addresses on the Client network (172.20.120.0) with a destination address on the server network (192.168.10.0):

```
config wanopt rule
  edit 5
```

```

set auto-detect active
set src-ip 172.20.120.0-172.20.120.255
set dst-ip 192.168.10.0-192.168.10.255
set port 21
set proto ftp
end

```

Accept default settings for `transparent (enable)`, `status (enable)`, `mode (full)`, `byte-caching (enable)`, `ssl (disable)`, `secure-tunnel (disable)`, `auth-group (null)`, `unknown-http-version (tunnel)`, and `tunnel-non-http (disable)`.



For improved privacy you can enable `secure-tunnel` and add an authentication group to both optimizing FortiGate units.

- 4 Add a rule to web cache HTTP traffic from IP addresses on the Client network (172.20.120.0) with any destination address:

```

config wanopt rule
edit 6
set mode webcache-only
set src-ip 172.20.120.0-172.20.120.255
set dst-ip 0.0.0.0
set port 80
set proto http
end

```

Accept default settings for `transparent (enable)`, `status (enable)`, `ssl (disable)`, `unknown-http-version (tunnel)`, and `tunnel-non-http (disable)`.

- 5 If required, use the `move` command to move the Web Cache Only rule below the full optimization HTTP and FTP rules in the list. The Web Cache Only rule should be below the full optimization rules because it will match all HTTP traffic and you need HTTP sessions with destination address 192.168.10.0 to match the full optimization HTTP rule

For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

Server-side configuration steps - CLI

This section describes the configuration steps required for the server-side FortiGate-620B unit to perform WAN optimization with the client-side FortiGate-311B unit and to send HTTP and FTP sessions to the server-side FortiGate-1000A cluster. This section also describes how to configure the FortiGate-1000A cluster to forward HTTP and FTP sessions from the client network to the server network.

The section breaks down the client-side configuration into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure routing for the FortiGate-620B unit.
- 2 Configure peers for the server-side FortiGate-620B unit.
- 3 Add a passive WAN optimization rule to the server-side FortiGate-620B unit.
- 4 Configure the FortiGate-1000A cluster to accept HTTP and FTP connections at port5 and forward them out port1 to the server network.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

To configure routing for the FortiGate-620B unit

- 1 Add a default route. The destination of the default route is the port16 interface. The gateway of the default route is the next hop router that the port16 interface connects with:

```
config router static
  edit 1
    set dst 0.0.0.0/0.0.0.0
    set device port16
    set gateway (next hop router IP address)
    set distance 10
  end
```

- 2 Add a route to send traffic for the server network out port1 to the port5 interface of the FortiGate-1000A cluster, which has the IP address 192.20.20.1:

```
config router static
  edit 2
    set dst 192.168.10.0/24
    set device port1
    set gateway 192.20.20.1
    set distance 10
  end
```

To configure peers for the server-side FortiGate-620B unit

- 1 Add the Local Host ID for the server-side FortiGate-620B unit:

```
config wanopt settings
  set host-id Server_Fgt
end
```

- 2 Add a Peer Host ID and the IP Address for the client-side FortiGate-311B unit:

```
config wanopt peer
  edit Client_Fgt
    set ip 10.10.10.2
  end
```

To add a passive WAN optimization rule to the server-side FortiGate-620B unit

You can add one passive WAN optimization rule to the server-side FortiGate-620B unit for both active rules on the FortiGate-311B unit. This rule can also allow the FortiGate-620B to perform WAN optimization with other client-side devices as long as the required Peer Host IDs are added to the FortiGate-620B configuration and to the client-side configurations.

- 1 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New* to add a passive rule that accepts any WAN optimization tunnel request:

```
config wanopt rule
  edit 5
    set auto-detect passive
    set src-ip 0.0.0.0
    set dst-ip 192.168.10.0-192.168.10.255
    set port 1-65535
    set webcache enable
  end
```

Accept default settings for `status` (enable) and `mode` (full).



You can also use a narrower `port` range such as 21-80 or add two rules, one with port set to 80 and one with port set to 21.

- 2 If required, use the `move` command to move the rule to a different position in the list so that the tunnel request from the client-side FortiGate unit matches with this rule. For more information, see [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate-1000A cluster to accept HTTP and FTP connections at port5 and forward them out port1 to the server network

- 1 Add a firewall address for the server network:

```
config firewall address
  edit Server_Net
    set type iprange
    set start-ip 192.168.10.0
    set end-ip 192.168.10.255
  end
```

- 2 Add a firewall address for the client network:

```
config firewall address
  edit Client_Net
    set type iprange
    set start-ip 172.20.120.0
    set end-ip 172.20.120.255
  end
```

- 3 Go to *Policy > Policy > Policy* and select *Create New* to add an security policy that accepts HTTP sessions at port5 destined for port1 and the server network:

```
config firewall policy
  edit 10
    set srcintf port5
    set dstintf port1
    set srcaddr Client_Net
    set dstaddr Server_Net
    set action accept
    set service HTTP
    set schedule always
  end
end
```



WAN optimization is operating in Transparent mode so the packets from the client network include their client network source IP addresses. To preserve these source IP addresses, the security policies on the FortiGate-1000A cluster that accept the sessions from the FortiGate- 620B unit should not apply NAT. If the policies were to apply NAT, the client network addresses would be replaced with the port1 IP address of the FortiGate-1000A cluster and the client network source IP addresses would be lost.

- 4 Go to *Policy > Policy > Policy* and select *Create New* to add an security policy that accepts FTP sessions at port5 destined for port1 and the server network:

```
config firewall policy
  edit 11
    set srcintf port5
    set dstintf port1
    set srcaddr Client_Net
    set dstaddr Server_Net
    set action accept
    set service FTP
    set schedule always
  end
end
```



As described above, selecting NAT would cause the loss of the client network source IP addresses.



SSL offloading for WAN optimization and web caching

WAN optimization SSL offloading uses a FortiGate unit's FortiASIC SSL encryption/decryption engine to accelerate SSL performance. Most commonly SSL offloading is used to accelerate the performance of a web server that secures communication using the HTTPS protocol. SSL offloading can be applied to WAN optimization configurations and reverse proxy web caching configurations.

In a WAN optimization configuration where clients on a client network use HTTPS to communicate with a web server over a WAN optimization tunnel, you can use SSL offloading to decrypt HTTPS traffic before sending it through the WAN optimization tunnel. Decrypting the HTTPS traffic before it enters the tunnel allows WAN optimization to apply optimization techniques to the in-the-clear HTTP traffic. These techniques that would not be as effective with encrypted HTTPS traffic. When the optimized traffic leaves the WAN optimization tunnel it can either be forwarded as clear text HTTP traffic or it can be re-encrypted as HTTPS traffic before being sent to its destination. To forward clear text traffic, SSL offloading must be configured in half mode. To forward encrypted traffic, SSL offloading must be configured in full mode.

In a half mode reverse proxy web caching SSL offloading configuration, the FortiGate unit intercepts HTTPS traffic from clients and decrypts it before sending it as HTTP clear text traffic to a web server. The HTTP clear text response from the web server is encrypted by the FortiGate unit and returned to the clients as encrypted HTTPS traffic. SSL offloading improves performance by offloading SSL processing from the web server.

In a full mode reverse proxy web cache SSL offloading configuration, the FortiGate unit decrypts HTTPS traffic, applies web caching, and then re-encrypts the HTTPS traffic before forwarding it to the web server.

You can also combine SSL offloading with other WAN optimization techniques such as HTTP protocol optimization, byte caching, and web caching to further enhance web server performance.

You enable SSL offloading by selecting *Enable SSL* in a WAN optimization rule. You must also add SSL servers to support SSL offloading by using the CLI command `config wanopt ssl-server`.

You must add one SSL server for each web server for which you are configuring SSL offloading. The SSL server configuration must include the destination IP address and port number of the HTTPS packets to be decrypted. The SSL server configuration also includes the SSL mode (full or half), the name of the HTTP server CA certificate used to decrypt and encrypt the traffic, the size of the Diffie-Hellman prime used in DHE-RSA negotiation, the relative strength of the encryption algorithms accepted during negotiation, the SSL/TLS versions to use, and the behavior regarding sending empty fragments to avoid a CBC IV attack.

You load HTTP server CA certificate into the FortiGate unit as a local certificate and then add it to the SSL server configuration using the `ssl-cert` keyword. The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

SSL offload does not perform address or port translation. If your configuration requires address translation this must be configured at the security policy level using firewall virtual IPs.

You can configure one WAN optimization rule to offload SSL encryption/decryption for multiple HTTP servers. To do this, you configure the WAN optimization rule source and destination addresses, so that the rule accepts packets destined for all of the HTTP servers for which you want offloading. Then you add one SSL server for each of the HTTP servers.

This chapter describes:

- [Example: SSL offloading for a WAN optimization tunnel](#)
- [Example: SSL offloading and reverse proxy web caching for an Internet web server using static one-to-one virtual IPs](#)
- [Example: SSL offloading and reverse proxy web caching for an Internet web server using a port forwarding virtual IP for HTTPS traffic](#)

About SSL server full and half mode

An SSL server configuration can operate in full mode or half mode. In full mode the SSL server performs both decryption and encryption of the HTTPS traffic. The full mode sequence is:

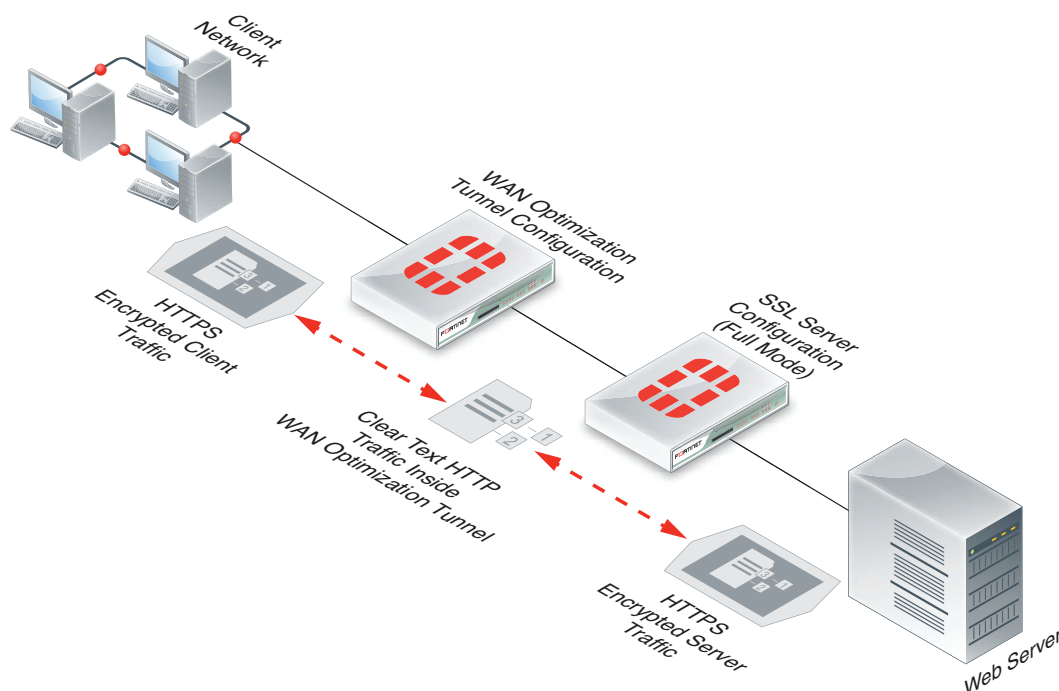
```
HTTPS -> HTTP -> HTTPS
HTTPS <- HTTP <- HTTPS
```

In half mode, the SSL server only performs one encryption or decryption action. If HTTP packets are received, the half mode SSL server encrypts them and converts them to HTTPS packets. If HTTPS packets are received, the SSL server decrypts them and converts them to HTTP packets. The half mode sequence is:

```
HTTPS -> HTTP
HTTPS <- HTTP
```

WAN optimization full mode SSL server configuration

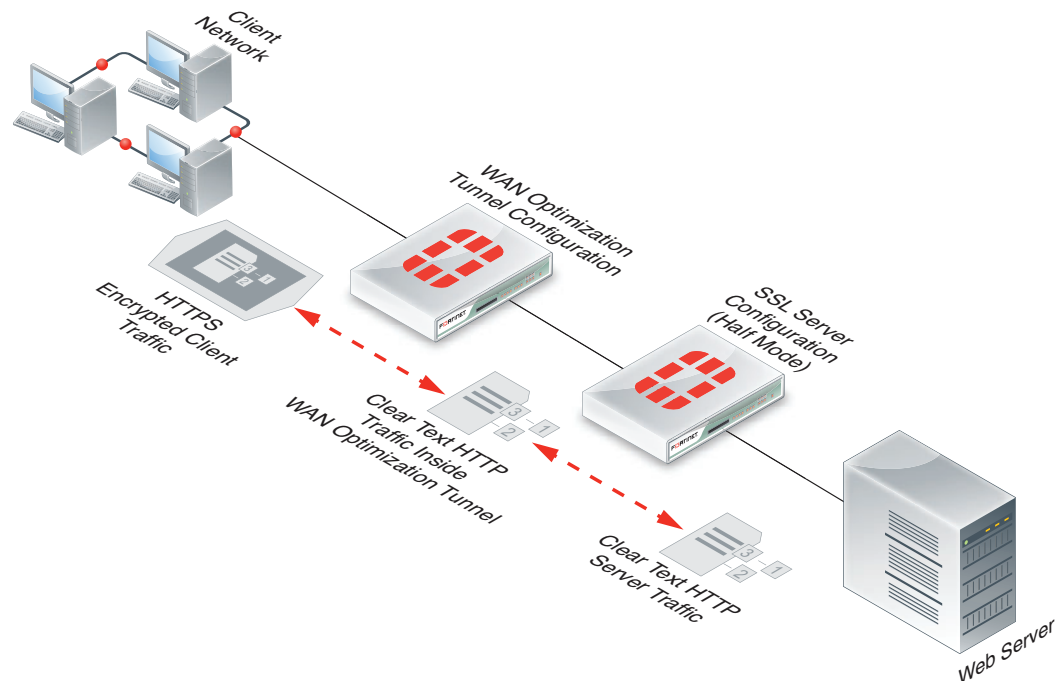
As shown in [Figure 330](#), in a WAN optimization full mode SSL server configuration, encrypted HTTPS packets received from a client network are decrypted into clear text HTTP packets that are sent over the WAN optimization tunnel to a server network. When the packets exit the tunnel they are re-encrypted and sent to the web server as HTTPS packets. HTTPS responses from the server are unencrypted before being sent across the tunnel and then re-encrypted before being returned to the clients.

Figure 330: WAN optimization full mode SSL server configuration

This full mode configuration is used when the traffic on both the client and server networks must be encrypted. This configuration enhances performance by allowing WAN optimization to apply HTTP optimization, byte caching, and web caching to the HTTP traffic in the WAN optimization tunnel.

WAN optimization half mode SSL server configuration


As shown in Figure 331, in a WAN optimization half mode SSL server configuration, encrypted HTTPS packets received from a client network are decrypted into clear text HTTP packets that are sent over the WAN optimization tunnel and forwarded to the server. HTTP responses from the server are sent across the tunnel and then re-encrypted before being returned to the clients.

Figure 331: WAN optimization half mode SSL server configuration

This half mode configuration is used when the traffic on the server network can be in clear text. This configuration enhances performance by accelerating SSL decryption and encryption and by allowing WAN optimization to apply HTTP optimization, byte caching, and web caching to the HTTP traffic in the WAN optimization tunnel.

Reverse proxy web cache full mode SSL server configuration


As shown in Figure 332, in a reverse proxy web cache full mode SSL server configuration, encrypted HTTPS packets received from a client network are decrypted into clear text HTTP packets that are cached and then re-encrypted and sent to the web server as HTTPS packets. HTTPS responses from the server are unencrypted and cached and then re-encrypted before being returned to the clients.

Figure 332: Reverse proxy web cache full mode SSL server configuration

This full mode configuration is used when the traffic on both the client and server networks must be encrypted. This configuration enhances performance by caching HTTP content.

Reverse proxy web cache half mode SSL server configuration

As shown in Figure 333, in a reverse proxy web cache half mode SSL server configuration, encrypted HTTPS packets received from a client network are decrypted into clear text HTTP packets that are cached and forwarded to the server. HTTP responses from the server are cached and then re-encrypted before being returned to the clients.

Figure 333: Reverse proxy web cache half mode SSL server configuration

This half mode configuration is used when the traffic on the server network can be in clear text. This configuration enhances performance by accelerating SSL decryption and encryption and by applying web caching to the HTTP traffic.

Example: SSL offloading for a WAN optimization tunnel

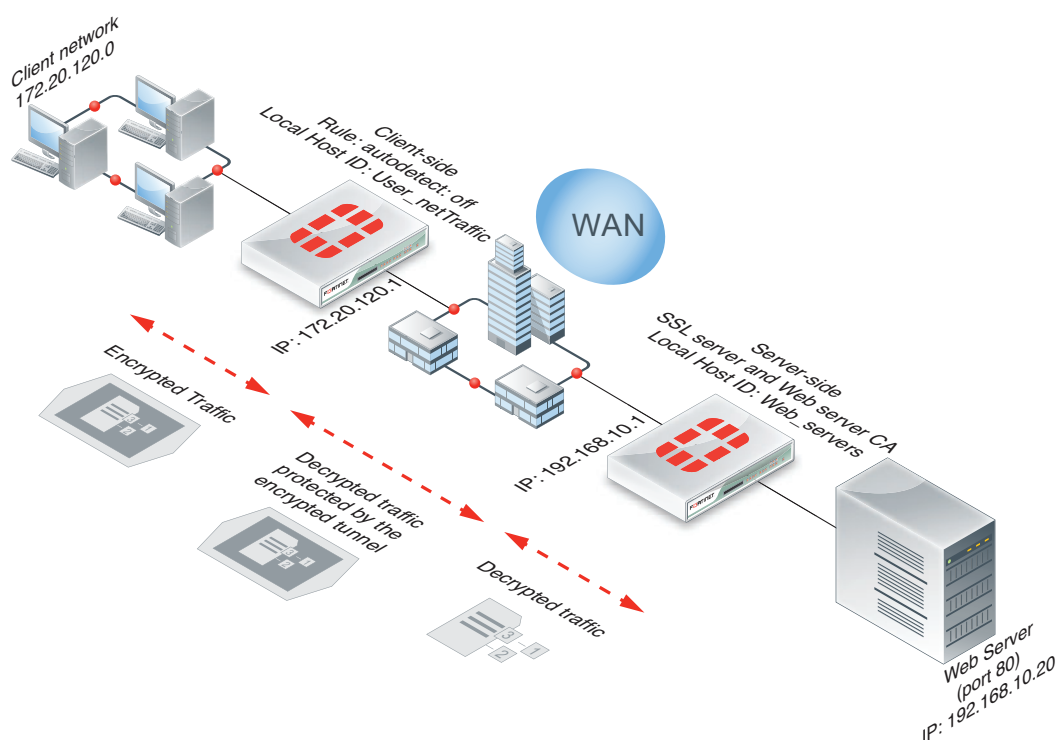
This example shows how to configure basic SSL offloading for a WAN optimization tunnel. This basic SSL offloading configuration can be applied to many network configurations.

Network topology and assumptions

In this example, clients on a client network use `https://192.168.10.20` to browse to a web server. A WAN optimization rule with *Auto-Detect* set to *Off* on the client-side FortiGate unit accepts sessions from the clients with source addresses on the 172.20.120.0 network and with a destination address of 192.168.10.0 and a destination port of 443. In this rule, *Enable Secure Tunnel* is selected so that the tunnel is encrypted. To support the encrypted tunnel, the configuration also includes an authentication group with a pre-shared key. Both FortiGate units must have the same authentication group with the same pre-shared key.

The server-side FortiGate unit includes an SSL server configuration with `ip` set to 192.168.10.20 and `port` to 443. The unit also includes the web server CA.

Figure 334: SSL offloading WAN optimization configuration



When the client-side FortiGate unit accepts an HTTPS connection for 192.168.10.20, the SSL server configuration provides the information that the client-side unit needs to decrypt the traffic and send it in clear text across a WAN optimization tunnel to the server-side unit. The server-side unit then forwards the clear text packets to the web server.

The web server CA is not downloaded from the server side to the client-side FortiGate unit. Instead, the client-side FortiGate unit proxies the SSL parameters from the client side to the server side, which returns an SSL key and other required information to the client-side unit so that it can decrypt and encrypt HTTPS traffic.



In this peer-to-peer configuration you do not need to add a WAN optimization rule to the server-side FortiGate unit as long as this server-side unit includes the peer host ID of the client-side FortiGate unit in its peer list. However, you can set *Auto-Detect* to *Active* on the client-side FortiGate unit and then add a passive rule to the server-side unit.

In this example, you do not require the secure tunnel and the authentication group configurations, but they are included to show how to protect the privacy of the WAN optimization tunnel. Alternatively, you could configure a route-based IPsec VPN between the FortiGate units and use IPsec to protect the privacy of the WAN optimization tunnel.

In this example, it is assumed that you have a local CA named `Web_Server_Cert_1.crt` stored in a file that you will import when you configure the server-side FortiGate unit.

General configuration steps

This example is divided into client-side and server-side steps, as configured through the web-based manager, and with CLI instructions provided for CLI-only steps. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

You also need access to the CLI to perform CLI-only steps.

Client-side configuration steps

To configure the client-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	User_net
----------------------	----------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the peer side FortiGate unit:

Peer Host ID	Web_servers
IP Address	192.168.10.1

- 4 Select *OK*.
- 5 Go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New* to add an authentication group named `SSL_auth_grp` to the client-side FortiGate unit.

The authentication group includes a pre-shared key and the peer added in step 3. An authentication group with the same name and the same pre-shared key must also be added to the server-side FortiGate unit. This authentication group is required for the secure tunnel:

Name	SSL_auth_grp
Authentication Method	Pre-shared key
Password	<pre-shared_key>
Peer Acceptance	Specify Peer: Web_servers

- 6 Select *OK*.
- 7 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New* to add the WAN optimization rule:

Mode	Full Optimization
Source	172.20.120.*
Destination	192.168.10.*
Port	443
Auto-Detect	Off
Protocol	HTTP
Peer	Web_servers
Transparent Mode	Select
Enable Byte Caching	Select
Enable SSL	Select
Enable Secure Tunnel	Select
Authentication Group	SSL_auth_grpa

- 8 Select *OK*.
The rule is added to the bottom of the WAN optimization list.
- 9 If required, move the rule to a different position in the list.
The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

Server-side configuration steps

To configure the server-side FortiGate unit

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > WAN Opt. Peer > Peer* and enter a *Local Host ID* for the server-side FortiGate unit:

Local Host ID	Web_servers
----------------------	-------------

- 2 Select *Apply* to save your setting.
- 3 Select *Create New* and add a *Peer Host ID* and the *IP Address* for the peer side FortiGate unit:

Peer Host ID	User_net
IP Address	172.20.120.1

- 4 Select *OK*.

- 5 Go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New* to add an authentication group named `SSL_auth_grp` to the server-side FortiGate unit.

The authentication group includes a pre-shared key and the peer added to the server-side FortiGate unit in step 3:

Name	SSL_auth_grp
Authentication Method	Pre-shared key
Password	<pre-shared_key>
Peer Acceptance	Specify Peer: User_net

- 6 Select *OK*.
- 7 Go to *System > Certificates > Local Certificates* and select *Import* to import the web server's CA.

For *Type*, select *Local Certificate*. Select the *Browse* button to locate the file, `Web_Server_Cert_1.crt`.

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

- 8 From the CLI, enter the following command to add the SSL server to the server-side FortiGate unit:

```
config wanopt ssl-server
  edit example_server
    set ip 192.168.10.20
    set port 443
    set ssl-mode half
    set ssl-cert Web_Server_Cert_1
  end
```

Configure other `ssl-server` settings that you may require for your configuration.

Example: SSL offloading and reverse proxy web caching for an Internet web server using static one-to-one virtual IPs

This section describes configuring SSL offloading for a reverse proxy Web Cache Only WAN optimization configuration using a static one-to-one firewall virtual IP (VIP). While the static one-to-one configuration described in this example is valid, it's also common to change the destination port of the unencrypted HTTPS traffic to a commonly used HTTP port such as 8080 using a port forwarding virtual IP. For an example of this configuration see [“Example: SSL offloading and reverse proxy web caching for an Internet web server using a port forwarding virtual IP for HTTPS traffic” on page 2795](#).

Network topology and assumptions

In this configuration, clients on the Internet use HTTP and HTTPS to browse to a web server. The FortiGate unit intercepts the HTTP traffic and a Web Cache Only WAN optimization rule forwards the HTTP traffic to the web server. The FortiGate unit intercepts the HTTPS traffic and a different Web Cache Only WAN optimization rule with SSL offloading enabled decrypts the traffic before sending it to the web server. The FortiGate unit also caches HTTP and HTTPS pages from the web server. Replies to HTTPS sessions from the web server are encrypted by the FortiGate unit before returning to the web browsing clients.

In the Web Cache Only rules transparent mode is enabled because the FortiGate unit is performing NAT between the Internet and the HTTP server and the web server network is not configured to route Internet traffic between the FortiGate unit and the web server.

In this configuration, the FortiGate unit is operating as a web cache in reverse proxy mode. Reverse proxy caches can be placed directly in front of a web server. Web caching on the FortiGate unit reduces the number of requests that the web server must handle, therefore leaving it free to process new requests that it has not serviced before.

Using a reverse proxy configuration:

- avoids the capital expense of additional web servers by increasing the capacity of existing servers
- serves more requests for static content from web servers
- serves more requests for dynamic content from web servers
- reduces operating expenses including the cost of bandwidth required to serve content
- accelerates the response time of web servers and of page download times to end users.

When planning a reverse proxy implementation, the web server's content should be written so that it is "cache aware" to take full advantage of the reverse proxy cache.

In reverse proxy mode, the FortiGate unit functions more like a web server for the clients it services. Unlike internal clients, external clients are not reconfigured to access the proxy server. Instead, the site URL routes the client to the FortiGate unit as if it were a web server. Replicated content is delivered from the proxy cache to the external client without exposing the web server or the private network residing safely behind the firewall.

In this example, the site URL translates to IP address 192.168.10.1, which is the port2 IP address of the FortiGate unit. The port2 interface is connected to the Internet.

This example also includes two Web Cache Only rules, one that accepts the HTTP traffic for web caching and one that accepts the HTTPS traffic for SSL offloading and web caching. You could instead add only one rule for both the HTTP and HTTPS traffic.

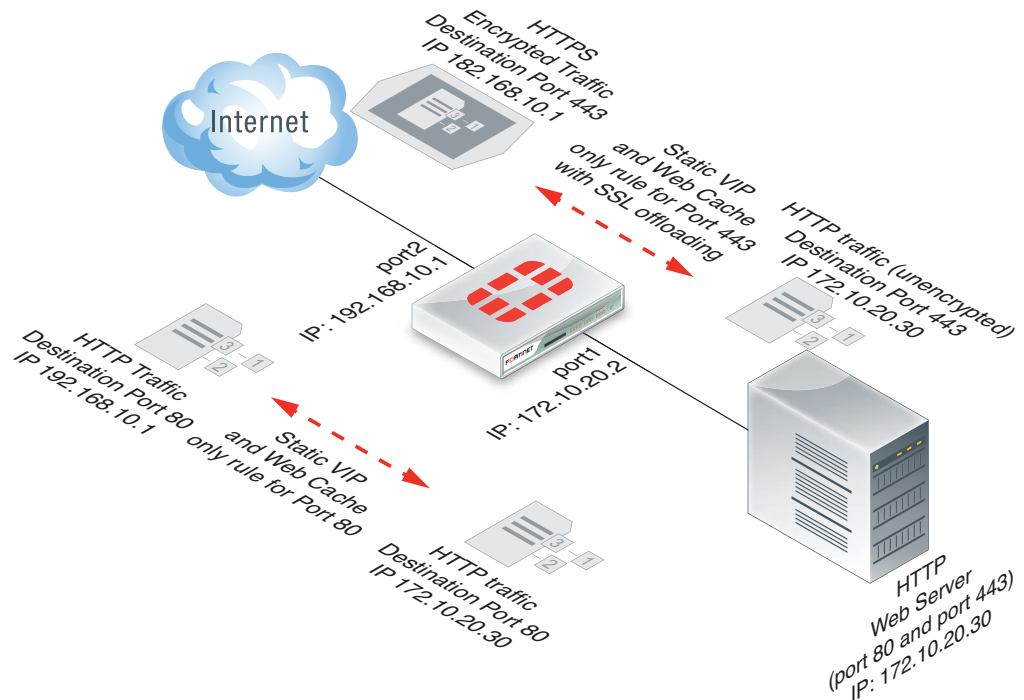
For this example, it is also assumed that all HTTP traffic uses port 80 and all HTTPS traffic uses port 443.

The FortiGate unit includes the web server CA and an SSL server configuration for IP address 172.10.20.30 and port to 443. The name of the file containing the CA is `Rev_Proxy_Cert_1.crt`.

The destination address of incoming web server requests are translated to the IP address of the web server using a static one-to-one virtual IP that performs destination address translation (DNAT) for the HTTP and HTTPS packets. The DNAT translates the destination address of the packets from 192.168.10.1 to 172.10.20.30 but does not change the destination port number.

WAN optimization receives the packets after the DNAT has occurred so the source address for the reverse proxy WAN optimization rules should be 172.10.20.30.

Figure 335: SSL offloading and reverse proxy web caching for an Internet web server using static one-to-one virtual IPs



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the FortiGate unit as a reverse proxy web cache server.
- 2 Configure the FortiGate unit for SSL offloading of HTTPS traffic.
- 3 Add an SSL server to offload SSL encryption and decryption for the web server.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuration steps - web-based manager

To configure the FortiGate unit as a reverse proxy web cache server

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New* to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

Name	Reverse_proxy_VIP
External Interface	port2
Type	Static NAT
Source Address Filter	Do not select.
External IP Address/Range	192.168.10.1
Mapped IP Address/Range	172.10.20.30
Port Forwarding	Do not select.

- 2 Select *OK* to save your settings.
- 3 Go to *Policy > Policy > Policy* and select *Create New* to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet.

Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	Reverse_proxy_VIP
Schedule	always
Service	HTTP and HTTPS Select <i>Multiple</i> to select more than one service.
Action	ACCEPT

- 4 Select *OK* to save your settings.
- 5 Go to *WAN Opt. & Cache > WAN Opt Rule > Rule* and select *Create New* to add a Web Cache Only WAN optimization rule.

Configure the rule to accept the HTTP traffic on port 80 accepted by the security policy:

Mode	Web Cache Only
Source	0.0.0.0
Destination	172.10.20.30 You need to set <i>Destination</i> to the translated (DNATed) IP address (172.10.20.30).
Port	80
Transparent Mode	Select
Enable SSL	Do not select

- 6 Select *OK*.
The rule is added to the bottom of the WAN optimization list.
- 7 If required, move the rule to a different position in the list.
The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate unit for SSL offloading of HTTPS traffic

The security policy added in the first procedure accepts HTTPS traffic so you do not have to add another one.

- 1 Go to *WAN Opt. & Cache > WAN Opt Rule > Rule* and select *Create New* to add a Web Cache Only WAN optimization rule.

Configure the rule to accept the HTTPS traffic on port 443 accepted by the security policy:

Mode	Web Cache Only
Source	0.0.0.0
Destination	172.10.20.30 You need to set <i>Destination</i> to the translated (DNATed) IP address (172.10.20.30).
Port	443
Transparent Mode	Select.
Enable SSL	Select.

- 2 Select *OK*.

The rule is added to the bottom of the WAN optimization list.

- 3 If required, move the rule to a different position in the list.

The HTTPS rule can be above or below the HTTP rule.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To add an SSL server to offload SSL encryption and decryption for the web server

- 1 Go to *System > Certificates > Local Certificates* and select *Import* to import the web server's CA.

For *Type*, select *Local Certificate*. Select the *Browse* button to locate the file *Rev_Proxy_Cert_1.crt*.

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

- 2 From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
  edit rev_proxy_server
    set ip 172.10.20.30
    set port 443
    set ssl-mode half
    set ssl-cert Rev_Proxy_Cert_1
  end
```

- 3 Configure other `ssl-server` settings that you may require for your configuration.

Configuration steps - CLI

To configure the FortiGate unit as a reverse proxy web cache server

- 1 Enter the following command to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

```
config firewall vip
  edit Reverse_proxy_VIP
    set extintf port2
    set type static-nat
    set extip 192.168.10.1
    set mappedip 172.10.20.30
  end
```

- 2 Enter the following command to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet.

Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
  edit 0
    set srcintf port2
    set srcaddr all
    set dstintf port1
    set dstaddr Reverse_proxy_VIP
    set schedule always
    set service HTTP HTTPS
    set action accept
  end
```

- 3 Enter the following command to add a Web Cache Only WAN optimization rule that accepts the HTTP traffic on port 80 accepted by the security policy.

Set `dst-ip` to the translated (DNATed) IP address (172.10.20.30).

```
config wanopt rule
  edit 0
    set mode webcache-only
    set src-ip 0.0.0.0
    set dst-ip 172.10.20.30
    set port 80
    set transparent enable
    set ssl disable
  end
```

- 4 If required, move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate unit for SSL offloading of HTTPS traffic

The security policy added in the first procedure accepts HTTPS traffic so you do not have to add another one.

- 1 Enter the following command to add a Web Cache Only WAN optimization rule that accepts the HTTPS traffic on port 443 accepted by the security policy.

Set `dst-ip` to the translated (DNATed) IP address (172.10.20.30).


```
config wanopt rule
edit 0
set mode webcache-only
set src-ip 0.0.0.0
set dst-ip 172.10.20.30
set port 443
set transparent enable
set ssl enable
end
```

- 2 If required, move the rule to a different position in the list.

The HTTPS rule can be above or below the HTTP rule.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To add an SSL server to offload SSL encryption and decryption for the web server

- 1 Place a copy of the web server’s CA (file name `Rev_Proxy_Cert_1.crt`) in the root folder of a TFTP server.
- 2 Enter the following command to import the web server’s CA from a TFTP server. The IP address of the TFTP server is 10.31.101.30:

```
execute vpn certificate local import tftp Rev_Proxy_Cert_1.crt
10.31.101.30
```

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

- 3 From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
edit rev_proxy_server
set ip 172.10.20.30
set port 443
set ssl-mode half
set ssl-cert Rev_Proxy_Cert_1
end
```

- 4 Configure other `ssl-server` settings that you may require for your configuration.

Example: SSL offloading and reverse proxy web caching for an Internet web server using a port forwarding virtual IP for HTTPS traffic

This section describes configuring SSL offloading for a reverse proxy Web Cache Only WAN optimization configuration using a static virtual IP (VIP) for HTTP traffic and a port forwarding virtual IP for HTTPS traffic. While the port forwarding configuration described in this example is commonly used, it is also possible to use a static virtual IP that does not change the destination port. For an example of this configuration see [“Example: SSL offloading and reverse proxy web caching for an Internet web server using static one-to-one virtual IPs” on page 2789](#).

Network topology and assumptions

In this configuration, clients on the Internet use HTTP (on port 80) and HTTPS (on port 9443) to browse to a web server. The FortiGate unit intercepts the HTTP traffic and a Web Cache Only WAN optimization rule forwards the HTTP traffic to the web server. The FortiGate unit intercepts the HTTPS traffic and a different Web Cache Only WAN optimization rule with SSL offloading enabled decrypts the traffic before sending it to the web server. The FortiGate unit also caches pages from the web server. Replies to HTTPS sessions from the web server are encrypted by the FortiGate unit before returning to the web browsing clients.

In the Web Cache Only WAN optimization rules, transparent mode is enabled because the FortiGate unit is performing NAT between the Internet and the HTTP server and the web server network is not configured to route Internet traffic between the FortiGate unit and the web server.

In this configuration, the FortiGate unit is operating as a web cache in reverse proxy mode. Reverse proxy caches can be placed directly in front of a web server. Web caching on the FortiGate unit reduces the number of requests that the web server must handle, therefore leaving it free to process new requests that it has not serviced before.

Using a reverse proxy configuration:

- avoids the capital expense of additional web servers by increasing the capacity of existing servers
- serves more requests for static content from web servers
- serves more requests for dynamic content from web servers
- reduces operating expenses including the cost of bandwidth required to serve content
- accelerates the response time of web servers and of page download times to end users.

When planning a reverse proxy implementation, the web server's content should be written so that it is "cache aware" to take full advantage of the reverse proxy cache.

In reverse proxy mode, the FortiGate unit functions more like a web server for the clients it services. Unlike internal clients, external clients are not reconfigured to access the proxy server. Instead, the site URL routes the client to the FortiGate unit as if it were a web server. Replicated content is delivered from the proxy cache to the external client without exposing the web server or the private network residing safely behind the firewall.

In this example, the site URL translates to IP address 192.168.10.1, which is the port2 IP address of the FortiGate unit. The port2 interface is connected to the Internet.

This example also includes two Web Cache Only rules, one that accepts the HTTP traffic for web caching and one that accepts the HTTPS traffic for SSL offloading and web caching.

For this example, it is also assumed that all HTTP traffic uses port 80 and all HTTPS traffic uses port 9443. (Since HTTPS traffic is received on the non-standard port of 9443 a security policy that accepts this traffic must have service set to ALL or you must add a custom service that matches traffic on port 9443. The example takes the custom service approach.)


The FortiGate unit includes the web server CA and an SSL server configuration for IP address 172.10.20.30 and port to 8080. The name of the file containing the CA is Rev_Proxy_Cert_1.crt.

The destination address of incoming HTTP requests is translated to the IP address of the web server using a static NAT virtual IP that performs destination address translation (DNAT). The DNAT translates the destination address of the HTTP packets from 192.168.10.1 to 172.10.20.30 but does not change the destination port of the HTTP traffic.

The destination address and port number of incoming HTTPS requests is translated to the IP address of the web server and to port 8080 using a port forwarding virtual IP that performs destination address translation (DNAT) and destination port translation (port forwarding). The DNAT translates the destination address of the packets from 192.168.10.1 to 172.10.20.30. Port forwarding changes the destination port number from 9443 to 8080.

WAN optimization receives the packets after the address and port translation has occurred so the source address for the reverse proxy WAN optimization rules should be 172.10.20.30 and the source port should be 80 for the HTTP traffic and 8080 for the HTTPS traffic.

Figure 336: SSL offloading and reverse proxy web caching for an Internet web server using a static virtual IP for HTTP traffic and a port forwarding virtual IP for HTTPS traffic



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Configure the FortiGate unit as a reverse proxy web cache server for the HTTP traffic (port 80, static DNAT).
- 2 Configure the FortiGate unit as a reverse proxy web cache server with SSL offloading for HTTPS traffic (DNAT, port forwarding (9443 to 8080), SSL offloading).

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuration steps - web-based manager

To configure the FortiGate unit as a reverse proxy web cache server for the HTTP traffic

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New* to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate the destination port):

Name	Reverse_proxy_http_VIP
External Interface	port2
Type	Static NAT
Source Address Filter	Do not select.
External IP Address/Range	192.168.10.1
Mapped IP Address/Range	172.10.20.30
Port Forwarding	Do not select.

- 2 Select *OK* to save your settings.
- 3 Go to *Policy > Policy > Policy* and select *Create New* to add a port2 to port1 security policy that accepts HTTP traffic from the Internet.
Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	Reverse_proxy_http_VIP
Schedule	always
Service	HTTP
Action	ACCEPT

- 4 Select *OK* to save your settings.
- 5 Go to *WAN Opt. & Cache > WAN Opt Rule > Rule* and select *Create New* to add a Web Cache Only WAN optimization rule.

Configure the rule to accept the HTTP traffic on port 80 accepted by the security policy:

Mode	Web Cache Only
Source	0.0.0.0
Destination	172.10.20.30 You need to set <i>Destination</i> to the translated (DNATed) IP address (172.10.20.30).
Port	80

Transparent Mode	Select
Enable SSL	Do not select

- 6 Select OK.

The rule is added to the bottom of the WAN optimization list.

- 7 If required, move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate unit as a reverse proxy web cache server for the HTTPS traffic

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New* to add a port forwarding virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 and translates the destination port from 9443 to 8080:

Name	Reverse_proxy_https_VIP
External Interface	port2
Type	Static NAT
Source Address Filter	Do not select.
External IP Address/Range	192.168.10.1
Mapped IP Address/Range	172.10.20.30
Port Forwarding	Select
Protocol	TCP
External Service Port	9443
Map to Port	8080

- 2 Select OK to save your settings.

- 3 Go to *Firewall Objects > Service > Custom* and select *Create New* to add a custom service for HTTPS on port 9443:

Name	HTTPS_9443
Protocol Type	TCP/UDP/SCTP
Protocol	TCP
Source Port (Low - High)	1 - 65535
Destination Port (Low - High)	9443 - 9443

- 4 Go to *Policy > Policy > Policy* and select *Create New* to add a port2 to port1 security policy that uses the custom service and accepts traffic from the Internet.

Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	Reverse_proxy_https_VIP

Schedule	always
Service	HTTPS_9443
Action	ACCEPT

5 Select *OK* to save your settings.

6 Go to *WAN Opt. & Cache > WAN Opt Rule > Rule* and select *Create New* to add a Web Cache Only WAN optimization rule.

Configure the rule to accept the HTTPS traffic on port 8080, the as translated by the port forwarding virtual IP:

Mode	Web Cache Only
Source	0.0.0.0
Destination	172.10.20.30 You need to set <i>Destination</i> to the translated (DNATed) IP address (172.10.20.30).
Port	8080
Transparent Mode	Select
Enable SSL	Select

7 Select *OK*.

The rule is added to the bottom of the WAN optimization list.

8 If required, move the rule to a different position in the list.

The HTTPS rule can be above or below the HTTP rule.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

9 Go to *System > Certificates > Local Certificates* and select *Import* to import the web server’s CA.

For *Type*, select *Local Certificate*. Select the *Browse* button to locate the file *Rev_Proxy_Cert_1.crt*.

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

10 From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic after being translated by the virtual IP (8080). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
edit rev_proxy_server
set ip 172.10.20.30
set port 8080
set ssl-mode half
set ssl-cert Rev_Proxy_Cert_1
end
```

11 Configure other `ssl-server` settings that you may require for your configuration.

Configuration steps - CLI

To configure the FortiGate unit as a reverse proxy web cache server for the HTTP traffic

- 1 Enter the following command to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

```
config firewall vip
  edit Reverse_proxy_http_VIP
    set extintf port2
    set type static-nat
    set extip 192.168.10.1
    set mappedip 172.10.20.30
  end
```

- 2 Enter the following command to add a port2 to port1 security policy that accepts HTTP traffic from the Internet.

Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
  edit 0
    set srcintf port2
    set srcaddr all
    set dstintf port1
    set dstaddr Reverse_proxy_http_VIP
    set schedule always
    set service HTTP
    set action accept
  end
```

- 3 Enter the following command to add a Web Cache Only WAN optimization rule that accepts the HTTP traffic on port 80 accepted by the security policy.

Set dst-ip to the translated (DNATed) IP address (172.10.20.30).

```
config wanopt rule
  edit 0
    set mode webcache-only
    set src-ip 0.0.0.0
    set dst-ip 172.10.20.30
    set port 80
    set transparent enable
    set ssl disable
  end
```

- 4 If required, move the rule to a different position in the list.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching” on page 2708](#) and [“Moving a rule to a different position in the rule list” on page 2709](#).

To configure the FortiGate unit as a reverse proxy web cache server for the HTTPS traffic

- 1 Enter the following command to add a port forwarding virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 and translates the destination port from 9443 to 8080:

```
config firewall vip
edit Reverse_proxy_https_VIP
set extintf port2
set type static-nat
set port-forward enable
set extip 192.168.10.1
set mappedip 172.10.20.30
set extport 9443
set mappedport 8080
end
```

- 2** Enter the following command to add a custom service for HTTPS on port 9443:

```
config firewall service custom
edit HTTPS_9443
set protocol TCP/UDP/SCTP
set tcp_portrange 9443
end
```

- 3** Enter the following command to add a port2 to port1 security policy that uses the custom service and accepts traffic from the Internet.

Do not select UTM features. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
edit 0
set srcintf port2
set srcaddr all
set dstintf port1
set dstaddr Reverse_proxy_https_VIP
set schedule always
set service HTTPS_9443
set action accept
end
```

- 4** Enter the following command to add a Web Cache Only WAN optimization rule that accepts the HTTPS traffic on port 443 accepted by the security policy.

Set `dst-ip` to the translated (DNATed) IP address (172.10.20.30) and set `port` to the translated port (8080).

```
config wanopt rule
edit 0
set mode webcache-only
set src-ip 0.0.0.0
set dst-ip 172.10.20.30
set port 8080
set transparent enable
set ssl enable
end
```

- 5** If required, move the rule to a different position in the list.

The HTTPS rule can be above or below the HTTP rule.

The order of the rules in the list significantly affects how the rules are applied. For more information, see [“How list order affects rule matching”](#) on page 2708 and [“Moving a rule to a different position in the rule list”](#) on page 2709.

- 6** Place a copy of the web server’s CA (file name `Rev_Proxy_Cert_1.crt`) in the root folder of a TFTP server.

- 7** Enter the following command to import the web server's CA from a TFTP server. The IP address of the TFTP server is 10.31.101.30:

```
execute vpn certificate local import tftp Rev_Proxy_Cert_1.crt
10.31.101.30
```


The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

- 8** From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic after being translated by the virtual IP (8080). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config wanopt ssl-server
edit rev_proxy_server
set ip 172.10.20.30
set port 8080
set ssl-mode half
set ssl-cert Rev_Proxy_Cert_1
end
```

- 9** Configure other `ssl-server` settings that you may require for your configuration.




FortiClient WAN optimization

FortiClient WAN optimization works together with WAN optimization on a FortiGate unit to accelerate network traffic between a PC running version 4.0 or greater of the FortiClient application and a network behind a FortiGate unit. When a user of a PC with FortiClient WAN optimization enabled attempts to connect to network resources behind a server-side FortiGate unit, the FortiClient application automatically detects if WAN optimization is enabled on the FortiGate unit. If WAN optimization is detected and the FortiClient application can successfully negotiate a WAN optimization tunnel with the FortiGate unit, a WAN optimization tunnel starts.

FortiClient WAN optimization includes protocol optimization settings selected in the FortiClient application and byte caching (byte caching is enabled by default in the FortiClient application and cannot be disabled). Web caching is applied if selected in the passive rule on the FortiGate unit that accepts FortiClient WAN optimization tunnel requests.

This chapter describes how to configure the FortiClient application for WAN optimization and how to configure a FortiGate unit to accept WAN optimization tunnel requests from the FortiClient application.

Figure 337: FortiClient WAN optimization topology



Configuring FortiClient WAN optimization

Configuring WAN optimization with the FortiClient application consists of enabling WAN optimization for the FortiClient application and configuring the FortiGate unit to accept WAN optimization tunnel requests from the FortiClient application.

FortiClient configuration steps

To configure WAN Optimization for the FortiClient application

- 1 From the FortiClient user interface, go to *Status > WAN Optimization*.
- 2 Select *Enable WAN Optimization*.
- 3 Enable the protocols to be optimized: *HTTP* (web browsing), *CIFS* (Windows file sharing), *MAPI* (Microsoft Exchange) and *FTP* (file transfers).
- 4 Set *Maximum Disk Cache* to 512, 1024, or 2048 MB.
The default is 512 MB. If the PC hard disk can accommodate a larger cache, better optimization performance is possible.
- 5 Select *Apply*.

FortiGate unit configuration steps

To configure FortiClient WAN Optimization on the FortiGate unit

Because PCs running the FortiClient application can have IP addresses that change often, it is usually not practical to add PCs running the FortiClient application to the WAN optimization peer list. Instead, a FortiGate unit that accepts WAN optimization tunnel requests from the FortiClient application should be configured to accept any peer (see [“Accepting any peers” on page 2697](#)) by adding an authentication group named *auth-fc* with *Peer acceptance* set to *Accept Any Peer*.

On the FortiGate unit, you also need to add a passive rule that includes source and destination addresses that will accept connections from the IP addresses of PCs running the FortiClient application. If these PCs can be anywhere on the Internet, the source address for this rule is 0.0.0.0. You can also use a more restrictive address range if the PCs running the FortiClient application have a restricted range of addresses.

You do not need to add security policies to the FortiGate unit because it is on the server side of the WAN optimization tunnel.

- 1 Go to *WAN Opt. & Cache > WAN Opt. Peer > Authentication Group* and select *Create New*.
- 2 Configure the authentication group:

Name	auth-fc
Authentication Method	Certificate
Certificate	Fortinet_Firmware
Peer Acceptance	Accept Any Peer

- 3 Select *OK*.
- 4 Go to *WAN Opt. & Cache > WAN Opt. Rule > Rule* and select *Create New*.
- 5 Configure a rule to accept FortiClient WAN optimization sessions:

Mode	Full Optimization
Source	0.0.0.0
Destination	0.0.0.0
Port	1-65535
Auto-Detect	Passive

- 6 Select *OK*.



The FortiGate explicit web proxy

You can use the FortiGate explicit web proxy to enable explicit HTTP, and HTTPS proxying on one or more FortiGate interfaces. The explicit web proxy also supports proxying FTP sessions from a web browser and proxy auto-config (PAC) to provide automatic proxy configurations for explicit web proxy users. From the CLI you can also configure the explicit web proxy to support SOCKS sessions from a web browser.

The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



Web proxies are configured for each VDOM when multiple VDOMs are enabled.


In most cases you would configure the explicit web proxy for users on a network by enabling the explicit web proxy on the FortiGate interface connected to that network. Users on the network would configure their web browsers to use a proxy server for HTTP and HTTPS, FTP, or SOCKS and set the proxy server IP address to the IP address of the FortiGate interface connected to their network. Users could also enter the PAC URL into their web browser PAC configuration to automate their web proxy configuration using a PAC file stored on the FortiGate unit.



Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate unit management IP address.

The web proxy receives web browser sessions to be proxied at FortiGate interfaces with the explicit web proxy enabled. The web proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit web proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address. For more information about explicit web proxy sessions, see [“Explicit proxy sessions and user limits” on page 2825](#).

Figure 338: Example explicit web proxy topology

To allow all explicit web proxy traffic to pass through the FortiGate unit you can set the explicit web proxy default firewall proxy action to accept. However, in most cases you would want to use security policies to control explicit web proxy traffic and apply security features such as access control/authentication, UTM, and traffic logging. You can do this by keeping the default explicit web proxy security policy action to deny and then adding web-proxy security policies.

You can also change the explicit web proxy default security policy action to accept and add explicit web proxy security policies. If you do this, sessions that match web-proxy security policies are processed according to the security policy settings. Connections to the explicit web proxy that do not match a web-proxy security policy are allowed with no restrictions or additional security processing. This configuration is not recommended and is not a best practice.

Web-proxy security policies can selectively allow or deny traffic, apply authentication using identity-based policies, enable traffic logging, and use UTM options to apply virus scanning, web filtering, and DLP to explicit web proxy traffic. There are some limitations to the UTM features that can be applied to explicit web proxy sessions. See [“UTM features and the explicit web proxy” on page 2817](#).

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit web proxy traffic. Security policies for the web proxy can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to *Any*. (On the web-based manager you must set the interface to *Any*. In the CLI you must `unset the associated-interface`.)

Authentication of explicit web proxy sessions uses HTTP authentication and can be based on the user's source IP address or on cookies from the user's web browser. For more information, see [“Explicit web proxy authentication” on page 2815](#).

To use the explicit web proxy, users must add the IP address of a FortiGate interface on which the explicit web proxy is enabled and the explicit web proxy port number (default 8080) to the proxy configuration settings of their web browsers.

On FortiGate units that support WAN optimization, you can also enable web caching for explicit web proxy sessions.

This section describes:

- [Explicit web proxy configuration overview](#)
- [Proxy chaining](#)
- [Explicit web proxy authentication](#)
- [UTM features and the explicit web proxy](#)

- [Web Proxy Services](#)
- [Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning](#)
- [Explicit proxy sessions and user limits](#)
- [Explicit web proxy configuration options](#)

Explicit web proxy configuration overview

You can use the following general steps to configure the explicit web proxy.

To enable the explicit web proxy - web-based manager

- 1 Go to *System > Network > Explicit Proxy*. Select *Enable Explicit Web Proxy* to turn on the explicit web proxy for HTTP and HTTPS traffic.

You can also select FTP to enable the web proxy for FTP over HTTP sessions in a web browser (not an FTP client) and PAC to enable automatic proxy configuration.

You can also optionally change the HTTP port that the proxy listens on (the default is 8080) and optionally specify different ports for HTTPS, FTP, and PAC.

- 2 Select *Apply*.

The default explicit web proxy configuration has *Default Firewall Policy Action* set to Deny and requires you to add a security policy to allow access to the explicit web proxy. This configuration is recommended as a best practice because you can use security policies to control access to the explicit web proxy and also apply security features such as logging, UTM, and authentication (by adding identity-based policies).

- 3 Go to *System > Network > Interface* and select one or more interfaces for which to enable the explicit web proxy. Edit the interface configuration and select *Enable Explicit Web Proxy*.



Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

- 4 Go to *Policy > Policy > Policy* and select *Create New* and set the *Source Interface/Zone* to *web-proxy*.

You can add multiple web-proxy security policies.

- 5 Configure the security policy as required to accept the traffic that you want to be processed by the explicit web proxy.

The source address of the policy should match client source IP addresses. The firewall address selected as the source address cannot be assigned to a FortiGate interface. The Interface field of the firewall address must be blank or it must be set to *Any*.

The destination address of the policy should match the IP addresses of web sites that clients are connecting to. Usually the destination address would be *all* if proxying Internet web browsing.

If *Default Firewall Policy Action* is set to Deny, traffic sent to the explicit web proxy that is not accepted by a web-proxy security policy is dropped. If *Default Firewall Policy Action* is set to Allow then all web-proxy sessions that don't match with a security policy are allowed.

For example the following security policy allows users on an internal network to access the Internet through the wan1 interface of a FortiGate unit.

Source Interface/Zone	web-proxy
Source Address	Internal_subnet
Destination Interface/Zone	wan1
Destination Address	all
Action	ACCEPT

- 6 You can select other security policy options as required.
For example, you can apply UTM protection to web proxy sessions and log allowed web proxy traffic.
- 7 You can also select *Enable Identity Based Policy* to apply authentication to explicit web proxy sessions.
- 8 You can add multiple identity based policies to apply different authentication for different user groups and also apply different UTM and logging settings for different user groups.

To enable the explicit web proxy - CLI



You must use the CLI to enable the explicit web proxy for VLAN interfaces.

- 1 Enter the following command to turn on the explicit web proxy for HTTP and HTTPS traffic.

```
config web-proxy explicit
  set status enable
end
```

You can also enter the following command to enable the web proxy for FTP sessions in a web browser.

```
config web-proxy explicit
  set ftp-over-http enable
end
```

The default explicit web proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit web proxy.

- 2** Enter the following command to enable the explicit web proxy for the internal interface.

```
config system interface
edit internal
set explicit-web-proxy enable
end
end
```

- 3** Use the following command to add a firewall address that matches the source address of users who connect to the explicit web proxy.

```
config firewall address
edit Internal_subnet
set type iprange
set start-ip 10.31.101.1
set end-ip 10.31.101.255
end
```

The source address for a web-proxy security policy cannot be assigned to a FortiGate unit interface.

- 4** Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall policy
edit 2
set srcintf web-proxy
set dstintf wan1
set scraddr Internal_subnet
set dstaddr all
set action accept
set identity-based enable
set schedule always
config identity-based-policy
edit 1
set groups Internal_users
set utm-status enable
set profile-protocol-options default
set av-profile Scan
set logtraffic enable
set schedule always
set service ANY
end
end
```

The firewall address selected as the source address cannot be assigned to a FortiGate unit interface. Either the field must be blank or it must be set to *Any*.

- 5** Use the following command to change global web proxy settings, for example to set the maximum request length for the explicit web proxy to 10:

```
config web-proxy global
set max-request-length 10
end
```

Proxy auto-config (PAC) configuration

A proxy auto-config (PAC) file defines how web browsers can choose a proxy server for receiving HTTP content. PAC files include the `FindProxyForURL(url, host)` JavaScript function that returns a string with one or more access method specifications. These specifications cause the web browser to use a particular proxy server or to connect directly.

To configure PAC for explicit web proxy users, you can use the port that PAC traffic from client web browsers use to connect to the explicit web proxy. explicit web proxy users must configure their web browser's PAC proxy settings to use the PAC port.

PAC File Content

You can edit the default PAC file from the web-based manager or use the following command to upload a custom PAC file:

```
config web-proxy explicit
    set pac-file-server-status enable
    set pac-file data <pac_file_str>
end
```

Where `<pac_file_str>` is the contents of the PAC file. Enter the contents of the PAC file. Enclose the PAC file text in quotes. You can copy the contents of a PAC text file and paste the contents into the CLI using this option. Enter the command followed by two sets of quotes then place the cursor between the quotes and paste the file content.

The maximum PAC file size is 256 kbytes. If your FortiGate unit is operating with multiple VDOMs each VDOM has its own PAC file. The total amount of FortiGate memory available to store all of these PAC files 2 MBytes. If this limit is reached you will not be able to load any additional PAC files.

You can use any PAC file syntax that is supported by your users's browsers. The FortiGate unit does not parse the PAC file.

To use PAC, users must add an automatic proxy configuration URL (or PAC URL) to their web browser proxy configuration. The default PAC file URL is:

```
http://<interface_ip>:<PAC_port_int>/<pac_file_str>
```

For example, if the interface with the explicit web proxy has IP address 172.20.120.122, the PAC port is the same as the default HTTP explicit web proxy port (8080) and the PAC file name is proxy.pac the PAC file URL would be:

```
http://172.20.120.122:8080/proxy.pac
```

From the CLI you can use the following command to display the PAC file urls:

```
get web-proxy explicit
```

Unknown HTTP version

You can select the action to take when the proxy server must handle an unknown HTTP version request or message. Set unknown HTTP version to Reject or Best Effort. Best Effort attempts to handle the HTTP traffic as best as it can. Reject treats known HTTP traffic as malformed and drops it. The Reject option is more secure.

Authentication realm

You can enter an authentication realm to identify the explicit web proxy. The realm can be any text string of up to 63 characters. If the realm includes spaces enclose it in quotes. When a user authenticates with the explicit web proxy the HTTP authentication dialog includes the realm so you can use the realm to identify the explicitly web proxy for your users.

Other explicit web proxy options

You can change the following explicit web proxy options as required by your configuration.

HTTP port, HTTPS port, FTP port, PAC port	The TCP port that web browsers use to connect to the explicit proxy for HTTP, HTTPS, FTP and PAC services. The default port is 8080 for all services. By default HTTPS, FTP, and PAC use the same port as HTTP. You can change any of these ports as required. Users configuring their web browsers to use the explicit web proxy should add the same port numbers to their browser configurations.
Proxy FQDN	Enter the fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server.
Max HTTP request length	Enter the maximum length of an HTTP request in Kbytes. Larger requests will be rejected.
Max HTTP message length	Enter the maximum length of an HTTP message in Kbytes. Larger messages will be rejected.
Add headers to Forwarded Requests	The web proxy server will forward HTTP requests to the internal network. You can include the following headers in those requests:
Client IP Header	Enable to include the Client IP Header from the original HTTP request.
Via Header	Enable to include the Via Header from the original HTTP request.
X-forwarded-for Header	Enable to include the X-Forwarded-For (XFF) HTTP header. The XFF HTTP header identifies the originating IP address of a web client or browser that is connecting through an HTTP proxy, and the remote addresses it passed through to this point.
Front-end HTTPS Header	Enable to include the Front-end HTTP Header from the original HTTPS request.

Proxy chaining

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with an already existing web proxy solution.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

FortiGate proxy chaining does not support authenticating with the remote forwarding server.

Adding a web proxy forwarding server

You add forwarding servers from the web-based manager by going to *System > Network > Explicit Proxy* and adding them. For each server you can define its address as a numeric IP address or fully qualified domain name. You can also specify the TCP port to use to connect to the proxy server.

Use the following CLI command to add a web proxy forwarding server named `fwd-srv` at address `proxy.example.com` and port 8080.

```
config web-proxy forward-server
edit fwd-srv
set addr-type fqdn
set fqdn proxy.example.com
set port 8080
end
```

Web proxy forwarding server monitoring and health checking

By default, a FortiGate unit monitors web proxy forwarding server by forwarding a connection to the remote server every 10 seconds. If the remote server does not respond it is assumed to be down. Checking continues and when the server does send a response the server is assumed to be back up. If you configure health checking, every 10 seconds the FortiGate unit attempts to get a response from a web server by connecting through the remote forwarding server.

You can configure health checking for each remote server and specify a different website to check for each one.

If the remote server is found to be down you can configure the FortiGate unit to block sessions until the server comes back up or to allow sessions to connect to their destination, bypassing the remote forwarding server. You cannot configure the FortiGate unit to fail over to another remote forwarding server.

Configure the server down action and enable health monitoring from the web-based manager by going to *System > Network > Explicit Proxy*, selecting a forwarding server, and changing the server down action and changing the health monitor settings.

Use the following CLI command to enable health checking for a web proxy forwarding server and set the server down option to bypass the forwarding server if it is down.

```
config web-proxy forward-server
edit fwd-srv
set healthcheck enable
set monitor http://example.com
set server-down-option pass
end
```

Adding proxy chaining to an explicit web proxy security policy

You enable proxy chaining for web proxy sessions by adding a web proxy forwarding server to an explicit web proxy security policy. In a policy you can select one web proxy forwarding server. All explicit web proxy traffic accepted by this security policy is forwarded to the specified web proxy forwarding server.

To add an explicit web proxy forwarding server - web-based manager

- 1 Go to *Policy > Policy > Policy* and select Create New.

2 Configure the security policy:

Source Interface/Zone	web-proxy
Source Address	Internal_subnet
Destination Interface/Zone	wan1
Destination Address	all
Schedule	always
Service	webproxy
Action	ACCEPT
Web Proxy Forwarding Server	Select, fwd-srv

3 Select OK to save the security policy.**To add an explicit web proxy forwarding server - CLI**

- 1 Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet. The policy forwards web proxy sessions to a remote forwarding server named fwd-srv

```
config firewall policy
edit 2
    set srcintf web-proxy
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr all
    set action accept
    set schedule always
    set service webproxy
    set webproxu-forward-server fwd-srv
end
```

Explicit web proxy authentication

You can add identity-based policies to apply authentication to explicit web proxy sessions. You can use authentication to control access to the explicit web proxy. You can also use identity-based policies to identify users and apply different UTM features to different users.

Authentication of web proxy sessions uses HTTP basic and digest authentication as described in [RFC 2617 \(HTTP Authentication: Basic and Digest Access Authentication\)](#) and prompts the user for credentials from the browser allowing individual users to be identified by their web browser instead of IP address. HTTP authentication allows the FortiGate unit to identify multiple users accessing services from a shared IP address. You can also select IP-based authentication to authenticate users according to their source IP address.

IP-Based authentication

IP-based authentication applies authentication by source IP address. Once a user authenticates, all sessions to the explicit web proxy from that IP address are assumed to be from that user and are accepted until the authentication timeout ends or the session times out.

This method of authentication is similar to standard (non-web proxy) firewall authentication and may not produce the desired results if multiple users share IP addresses (such as in a network that uses virtualization solutions or includes a NAT device between the users and the explicit web proxy).

To configure IP-based authentication, add a security policy for the explicit web proxy, set the source interface/zone to web-proxy, select Enable Identify Based Policy, and make sure IP Based is selected before adding identity-based policies. You can also set the authentication method to basic, digest, NTLM or FSAE.

Use the following CLI command to add IP-based authentication to a web proxy security policy. IP-based authentication is selected by setting `ip-based` to `enable`.

```
config firewall policy
  edit 3
    set srcintf web-proxy
    set dstintf port1
    set scraddr User_network
    set dstaddr all
    set action accept
    set identity-based enable
    set ip-based enable
    config identity-based-policy
      edit 1
        set groups Internal_users
        set service ANY
        set schedule always
      end
    end
end
```

Per session authentication

If you don't select *IP Based* the FortiGate unit applies HTTP authentication per session. This authentication is browser-based (see [Figure 339 on page 2817](#)). When a user enters a user name and password in their browser to authenticate with the explicit web proxy, this information is stored by the browser in a session cookie. Each new session started by the same web browser uses the session cookie for authentication. When the session cookie expires the user has to re-authenticate. If the user starts another browser on the same PC or closes and then re-opens their browser they have to authenticate again.

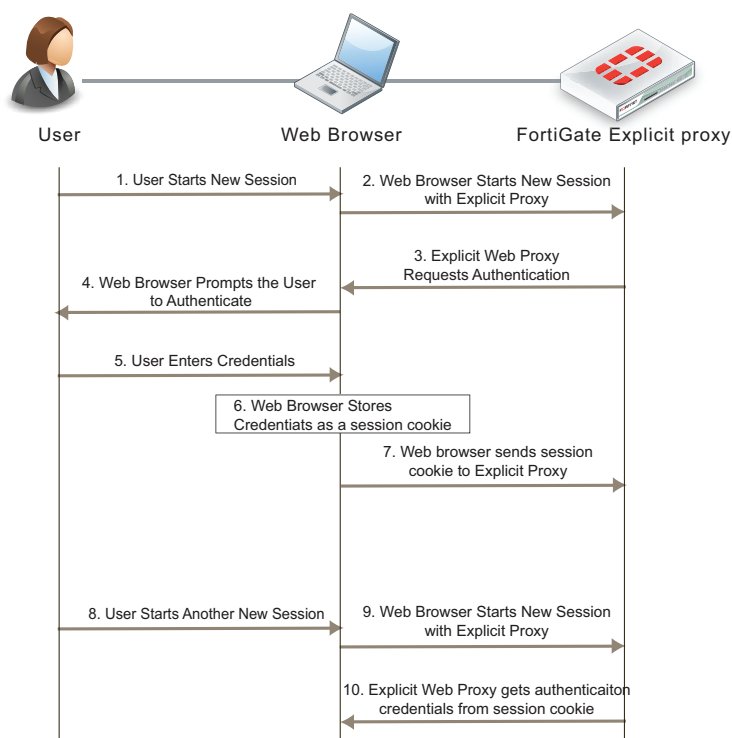
Since the authentication is browser-based, multiple clients with the same IP address can authenticate with the proxy using their own credentials. HTTP authentication provides authentication for multiple user sessions from the same source IP address. This can happen if there is a NAT device between the users and the FortiGate unit. HTTP authentication also supports authentication for other configurations that share one IP address among multiple users. These include Citrix products and Windows Terminal Server and other similar virtualization solutions.

To configure per session authentication, add a security policy for the explicit web proxy, set the source interface/zone to web-proxy, select Enable Identify Based Policy, and make sure IP Based is not selected before adding identity-based policies. You can also set the authentication method to basic, digest, NTLM or FSAE.

Use the following CLI command to add per session authentication to a security policy. Per session authentication is selected by setting `ip-based` to `disable`.

```
config firewall policy
edit 5
set srcintf web-proxy
set dstintf port1
set scraddr User_network
set dstaddr all
set action accept
set identity-based enable
set ip-based disable
config identity-based-policy
edit 1
set groups Internal_users
set service ANY
set schedule always
end
end
```

Figure 339: Per session HTTP authentication



UTM features and the explicit web proxy

You can apply protocol options, antivirus, web filtering, FortiGuard Web Filtering and data leak prevention (DLP) including DLP archiving to explicit web proxy sessions. UTM features are applied by selecting them in a web proxy security policy or an identity based policy in a web proxy security policy. You cannot apply intrusion protection (IPS), email filtering, application control, or VoIP UTM features to explicit web proxy sessions.

To apply intrusion protection to explicit web proxy traffic you can add DoS policies to the FortiGate interfaces that receive and send explicit web proxy traffic. However, you cannot apply application control to explicit web proxy traffic, so you cannot filter explicit web proxy traffic by application and explicit web proxy traffic does not contribute to application control monitoring or reporting.

Explicit web proxy sessions and flow-based scanning

Flow-based scanning cannot be applied to explicit web proxy sessions. This includes:

- IPS
- Application control
- Flow-based virus scanning
- Flow-based web filtering
- Flow-based FortiGuard web filtering
- Flow-based Data leak prevention (DLP)

Explicit web proxy sessions and protocol options

Since the traffic accepted by the explicit web proxy is known to be either HTTP, HTTPS, or FTP over HTTP and since the ports are already known by the proxy, the explicit web proxy does not use the HTTP or HTTPS port protocol options settings.

When adding UTM features to a web proxy security policy, you must select a protocol options profile. In most cases you can select the default protocol options profile. You could also create a custom protocol options profile.

The explicit web proxy supports the following protocol options:

- Enable chunked bypass
- HTTP oversized file action and threshold

The explicit web proxy does not support the following protocol options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit web proxy users are not added to dashboard usage and log and archive statistics widgets.

Explicit web proxy sessions web filtering and FortiGuard web filtering

For explicit web proxy sessions, the FortiGate unit applies web filtering to an HTTP request when it receives the headers of the request. If web filtering allows the HTTP request, it is forwarded to the web server. If web filtering blocks the HTTP request, the request is dropped and a blocking HTTP response is generated by the FortiGate unit and returned to the client web browser.

The explicit web proxy completely supports all web filtering and FortiGuard web filtering options and their configuration settings.

The web page displayed when FortiGuard Web Filtering blocks a web page through the explicit web proxy may be different than the page displayed through a normal firewall session.

Explicit web proxy sessions and HTTPS deep scanning

HTTPS explicit web proxy sessions are subject to web filtering and FortiGuard web filtering if you select *HTTPS scanning* in the web filtering profile applied to the explicit web proxy HTTPS sessions.

However, HTTPS deep scanning (also called SSL content scanning and inspection) is not supported for explicit web proxy HTTPS sessions. This means that web filtering of HTTPS sessions will only be done based on the URL in the session certificate.

The following features are also not supported for explicit web proxy HTTPS sessions:

- Virus scanning
- Data leak prevention (DLP)
- Enabling deep scanning in a protocol options profile

Explicit web proxy sessions and antivirus

For explicit web proxy sessions, the FortiGate unit applies antivirus scanning to HTTP POST requests and HTTP responses. The FortiGate unit starts virus scanning a file in an HTML session when it receives a file in the body of an HTML request. The explicit web proxy can receive HTTP responses from either the originating web server or the FortiGate web cache module.

Flow-based virus scanning is not available for explicit web proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit web proxy sessions use the regular virus database.

Web Proxy Services

Use the Web Proxy Service menu to configure multiple web proxy services for the web proxy feature. Web proxy services can only be selected in a security policy when *web-proxy* is selected as the source interface.

Web proxy services are similar to standard firewall services. You can configure web proxy services to define one or more protocols and port numbers that are associated with each web proxy service. Web proxy services can also be grouped into web proxy service groups.

The following are web proxy service configuration settings in *Firewall Objects > Service > Web Proxy Service*.

Name	Enter a name for the web proxy service.
Protocol	Select a protocol from the drop-down list.
Source Port	Enter the low and high source ports.
Low	Enter the lowest source port.
High	Enter the highest source port.
Destination Port	Enter the low and high destination ports.
Low	Enter the lowest destination port.
High	Enter the highest destination port.

Web Proxy Service Groups


In *Firewall Objects > Service > Web Proxy Service Group*, you can configure groups of web proxy services. By using service groups, you can efficiently apply several web proxy services to a security policy. This feature is applicable only when *web-proxy* is selected as the source interface in a security policy.

Group Name	Enter the name of the group.
Available Services	The web proxy services that can be members of the group. Select a web proxy service and use the -> arrow to move it to <i>Members</i> . Repeat for each service that you want to be a member in the group.
Members	The members of a group. To remove a member from <i>Members</i> , select the service and then use the <- arrow to move it back to <i>Available Services</i> .

Example: users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering and virus scanning

This example describes how to configure the explicit web proxy for the example network shown in [Figure 340](#). In this example, users on the internal network connect to the explicit web proxy through the Internal interface of the FortiGate-51B unit. The explicit web proxy is configured to use port 8888 so users must configure their web browser proxy settings to use port 8888 and IP address 10.31.101.100.

Figure 340: Example explicit web proxy network topology



In this example, explicit web proxy users must authenticate with a RADIUS server before getting access to the proxy. To apply authentication, the security policy that accepts explicit web proxy traffic includes an identity based policy that applies per session authentication to explicit web proxy users and includes a user group with the RADIUS server in it. The identity based policy also applies UTM web filtering and virus scanning.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Enable the explicit web proxy for HTTP and HTTPS and change the HTTP and HTTPS ports to 8888.
- 2 Enable the explicit web proxy on the internal interface.
- 3 Add a RADIUS server and user group for the explicit web proxy.
- 4 Add web filtering and antivirus profiles for the explicit web proxy.
- 5 Add a security policy for the explicit web proxy.

Enable web caching in the security policy.

Add an identity based policy to the security policy to support authentication. Enable antivirus, web filter, and DLP UTM features for the identity-based policy.

Configuring the explicit web proxy - web-based manager

Use the following steps to configure the explicit web proxy from FortiGate web-based manager.

To enable and configure the explicit web proxy - web-based manager

- 1 Go to *System > Network > Explicit Proxy* and change the following settings:

Enable Explicit Web Proxy	Select <i>HTTP/HTTPS</i> .
Listen on Interfaces	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
HTTP Port	8888
HTTPS Port	8888
Realm	You are authenticating with the explicit web proxy.
Default Firewall Policy Action	Deny

- 2 Select *Apply*.

To enable the explicit web proxy on the Internal interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Edit the internal interface.
- 3 Select *Enable Explicit Web Proxy*.
- 4 Select *OK*.

To add a RADIUS server and user group for the explicit web proxy - web-based manager

- 1 Go to *User > Remote > RADIUS*.

- 2 Select Create New to add a new RADIUS server:

Name	RADIUS_1
Type	Query
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

- 3 Go to *User > User Group > User Group* and select *Create New*.

Name	Explicit_proxy_user_group
Type	Firewall
Members	RADIUS_1

- 4 Select OK.

To add a security policy for the explicit web proxy - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.

- 2 Add a firewall address for the internal network:

Address Name	Internal_subnet
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.[1-255]
Interface	Any

- 3 Go to *Policy > Policy > Policy* and select *Create New*.

- 4 Configure the explicit web proxy security policy.

Source Interface/Zone	web-proxy
Source Address	Internal_subnet
Destination Interface/Zone	wan1
Destination Address	all
Action	ACCEPT

- 5 Select *Enable web cache* to enable web caching for the explicit web proxy.
- 6 Select *Enable Identity Based Policy*, make sure *IP Based* is not selected and *Auth Method* is set to *Basic*.
- 7 Select Add and configure the following settings for the identity based policy:

User Group	Explicit_policy
Service	webproxy
UTM	Select
Enable Antivirus	default
Enable Web Filter	default
Enable DLP Sensor	default
Protocol Options	default

You can also create custom antivirus, web filter, and DLP profiles instead of using the defaults.

- 8 Select OK.

Configuring the explicit web proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable the explicit web proxy on the Internal interface - CLI

- 1 Enter the following command to enable the explicit web proxy on the internal interface.

```
config system interface
  edit internal
    set explicit-web-proxy enable
  end
```

To enable and configure the explicit web proxy - CLI

- 1 Enter the following command to enable the explicit web proxy and set the TCP port that proxy accepts HTTP and HTTPS connections on to 8888.

```
config web-proxy explicit
  set status enable
  set http-incoming-port 8888
  set https-incoming-port 8888
  set realm "You are authenticating with the explicit web proxy"
  set sec-default-action deny
end
```

To add a RADIUS server and user group for the explicit web proxy - CLI

- 1 Enter the following command to add a RADIUS server:

```
config user radius
  edit RADIUS_1
    set server 10.31.101.200
    set secret RADIUS_server_secret
  end
```

- 2 Enter the following command to add a user group for the RADIUS server.

```
config user group
  edit Explicit_proxy_user_group
    set group-type firewall
    set member RADIUS_1
  end
```

To add a security policy for the explicit web proxy - CLI

- 1 Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
  edit Internal_subnet
    set type iprange
    set start-ip 10.31.101.1
    set end-ip 10.31.101.255
  end
```

- 2 Enter the following command to add the explicit web proxy security policy:

```
config firewall policy
  edit 0
    set srcintf web-proxy
```

```
set dstintf wan1
set srcaddr Internal_subnet
set dstaddr all
set action accept
set schedule always
set webcache enable
set identity-based enable
set ipbased disable
set auth-method basic
config identity-based-policy
edit 1
    set groups Explicit_Proxy_user_group
    set schedule always
    set service explicit-web
    set utm-status enable
    set av-profile Explicit_av_pro
    set webfilter-profile Explicit_wf_pro
    set dlp-sensor default
    set profile-protocol-options default
end
end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit web proxy configuration is working as expected:

To test the explicit web proxy configuration

- 1 Configure a web browser on the internal subnet to use a web proxy server at IP address 10.31.101.100 and port 8888.
- 2 Browse to an Internet web page.
The web browser should pop up an authentication window that includes the phrase that you added to the Realm option.
- 3 Enter the username and password for an account on the RADIUS server.
If the account is valid you should be allowed to browse web pages on the Internet.
- 4 Close the browser and clear its cache and cookies.
- 5 Restart the browser and connect to the Internet.
You could also start a second web browser on the same PC. Or you could start a new instance of the same browser as long as the browser asks for a user name and password again.
You should have to authenticate again because identity-based policies are set to session-based authentication.
- 6 If this basic functionality does not work, check your FortiGate and web browser configuration settings.
- 7 Browse to a URL on the URL filter list and confirm that the web page is blocked.
- 8 Browse to <http://eicar.org> and attempt to download an anti-malware test file.
The antivirus configuration should block the file.
Sessions for web-proxy security policies do not appear on the Top Sessions dashboard widget and the count column for security policies does not display a count for explicit web proxy security policies.

9 You can use the following command to display explicit web proxy sessions

```
get test wad 60
IP based users:

Session based users:
    user:0x9c20778, username:User1, vf_id:0, ref_cnt:9

Total allocated user:1

Total user count:3, shared user quota:50, shared user count:3
This command output shows one explicit proxy user with user name User1
authenticated using session-based authentication.
```

Explicit proxy sessions and user limits

Web browsers and web servers open and close multiple sessions with the explicit web proxy. Some sessions open and close very quickly. HTTP 1.1 keepalive sessions are persistent and can remain open for long periods of time. Sessions can remain on the explicit web proxy session list after a user has stopped using the proxy (and has, for example, closed their browser). If an explicit web proxy session is idle for more than 3600 seconds it is torn down by the explicit web proxy. See [RFC 2616](#) for information about HTTP keepalive/persistent HTTP sessions.

This section describes proxy sessions and user limits for both the explicit web proxy and the explicit FTP proxy. Session and user limits for the two proxies are counted and calculated together. However, in most cases if both proxies are active there will be many more web proxy sessions than FTP proxy sessions.

The FortiGate unit adds two sessions to its session table for every explicit proxy session started by a web browser and every FTP session started by an FTP client. An entry is added to the session table for the session from the web browser or client to the explicit proxy. All of these sessions have the same destination port as the explicit web proxy port (usually 8080 for HTTP and 21 for FTP). An entry is also added to the session table for the session between the exiting FortiGate interface and the web or FTP server destination of the session. All of these sessions have a FortiGate interface IP address and the source address of the session and usually have a destination port of 80 for HTTP and 21 for FTP.

Proxy sessions that appear in the Top sessions dashboard widget do not include the Policy ID of the web-proxy or ftp-proxy security policy that accepted them. However, the explicit proxy sessions appear in the Top Sessions dashboard widget with a destination port that matches the explicit proxy port number (usually 8080 for the web proxy and 21 for the FTP proxy). The proxied sessions from the FortiGate unit have their source address set to the IP address of the FortiGate unit interface that the sessions use to connect to their destinations (for example, for connections to the Internet the source address would be the IP address of the FortiGate interface connected to the Internet).

FortiOS limits the number of explicit proxy users. This includes both explicit FTP proxy and explicit web proxy users. The number of users varies by FortiGate model from as low as 10 to up to 18000 for high end models. You can use the following command to display the limit on the number of explicit web proxy users for a FortiGate unit:

```
get test wad 62

Total user count:1, shared user quota:500, shared user count:1
form_auth_keepalive=0 vd=root max=0 guarantee=0 used=1
```

This command output shows that the explicit proxy user limit (the shared user quota) for this FortiGate unit is 500 users.

You cannot change this limit. If your FortiGate unit is configured for multiple VDOMs this limit must be shared by all VDOMs. You can also use VDOM resource limiting to limit the number of explicit proxy users for the FortiGate unit and for each VDOM. To limit the number of explicit proxy users for the FortiGate unit from the web-based manager enable multiple VDOMs and go to *System > VDOM > Global Resources* set the number of Concurrent explicit proxy users or use the following command:

```
config global
  config system resource-limits
    set proxy 50
  end
end
```

To limit the number of explicit proxy users for a VDOM, from the web-based manager enable multiple VDOMs and go to *System > VDOM > VDOM* and edit a VDOM or use the following command to change the number of explicit web proxy users for VDOM_1:

```
config global
  config system vdom-property
    edit VDOM_1
      set proxy 25
    end
  end
```

The VDOM resource limit pages on the web-based manager also display the current number of explicit web proxy users. You can also use the `get test wad 60` CLI command to view the number of explicit web proxy users. For example:

```
get test wad 60
IP based users:
  user:0x9ab8350 username:User1, vf_id:0,
ip_addr:10.31.101.10, ref_cnt:9

Session based users:
  user:0x9ac3c40, username:User2, vf_id:0, ref_cnt:3
  user:0x9ab94f0, username:User3, vf_id:0, ref_cnt:1

Total allocated user:3
```

```
Total user count:3, shared user quota:50, shared user count:3
```

Users may be displayed with this command even if they are no longer actively using the proxy. All idle sessions time out after 3600 seconds.

The command output shows three explicit proxy users. The user named User1 has authenticated with a security policy that includes IP-based authentication and the user's source IP address is 10.31.101.10. The users named User2 and User3 have authenticated with a security policy that includes session-based authentication.

You can use the following command to flush all current explicit proxy users. This means delete information about all users and force them re-authenticate.

```
get test wad 61
```



Users that authenticate with explicit web-proxy or ftp-proxy security policies do not appear in the *User > Monitor > Firewall* list and selecting *De-authenticate All Users* has no effect on explicit proxy users.

How the number of concurrent explicit proxy users is determined depends on their authentication method:

- For session-based authenticated users, each authenticated user is counted as a single user. Since multiple users can have the same user name, the proxy attempts to identify users according to their authentication membership (based upon whether they were authenticated using RADIUS, LDAP, FSAE, local database etc.). If a user of one session has the same name and membership as a user of another session, the explicit proxy assumes this is one user.
- For IP Based authentication, or no authentication, or if no web-proxy security policy has been added, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

The explicit proxy does not limit the number of active sessions for each user. As a result the actual explicit proxy session count is usually much higher than the number of explicit web proxy users. If an excessive number of explicit web proxy sessions is compromising system performance you can limit the amount of users if the FortiGate unit is operating with multiple VDOMs.

Explicit web proxy configuration options

Use the explicit web proxy to enable explicit HTTP and HTTPS proxying on one or more FortiGate interfaces. The explicit web proxy also supports proxying FTP sessions sent from a web browser and proxy auto-config (PAC) to provide automatic proxy configurations for explicit web proxy users. From the CLI, you can also configure the explicit web proxy to support SOCKS sessions sent from a web browser. See [“Explicit web proxy configuration overview” on page 2809](#).

To configure the explicit web proxy, go to *System > Network > Explicit Proxy > Explicit Web Proxy Options*:



For explicit FTP proxy options, see [“Explicit FTP proxy options” on page 2844](#).

Explicit Web Proxy Options

Use the following options to configure the explicit web proxy.

Enable Explicit Web Proxy	Enable the explicit web proxy server for HTTP/ HTTPS, FTP and proxy auto-config PAC sessions. You must select this option for the explicit web proxy to accept and forward packets. FTP and PAC is only supported from a web browser and not a standalone client.
Listen on Interfaces	Displays the interfaces that are being monitored by the explicit web proxy. If VDOMs are enabled, only interfaces that belong to the current VDOM and have explicit web proxy enabled will be displayed. If you enable the web proxy on an interface that has VLANs on it, the explicit web proxy will not be enabled for those VLAN interfaces. You must use the CLI to enable the explicit proxy for VLAN interfaces.

HTTP Port HTTPS Port FTP Port PAC Port	<p>Enter the port number that traffic from client web browsers use to connect to the explicit proxy for the specific protocol. Explicit proxy users must configure their web browser's protocols proxy settings to use this port.</p> <p>The default value of 0 means use the same port as the configured HTTP port.</p>
PAC File Content	<p>Select the <i>Edit</i> icon to change the contents in a PAC file, or import a PAC file.</p> <p>The maximum PAC file size is 8192 bytes.</p> <p>You can use any PAC file syntax that is supported by your users's browsers. The FortiGate unit does not parse the PAC file.</p> <p>To use PAC, users must add an automatic proxy configuration URL (or PAC URL) to their web browser proxy configuration. The default PAC file URL is:</p> <p>http://<interface_ip>:<PAC_port_int>/<pac_file_str></p> <p>For example, if the interface with the explicit web proxy has IP address 172.20.120.122, the PAC port is the same as the default HTTP explicit proxy port (8080) and the PAC file name is proxy.pac the PAC file URL would be:</p> <p>http://172.20.120.122:8080/proxy.pac</p> <p>From the CLI you can use the following command to display the PAC file url:</p> <p>get web-proxy explicit</p>
Proxy FQDN	<p>Enter the fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server.</p>
Max HTTP request length	<p>Enter the maximum length of an HTTP request. Larger requests will be rejected.</p>
Max HTTP message length	<p>Enter the maximum length of an HTTP message. Larger messages will be rejected.</p>
Unknown HTTP version	<p>Select the action to take when the proxy server must handle an unknown HTTP version request or message.</p> <p><i>Best Effort</i> attempts to handle the HTTP traffic as best as it can.</p> <p><i>Reject</i> treats known HTTP traffic as malformed and drops it.</p>

Realm	<p>Enter an authentication realm to identify the explicit web proxy. The realm can be any text string of up to 63 characters. If the realm includes spaces enclose the name in quotes.</p> <p>When a user authenticates with the explicit proxy the HTTP authentication dialog includes the realm so you can use the realm to identify the explicitly web proxy for your users.</p>
Default Firewall Policy Action	<p>Configure the explicit web proxy to block (deny) or accept sessions if security policies have not been added for the explicit web proxy. To add security policies for the explicit web proxy add a security policy and set the source interface to <i>web-proxy</i>.</p> <p>The default setting (or <i>Deny</i>) blocks access to the explicit web proxy before adding a security policy. If you set this option to <i>Accept</i> the explicit web proxy server accepts sessions even if you haven't defined a security policy.</p>

Web Proxy Forwarding Servers Options

Use the following options to configure web proxy forwarding servers.

Create New	Select to create a new forwarding server.
Service Name	The name of the forwarding server.
Address	The IP address of the forwarding server.
Port	The port number of the forwarding server.
Health Check	Indicates whether the health check is disabled or enabled for that forwarding server. A green checkmark indicates health check is enabled; a gray x indicates that health check is disabled.
Server Down	The action that the FortiGate unit will take when the server is down.
Ref.	Displays the number of times the object is referenced to other objects.

Adding Web Proxy Forwarding Servers

To add a forwarding server, select *Create New* in the *Web Proxy Forwarding Servers* section of the *Explicit Proxy* page by going to *System > Network > Explicit Proxy*.

Server Name	Enter the name of the forwarding server.
Proxy Address	Enter the IP address of the forwarding server.
Proxy Address Type	Select the type of IP address of the forwarding server. A forwarding server can have an FQDN or IP address.
Port	Enter the port number.
Server Down action	Select what action the FortiGate unit will take if the forwarding server is down.
Enable Health Monitor	Select to enable health check monitoring.
Health Check Monitor Site	Enter the URL address of the health check monitoring site.

Restricting the IP address of the explicit web proxy

You can use the following command to restrict access to the explicit web proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require users to connect to the IP address 10.31.101.100 to connect to the explicit HTTP proxy:


```
config web-proxy explicit
  set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config http-proxy explicit
  set outgoing-ip 172.20.120.100
end
```



The FortiGate explicit FTP proxy

You can use the FortiGate explicit FTP proxy to enable explicit FTP proxying on one or more FortiGate interfaces. The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



Explicit FTP proxies are configured for each VDOM when multiple VDOMs are enabled.

In most cases you would configure the explicit FTP proxy for users on a network by enabling the explicit FTP proxy on the FortiGate interface connected to that network. Users on the network would connect to and authenticate with the explicit FTP proxy before connecting to an FTP server. In this case the IP address of the explicit FTP proxy is the IP address of the FortiGate interface on which the explicit FTP proxy is enabled.




Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate unit management IP address.

The FTP proxy receives FTP sessions to be proxied at FortiGate interfaces with the explicit FTP proxy enabled. The FTP proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit FTP proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address.

Figure 341: Example explicit FTP proxy topology



To allow anyone to anonymously log into explicit FTP proxy and connect to any FTP server you can set the explicit FTP proxy default firewall proxy action to accept. When you do this, users can log into the explicit FTP proxy with any username and password.

In most cases you would want to use security policies to control explicit FTP proxy traffic and apply security features such as access control/authentication, UTM, and traffic logging. You can do this by keeping the default explicit FTP proxy firewall policy action to deny and then adding ftp-proxy security policies. In most cases you would also want users to authenticate with the explicit FTP proxy. By default an anonymous FTP login is required. Usually you would add authentication, in the form of identity based policies, to ftp-proxy security policies. Users can then authenticate with the explicit FTP proxy according to user groups added to the identity based policies. User groups added to FTP proxy identity based policies can use any authentication method supported by FortiOS including the local user database and RADIUS and other remote servers.

If you leave the default firewall policy action set to deny and add ftp-proxy security policies, all connections to the explicit FTP proxy must match an ftp-proxy security policy or else they will be dropped. Sessions that are accepted are processed according to the ftp-proxy security policy settings.

You can also change the explicit FTP proxy default firewall policy action to accept and add explicit FTP proxy security policies. If you do this, sessions that match ftp-proxy security policies are processed according to the security policy settings. Connections to the explicit FTP proxy that do not match an ftp-proxy security policy are allowed and the users can authenticate with the proxy anonymously user any username and password.

There are some limitations to the UTM features that can be applied to explicit web proxy sessions. See [“UTM features and the explicit FTP proxy” on page 2837](#).

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit FTP proxy traffic. Security policies for the FTP proxy can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to *any*. (On the web-based manager you must set the interface to *Any*. In the CLI you must `unset the associated-interface`.)

This section describes:

- [How to use the explicit FTP proxy to connect to an FTP server](#)
- [Explicit FTP proxy configuration overview](#)
- [UTM features and the explicit FTP proxy](#)
- [Example: users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning](#)
- [Explicit FTP proxy sessions and user limits](#)
- [Explicit FTP proxy options](#)

How to use the explicit FTP proxy to connect to an FTP server

To connect to an FTP server using the explicit FTP proxy, users must run an FTP client and connect to the IP address of a FortiGate interface on which the explicit FTP proxy is enabled. This connection attempt must use the configured explicit FTP proxy port number (default 21).

The explicit FTP proxy is not compatible with using a web browser as an FTP client. To use web browsers as FTP clients configure the explicit web proxy to accept FTP sessions.

The following steps occur when a user starts an FTP client to connect to an FTP server using the explicit FTP proxy. Any RFC-compliant FTP client can be used. This example describes using a command-line FTP client. Some FTP clients may require a custom FTP proxy connection script.

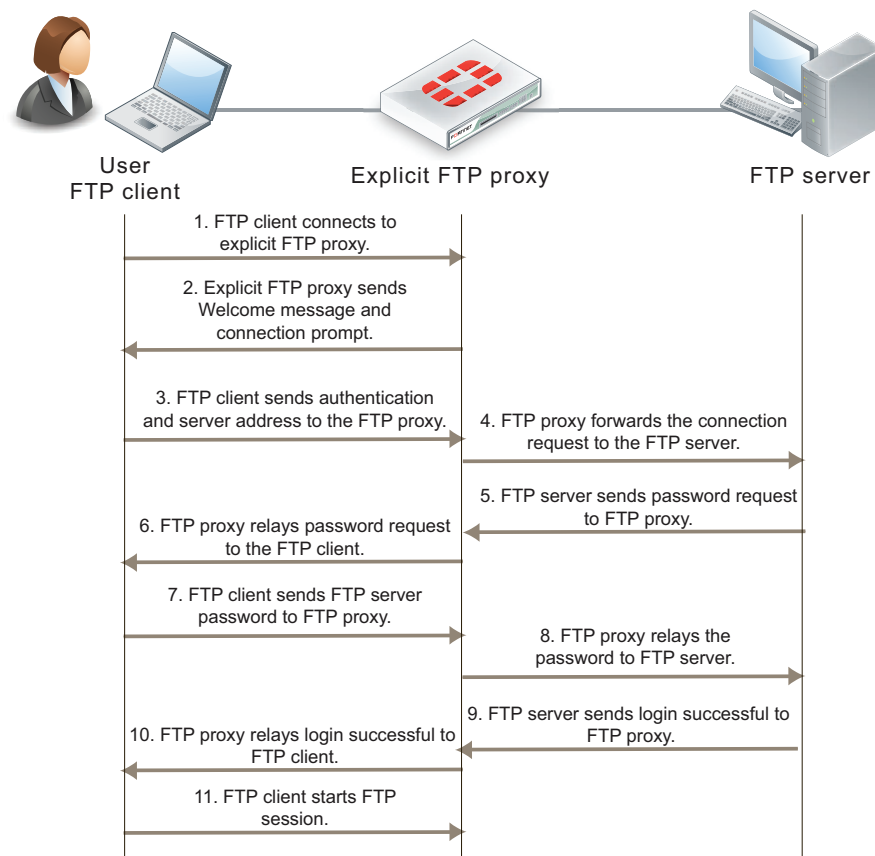
- 1 The user enters a command on the FTP client to connect to the explicit FTP proxy.
For example, if the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100, enter:
`ftp 10.31.101.100`
- 2 The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:
`Connected to 10.31.101.100.`
`220 Welcome to Fortigate FTP proxy`
`Name (10.31.101.100:user):`
You can change the message by editing the FTP Proxy replacement message.
- 3 At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:
`<proxy-user>:<proxy-password>:<server-user>@<server-address>`
For example, if the proxy username and password are p-name and p-pass and a valid username for the FTP server is s-name and the server's IP address is ftp.example.com the syntax would be:
`p-name:p-pass:s-name@ftp.example.com`



If the FTP proxy accepts anonymous logins p-name and p-pass can be any characters.

- 4 The FTP proxy forwards the connection request, including the user name, to the FTP server.
- 5 If the user name is valid for the FTP server it responds with a password request prompt.
- 6 The FTP proxy relays the password request to the FTP client.
- 7 The user enters the FTP server password and the client sends the password to the FTP proxy.
- 8 The FTP proxy relays the password to the FTP server.
- 9 The FTP server sends a login successful message to the FTP proxy.
- 10 The FTP proxy relays the login successful message to the FTP client.
- 11 The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

Figure 342: Explicit FTP proxy session

From a simple command line FTP client connecting to an the previous sequence could appear as follows:

```

ftp 10.31.101.100 21
Connected to 10.31.101.100.
220 Welcome to Fortigate FTP proxy
Name (10.31.101.100:user): p-name:p-pass:s-name@ftp.example.com
331 Please specify the password.
Password: s-pass
230 Login successful.
Remote system type is UNIX
Using binary mode to transfer files.
ftp>
  
```

Explicit FTP proxy configuration overview

You can use the following general steps to configure the explicit FTP proxy.

To enable the explicit FTP proxy - web-based manager

- 1 Go to *System > Network > Explicit Proxy > Explicit FTP Proxy Options*. Select *Enable Explicit FTP Proxy* to turn on the explicit FTP proxy.

2 Select Apply.

The default explicit FTP proxy configuration has *Default Firewall Policy Action* set to *Deny* and requires you to add a security policy to allow access to the explicit FTP proxy. This configuration is recommended and is a best practice because you can use security policies to control access to the explicit web proxy and also apply security features such as logging, UTM, and authentication (by adding identity-based policies).

3 Go to *System > Network > Interface* and select one or more interfaces for which to enable the explicit web proxy. Edit the interface configuration and select *Enable Explicit FTP Proxy*.

Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

4 Go to *Policy > Policy > Policy* and select *Create New* and set the *Source Interface/Zone* to *ftp-proxy*.

You can add multiple ftp-proxy security policies.

5 Configure the security policy as required to accept the traffic that you want to be processed by the explicit web proxy.

The source address of the policy should match client source IP addresses. The firewall address selected as the source address cannot be assigned to a FortiGate interface. The Interface field of the firewall address must be blank or it must be set to *Any*.

The destination address of the policy should match the IP addresses of FTP servers that clients are connecting to. The destination address could be *all* to allow connections to any FTP server.

If *Default Firewall Policy Action* is set to *Deny*, traffic sent to the explicit FTP proxy that is not accepted by an ftp-proxy security policy is dropped. If *Default Firewall Policy Action* is set to *Allow* then all web-proxy sessions that don't match with a security policy are allowed.

For example the following security policy allows users on an internal network to access FTP servers on the Internet through the wan1 interface of a FortiGate unit.

Source Interface/Zone	ftp-proxy
Source Address	Internal_subnet
Destination Interface/Zone	wan1
Destination Address	all
Action	ACCEPT

6 You can select other security policy options as required.

For example, you can apply UTM protection to web proxy sessions and log allowed ftp proxy traffic.

7 You can also select *Enable Identity Based Policy* to apply authentication to explicit web proxy sessions.

To add identity based policies you must first add user groups to the FortiGate authentication configuration.

- 8 You can add multiple identity based policies to apply different authentication for different user groups and also apply different UTM and logging settings for different user groups.

To enable the explicit web proxy - CLI

- 1 Enter the following command to turn on the explicit FTP proxy. This command also changes the explicit FTP proxy port to 2121.

```
config ftp-proxy explicit
  set status enable
  set incoming-port 2121
end
```

The default explicit FTP proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit FTP proxy.

- 1 Enter the following command to enable the explicit FTP proxy for the internal interface.

```
config system interface
  edit internal
    set explicit-ftp-proxy enable
  end
end
```

- 2 Use the following command to add a firewall address that matches the source address of users who connect to the explicit FTP proxy.

```
config firewall address
  edit Internal_subnet
    set type iprange
    set start-ip 10.31.101.1
    set end-ip 10.31.101.255
  end
```

The source address for a ftp-proxy security policy cannot be assigned to a FortiGate unit interface.

- 3 Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall policy
  edit 2
    set srcintf ftp-proxy
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr all
    set action accept
    set identity-based enable
    set schedule always
    config identity-based-policy
      edit 1
        set groups Internal_users
        set utm-status enable
        set profile-protocol-options default
        set av-profile Scan
        set logtraffic enable
        set schedule always
        set service ANY
      end
    end
```

```
end
end
```

The firewall address selected as the source address cannot be assigned to a FortiGate unit interface. Either the field must be blank or it must be set to *Any*.

Restricting the IP address of the explicit FTP proxy

You can use the following command to restrict access to the explicit FTP proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require users to connect to the IP address 10.31.101.100 to connect to the explicit FTP proxy:

```
config ftp-proxy explicit
  set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the explicit FTP proxy

You can use the following command to restrict the source address of outgoing FTP proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config ftp-proxy explicit
  set outgoing-ip 172.20.120.100
end
```

UTM features and the explicit FTP proxy

You can apply protocol options, antivirus, and data leak prevention (DLP) including DLP archiving to explicit FTP proxy sessions. UTM features are applied by selecting them in a ftp proxy security policy or an identity based policy in a FTP proxy security policy. You cannot apply intrusion protection (IPS), or application control to explicit FTP proxy sessions.

To apply intrusion protection to explicit FTP proxy traffic you can add DoS policies to the FortiGate interfaces that receive and send explicit FTP proxy traffic. However, you cannot apply application control to explicit FTP proxy traffic, so you cannot filter explicit FTP proxy traffic by application and explicit FTP proxy traffic does not contribute to application control monitoring or reporting.

Explicit FTP proxy sessions and protocol options

Since the traffic accepted by the explicit FTP proxy is known to be FTP and since the ports are already known by the proxy, the explicit FTP proxy does not use the FTP port protocol options settings.

When adding UTM features to an FTP proxy security policy, you must select a protocol options profile. In most cases you can select the default protocol options profile. You could also create a custom protocol options profile.

The explicit FTP proxy supports the following protocol options:

- FTP oversized file action and threshold

The explicit FTP proxy does not support the following protocol options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit FTP proxy users are not added to dashboard usage and log and archive statistics widgets.

Explicit FTP proxy sessions and antivirus


For explicit FTP proxy sessions, the FortiGate unit applies antivirus scanning to FTP file GET and PUT requests. The FortiGate unit starts virus scanning a file in an FTP session when it receives a file in the body of an FTP request.

Flow-based virus scanning is not available for explicit FTP proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit FTP proxy sessions use the regular virus database.

Example: users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning

This example describes how to configure the explicit FTP proxy for the example network shown in Figure 343. In this example, users on the internal network connect to the explicit FTP proxy through the Internal interface with IP address 10.31.101.100 of the FortiGate-51B unit. The explicit web proxy is configured to use port 2121 so to connect to an FTP server on the Internet users must first connect to the explicit FTP proxy using IP address 10.31.101.100 and port 2121.

Figure 343: Example explicit FTP proxy network topology



In this example, explicit FTP proxy users must authenticate with a RADIUS server before getting access to the proxy. To apply authentication, the security policy that accepts explicit FTP proxy traffic includes an identity based policy that applies per session authentication to explicit FTP proxy users and includes a user group with the RADIUS server in it. The identity based policy also applies UTM virus scanning and DLP.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

- 1 Enable the explicit FTP proxy and change the FTP port to 2121.
- 2 Enable the explicit FTP proxy on the internal interface.
- 3 Add a RADIUS server and user group for the explicit FTP proxy.
- 4 Add a security policy for the explicit FTP proxy.

Add an identity based policy to the security policy to support authentication. Enable antivirus and DLP UTM features for the identity-based policy.

Configuring the explicit FTP proxy - web-based manager

Use the following steps to configure the explicit FTP proxy from FortiGate web-based manager.

To enable and configure the explicit FTP proxy - web-based manager

- 1 Go to *System > Network > Explicit Proxy > Explicit FTP Proxy Options* and change the following settings:

Enable Explicit FTP Proxy	Select.
Listen on Interface	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
FTP Port	2121
Default Firewall Policy Action	Deny

- 2 Select *Apply*.

To enable the explicit FTP proxy on the Internal interface - web-based manager

- 1 Go to *System > Network > Interface*.
- 2 Edit the internal interface.
- 3 Select *Enable Explicit FTP Proxy*.
- 4 Select *OK*.

To add a RADIUS server and user group for the explicit FTP proxy - web-based manager

- 1 Go to *User > Remote > RADIUS*.
- 2 Select *Create New* to add a new RADIUS server:

Name	RADIUS_1
Type	Query
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

- 3 Go to *User > User Group > User Group* and select *Create New*.

Name	Explicit_proxy_user_group
Type	Firewall
Members	RADIUS_1

- 4 Select OK.

To add a security policy for the explicit FTP proxy - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select Create New.
- 2 Add a firewall address for the internal network:

Address Name	Internal_subnet
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.[1-255]
Interface	Any

- 3 Go to *Policy > Policy > Policy* and select Create New.
- 4 Configure the explicit FTP proxy security policy.

Source Interface/Zone	ftp-proxy
Source Address	Internal_subnet
Destination Interface/Zone	wan1
Destination Address	all
Action	ACCEPT

- 5 Select Enable Identity Based Policy, make sure *IP Based* is not selected and *Auth Method* is set to *Basic*.
- 6 Select Add and configure the following settings for the identity based policy:

User Group	Explicit_policy
UTM	Select
Enable Antivirus	default
Enable DLP Sensor	default
Protocol Options	default

- 7 Select OK.

Configuring the explicit FTP proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable and configure the explicit FTP proxy - CLI

- 1 Enter the following command to enable the explicit FTP proxy and set the TCP port that proxy accepts FTP connections on to 2121.

```
config ftp-proxy explicit
set status enable
set incoming-port 2121
set sec-default-action deny
end
```

To enable the explicit FTP proxy on the Internal interface - CLI

- 1 Enter the following command to enable the explicit FTP proxy on the internal interface.

```
config system interface
```

```
edit internal
set explicit-ftp-proxy enable
end
```

To add a RADIUS server and user group for the explicit FTP proxy - CLI

- 1 Enter the following command to add a RADIUS server:

```
config user radius
edit RADIUS_1
set server 10.31.101.200
set secret RADIUS_server_secret
end
```

- 2 Enter the following command to add a user group for the RADIUS server.

```
config user group
edit Explicit_proxy_user_group
set group-type firewall
set member RADIUS_1
end
```

To add a security policy for the explicit FTP proxy - CLI

- 1 Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
edit Internal_subnet
set type iprange
set start-ip 10.31.101.1
set end-ip 10.31.101.255
end
```

- 2 Enter the following command to add the explicit FTP proxy security policy:

```
config firewall policy
edit 0
set srcintf web-proxy
set dstintf wan1
set srcaddr Internal_subnet
set dstaddr all
set action accept
set schedule always
set identity-based enable
set ipbased disable
set auth-method basic
config identity-based-policy
edit 1
set groups Explicit_Proxy_user_group
set schedule always
set utm-status enable
set av-profile default
set dlp-sensor default
set profile-protocol-options default
end
end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit web proxy configuration is working as expected:

- 1 The user enters a command on the FTP client to connect to the explicit FTP proxy.
For example, if the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100, enter:
- 2 The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:

```
ftp 10.31.101.100  
Connected to 10.31.101.100.  
220 Welcome to Fortigate FTP proxy  
Name (10.31.101.100:user):
```

You can change the message by editing the FTP Proxy replacement message.

- 3 At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:

```
<proxy-user>:<proxy-password>:<server-user>@<server-address>
```

For example, if the proxy username and password are `p-name` and `p-pass` and a valid username for the FTP server is `s-name` and the server's IP address is `ftp.example.com` the syntax would be:

```
p-name:p-pass:s-name@ftp.example.com
```



If the FTP proxy accepts anonymous logins `p-name` and `p-pass` can be any characters.

- 4 The FTP proxy forwards the connection request, including the user name, to the FTP server.
- 5 If the user name is valid for the FTP server it responds with a password request prompt.
- 6 The FTP proxy relays the password request to the FTP client.
- 7 The user enters the FTP server password and the client sends the password to the FTP proxy.
- 8 The FTP proxy relays the password to the FTP server.
- 9 The FTP server sends a login successful message to the FTP proxy.
- 10 The FTP proxy relays the login successful message to the FTP client.
- 11 The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

To test the explicit web proxy configuration

- 1 From a system on the internal network start an FTP client and enter the following command to connect to the FTP proxy:

```
ftp 10.31.101.100
```

The explicit FTP proxy should respond with a message similar to the following:

```
Connected to 10.31.101.100.
```



```
220 Welcome to Fortigate FTP proxy
Name (10.31.101.100:user):
```

- 2 At the prompt enter a valid username and password for the RADIUS server followed by a user name for an FTP server on the Internet and the address of the FTP server. For example, if a valid username and password on the RADIUS server is ex_name and ex_pass and you attempt to connect to an FTP server at ftp.example.com with user name s_name, enter the following at the prompt:

```
Name
(10.31.101.100:user):ex_name:ex_pass:s_name@ftp.example.com
```

- 3 You should be prompted for the password for the account on the FTP server.
- 4 Enter the password and you should be able to connect to the FTP server.
- 5 Attempt to explore the FTP server file system and download or upload files.
- 6 To test UTM functionality, attempt to upload or download an ECAR test file. Or upload or download a tex file containing text that would be matched by the DLP sensor.

For eicar test files, go to <http://eicar.org>.

- 7 You can use the following command to display explicit web proxy sessions. The following output shows one explicit web proxy user with user name user_1 authenticated using IP-based authentication.

```
get test wad 60
users:
user:user_1@10.31.101.20(0x40d1b170), type:IP, vfid:0,
  ref_cnt:2, user:1(0x40f18020), user_ip:1(0x40db70f0),
  timeout:alive, id: 1

Total allocated user:1

Total user count:1, shared user quota:500, shared user count:1
  form_auth_keepalive=0
```

```
Explicit proxy authentication timeout: 300 sec, timeout
precision: 30000 msec
```

The following output shows one explicit web proxy user with user name user2 authenticated using IP-based authentication.

```
get test wad 60
users:
user:user2@10.31.101.20(0x40d1b170), type:SESSION, vfid:0,
  ref_cnt:2, user:1(0x40f18020), user_ip:1(0x40db70f0),
  timeout:alive, id: 2

Total allocated user:1

Total user count:1, shared user quota:500, shared user count:1
  form_auth_keepalive=0
```

```
Explicit proxy authentication timeout: 300 sec, timeout precision:
30000 msec
```

Explicit FTP proxy sessions and user limits

FTP clients do not open large numbers of sessions with the explicit FTP proxy. Most sessions stay open for a short while depending on how long a user is connected to an FTP server and how large the file uploads or downloads are. So unless you have large numbers of FTP users, the explicit FTP proxy should not be adding large numbers of sessions to the session table.

Explicit FTP proxy sessions and user limits are combined with explicit web proxy session and user limits. For information about explicit proxy session and user limits, see [“Explicit proxy sessions and user limits” on page 2825](#).

Explicit FTP proxy options

To configure the explicit FTP proxy, go to *System > Network > Explicit Proxy > Explicit FTP Proxy Options* and configure explicit FTP proxy options. For more information about the explicit FTP proxy, see [“Explicit FTP proxy configuration overview” on page 2834](#).

Enable Explicit FTP Proxy	Select to enable the explicit FTP proxy.
Listen on Interface	Select <i>Edit</i> to select the interface that the FortiGate unit will use to listen on. If <i>None</i> displays, you need to go to the interface that you want to use to listen on and select <i>Enable FTP Proxy</i> .
FTP Port	Enter the port number for the FTP server.
Default Firewall Policy Action	apply the action that the default security policy will take with regards to the explicit FTP proxy activity.



FortiGate WCCP

The Web Cache Communication Protocol (WCCP) can be used to provide web caching with load balancing and fault tolerance. In a WCCP configuration, a WCCP server receives HTTP requests from user's web browsers and redirects the requests to one or more WCCP clients. The clients either return cached content or request new content from the destination web servers before caching it and returning it to the server which in turn returns the content to the original requestor. If a WCCP configuration includes multiple WCCP clients, the WCCP server load balances traffic among the clients and can detect when a client fails and failover sessions to still operating clients. WCCP is described by the [Web Cache Communication Protocol internet draft](#).

The sessions that are cached by WCCP depend on the configuration of the WCCP clients. If the client is a FortiGate unit, you can configure the port numbers and protocol number of the sessions to be cached. For example, to cache HTTPS traffic on port 443 the WCCP client port must be set to 443 and protocol must be set to 6. If the WCCP client should also cache HTTPS traffic on port 993 the client ports option should include both port 443 and 993.

WCCP sessions are accepted by a security policy before being cached. If the security policy that accepts sessions that do not match the port and protocol settings in the WCCP clients the traffic is dropped.

WCCP is configured per-VDOM. A single VDOM can operate as a WCCP server or client (not both at the same time). FortiGate units are compatible with third-party WCCP clients and servers. If a FortiGate unit is operating as an Internet firewall for a private network, you can configure it to cache and serve some or all of the web traffic on the private network using WCCP by adding one or more WCCP clients, configuring WCCP server settings on the FortiGate unit and adding WCCP to security policies that accept HTTP session from the private network.

FortiGate units support WCCPv1 and WCCPv2. A FortiGate unit in NAT/Route or transparent mode can operate as a WCCP server. To operate as a WCCP client a FortiGate unit must be in NAT/Route mode. FortiGate units communicate between WCCP servers and clients uses UDP port 2048. This communication can be encapsulated in a GRE tunnel or just use layer 2 forwarding.



A WCCP server can also be called a WCCP router. A WCCP client can also be called a WCCP cache engine.

This section describes:

- [WCCP service groups, service numbers, service IDs and well known services](#)
- [WCCP configuration overview](#)
- [Example: caching HTTP sessions on port 80 using WCCP](#)
- [Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP](#)
- [WCCP packet flow](#)
- [Configuring the forward and return methods and adding authentication](#)

- [WCCP Messages](#)
- [Troubleshooting WCCP](#)

WCCP service groups, service numbers, service IDs and well known services

A FortiGate unit configured as a WCCP server or client can include multiple server or client configurations. Each of these configurations is called a WCCP service group. A service group consists of one or more WCCP servers (or routers) and one or more WCCP clients working together to cache a specific type of traffic. The service group configuration includes information about the type of traffic to be cached, the addresses of the WCCP clients and servers and other information about the service.

A service group is identified with a numeric WCCP service ID (or service number) in the range 0 to 255. All of the servers and clients in the same WCCP service group must have service group configurations with the same WCCP service ID.

The value of the service ID provides some information about the type of traffic to be cached by the service group. Service IDs in the range 0 to 50 are reserved for well known services. A well known service is any service that is defined by the WCCP standard as being well known. Since the service is well known, just the service ID is required to identify the traffic to be cached.

Even though the well known service ID range is 0 to 50, at this time only one well known service has been defined. Its service ID 0, which is used for caching HTTP (web) traffic.

So to configure WCCP to cache HTTP sessions you can add a service group to the WCCP router and WCCP clients with a service ID of 0. No other information about the type of traffic to cache needs to be added to the service group.

Since service IDs 1 to 50 are reserved for well know services and since these services are not defined yet, you should not add service groups with IDs in the range 1 to 50.



FortiOS does allow you to add service groups with IDs between 1 and 50. Since these service groups have not been assigned well known services; however, they will not cache any sessions. Service groups with IDs 51 to 255 allow you to set the port numbers and protocol number of the traffic to be cached. So you can use service groups with IDs 51 to 255 to cache different kinds of traffic based on port numbers and protocol number of the traffic. Service groups 1 to 50; however, do not allow you to set port numbers or protocol numbers so cannot be used to cache any traffic.

To cache traffic other than HTTP traffic you must add service groups with IDs in the range 51 to 255. These service group configurations must include the port numbers and protocol number of the traffic to be cached. It is the port and protocol number configuration in the service group that determines what traffic will be cached by WCCP.

Example WCCP server and client configuration for caching HTTP sessions (service ID = 0)

Enter the following command to add a WCCP service group to a WCCP server that caches HTTP sessions. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service

ID of this service group is 0.

```
config system wccp
edit 0
set router-id 10.31.101.100
```

```

    set server-list 10.31.101.0 255.255.255.0
end

```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures the client to cache HTTP sessions. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group is 0.

```

config system settings
    set wccp-cache-engine enable
end

config system wccp
    edit 0
        set cache-id 10.31.101.1
        set router-list 10.31.101.100
    end

```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

Example WCCP server and client configuration for caching HTTPS sessions

Enter the following command to add a service group to a WCCP server that caches HTTPS content on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 80.

```

config system wccp
    edit 80
        set router-id 10.31.101.100
        set server-list 10.31.101.0 255.255.255.0
        set ports 443
        set protocol 6
    end

```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTPS sessions on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 80 to match the service ID added to the server.

```

config system settings
    set wccp-cache-engine enable
end

config system wccp
    edit 80
        set cache-id 10.31.101.1
        set router-list 10.31.101.100
        set ports 443
        set protocol 6
    end

```

Example WCCP server and client configuration for caching HTTP and HTTPS sessions

You could do this by configuring two WCCP service groups as described in the previous examples. Or you could use the following commands to configure one service group for both types of traffic. The example also caches HTTP sessions on port 8080.

Enter the following command to add a service group to a WCCP server that caches HTTP sessions on ports 80 and 8080 and HTTPS sessions on port 443. Both of these protocols use protocol number 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 90.

```
config system wccp
  edit 90
    set router-id 10.31.101.100
    set server-list 10.31.101.0 255.255.255.0
    set ports 443 80 8080
    set protocol 6
  end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTP sessions on port 80 and 8080 and HTTPS sessions on port 443. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 90 to match the service ID added to the server.

```
config system settings
  set wccp-cache-engine enable
end

config system wccp
  edit 90
    set cache-id 10.31.101.1
    set router-list 10.31.101.100
    set ports 443 80 8080
    set protocol 6
  end
```

Other WCCP service group options

In addition to using WCCP service groups to define the types of traffic to be cached by WCCP the following options are available for servers and clients.

Server configuration options

The server configuration must include the `router-id`, which is the WCCP server IP address. This is the IP address of the interface that the server uses to communicate with WCCP clients.

The `group-address` is used for multicast WCCP configurations to specify the multicast addresses of the clients.

The `server-list` defines the IP addresses of the WCCP clients that the server can connect to. Often the server list can be the address of the subnet that contains the WCCP clients.

The `authentication` option enables or disables authentication for the WCCP service group. Authentication must be enabled on all servers and clients in a service group and members of the group must have the same `password`.

The `forward-method` option specifies the protocol used for communication between the server and clients. The default forwarding method is GRE encapsulation. If required by your network you can also select to use unencapsulated layer-2 packets instead of GRE or select any to allow both. The `return-method` allows you to specify the communication method from the client to the server. Both GRE and layer-2 are supported.

The `assignment-method` determines how the server load balances sessions to the clients if there are multiple clients. Load balancing can be done using hashing or masking.

Client configuration options

The client configuration includes the `cache-id` which is the IP address of the FortiGate interface of the client that communicates with WCCP server. The `router-list` option is the list of IP addresses of the WCCP servers in the WCCP service group.

The `ports` option lists the port numbers of the sessions to be cached by the client and the `protocol` sets the protocol number of the sessions to be cached. For TCP sessions the protocol is 6.

The `service-type` option can be auto, dynamic or standard. Usually you would not change this setting.

The client configuration also includes options to influence load balancing including the `primary-hash`, `priority`, `assignment-weight` and `assignment-bucket-format`.

WCCP configuration overview

To configure WCCP you must create a service group that includes WCCP servers and clients. WCCP servers intercept sessions to be cached (for example, sessions from users browsing the web from a private network). To intercept sessions to be cached the WCCP server must include a security policy that accepts sessions to be cached and WCCP must be enabled in this security policy.

The server must have an interface configured for WCCP communication with WCCP clients. That interface sends and receives encapsulated GRE traffic to and from WCCP clients. The server must also include a WCCP service group that includes a service ID and the addresses of the WCCP clients as well as other WCCP configuration options.

To use a FortiGate unit as a WCCP client, the FortiGate unit must be set to be a WCCP client (or cache engine). You must also configure an interface on the client for WCCP communication. The client sends and receives encapsulated GRE traffic to and from the WCCP server using this interface.

The client must also include a WCCP service group with a service ID that matches a service ID on the server. The client service group also includes the IP address of the servers in the service group and specifies the port numbers and protocol number of the sessions that will be cached on the client.

When the client receives sessions from the server on its WCCP interface, it either returns cached content over the WCCP interface or connects to the destination web servers using the appropriate interface depending on the client routing configuration. Content received from web servers is then cached by the client and returned to the WCCP server over the WCCP link. The server then returns the received content to the initial requesting user web browser.

Finally you may also need to configure routing on the server and client FortiGate units and additional security policies may have to be added to the server to accept sessions not cached by WCCP.

Example: caching HTTP sessions on port 80 using WCCP

In this example configuration (shown in [Figure 344](#)), a FortiGate unit with host name WCCP_srv is operating as an Internet firewall for a private network is also configured as a WCCP server. The port1 interface of WCCP_srv is connected to the Internet and the port2 interface is connected to the internal network.

All HTTP traffic on port 80 that is received at the port2 interface of WCCP_srv is accepted by a port2 to port1 security policy with WCCP enabled. All other traffic received at the port2 interface is allowed to connect to the Internet by adding a general port2 to port1 security policy below the HTTP on port 80 security policy.


A WCCP service group is added to WCCP_srv with a service ID of 0 for caching HTTP traffic on port 80. The port5 interface of WCCP_srv is configured for WCCP communication.

A second FortiGate unit with host name WCCP_client is operating as a WCCP client. The port1 interface of WCCP_client is connected to port5 of WCCP_srv and is configured for WCCP communication.

WCCP_client is configured to cache HTTP traffic because it also has a WCCP service group with a service ID of 0.

WCCP_client connects to the Internet through WCCP_srv. To allow this, a port5 to port1 security policy is added to WCCP_srv.

Figure 344: FortiGate WCCP server and client configuration



Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

- 1 Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and is configured for WCCP:

```
config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service HTTP
    set wccp enable
    set nat enable
  end
```

- 2 Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
.config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
```

- 3 Move this policy below the WCCP policy in the port2 to port1 policy list.

- 4 Enable WCCP on the port5 interface.

```
config system interface
  edit port5
    set wccp enable
  end
```

- 5 Add a WCCP service group with service ID 0.

```
config system wccp
  edit 0
    set router-id 10.51.101.100
    set server-list 10.51.101.0 255.255.255.0
  end
```

- 6 Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```
config firewall address
  edit WCCP_client_addr
    set subnet 10.51.101.10
  end
```

```

config firewall policy
  edit 0
    set srtintf port5
    set dstintf port1
    set srcaddr WCCP_client_addr
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end

```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

- 1 Configure WCCP_client to operate as a WCCP client.

```

config system settings
  set wccp-cache-engine enable
end

```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

- 2 Enable WCCP on the port1 interface.

```

config system interface
  edit port1
    set wccp enable
  end

```

- 3 Add a WCCP service group with service ID 0.

```

config system wccp
  edit 0
    set cache-id 10.51.101.10
    set router-list 10.51.101.100
  end

```

Example: caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP

This example configuration is the same as that shown in [Figure 344](#) and described in “[Example: caching HTTP sessions on port 80 using WCCP](#)” on page 2850 except that WCCP now also cached HTTPS traffic on port 443. To cache HTTP and HTTPS traffic the WCCP service group must have a service ID in the range 51 to 255 and you must specify port 80 and 443 and protocol 6 in the service group configuration of the WCCP client.

Also the security policy on the WCCP_srv that accepts sessions from the internal network to be cached must accept HTTP and HTTPS sessions.

Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

- 1 Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and HTTPS traffic on port 443 and is configured for WCCP:

```
config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service HTTP HTTPS
    set wccp enable
    set nat enable
  end
```

- 2 Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
.config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
```

- 3 Move this policy below the WCCP policy in the port2 to port1 policy list.

- 4 Enable WCCP on the port5 interface.

```
config system interface
  edit port5
    set wccp enable
  end
```

- 5 Add a WCCP service group with service ID 90 (can be any number between 51 and 255).

```
config system wccp
  edit 90
    set router-id 10.51.101.100
    set server-list 10.51.101.0 255.255.255.0
  end
```

- 6 Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```
config firewall address
  edit WCCP_client_addr
    set subnet 10.51.101.10
```

```

end

.config firewall policy
edit 0
    set srtintf port5
    set dstintf port1
    set srcaddr WCCP_client_addr
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
end

```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

- 1 Configure WCCP_client to operate as a WCCP client.

```

config system settings
    set wccp-cache-engine enable
end

```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

- 2 Enable WCCP on the port1 interface.

```

config system interface
    edit port1
        set wccp enable
    end
end

```

- 3 Add a WCCP service group with service ID 90. This service group also specifies to cache sessions on ports 80 and 443 (for HTTP and HTTPS) and protocol number 6.

```

config system wccp
    edit 90
        set cache-id 10.51.101.10
        set router-list 10.51.101.100
        ports 80 443
        set protocol 6
    end
end

```

WCCP packet flow

The following packet flow sequence assumes you have configured a FortiGate unit to be a WCCP server and one or more FortiGate units to be WCCP clients.

- 1 A user's web browser sends a request for web content.

- 2 The FortiGate unit configured as a WCCP server includes a security policy that intercepts the request and forwards it to a WCCP client.
The security policy can apply UTM features to traffic accepted by the policy.
- 3 The WCCP client receives the WCCP session.
- 4 The client either returns requested content to the WCCP server if it is already cached, or connects to the destination web server, receives and caches the content and then returns it to the WCCP server.
- 5 The WCCP server returns the requested content to the user's web browser.
- 6 The WCCP router returns the request to the client web browser.
The client web browser is not aware that all this is taking place and does not have to be configured to use a web proxy.

Configuring the forward and return methods and adding authentication

The WCCP forwarding method determines how intercepted traffic is transmitted from the WCCP router to the WCCP cache engine. There are two different forwarding methods:

- GRE forwarding (the default) encapsulates the intercepted packet in an IP GRE header with a source IP address of the WCCP router and a destination IP address of the target WCCP cache engine. The result is a tunnel that allows the WCCP router to be multiple hops away from the WCCP cache server.
- L2 forwarding rewrites the destination MAC address of the intercepted packet to match the MAC address of the target WCCP cache engine. L2 forwarding requires that the WCCP router is Layer 2 adjacent to the WCCP client.

You can use the following command on a FortiGate unit configured as a WCCP router to change the forward and return methods to L2:

```
config system wccp
  edit 1
    set forward-method L2
    set return-method L2
  end
```

You can also set the forward and return methods to any in order to match the cache server configuration.

By default the WCCP communication between the router and cache servers is unencrypted. If you are concerned about attackers sniffing the information in the WCCP stream you can use the following command to enable hash-based authentication of the WCCP traffic. You must enable authentication on the router and the cache engines and all must have the same password.

```
config system wccp
  edit 1
    set authentication enable
    set password <password>
  end
```

WCCP Messages

When the WCCP service is active on a web cache server it periodically sends a WCCP HERE I AM broadcast or unicast message to the FortiGate unit operating as a WCCP router. This message contains the following information:

- Web cache identity (the IP address of the web cache server).
- Service info (the service group to join).

If the information received in the previous message matches what is expected, the FortiGate unit replies with a WCCP I SEE YOU message that contains the following details:

- Router identity (the FortiGate unit's IP address).
- Sent to IP (the web cache IP addresses to which the packets are addressed)

When both ends receive these two messages the connection is established, the service group is formed and the designated web cache is elected.

Troubleshooting WCCP

Two types of debug commands are available for debugging or troubleshooting a WCCP connection between a FortiGate unit operating as a WCCP router and its WCCP cache engines.

Real time debugging

The following commands can capture live WCCP messages:

```
diag debug en
diag debug application wccpd <debug level>
```

Application debugging

The following commands display information about WCCP operations:

```
get test wccpd <integer>
diag test application wccpd <integer>
```

Where <integer> is a value between 1 and 5:

- 1 Display WCCP stats
- 2 Display WCCP config
- 3 Display WCCP cache servers
- 4 Display WCCP services
- 5 Display WCCP assignment

Enter the following command to view debugging output:


```
diag test application wccpd 3
```

Sample output from a successful WCCP connection:

```
service-0 in vdom-root: num=1, usable=1
cache server ID:
len=44, addr=172.16.78.8, weight=4135, status=0
rcv_id=6547, usable=1, fm=1, nq=0, dev=3(k3),
to=192.168.11.55
ch_no=0, num_router=1:
192.168.11.55
```

Sample output from the same command from an unsuccessful WCCP connection (because of a service group password mismatch):

```
service-0 in vdom-root: num=0, usable=0
diag debug application wccpd -1
Sample output:
wccp_on_recv()-98: vdom-root recv: num=160, dev=3(3),
172.16.78.8->192.168.11.55
wccp2_receive_pkt()-1124: len=160, type=10, ver=0200,
length=152
wccp2_receive_pkt()-1150: found component:t=0, len=20
wccp2_receive_pkt()-1150: found component:t=1, len=24
wccp2_receive_pkt()-1150: found component:t=3, len=44
wccp2_receive_pkt()-1150: found component:t=5, len=20
wccp2_receive_pkt()-1150: found component:t=8, len=24
wccp2_check_security_info()-326: MD5 check failed
```

WAN optimization, web cache, explicit proxy and WCCP get and diagnose commands

The following get and diagnose commands are available for troubleshooting WAN optimization, web cache, explicit proxy and WCCP.

- `get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>`
- `diagnose wad`
- `diagnose wacs`
- `diagnose wadbd`
- `diagnose debug application {wa_cs | wa_dbd | wad | wad_diskd | wccpd} [<debug_level>]`

`get test {wa_cs | wa_dbd | wad | wad_diskd | wccpd} <test_level>`

Display usage information about WAN optimization and web-cache-related applications. Use `<test_level>` to display different information.

```
get test wa_cs <test_level>
get test wa_dbd <test_level>
get test wad <test_level>
get test wad_diskd <test_level>
get test wccpd <test_level>
```

Variable	Description
wad	Display information about WAN optimization, web caching, the explicit web proxy, and the explicit FTP proxy.
wa_cs	Display information about the WAN optimization web cache server.
wa_dbd	Display information about the WAN optimization storage server application.
wad_diskd	Display information about the WAN optimization disk access daemon application.
wccpd	Display information about the WCCP application.

Examples

Enter the following command to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 11
wad tunnel protocol stats:
http tunnel
```

```
bytes_in=1751767 bytes_out=325468
ftp tunnel
bytes_in=0 bytes_out=0
cifs tunnel
bytes_in=0 bytes_out=0
mapi tunnel
bytes_in=0 bytes_out=0
tcp tunnel
bytes_in=3182253 bytes_out=200702
maintenance tunnel
bytes_in=11800 bytes_out=15052
```

Enter the following command to display the current WAN optimization peers. You can use this command to make sure all peers are configured correctly. The command output shows one peer with IP address 172.20.120.141, peer name Web_servers, with 10 active tunnels.

```
get test wad 26
peer name=Web_servers ip=172.20.120.141 vd=0 version=1
tunnels(active/connecting/failover)=10/0/0
sessions=0 n_retries=0 version_valid=true
```

Enter the following command to restart the WAN optimization web cache server.

```
get test wa_cs 99
```

Enter the following command to display all test options:

```
get test wad
```

WAD Test Usage

- 1: display total memory usage
- 3: display proxy status
- 4: display all stats and connections
- 5: toggle AV conserve mode(for debug purpose).
- 8: display all fix-sized advanced memory stats
- 10: toggle cifs read-ahead
- 11: display tunnel protocol stats
- 12: flush tunnel protocol stats
- 13: display http protocol stats
- 14: flush http protocol stats
- 15: display cifs protocol stats
- 16: flush cifs protocol stats
- 17: display ftp protocol stats
- 18: flush ftp protocol stats
- 19: display mapi protocol stats
- 20: flush mapi protocol stats
- 21: display tcp protocol stats
- 22: flush tcp protocol stats
- 23: display all protocols stats
- 24: flush all protocols stats
- 25: display all listeners
- 26: display all peers
- 27: display DNS stats
- 28: display security profile mapping for regular firewall policy
- 30: display Byte Cache DB state
- 31: flush Byte Cache DB stats

```
32: display Web Cache DB state
33: flush Web Cache DB stats
35: display tunnel compressor state
36: flush tunnel compressor stats
37: discard all wad debug info that is currently pending
38: display rules
39: display video cache rules (patterns)
40: display cache state
41: flush cache stats
42: display all fix-sized advanced memory stats in details
45: display memory cache state
46: flush memory cache stats
47: display SSL stats
48: flush SSL stats
49: display SSL mem stats
50: display Web Cache stats
51: flush Web Cache stats
52: flush idle Web cache objects
53: display firewall policies
54: display WAD tunnel stats.
55: display WAD fsae state.
56yxxx: set xxx concurrent Web Cache session for object
      storage y.
57yxxx: set xxxK(32K, 64K,...) unconfirmed write/read size per
      Web Cache object for object storage y.
58yxxxx: set xxxxK maximum ouput buffer size for object
      storage y.
59yxx: set lookup lowmark(only if more to define busy status)
      to be xx for object storage y.
60: display current web proxy users
61: flush current web proxy users
62: display current web proxy user summary
63: display web cache cache sessions
65: display cache exemption patterns
66: toggle dumping URL when daemon crashes.
67: list all used fqdns.
68: list all current ftpproxy sessions.
69: display ftpproxy stats.
70: clear ftpproxy stats.
600000..699999 cmem bucket stats (699999 for usage)
70yxxx: set xxxK maximum ouput buffer size for byte storage y.
71yxxx: set number of buffered add requests to be xxx for byte
      storage y.
72yxxxx: set number of buffered query requests to be xxxx for
      byte storage y.
73yxxxxx: set number of concurrent query requests to be xxxxx
      for byte storage y.
79xxxx: set xxxxMiB maximum AV memory.(0: set to default.
80: display av memory usage
81: toggle av memory protection
800..899: mem_diag commands (800 for help & usage)
800000..899999: mem_diag commands with 1 arg (800 for help &
      usage)
```

```

800000000..899999999: mem_diag commands with 2 args (800 for
  help & usage)
90: set to test disk failure
91: unset to test disk failure
92: trigger a disk failure event
98: gracefully stopping wad proxy
99: restart proxy

```

diagnose wad

Display diagnostic information about the WAN optimization daemon (wad).

```

diagnose wad console-log {disable | enable}
diagnose wad filter {clear | dport | dst | list | negate |
  protocol | sport | src | vd}
diagnose wad history
diagnose wad session
diagnose wad stats {cache | cifs | clear | crypto | ftp | http |
  list | mapi | mem | scan | scripts | summary | tcp | tunnel}
diagnose wad tunnel {clear | list}

```

Variable	Description
console-log	Enable or disable displaying WAN optimization log messages on the CLI console.
filter	Set a filter for listing WAN optimization daemon sessions or tunnels. clear reset or clear the current log filter settings. dport enter the destination port range to filter by. dst enter the destination address range to filter by. list display the current log filter settings
history	Display statistics for one or more WAN optimization protocols for a specified period of time (the last 10 minutes, hour, day or 30 days).
session	Display diagnostics for WAN optimization sessions or clear active sessions.
stats	Display statistics for various parts of WAN optimization such as cache statistics, CIFS statistics, MAPI statistics, HTTP statistics, tunnel statistics etc. You can also clear WAN optimization statistics and display a summary.
tunnel	Display diagnostic information for one or all active WAN optimization tunnels. Clear all active tunnels. Clear all active tunnels.

Examples

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```

diagnose wad tunnel list

Tunnel: id=100 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=100 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=

```

```
bytes_in=348 bytes_out=384

Tunnel: id=99 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=99 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=98 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=98 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=39 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=39 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1068 bytes_out=1104

Tunnel: id=7 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=7 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=8 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=8 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=5 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=5 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=4 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=1 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=1 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web_servers id=2 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
```

```
bytes_in=1228 bytes_out=1264
```

```
Tunnels total=10 manual=10 auto=0
```

diagnose wacs

Display diagnostic information for the web cache database daemon (wacs).

```
diagnose wacs clear
diagnose wacs reents
diagnose wacs restart
diagnose wacs stats
```

Variable	Description
clear	Remove all entries from the web cache database.
reents	Display recent web cache database activity.
restart	Restart the web cache daemon and reset statistics.
stats	Display web cache statistics.

diagnose wadbd

Display diagnostic information for the WAN optimization database daemon (wadbd).

```
diagnose wadbd {check | clear | reents | restart | stats}
```


Variable	Description
check	Check WAN optimization database integrity.
clear	Remove all entries from the WAN optimization database.
reents	Display recent WAN optimization database activity.
restart	Restart the WAN optimization daemon and reset statistics.
stats	Display WAN optimization statistics.

diagnose debug application {wa_cs | wa_dbd | wad | wad_diskd | wccpd} [<debug_level>]

View or set the debug level for displaying WAN optimization and web cache-related daemon debug messages. Include a <debug_level> to change the debug level. Leave the <debug_level> out to display the current debug level. Default debug level is 0.

```
diagnose debug application wa_cs [<debug_level>]
diagnose debug application wa_dbd [<debug_level>]
diagnose debug application wad [<debug_level>]
diagnose debug application wccpd [<debug_level>]
```

Variable	Description
wa_cs	Set the debug level for the web cache server.
wa_dbd	Set the debug level for the WAN optimization database server.
wad	Set the debug level for the WAN optimization daemon.
wccpd	Set the debug level for the WCCP daemon.




Chapter 18 Load Balancing

This FortiOS Handbook chapter contains the following sections:

[Configuring load balancing](#) describes FortiGate firewall load balancing.

[Load balancing configuration examples](#) describes includes basic and advanced load balancing configurations.



Configuring load balancing

This section describes how to use the FortiGate firewall load balancing configuration to load balance traffic to multiple backend servers.


This section describes:

- [Load balancing overview](#)
- [Basic load balancing configuration example](#)
- [HTTP and HTTPS load balancing, multiplexing, and persistence](#)
- [SSL/TLS load balancing](#)
- [IP, TCP, and UDP load balancing](#)

Load balancing overview

You can configure FortiOS load balancing to intercept incoming traffic with a virtual server and share it among one or more backend real servers. By doing so, the FortiGate unit enables multiple real servers to respond as if they were a single device or virtual server. This in turn means that more simultaneous requests can be handled.

Figure 345: Load balancing configuration



Traffic can be balanced across multiple backend real servers based on a selection of load balancing methods including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time, number of connections. The load balancer can balance layer 7 HTTP, HTTPS, SSL, generic layer 4 TCP, UDP and generic layer 3 IP protocols. Session persistence is supported based on injected HTTP/HTTPS cookies or the SSL session ID.

You can bind up to 8 real servers can to one virtual server. The real server topology is transparent to end users, and the users interact with the system as if it were only a single server with the IP address and port number of the virtual server. The real servers may be interconnected by high-speed LAN or by geographically dispersed WAN. The FortiGate unit schedules requests to the real servers and makes parallel services of the virtual server to appear to involve a single IP address.

There are additional benefits to load balancing. First, because the load is distributed across multiple servers, the service being provided can be highly available. If one of the servers breaks down, the load can still be handled by the other servers. Secondly, this increases scalability. If the load increases substantially, more servers can be added behind the FortiGate unit in order to cope with the increased load.

Load balancing, UTM, authentication, and other FortiOS features

Flow-based and proxy-based UTM features such as virus scanning, IPS, DLP, application control, and web filtering can be applied to sessions that are to be load balanced. This includes SSL offloading and multiplexing. Applying these UTM features to load balancing traffic may reduce load balancing performance.

Authentication and dynamic profiles are not supported for load balancing sessions. Usually FortiGate load balancing is used to allow public access to services on servers protected by a FortiGate unit. Authentication is not generally not required for this kind of configuration.

Features such web proxying, web caching, and WAN optimization also do not work with load balanced sessions. However, most other features that can be applied by a security policy are supported.

Configuring load balancing virtual servers

A virtual server is a specialized firewall virtual IP that performs server load balancing. From the web-based manager you add load balancing virtual server by going to *Firewall Objects > Load Balance > Virtual Server*.

Name	Enter the name for the virtual server.
Color	Select <i>Change</i> beside the icon to change the color of the icon. When you select <i>Change</i> , a color palette window appears; select a color from the palette window.

Type	<p>Select the protocol to be load balanced by the virtual server. If you select a general protocol such as <i>IP</i>, <i>TCP</i>, or <i>UDP</i> the virtual server load balances all IP, TCP, or UDP sessions. If you select specific protocols such as <i>HTTP</i>, <i>HTTPS</i>, or <i>SSL</i> you can apply additional server load balancing features such as <i>Persistence</i> and <i>HTTP Multiplexing</i>.</p> <ul style="list-style-type: none"> • Select <i>HTTP</i> to load balance only HTTP sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 80 for HTTP sessions). You can also select <i>HTTP Multiplex</i>. You can also set <i>Persistence</i> to <i>HTTP Cookie</i> to select cookie-based persistence. • Select <i>HTTPS</i> to load balance only HTTPS sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced (usually port 443 for HTTPS sessions). You can also select <i>HTTP Multiplex</i>. You can also set <i>Persistence</i> to <i>HTTP Cookie</i> to select cookie-based persistence. You can also set <i>Persistence</i> to <i>SSL Session ID</i>. • Select <i>IP</i> to load balance all sessions accepted by the security policy that contains this virtual server. • Select <i>SSL</i> to load balance only SSL sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced. • Select <i>TCP</i> to load balance only TCP sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced. • Select <i>UDP</i> to load balance only UDP sessions with destination port number that matches the <i>Virtual Server Port</i> setting. Change <i>Virtual Server Port</i> to match the destination port of the sessions to be load balanced.
Interface	Select the virtual server external interface from the list. The external interface is connected to the source network and receives the packets to be forwarded to the destination network.
Virtual Server IP	The IP address of the virtual server. This is an IP address on the external interface that you want to map to an address on the destination network.
Virtual Server Port	Enter the external port number that you want to map to a port number on the destination network. Sessions with this destination port are load balanced by this virtual server.
Load Balance Method	Select the load balancing method used by the virtual server. See “Load balancing methods” on page 2873 .
Persistence	Configure persistence to make sure that a user is connected to the same server every time they make a request that is part of the same session. See “Session persistence” on page 2874 . For HTTP and HTTPS sessions, see “HTTP and HTTPS persistence” on page 2886 .

HTTP Multiplexing	Select to use the FortiGate unit to multiplex multiple client connections into a few connections between the FortiGate unit and the real server. See “HTTP and HTTPS multiplexing” on page 2885 .
Preserve Client IP	Select to preserve the IP address of the client in the X-Forwarded-For HTTP header. This can be useful if you want log messages on the real servers to the client’s original IP address. If this option is not selected, the header will contain the IP address of the FortiGate unit. This option appears only if <i>HTTP</i> or <i>HTTPS</i> are selected for <i>Type</i> , and is available only if <i>HTTP Multiplexing</i> is selected.
SSL Offloading	Select to accelerate clients’ SSL connections to the server by using the FortiGate FortiGate unit to perform SSL operations, then select which segments of the connection will receive SSL offloading. See “SSL offloading” on page 2890
Certificate	Select the certificate to use with <i>SSL Offloading</i> . The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported. This option appears only if <i>HTTPS</i> or <i>SSL</i> are selected for <i>Type</i> , and is available only if <i>SSL Offloading</i> is selected.
Health Check	Select which health check monitor configuration will be used to determine a server’s connectivity status. See “Health check monitoring” on page 2876 .

From the CLI you configure a virtual server by added a firewall virtual IP and setting the virtual IP type to server load balance:

```
config firewall vip
edit Vserver-HTTP-1
set type server-load-balance
...
```

A virtual server includes a virtual server IP address bound to an interface. The virtual server IP address is the destination address incoming packets to be load balanced and the virtual server is bound to the interface that receives the packets to be load balanced.

For example, if you want to load balance incoming HTTP traffic from the Internet to a group of web servers on a DMZ network, the virtual server IP address is the known Internet IP address of the web servers and the virtual server binds this IP address to the FortiGate interface connected to the Internet.

When you bind the virtual server’s external IP address to a FortiGate unit interface, by default, the network interface responds to ARP requests for the bound IP address. Virtual servers use proxy ARP, as defined in [RFC 1027](#), so that the FortiGate unit can respond to ARP requests on a network for a real server that is actually installed on another network. In some cases you may not want the network interface sending ARP replies. You can use the `arp-reply` option disable sending ARP replies:

```
config firewall vip
edit Vserver-HTTP-1
set type server-load-balance
set arp-reply disable
...
```

The load balancing virtual server configuration also includes the virtual server port. This is the TCP port on the bound interface that the virtual server listens for traffic to be load balanced on. The virtual server can listen on any port.

Load balancing methods

The load balancing method defines how sessions are load balanced to real servers. A number of load balancing methods are available as listed in [Table 150](#).

All load balancing methods will not send traffic to real servers that are down or not responding. However, the FortiGate unit can only determine if a real server is not responding by using a health check monitor. You should always add at least one health check monitor to a virtual server or to individual real servers, or load balancing methods may attempt to distribute sessions to real servers that are not functioning.

Table 150: Load balancing methods

Method	Description
Source IP Hash	The traffic load is statically spread evenly across all real servers. However, sessions are not assigned according to how busy individual real servers are. This load balancing method provides some persistence because all sessions from the same source address always go to the same real server. However, the distribution is stateless, so if a real server is added or removed (or goes up or down) the distribution is changed and persistence could be lost.
Round Robin	Directs new requests to the next real server, and treats all real servers as equals regardless of response time or number of connections. Dead real servers or non responsive real servers are avoided.
Weighted	Real servers with a higher weight value receive a larger percentage of connections. Set the real server weight when adding a real server.
First Alive	Always directs sessions to the first alive real server. This load balancing schedule provides real server failover protection by sending all sessions to the first alive real server and if that real server fails, sending all sessions to the next alive real server. Sessions are not distributed to all real servers so all sessions are processed by the “first” real server only. First refers to the order of the real servers in the virtual server configuration. For example, if you add real servers A, B and C in that order, then all sessions always go to A as long as it is alive. If A goes down then sessions go to B and if B goes down sessions go to C. If A comes back up sessions go back to A. Real servers are ordered in the virtual server configuration in the order in which you add them, with the most recently added real server last. If you want to change the order you must delete and re-add real servers in the required order.
Least RTT	Directs sessions to the real server with the least round trip time. The round trip time is determined by a Ping health check monitor and is defaulted to 0 if no Ping health check monitors are added to the virtual server.
Least Session	Directs requests to the real server that has the least number of current connections. This method works best in environments where the real servers or other equipment you are load balancing all have similar capabilities. This load balancing method uses the FortiGate session table to track the number of sessions being processed by each real server. The FortiGate unit cannot detect the number of sessions actually being processed by a real server.

Table 150: Load balancing methods

Method	Description
HTTP Host	Load balances HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server.

Session persistence

Use persistence to make sure that a user is connected to the same real server every time they make an HTTP, HTTPS, or SSL request that is part of the same user session. For example, if you are load balancing HTTP and HTTPS sessions to a collection of eCommerce web servers, when a user is making a purchase they will be starting multiple sessions as they navigate the eCommerce site. In most cases all of the sessions started by this user during an eCommerce session should be processed by the same real server. Typically, the HTTP protocol keeps track of these related sessions using cookies. HTTP cookie persistence makes sure that all sessions that are part of the same user session are processed by the same real server.

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the load balance method. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server. For more information about HTTP and HTTPS persistence, see [“HTTP and HTTPS persistence” on page 2886](#).

Real servers

Add real servers to a load balancing virtual server to provide the information the virtual server requires to be able to send sessions to the server. A real server configuration includes the IP address of the real server and port number that the real server receives sessions on. The FortiGate unit sends sessions to the real server's IP address using the destination port number in the real server configuration.

When configuring a real server you can also specify the weight (used if the load balance method is set to weighted) and you can limit the maximum number of open connections between the FortiGate unit and the real server. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests to other real servers until the connection number drops below the specified limit. Setting Maximum Connections to 0 means that the FortiGate unit does not limit the number of connections to the real server.

Real server active, standby, and disabled modes

By default the real server mode setting is active indicating that the real server is available to receive connections. If the real server is removed from the network (for example, for routine maintenance or because of a hardware or software failure) you can change the mode to standby or disabled. In disabled mode the FortiGate unit no longer sends sessions to the real server.

If a real server is in standby mode the FortiGate also does not send sessions to it unless other real servers added to the same virtual server become unavailable. For example:

- A virtual server that includes two real servers one in active mode and one in standby mode. If the real server in active mode fails, the real server in standby mode is changed to active mode and all sessions are sent to this real server.

- A virtual server includes three real servers, two in active mode and one in standby mode, if one of the real servers in active mode fails, the real server in standby mode is changed to active mode and sessions are load balanced between it and still operating real server. If both real servers in active mode fail, all sessions are sent to the real server in standby mode.

Adding real servers

To add a real server from the web-based manager go to *Firewall Objects > Load Balance > Real Server*.

Virtual Server	Select the virtual server that will send sessions to this real server.
IP Address	Enter the IP address of the real server.
Port	Enter the port number on the destination network to which the external port number is mapped.
Weight	Enter the weight value of the real server. The higher the weight value, the higher the percentage of connections the server will handle. A range of 1-255 can be used. This option is available only if the associated virtual server's load balance method is <i>Weighted</i> .
Max Connections	Enter the limit on the number of active connections directed to a real server. A range of 1-99999 can be used. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests to another server until the connection number drops below the specified limit. Setting <i>Maximum Connections</i> to 0 means that the FortiGate unit does not limit the number of connections to the real server.
HTTP Host	Enter the HTTP header for load balancing across multiple real servers. This feature is used for load balancing HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server, providing better load balancing for those specific connections.
Mode	Select a mode for the real server.

To add a real server from the CLI you configure a virtual server and add real servers to it. For example, to add three real servers to a virtual server that load balances UDP sessions on port 8190 using weighted load balancing. For each real server the port is not changed. The default real server port is 0 resulting in the traffic being sent the real server with destination port 8190. Each real sever is given a different weight. Servers with higher weights have a max-connections limit to prevent too many sessions from being sent to them.

```
config firewall vip
edit Vserver-UDP-1
set type server-load-balance
set server-type udp
set ldb-method weighted
set extip 172.20.120.30
set extintf wan1
set extport 8190
set monitor ping-mon-1
config realservers
```

```
edit 1
    set ip 10.31.101.30
    set weight 100
    set max-connections 10000
next
edit 2
    set ip 10.31.101.40
    set weight 100
    set max-connections 10000
next
edit 3
    set ip 10.31.101.50
    set weight 10
end
end
```

Health check monitoring

From the FortiGate web-based manager you can go to *Firewall Objects > Load Balance > Health Check Monitor* and configure health check monitoring so that the FortiGate unit can verify that real servers are able respond to network connection attempts. If a real server responds to connection attempts the load balancer continues to send sessions to it. If a real server stops responding to connection attempts the load balancer assumes that the server is down and does not send sessions to it. The health check monitor configuration determines how the load balancer tests the real servers. You can use a single health check monitor for multiple load balancing configurations.

You can configure TCP, HTTP and Ping health check monitors. Usually you would want the health check monitor to use the same protocol for checking the health of the server as the traffic being load balanced to it. For example, for an HTTP load balancing configuration you would normally use an HTTP health check monitor.

For the TCP and HTTP health check monitors you can specify the destination port to use to connect to the real servers. If you set the port to 0, the health check monitor uses the port defined in the real server. This allows you to use the same health check monitor for multiple real servers using different ports. You can also configure the interval, timeout and retry. A health check occurs every number of seconds indicated by the interval. If a reply is not received within the timeout period the health check is repeated. If no response is received after the number of configured retries, the virtual server is considered unresponsive, and load balancing will disabling traffic to that real server. The health check monitor will continue to contact the real server and if successful, the load balancer can resume sending sessions to the recovered real server.

For HTTP health check monitors, you can add URL that the FortiGate unit connects to when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.

The URL would not usually include an IP address or domain name. Instead it should start with a "/" and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.31.101.30, the URL "/test_page.htm" causes the FortiGate unit to send an HTTP get request to "http://10.31.101.30/test_page.htm".

For HTTP health check monitors, you can also add a matched content phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the URL option. If the URL returns a web page, the matched content should exactly match some of the text on the web page. You can use the URL and Matched Content options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.

For example, you can set matched content to “server test page” if the real HTTP server page defined by the URL option contains the phrase “server test page”. When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the matched content phrase.

Health Check page	
Lists each individual health check monitor that you created. On this page, you can edit, delete and create a new health check monitor.	
Create New	Creates a new health check monitor. When you select <i>Create New</i> , you are automatically redirected to the Add New Health Check Monitor page.
Edit	Modifies settings within the health check monitor configuration. When you select <i>Edit</i> , you are automatically redirected to the Edit Health Check Monitor page.
Delete	<p>Removes a health check monitor from the list on the Health Check Monitor page. This option appears only if the health check monitor configuration is not currently being used by a virtual server configuration.</p> <p>To remove multiple health check monitors, on the Health Check Monitor page, in each of the rows of the monitors you want removed, select the check box and then select <i>Delete</i>.</p> <p>To remove all health check monitors, on the Health Check Monitor page, select the check box in the check box column, and then select <i>Delete</i>.</p>
Name	The name of the health check monitor configuration. The names are grouped by the health check monitor types.
Details	<p>The details of the health check monitor configuration, which vary by the type of the health check monitor, and do not include the interval, timeout, or retry, which are settings common to all types.</p> <p>This field is empty if the type of the health check monitor is PING.</p>
Ref.	<p>Displays the number of virtual servers the health check monitor has been added to.</p> <p>To view the location of the referenced object, select the number in <i>Ref.</i>, and the Object Usage window appears displaying the list of virtual servers that the health check monitor has been added to.</p>
Add New Health Check Monitor	
Provides settings for configuring a health check monitor.	
Name	Enter the name of the health check monitor configuration.

Type	<p>Select the protocol used to perform the health check.</p> <ul style="list-style-type: none"> • TCP • HTTP • PING
Port	<p>Enter the port number used to perform the health check. If you set the <i>Port</i> to 0, the health check monitor uses the port defined in the real server. This way you can use a single health check monitor for different real servers.</p> <p>This option does not appear if the <i>Type</i> is <i>PING</i>.</p>
Interval	Enter the number of seconds between each server health check.
URL	<p>For HTTP health check monitors, add a URL that the FortiGate unit uses when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.</p> <p>The URL would not usually include an IP address or domain name. Instead it should start with a "/" and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.10.10.1, the <i>URL</i> "/test_page.htm" causes the FortiGate unit to send an HTTP get request to "http://10.10.10.1/test_page.htm".</p> <p>This option appears only if <i>Type</i> is <i>HTTP</i>.</p>
Matched Content	<p>For HTTP health check monitors, add a phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the <i>URL</i> option. If the <i>URL</i> returns a web page, the <i>Matched Content</i> should exactly match some of the text on the web page. You can use the <i>URL</i> and <i>Matched Content</i> options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.</p> <p>For example, you can set <i>Matched Content</i> to "server test page" if the real HTTP server page defined by the URL option contains the phrase "server test page". When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the <i>Matched Content</i> phrase.</p> <p>This option appears only if <i>Type</i> is <i>HTTP</i>.</p>
Timeout	Enter the number of seconds which must pass after the server health check to indicate a failed health check.
Retry	Enter the number of times, if any, a failed health check will be retried before the server is determined to be inaccessible.

Virtual IP, load balance virtual server and load balance real server limitations

The following limitations apply when adding virtual IPs, Load balancing virtual servers, and load balancing real servers. Load balancing virtual servers are actually server load balancing virtual IPs. You can add server load balance virtual IPs from the CLI.

- Virtual IP *External IP Address/Range* entries or ranges cannot overlap with each other or with load balancing virtual server *Virtual Server IP* entries.
- A virtual IP *Mapped IP Address/Range* cannot be 0.0.0.0 or 255.255.255.255.
- A real server *IP* cannot be 0.0.0.0 or 255.255.255.255.
- If a static NAT virtual IP *External IP Address/Range* is 0.0.0.0, the *Mapped IP Address/Range* must be a single IP address.
- If a load balance virtual IP *External IP Address/Range* is 0.0.0.0, the *Mapped IP Address/Range* can be an address range.
- When port forwarding, the count of mapped port numbers and external port numbers must be the same. The web-based manager does this automatically but the CLI does not.
- Virtual IP and virtual server names must be different from firewall address or address group names.

Monitoring load balancing

From the web-based manager you can go to *Firewall Objects > Monitor > Load Balance Monitor* to monitor the status of configured virtual servers and real server and start or stop the real servers. You can also use the `get test ipldb` command from the CLI to display similar information.

For each real server the monitor displays health status (up or down), active sessions, round trip time and the amount of bytes of data processed. From the monitor page you can also stop sending new sessions to any real server. When you select to stop sending sessions the FortiGate unit performs a graceful stop by continuing to send data for sessions that were established or persistent before you selected stop. However, no new sessions are started.

Virtual Server	The IP addresses of the existing virtual servers.
Real Server	The IP addresses of the existing real servers.
Health Status	Displays the health status according to the health check results for each real server. A green arrow means the server is up. A red arrow means the server is down.
Monitor Events	Display each real server's up and down times.
Active Sessions	Display each real server's active sessions.
RTT (ms)	Displays the Round Trip Time (RTT) of each real server. By default, the RTT is "<1". This value will change only when ping monitoring is enabled on a real server.
Bytes Processed	Displays the traffic processed by each real server.
Graceful Stop/Start	Select to start or stop real servers. When stopping a server, the FortiGate unit will not accept new sessions but will wait for the active sessions to finish.

Load balancing get command

The following get command is available to display testing and debug information for the FortiGate virtual server process:

```
get test vs <test-level_int>
```

Where <test-level_int> can be:

3 to display the virtual server process id.

8 to display the virtual server log configuration.

30 to display the virtual server configuration statistics.

99 to restart the virtual server process.

Load balancing diagnose commands

You can also use the following diagnose commands to view status information for load balancing virtual servers and real servers:

```
diagnose firewall vip realserver {down | flush | healthcheck |
  list | up}
diagnose firewall vip virtual-server {filter | log | real-server
  | session | stats}
```

For example, the following command lists and displays status information for all real servers:

```
diagnose firewall vip virtual-server real-server

vd root/0  vs vs/2  addr 10.31.101.30:80  status 1/1
  conn: max 0  active 0  attempts 0  success 0  drop 0  fail 0

vd root/0  vs vs/2  addr 10.31.101.20:80  status 1/1
  conn: max 0  active 0  attempts 0  success 0  drop 0  fail 0
```

Many of the diagnostic commands involve retrieving information about one or more virtual servers. To control which servers are queried you can define a filter:

```
diagnose firewall vip virtual-server filter <filter_str>
```

Where <filter_str> can be:

clear erase the current filter

dst the destination address range to filter by

dst-port the destination port range to filter by

list display the current filter

name the vip name to filter by

negate negate the specified filter parameter

src the source address range to filter by

src-port the source port range to filter by

vd index of virtual domain. -1 matches all

The default filter is empty so no filtering is done.

Logging Diagnostics

The logging diagnostics provide information about two separate features:

```
diagnose firewall vip virtual-server log {console | filter}
```

Where

`console {disable | enable}` enables or disables displaying the event log messages generated by virtual server traffic on the console to simplify debugging.

`filter` sets a filter for the virtual server debug log

The filter option controls what entries the virtual server daemon will log to the console if `diagnose debug application vs` level is non-zero. The filtering can be done on source, destination, virtual-server name, virtual domain, and so on:

```
diagnose firewall vip virtual-server log filter <filter_str>
```

where `<filter_str>` can be

`clear` erase the current filter

`dst` the destination address range to filter by

`dst-port` the destination port range to filter by

`list` display the current filter

`name` the virtual-server name to filter by

`negate` negate the specified filter parameter

`src` the source address range to filter by

`src-port` the source port range to filter by

`vd` index of virtual domain. -1 matches all

The default filter is empty so no filtering is done.

Real server diagnostics

Enter the following command to list all the real servers:

```
diag firewall vip virtual-server real-server list
```

In the following example there is only one virtual server called `slb` and it has two real-servers:

```
diag firewall vip virtual-server server
vd root/0  vs slb/2  addr 172.16.67.191:80  status 1/1
  conn: max 10  active 0  attempts 0  success 0  drop 0  fail 0
  http: available 0  total 0

vd root/0  vs slb/2  addr 172.16.67.192:80  status 1/1
  conn: max 10  active 1  attempts 4  success 4  drop 0  fail 0
  http: available 1  total 1
```

The `status` indicates the administrative and operational status of the real-server.

`max` indicates that the real-server will only allow 10 concurrent connections.

`active` is the number of current connections to the server `attempts` is the total number of connections attempted `success` is the total number of connections that were successful.

`drop` is the total number of connections that were dropped because the active count hit `max`.

`fail` is the total number of connections that failed to complete due to some internal problem (for example, lack of memory).


If the virtual server has HTTP multiplexing enabled then the HTTP section indicates how many established connections to the real-server are available to service a HTTP request and also the total number of connections.

Basic load balancing configuration example

This section describes the steps required to configure the load balancing configuration shown in Figure 346. In this configuration a FortiGate-51B unit is load balancing HTTP traffic from the Internet to three HTTP servers on the Internal network. HTTP sessions are accepted at the wan1 interface with destination IP address 172.20.120.121 on TCP port 8080 and forwarded from the internal interface to the web servers. When forwarded the destination address of the sessions is translated to the IP address of one of the web servers.

The load balancing configuration also includes session persistence using HTTP cookies, round-robin load balancing, and TCP health monitoring for the real servers. Ping health monitoring consists of the FortiGate unit using ICMP ping to make sure the web servers can respond to network traffic.

Figure 346: Virtual server and real servers setup



To configure the example load balancing configuration - general configuration steps

- 1 Add a load balance ping health check monitor
A ping health check monitor causes the FortiGate unit to ping the real servers every 10 seconds. If one of the servers does not respond within 2 seconds, the FortiGate unit will retry the ping 3 times before assuming that the HTTP server is not responding.
- 2 Add a load balance virtual server.
- 3 Add the three load balance real servers. Include the virtual server in each real server configuration.
- 4 Add a security policy that includes the load balance virtual server as the destination address.

To configure the example load balancing configuration - web-based manager

- 1 Go to *Firewall Objects > Load Balance > Health Check* and add the following health check monitor.

Name	Ping-mon-1
Type	Ping
Interval	10 seconds
Timeout	2 seconds
Retry	3

- 2 Go to *Firewall Objects > Load Balance > Virtual Server* and add virtual server that accepts the traffic to be load balanced.

Name	Vserver-HTTP-1
Type	HTTP
Interface	wan1
Virtual Server IP	172.20.120.121
Virtual Server Port	8080
Load Balance Method	Round Robin
Persistence	HTTP Cookie
HTTP Multiplexing	Do not select
Health Check	Move Ping-mon-1 to the Selected list.

- 3 Go to *Firewall Objects > Load Balance > Real Server* and add the real servers.

Virtual Server	Vserver-HTTP-1
IP Address	10.31.101.30
Port	80
Weight	n/a
Max Connections	0
Mode	Active

Virtual Server	Vserver-HTTP-1
IP Address	10.31.101.40
Port	80
Weight	n/a
Max Connections	0
Mode	Active

Virtual Server	Vserver-HTTP-1
IP Address	10.31.101.50
Port	80

Weight	n/a
Max Connections	0
Mode	Active

- 4 Go to *Policy > Policy > Policy* and add a wan1 to internal security policy that includes the virtual server. This policy also applies an Antivirus profile to the load balanced sessions.

Source Interface/Zone	wan1
Source Address	all
Destination Interface/Zone	internal
Destination Address	Vserver-HTTP-1
Schedule	always
Service	ANY
Action	ACCEPT
NAT	Enable NAT
UTM	Select
Protocol Options	Select and select a protocol options profile.
Enable AntiVirus	Select and select an antivirus profile.

To configure the example load balancing configuration- CLI

- 1 Use the following command to add a Ping health check monitor.

```
config firewall ldb-monitor
  edit ping-mon-1
    set type ping
    set interval 10
    set timeout 2
    set retry 3
  end
```

- 2 Use the following command to add the virtual server that accepts HTTP sessions on port 8080 at the wan1 interface and load balances the traffic to three real servers.

```
config firewall vip
  edit Vserver-HTTP-1
    set type server-load-balance
    set server-type http
    set ldb-method round-robin
    set extip 172.20.120.30
    set extintf wan1
    set extport 8080
    set persistence http-cookie
    set monitor tcp-mon-1
    config realservers
      edit 1
        set ip 10.31.101.30
        set port 80
      next
```

```
edit 2
    set ip 10.31.101.40
    set port 80
end
end
```

- 3 Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```
config firewall policy
edit 0
    set srcintf wan1
    set srcaddr all
    set dstintf internal
    set dstaddr Vserver-HTTP-1
    set action accept
    set schedule always
    set service ANY
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set av-profile scan
end
```

HTTP and HTTPS load balancing, multiplexing, and persistence

In a firewall load balancing virtual server configuration, you can select HTTP to load balance only HTTP sessions. The virtual server will load balance HTTP sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. The default virtual server port for HTTP load balancing is 80, but you can change this to any port number. Similarly for HTTPS load balancing, set the virtual server type to HTTPS and then select the interface, virtual server IP, and virtual server port that matches the HTTPS traffic to be load balanced. Usually HTTPS traffic uses port 443.

You can also configure load balancing to offload SSL processing for HTTPS and SSL traffic. See [“SSL offloading” on page 2890](#) for more information.

HTTP and HTTPS multiplexing

For both HTTP and HTTPS load balancing you can multiplex HTTP requests and responses over a single TCP connection. HTTP multiplexing is a performance saving feature of HTTP/1.1 compliant web servers that provides the ability to pipeline many unrelated HTTP or HTTPS requests on the same connection. This allows a single HTTPD process on the server to interleave and serve multiple requests. The result is fewer idle sessions on the web server so server resources are used more efficiently. HTTP multiplexing can take multiple separate inbound sessions and multiplex them over the same internal session. This may reduce the load on the backend server and increase the overall performance.

HTTP multiplexing may improve performance in some cases. For example, if users web browsers are only compatible with HTTP 1.0. HTTP multiplexing can also improve performance between a web server and the FortiGate unit if the FortiGate unit is performing SSL acceleration. However, in most cases HTTP multiplexing should only be used if enabling it leads to a measurable improvement in performance.

To enable HTTP multiplexing from the web-based manager, select multiplex HTTP requests/responses over a single TCP connection. To enable HTTP multiplexing from the CLI enable the `http-multiplex` option.

Preserving the client IP address

Select preserve client IP from the web-based manager or enable the `http-ip-header` option from the CLI to preserve the IP address of the client in the X-Forwarded-For HTTP header. This can be useful in an HTTP multiplexing configuration if you want log messages on the real servers to the client's original IP address. If this option is not selected, the header will contain the IP address of the FortiGate unit.

HTTP and HTTPS persistence

Configure load balancing persistence for HTTP or HTTPS to make sure that a user is connected to the same server every time they make a request that is part of the same session. HTTP cookie persistence uses injected cookies to enable persistence.

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the *Load Balance Method*. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server.

The following example shows how to enable cookie persistence and set the cookie domain to `.example.org`.

```
config firewall vip
  edit HTTP_Load_Balance
    set type server-load-balance
    set server-type http
    set extport 8080
    set extintf port2
    set extip 192.168.20.20
    set persistence http-cookie
    set http-cookie-domain .example.org
  config realservers
    edit 1
      set ip 10.10.10.1
    next
    edit 2
      set ip 10.10.10.2
    next
    edit 3
      set ip 10.10.10.3
    end
  end
```

How HTTP cookie persistence options work

The following options are available for the `config firewall vip` command when type is set to `server-load-balance`, server-type is set to `http` or `https` and persistence is set to `http-cookie`:

```
http-cookie-domain
http-cookie-path
http-cookie-generation
http-cookie-age
http-cookie-share
```

`https-cookie-share` (appears when `server-type` is set to `https`)

When HTTP cookie persistence is enabled the FortiGate unit inserts a header of the following form into each HTTP response unless the corresponding HTTP request already contains a `FGTServer` cookie:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;  
Version=1; Max-Age=3600
```

The value of the `FGTServer` cookie encodes the server that traffic should be directed to. The value is encoded so as to not leak information about the internal network.

Use `http-cookie-domain` to restrict the domain that the cookie should apply to. For example, to restrict the cookie to `.server.com`, enter:

```
set http-cookie-domain .server.com
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;  
Version=1; Domain=.server.com; Max-Age=3600
```

Use `http-cookie-path` to limit the cookies to a particular path. For example, to limit cookies to the path `/sales`, enter:

```
set http-cookie-path /sales
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;  
Version=1; Domain=.server.com; Path=/sales; Max-Age=3600
```

Use `http-cookie-age` to change how long the browser caches the cookie. You can enter an age in minutes or set the age to 0 to make the browser keep the cookie indefinitely:

```
set http-cookie-age 0
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;  
Version=1; Domain=.server.com; Path=/sales
```

Use `http-cookie-generation` to invalidate all cookies that have already been generated. The exact value of the generation is not important, only that it is different from any generation that has already been used for cookies in this domain. The simplest approach is to increment the generation by one each time invalidation is required. Since the default is 0, enter the following to invalidate all existing cookies:

```
set http-cookie-generation 1
```

Use `http-cookie-share {disable | same-ip}` to control the sharing of cookies across virtual servers in the same virtual domain. The default setting `same-ip` means that any `FGTServer` cookie generated by one virtual server can be used by another virtual server in the same virtual domain. For example, if you have an application that starts on HTTP and then changes to HTTPS and you want to make sure that the same server is used for the HTTP and HTTPS traffic then you can create two virtual servers, one for port 80 (for HTTP) and one for port 443 (for HTTPS). As long as you add the same real servers to both of these virtual servers (and as long as both virtual servers have the same number of real servers with the same IP addresses), then cookies generated by accessing the HTTP server are reused when the application changes to the HTTPS server.

If for any reason you do not want this sharing to occur then select `disable` to make sure that a cookie generated for a virtual server cannot be used by other virtual servers.

Use `https-cookie-secure` to enable or disable using secure cookies. Secure cookies are disabled by default because secure cookies can interfere with cookie sharing across HTTP and HTTPS virtual servers. If enabled, then the `Secure` tag is added to the cookie inserted by the FortiGate unit:

```
Set-Cookie: FGTSer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158;  
Version=1; Max-Age=3600; Secure
```

HTTP host-based load balancing

When configuring HTTP or HTTPS load balancing you can select HTTP host load balancing to load balance HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server. HTTP 1.1 includes the concept of a virtual server which allows a HTTP or HTTPS server with a single external IP address to serve requests for multiple DNS domains by using the mandatory `Host:` header in a HTTP request to indicate which DNS domain the request is destined for.

FortiOS can load-balance HTTP and HTTPS connections among multiple real servers using the `Host:` header to guide the connection to the correct real server. The host load balancing method allows a real server to specify a `http-host` attribute which is the domain name of the traffic for that real server. Each real server can only specify a single domain name. The same domain name can appear in more than one real server but only the first one that is up will be used, any others are purely for redundancy. If the `Host:` header contains a domain that does not match any `http-host` entry then the connection will be dropped. A real server with no `http-host` can be matched by any `Host: domain`.

For example, consider a FortiGate unit that is load-balancing traffic to three real servers. Traffic for `www.example1.com` should go to `192.168.2.1`, traffic for `www.example2.com` should go to `192.168.2.2` and traffic to any other domain should go to `192.168.2.3`. To enable this configuration you would add a virtual server and set the load balance method to HTTP host. Then you would add three real servers and set the HTTP host of the real server with IP address `192.168.2.1` to `www.example1.com`, the HTTP host of the real server with IP address `192.168.2.2` to `www.example2.com` and you would not specify an HTTP host for the third real server.

The configuration of a virtual IP to achieve this result would be:

```
config firewall vip  
  edit "http-host-ldb"  
    set type server-load-balance  
    set extip 172.16.67.195  
    set extintf "lan"  
    set server-type http  
    set ldb-method http-host  
    set extport 80  
  config realservers  
    edit 1  
      set http-host "www.example1.com"  
      set ip 192.168.2.1  
      set port 80  
    next  
    edit 2  
      set http-host "www.example2.com"  
      set ip 192.168.2.2  
      set port 80  
    next  
    edit 3
```

```
        set ip 192.168.2.3
        set port 80
    next
end
end
```

Host load balancing and HTTP cookie persistence

In an HTTP host-based load balancing configuration with HTTP cookie persistence enabled you can optionally configure cookie persistence to use the domain set in the host header as the cookie domain. You can do this by enabling the `http-cookie-domain-from-host` option, for example:

```
config firewall vip
  edit "http-host-ldb"
    set type server-load-balance
    set extip 172.16.67.195
    set extintf "lan"
    set server-type http
    set ldb-method http-host
    set extport 80
    set persistence http-cookie
    set http-cookie-domain-from-host enable
  config realservers
    edit 1
      set http-host "www.example1.com"
      set ip 192.168.2.1
      set port 80
    next
    edit 2
      set http-host "www.example2.com"
      set ip 192.168.2.2
      set port 80
    next
    edit 3
      set ip 192.168.2.3
      set port 80
    next
  end
end
```

SSL/TLS load balancing

In a firewall load balancing virtual server configuration, you can select SSL to load balance only SSL and TLS sessions. The virtual server will load balance SSL and TLS sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. Change this port to match the destination port of the sessions to be load balanced.

For SSL load balancing you can also set persistence to SSL session ID. Persistence is achieved by the FortiGate unit sending all sessions with the same SSL session ID to the same real server. When you configure persistence, the FortiGate unit load balances a new session to a real server according to the *Load Balance Method*. If the session has an SSL session ID, the FortiGate unit sends all subsequent sessions with the same SSL session ID to the same real server.


SSL offloading

Use SSL offloading to accelerate clients' SSL or HTTPS connections to real servers by using the FortiGate unit to perform SSL operations (offloading them from the real servers using the FortiGate unit's SSL acceleration hardware). FortiGate units can offload SSL 3.0 and TLS 1.0. SSL offloading is available on FortiGate units that support SSL acceleration.

To configure SSL offloading from the web-based manager go to *Firewall Objects > Load Balance > Virtual Server*. Add a virtual server and set the type to HTTPS or SSL and select the SSL offloading type (Client <-> FortiGate or Client <-> FortiGate <-> Server).

Select Client <-> FortiGate to apply hardware accelerated SSL processing only to the part of the connection between the client and the FortiGate unit. This mode is called half mode SSL offloading. The segment between the FortiGate unit and the server will use clear text communications. This results in best performance, but cannot be used in failover configurations where the failover path does not have an SSL accelerator.

Select Client <-> FortiGate <-> Server to apply hardware accelerated SSL processing to both parts of the connection: the segment between client and the FortiGate unit, and the segment between the FortiGate unit and the server. This mode is called full mode SSL offloading. The segment between the FortiGate unit and the server will use encrypted communications, but the handshakes will be abbreviated. This results in performance which is less than the other option, but still improved over communications without SSL acceleration, and can be used in failover configurations where the failover path does not have an SSL accelerator. If the server is already configured to use SSL, this also enables SSL acceleration without requiring changes to the server's configuration.

Figure 347: SSL Offloading modes

Configuring SSL offloading also requires selecting a certificate to use for the SSL offloading sessions. The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

The following CLI command shows an example half mode HTTPS SSL offloading configuration. In the example the `ssl-mode` option sets the SSL offload mode to `half` (which is the default mode).

```
config firewall vip
  edit Vserver-ssl-offload
    set type server-load-balance
    set server-type https
    set ldb-method round-robin
    set extip 172.20.120.30
    set extintf wan1
    set extport 443
    set persistence ssl-session-id
    set ssl-mode half
    set ssl-certificate my-cert
    set monitor tcp-mon-1
    config realservers
      edit 1
        set ip 10.31.101.30
```

```

        set port 443
    next
    edit 2
        set ip 10.31.101.40
        set port 443
    end
end

```

Additional SSL load balancing options

The following SSL load balancing and SSL offloading options are only available from the CLI:

```
ssl-client-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-type {both | client | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the client and the FortiGate unit.

- **both:** Select to expire SSL session states when either `ssl-client-session-state-max` or `ssl-client-session-state-timeout` is exceeded, regardless of which occurs first.
- **count:** Select to expire SSL session states when `ssl-client-session-state-max` is exceeded.
- **disable:** Select to keep no SSL session states.
- **time:** Select to expire SSL session states when `ssl-client-session-state-timeout` is exceeded.

```
ssl-dh-bits <bits_int>
```

Enter the number of bits of the prime number used in the Diffie-Hellman exchange for RSA encryption of the SSL connection. Larger prime numbers are associated with greater cryptographic strength.

```
ssl-http-location-conversion {enable | disable}
```

Select to replace `http` with `https` in the reply's `Location` HTTP header field. For example, in the reply, `Location: http://example.com/` would be converted to `Location: https://example.com/`

```
ssl-http-match-host {enable | disable}
```

Select to apply `Location` conversion to the reply's HTTP header only if the host name portion of `Location` matches the request's `Host` field, or, if the `Host` field does not exist, the host name portion of the request's URI. If disabled, conversion occurs regardless of whether the host names in the request and the reply match.

For example, if host matching is enabled, and a request contains `Host: example.com` and the reply contains `Location: http://example.cc/`, the `Location` field does not match the host of the original request and the reply's `Location` field remains unchanged. If the reply contains `Location: http://example.com/`, however, then the FortiGate unit detects the matching host name and converts the reply field to `Location: https://example.com/`.

This option appears only if `ssl-http-location-conversion` is `enable`.

```
ssl-max-version {ssl-3.0 | tls-1.0}
```

Enter the maximum version of SSL/TLS to accept in negotiation.

```
ssl-min-version {ssl-3.0 | tls-1.0}
```

Enter the minimum version of SSL/TLS to accept in negotiation.

```
ssl-send-empty-frags {enable | disable}
```

Select to precede the record with empty fragments to thwart attacks on CBC IV. You might disable this option if SSL acceleration will be used with an old or buggy SSL implementation which cannot properly handle empty fragments.

```
ssl-server-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the server and the FortiGate unit.

```
ssl-server-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

```
ssl-server-session-state-type {both | count | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

- **both:** Select to expire SSL session states when either `ssl-server-session-state-max` or `ssl-server-session-state-timeout` is exceeded, regardless of which occurs first.
- **count:** Select to expire SSL session states when `ssl-server-session-state-max` is exceeded.
- **disable:** Select to keep no SSL session states.
- **time:** Select to expire SSL session states when `ssl-server-session-state-timeout` is exceeded.

SSL offloading support or Internet Explorer 6

In some cases the Internet Explorer 6 web browser may be able to access real servers. To resolve this issue, disable the `ssl-send-empty-frags` option:

```
config firewall vip
  edit vip_name
    set ssl-send-empty-frags disable
  end
```

You can disable this option if SSL acceleration will be used with an old or buggy SSL implementation that cannot properly handle empty fragments.

Disabling SSL/TLS re-negotiation

The vulnerability [CVE-2009-3555](#) affects all SSL/TLS servers that support re-negotiation. FortiOS when configured for SSL/TLS offloading is operating as a SSL/TLS server. The IETF is working on a TLS protocol change that will fix the problem identified by CVE-2009-3555 while still supporting re-negotiation. Until that protocol change is available, you can use the `ssl-client-renegotiation` option to disable support for SSL/TLS re-negotiation. The default value of this option is `allow`, which allows an SSL client to renegotiate. You can change the setting to `deny` to abort any attempts by an SSL client to renegotiate. If you select `deny` as soon as a `ClientHello` message indicating a re-negotiation is received from the client FortiOS terminates the TCP connection.

Since SSL offloading does not support requesting client certificates the only circumstance in which a re-negotiation is required is when more than 2^{32} bytes of data are exchanged over a single handshake. If you are sure that this volume of traffic will not occur then you can disable re-negotiation and avoid any possibility of the attack described in CVE-2009-3555.

The re-negotiation behavior can be tested using OpenSSL. The OpenSSL `s_client` application has the feature that the user can request that it do renegotiation by typing "R". For example, the following shows a successful re-negotiation against a FortiGate unit configured with a VIP for 192.168.2.100:443:

```
$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1
/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
---
Certificate chain
0
s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/emailAddress=support@fortinet.com
i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
---
Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----

subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/emailAddress=support@fortinet.com
issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
---
No client certificate CA names sent
---
SSL handshake has read 2370 bytes and written 316 bytes
---
New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
    Protocol : TLSv1
    Cipher   : DHE-RSA-AES256-SHA
    Session-ID:

02781E1E368DCCE97A95396FAA82E8F740F5BBA96CF022F6FEC3597B0CC88095
```

```

Session-ID-ctx:
Master-Key:

A6BBBD8477A2422D56E57C1792A4EA9C86F37D731E67D0A66E5CDB2B5C76650780
C0E7F01CFF851EC4466186F4C48397
Key-Arg      : None
Start Time: 1264453027
Timeout      : 300 (sec)
Verify return code: 19 (self signed certificate in
certificate
chain)
---
GET /main.c HTTP/1.0
R
RENEGOTIATING
depth=1
/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
HTTP/1.0 200 ok
Content-type: text/plain

/*
 * Copyright (C) 2004-2007 Fortinet
 */

#include <stdio.h>
#include "vsd_ui.h"

int main(int argc, char **argv)
{
    return vsd_ui_main(argc, argv);
}
closed
$

```

The following is the same test, but this time with the VIP configuration changed to `ssl-client-renegotiation deny`:

```

$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1
/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
---
Certificate chain
0

s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80C
M3909604325/emailAddress=support@fortinet.com
i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate

```

```

    Authority/CN=support/emailAddress=support@fortinet.com
1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
    Authority/CN=support/emailAddress=support@fortinet.com
    i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
    Authority/CN=support/emailAddress=support@fortinet.com
---
Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----

subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN
=FW80CM3909604325/emailAddress=support@fortinet.com
    issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
    Authority/CN=support/emailAddress=support@fortinet.com
---
No client certificate CA names sent
---
SSL handshake has read 2370 bytes and written 316 bytes
---
New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
    Protocol    : TLSv1
    Cipher      : DHE-RSA-AES256-SHA
    Session-ID:

8253331D266DDE38E4D8A04AFCA9CBDED5B1134932CE1718EED6469C1FBC7474
    Session-ID-ctx:
    Master-Key:

ED05A3EF168AF2D06A486362FE91F1D6CAA55CEFC38A3C36FB8BD74236BF2657D4
701B6C1456CEB5BB5EF7619EF12D
    Key-Arg     : None
    Start Time: 1264452957
    Timeout     : 300 (sec)
    Verify return code: 19 (self signed certificate in
certificate
    chain)
---
GET /main.c HTTP/1.0
R
RENEGOTIATING
19916:error:1409E0E5:SSL routines:SSL3_WRITE_BYTES:ssl handshake
failure:s3_pkt.c:530:
Use the following command to check the SSL stats to see that the renegotiations
blocked counter is now 1:
firewall vip virtual-server stats ssl
ssl
    client
        connections total 0 active 0 max 0

```

```

handshakes total 4 active 0 max 0 completed 4 abbreviated 0
session states total 4 active 4 max 4
cipher-suite failures 0
embryonics total 0 active 0 max 0 terminated 0
renegotiations blocked 1
server
connections total 0 active 0 max 0
handshakes total 3 active 0 max 0 completed 2 abbreviated 1
session states total 1 active 1 max 1
cipher-suite failures 0
internal error 0
bad handshake length 0
bad change cipher spec length 0
pubkey too big 0
persistence
find 0 found 0 clash 0 addr 0 error 0

```

If the virtual server debug log is examined (diag debug appl vs -1) then at the point the renegotiation is blocked there is a log:

```

vs ssl 12 handshake recv ClientHello
vs ssl 12 handshake recv 1
(0100005403014b5e056c7f573a563bebe0258c3254bbaff7046a461164f34f94f
4f3d019c41800002600390038003500160013000a00330032002f0005000400150
012000900140011000800060003020100000400230000)
vs ssl 12 client renegotiation attempted rejected, abort
vs ssl 12 closing 0 up
vs src 12 close 0 in
vs src 12 error closing
vs dst 14 error closing
vs dst 14 closed
vs ssl 14 close
vs sock 14 free
vs src 12 closed
vs ssl 12 close
vs sock 12 free

```

IP, TCP, and UDP load balancing

You can load balance all IP, TCP or UDP sessions accepted by the security policy that includes a load balancing virtual server with the type set to IP, TCP, or UDP. Traffic with destination IP and port that matches the virtual server IP and port is load balanced. For these protocol-level load balancing virtual servers you can select a load balance method and add real servers and health checking. However, you can't configure persistence, HTTP multiplexing and SSL offloading.



Load balancing configuration examples


This chapter includes the following examples:

- Example: HTTP load balancing to three real web servers
- Example: Basic IP load balancing configuration
- Example: Adding a server load balance port forwarding virtual IP
- Example: Weighted load balancing configuration
- Example: HTTP and HTTPS persistence configuration

Example: HTTP load balancing to three real web servers

In this example, the virtual web server IP address 192.168.37.4 on the Internet, is mapped to three real web servers connected to the FortiGate unit dmz1 interface. The real servers have IP addresses 10.10.123.42, 10.10.123.43, and 10.10.123.44. The virtual server uses the *First Alive* load balancing method. The configuration also includes an HTTP health check monitor that includes a URL used by the FortiGate unit for get requests to monitor the health of the real servers.

Connections to the virtual web server at IP address 192.168.37.4 from the Internet are translated and load balanced to the real servers by the FortiGate unit. First alive load balancing directs all sessions to the first real server. The computers on the Internet are unaware of this translation and load balancing and see a single virtual server at IP address 192.168.37.4 rather than the three real servers behind the FortiGate unit.

Figure 348: Virtual server configuration example

Web-based manager configuration

Use the following procedures to configure this load balancing setup from the web-based manager.

To add an HTTP health check monitor

In this example, the HTTP health check monitor includes the URL “/index.html” and the Matched Phrase “Fortinet products”.

- 1 Go to *Firewall Objects > Load Balance > Health Check*.
- 2 Select *Create New*.
- 3 Add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase “Fortinet products”.

Name	HTTP_health_chk_1
Type	HTTP
Port	80
URL	/index.html
Matched Content	Fortinet products
Interval	10 seconds
Timeout	2 seconds
Retry	3

- 4 Select *OK*.

To add the HTTP virtual server

- 1 Go to *Firewall Objects > Load Balance > Virtual Server*.

- 2 Select *Create New*.
- 3 Add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

Name	Load_Bal_VS1
Type	HTTP
Interface	wan1
Virtual Server IP	192.168.37.4 The public IP address of the web server. The virtual server IP address is usually a static IP address obtained from your ISP for your web server. This address must be a unique IP address that is not used by another host and cannot be the same as the IP address of the external interface the virtual IP will be using. However, the external IP address must be routed to the selected interface. The virtual IP address and the external IP address can be on different subnets. When you add the virtual IP, the external interface responds to ARP requests for the external IP address.
Virtual Server Port	80
Load Balance Method	First Alive
Persistence	HTTP cookie
HTTP Multiplexing	Select. The FortiGate unit multiplexes multiple client into a few connections between the FortiGate unit and each real HTTP server. This can improve performance by reducing server overhead associated with establishing multiple connections.
Preserve Client IP	Select The FortiGate unit preserves the IP address of the client in the X-Forwarded-For HTTP header.
Health Check	Move the HTTP_health_chk_1 health check monitor to the <i>Selected</i> list.

- 4 Select *OK*.

To add the real servers and associate them with the virtual server

- 1 Go to *Firewall Objects > Load Balance > Real Server*.
- 2 Select *Create New*.

- 3 Configure three real servers that include the virtual server Load_Bal_VS1. Each real server must include the IP address of a real server on the internal network.

Configuration for the first real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.42
Port	80
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

Configuration for the second real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.43
Port	80
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

Configuration for the third real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.44
Port	80

Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

To add the virtual server to a security policy

Add a wan1 to dmz1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Configure the security policy:

Source Interface/Zone	wan1
Source Address	all (or a more specific address)
Destination Interface/Zone	dmz1
Destination Address	Load_Bal_VS1
Schedule	always
Service	HTTP
Action	ACCEPT
Log Allowed Traffic	Select to log virtual server traffic
NAT	Enable NAT

- 4 Select other security policy options as required.
- 5 Select *OK*.

CLI configuration

Use the following procedure to configure this load balancing setup from the CLI.

To configure HTTP load balancing

- 1 Use the following command to add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase "Fortinet products".

```
config firewall ldb-monitor
edit HTTP_health_chk_1
set type http
set port 80
```

```
set http-get /index.html
set http-match "Fortinet products"
set interval 10
set timeout 2
set retry 3
end
```

- 2** Use the following command to add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

```
config firewall vip
edit Load-Bal_VS1
set type server-load-balance
set server-type http
set ldb-method first-alive
set http-multiplex enable
set http-ip-header enable
set extip 192.168.37.4
set extintf wan1
set extport 80
set persistence http-cookie
set monitor HTTP_health_chk_1
config realservers
edit 1
set ip 10.10.10.42
set port 80
next
edit 2
set ip 10.10.10.43
set port 80
next
edit 3
set ip 10.10.10.44
set port 80
end
end
```

- 3** Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```
config firewall policy
edit 0
set srcintf wan1
set srcaddr all
set dstintf dmz1
set dstaddr Load-Bal_VS1
set action accept
set schedule always
set service ANY
set nat enable
end
```

Configure other security policy settings as required.

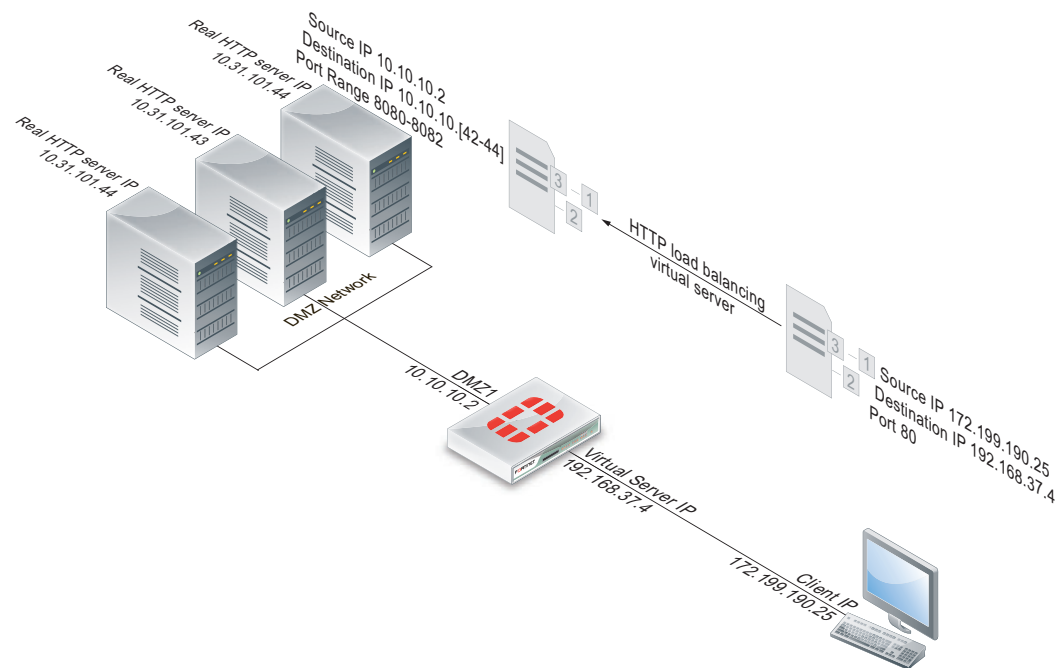
Example: Basic IP load balancing configuration

This example shows how to add a server load balancing virtual IP that load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is 192.168.20.20. The load balancing method is `weighted`. The IP addresses of the real servers are 10.10.10.1, 10.10.10.2, and 10.10.10.3. The weights for the real servers are 1, 2, and 3. The default weight is 1 and does not have to be changed for the first real server.

```
config firewall vip
  edit All_Load_Balance
    set type server-load-balance
    set server-type ip
    set extintf port2
    set extip 192.168.20.20
    set ldb-method weighted
    config realservers
      edit 1
        set ip 10.10.10.1
      next
      edit 2
        set ip 10.10.10.2
        set weight 2
      next
      edit 3
        set ip 10.10.10.3
        set weight 3
      end
    end
end
```

Example: Adding a server load balance port forwarding virtual IP

This example is the same as the example described in [“Example: HTTP load balancing to three real web servers” on page 2899](#) except that each real server accepts HTTP connections on a different port number. The first real server accepts connections on port 8080, the second on port 8081, and the third on 8082.

Figure 349: Server load balance virtual IP port forwarding

To complete this configuration, all of the steps would be the same as in [“Example: HTTP load balancing to three real web servers”](#) on page 2899 except for configuring the real servers.

To add the real servers and associate them with the virtual server

Use the following steps to configure the FortiGate unit to port forward HTTP packets to the three real servers on ports 8080, 8081, and 8082.

- 1 Go to *Firewall Objects > Load Balance > Real Server*.
- 2 Select *Create New*.

- 3 Configure three real servers that include the virtual server Load_Bal_VS1. Each real server must include the IP address of a real server on the internal network and have a different port number.

Configuration for the first real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.42
Port	8080
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the second real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.43
Port	8081
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the third real server.

Virtual Server	Load_Bal_VS1
IP	10.10.10.44
Port	8082
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Example: Weighted load balancing configuration

This example shows how to using firewall load balancing to load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is 192.168.20.20. The load balancing method is `weighted`. The IP addresses of the real servers are 10.10.10.1, 10.10.10.2, and 10.10.10.3. The weights for the real servers are 1, 2, and 3.

This configuration does not include an health check monitor.

Web-based manager configuration

Use the following procedures to configure this load balancing setup from the web-based manager.

To add the HTTP virtual server

- 1 Go to *Firewall Objects > Load Balance > Virtual Server*.
- 2 Select *Create New*.

- 3 Add an IP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate port2 interface is connected to the Internet.

Name	HTTP_weghted_LB
Type	IP
Interface	port2
Virtual Server IP	192.168.20.20
Load Balance Method	Weighted

All other virtual server settings are not required or cannot be changed.

- 4 Select *OK*.

To add the real servers and associate them with the virtual server

- 1 Go to *Firewall Objects > Load Balance > Real Server*.
- 2 Select *Create New*.

- 3 Configure three real servers that include the virtual server All_Load_Balance. Because the *Load Balancing Method* is *Weighted*, each real server includes a weight. Servers with a greater weight receive a greater proportion of forwarded connections, Configuration for the first real server.

Virtual Server	HTTP_weghted_LB
IP	10.10.10.1
Port	Cannot be configured because the virtual server is an IP server.
Weight	1
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

Configuration for the second real server.

Virtual Server	HTTP_weghted_LB
IP	10.10.10.2
Port	Cannot be configured because the virtual server is an IP server.
Weight	2
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

Configuration for the third real server.

Virtual Server	HTTP_weghted_LB
IP	10.10.10.3
Port	Cannot be configured because the virtual server is an IP server.

Weight	3
Maximum Connections	0 Setting <i>Maximum Connections</i> to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses <i>First Alive</i> load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the <i>Maximum Connections</i> is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.

To add the virtual server to a security policy

Add a prot2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Configure the security policy:

Source Interface/Zone	port2
Source Address	all (or a more specific address)
Destination Interface/Zone	port1
Destination Address	HTTP_weghted_LB
Schedule	always
Service	ANY
Action	ACCEPT
NAT	Select

- 4 Select other security policy options as required.
- 5 Select *OK*.

CLI configuration

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`. The default weight is 1 and does not have to be changed for the first real server.

Use the following command to add the virtual server and the three weighted real servers.

```
config firewall vip
  edit HTTP_weghted_LB
    set type server-load-balance
    set server-type ip
    set extintf port2
    set extip 192.168.20.20
    set ldb-method weighted
  config realservers
    edit 1
```

```

        set ip 10.10.10.1
    next
    edit 2
        set ip 10.10.10.2
        set weight 2
    next
    edit 3
        set ip 10.10.10.3
        set weight 3
    end
end

```

Example: HTTP and HTTPS persistence configuration

This example shows how to add a virtual server named *Http_Load_Balance* that load balances HTTP traffic using port 80 and a second virtual server named *Https_Load_Balance* that load balances HTTPS traffic using port 443. The Internet is connected to port2 and the virtual IP address of the virtual server is 192.168.20.20. Both server load balancing virtual IPs load balance sessions to the same three real servers with IP addresses 10.10.10.2, 10.10.10.2, and 10.10.10.3. The real servers provide HTTP and HTTPS services.

For both virtual servers, persistence is set to *HTTP Cookie* to enable HTTP cookie persistence.

To add the HTTP and HTTPS virtual servers

- 1 Go to *Firewall Objects > Load Balance > Virtual Server*.
- 2 Add the HTTP virtual server that includes HTTP Cookie persistence.

Name	HTTP_Load_Balance
Type	HTTP
Interface	port2
Virtual Server IP	192.168.20.20
Virtual Server Port	80 In this example the virtual server uses port 8080 for HTTP sessions instead of port 80.
Load Balance Method	Static
Persistence	HTTP cookie

- 3 Select *OK*.
- 4 Select *Create New*.
- 5 Add the HTTPS virtual server that also includes HTTP Cookie persistence.

Name	HTTPS_Load_Balance
Type	HTTPS
Interface	port2
Virtual Server IP	192.168.20.20
Virtual Server Port	443

Load Balance Method	Static
Persistence	HTTP cookie

6 Select *OK*.

To add the real servers and associate them with the virtual servers

- 1 Go to *Firewall Objects > Load Balance > Real Server*.
- 2 Select *Create New*.
- 3 Configure three real servers for HTTP that include the virtual server HTTP_Load_Balance.

Configuration for the first HTTP real server.

Virtual Server	HTTP_Load_Balance
IP	10.10.10.1
Port	80
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the second HTTP real server.

Virtual Server	HTTP_Load_Balance
IP	10.10.10.2
Port	80
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the third HTTP real server.

Virtual Server	HTTP_Load_Balance
IP	10.10.10.3
Port	80
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

- 4 Configure three real servers for HTTPS that include the virtual server HTTPS_Load_Balance.

Configuration for the first HTTPS real server.

Virtual Server	HTTP_Load_Balance
IP	10.10.10.1
Port	443
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the second HTTPS real server.

Virtual Server	HTTP_Load_Balance
IP	10.10.10.2
Port	443
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

Configuration for the third HTTPS real server.

Virtual Server	HTTPS_Load_Balance
IP	10.10.10.3
Port	443
Weight	Cannot be configured because the virtual server does not include weighted load balancing.
Maximum Connections	0

To add the virtual servers to security policies

Add a port2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

- 1 Go to *Policy > Policy > Policy*.
- 2 Select *Create New*.
- 3 Configure the HTTP security policy:

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	HTTP_Load_Balance
Schedule	always
Service	HTTP

Action	ACCEPT
NAT	Enable NAT

- 4 Select other security policy options as required.
- 5 Select *OK*.
- 6 Select *Create New*.
- 7 Configure the HTTP security policy:

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	HTTPS_Load_Balance
Schedule	always
Service	HTTPS
Action	ACCEPT
NAT	Enable NAT

- 8 Select other security policy options as required.
- 9 Select *OK*.

CLI configuration: adding persistence for a specific domain

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`.

For the CLI configuration, both virtual servers include setting `http-cookie-domain` to `.example.org` because HTTP cookie persistence is just required for the `example.org` domain.

First, the configuration for the HTTP virtual IP:

```
config firewall vip
  edit HTTP_Load_Balance
    set type server-load-balance
    set server-type http
    set extport 8080
    set extintf port2
    set extip 192.168.20.20
    set persistence http-cookie
    set http-cookie-domain .example.org
  config realservers
    edit 1
      set ip 10.10.10.1
    next
    edit 2
      set ip 10.10.10.2
    next
    edit 3
      set ip 10.10.10.3
    end
  end
```


Second, the configuration for the HTTPS virtual IP. In this configuration you don't have to set extport to 443 because extport is automatically set to 443 when server-type is set to https.

```
config firewall vip
edit HTTPS_Load_Balance
set type server-load-balance
set server-type https
set extport 443
set extintf port2
set extip 192.168.20.20
set persistence http-cookie
set http-cookie-domain .example.org
config realservers
edit 1
set ip 10.10.10.1
next
edit 2
set ip 10.10.10.2
next
edit 3
set ip 10.10.10.3
end
end
```




Chapter 19 Hardware

This FortiOS Handbook chapter contains the following sections:

[FortiGate installation](#) describes installing your FortiGate unit, and how to mount the FortiGate in a rack, if applicable.

[FortiGate hardware accelerated processing](#) some FortiGate models incorporate network processors in the main unit, others support the addition of AMC (Advanced Mezzanine Card) modules. The FortiGate-5000 series supports rear transition modules (RTMs) that incorporate network processors. This chapter describes how hardware acceleration works as well as how to take full advantage of its benefits.

[Configuring RAID](#) some FortiGate models have two or more hard disks configured in a RAID array to store log messages locally on the FortiGate unit. A RAID array can provide faster disk access, redundancy in case of partial failure, or both depending on the RAID level you select. This section describes how to configure RAID on FortiGate units that support it.

[FortiBridge installation and operation](#) describes a typical transparent mode FortiGate network and how to add a FortiBridge unit to provide fail open protection. In addition, detailed information about how FortiBridge units operate, a description of to add a FortiBridge unit to an HA cluster, and connecting a FortiBridge unit other FortiGate interfaces is included.



FortiGate installation

This chapter describes installing your FortiGate unit, environmental specifications, and how to mount the FortiGate unit.

This chapter contains the following topics:

- [Mounting the FortiGate unit](#)
- [Plugging in the FortiGate unit](#)
- [Turning off the FortiGate unit](#)
- [Further configuration](#)

Mounting the FortiGate unit

Most FortiGate units can be either rack mounted, or placed on a desk or table. Only the smallest units have no rack mounting hardware. The largest units are designed for rack mounting.

Desk or table mounting

Attach the provided rubber feet to the bottom of the FortiGate unit if they are not already attached.

Place the FortiGate unit on any flat, stable surface, ensure the unit has at least 1.5 inches (3.75 cm) of clearance on each side to ensure adequate airflow for cooling.

Rack mounting

If you are placing a 1U or 2U FortiGate unit into a rack, remove the rubber feet from the bottom of the FortiGate unit.

For rack mounting, use the mounting brackets and screws included with the FortiGate unit. The 3U 3900-series FortiGate units can be rack-mounted using either slide rails or middle-mount brackets and both procedures are covered below.

Rack mount considerations

Elevated operating ambient — If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.

Reduced air flow — Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.

Mechanical loading — Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.

Circuit overloading — Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.

Reliable ground — Reliable electrical grounding of rack-mounted equipment should be maintained.

Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips).



Depending on the size of your FortiGate unit, you may require two or more people to safely install the unit in the rack.

To install a 1U or 2U FortiGate unit into a rack

- 1 Attach the mounting brackets to the side to the unit so that the brackets are on the front portion of the FortiGate unit. Ensure that the screws are tight.

The following photos illustrate how the brackets should be mounted. Note that the screw configuration may vary depending on your FortiGate unit.

Figure 350: Installed 1U mounting brackets



Figure 351: Installed 2U mounting brackets




- 2 Position the FortiGate unit in the rack to allow for sufficient air flow.
- 3 Line up the mounting bracket holes to the holes on the rack, ensuring the FortiGate unit is level.
- 4 Finger tighten the screws to attach the FortiGate unit to the rack.
- 5 Once you verify the spacing of the FortiGate unit and that it is level, tighten the screws with a screwdriver. Ensure that the screws are tight.

The following photos illustrate how the mounting brackets and FortiGate unit should be attached to the rack.

Figure 352: Mounting a 1U FortiGate unit in a rack**Figure 353: Mounting a 2U FortiGate unit in a rack****To install a 3U 3900-series FortiGate using slide rails**

- 1 Before you start, confirm that you have the two slide rails and two front handles.

Figure 354: Slide rails and front handles.

- 2 Attach the internal rails to each side of the unit. The rail should snap on and slide over until you hear a click from the rear clip.

Figure 355: Locking rear clip on unit.

- 3 Optionally, you can add a screw to make the rail more secure.

Figure 356: Optional screw hole for additional support.




Release tab

Insert screw here to provide additional support to the internal rail


- 4 Attach the front handles to each side at the front of the unit with three screws. Note that the front handles are not used as rack mounts. Use only as handles to slide the unit in and out of the rack.

Figure 357: Attaching front handle to unit.




Attach front handles to unit with three screws.

- 5 Orient the external rail on the rack. Ensure that the ball bearing track is forward. The front of the rail is labelled "Front" and the end of the rail is labelled "Rear".
- 6 Extend the external rail to fit the rack. Use the locking mechanism on the front and back of the rail to lock into place.

Figure 358: Front locking mechanism.**Figure 359: Rear locking mechanism.**

- 7 Use at least two people to lift the unit and insert the system approximately halfway onto the rack by sliding the external rails over the internal rails.
- 8 Slide the release tabs on both sides of the internal rails and push the system into the rack. Move your fingers away from the release tabs once the system is in motion.

Figure 360: FortiGate unit halfway on rack showing release tabs.

- 9 Lock the system into place by squeezing together the buttons at the front of the rail.

Figure 361: FortiGate unit on rack.



Squeeze buttons to lock unit into place

10 Optionally, you can add a screw through the front handles for more security.

Figure 362: Location of locking mechanism on rail and screw hole in front handle.




Hole in front handle allows you to screw the unit to the rack


Locking mechanism on the rail.

To install a 3U 3900-series FortiGate using the middle rack mount brackets

1 Before you start, confirm that you have the two middle rack mount brackets.

Figure 363: Middle rack mount brackets.


- 2 Attach the middle rack mount brackets to each side of the unit using five screws for each mount. Ensure the middle piece faces outwards.

Figure 364: Attaching the middle rack mount brackets to the sides of the unit.

Attach the middle rack mount ears to the side of the unit using five screws

Middle piece should face outwards

- 3 Use at least two people to lift the unit and insert the system halfway onto the rack until the middle rack mount brackets meet the stand-alone rack.
- 4 While the two people hold the unit, use another person to attach the middle piece of the middle rack mount brackets to the stand-alone rack using two screws.

Figure 365: Attaching the middle rack mount brackets to the stand-alone rack.

Attach the middle rack mount ears to the stand-alone rail using two screws

- 5 Ensure you attach both middle rack mount brackets to both sides of stand-alone rack.

Figure 366: Fortinet unit on the stand-alone rack.



Plugging in the FortiGate unit

Most FortiGate unit do not have an on/off switch. Check the quick-start guide included with your FortiGate unit to see if your model has an on/off switch.

To power on the FortiGate unit

- 1 Connect the power cable to the power connection on the back of the FortiGate unit. If your model has multiple power connections, connect cables to all the connections.
- 2 Connect the power cable or cables to power outlets.

Each power cable should be connected to a different power source. If one power source fails, the other may still be operative.

The FortiGate unit starts and the Power and Status (if available) LEDs light up. The Status LED (if available) flashes while the FortiGate unit starts, and remains lit when the system is running.



If the FortiGate unit has two power supplies and only one is connected, an audible alarm sounds to indicate a failed power supply. Press the red alarm cancel button on the rear panel next to the power supply to stop the alarm.

Connecting to the network

Using the supplied Ethernet cable, connect one end of the cable to your router or modem, whichever is the connection to the Internet. Connect the other end to the FortiGate unit. Connect it to either the External, WAN port, or port 1 interface. Use additional cables to connect the Internal port or port 2 to your internal hub or switch.

Turning off the FortiGate unit

Always shut down the FortiGate operating system properly before turning off the power switch or unplugging the unit to avoid potential hardware problems.

To power off the FortiGate unit

- 1 From the web-based manager, go to *System > Dashboard > Status*.
- 2 In the *Unit Operation* display, select *Shutdown*, or from the CLI enter:
`execute shutdown`
- 3 Wait a moment for the shutdown operation to finish.
- 4 Disconnect the power cables from the power supply.

Further configuration

Further configuration is beyond the scope of this installation guide.

“[Basic setup](#)” on [page 335](#) describes how to configure the operating mode, interface addresses, DNS server, the default gateway, and firewall policies.



Management interfaces do not support VLANs



AMC module configuration

This section explains how to configure AMC modules on the FortiGate unit. This includes auto-bypass and recovery for AMC bridge modules.

The following topics are included in this section:

- [Configuring AMC modules](#)
- [Auto-bypass and recovery for AMC bridge module](#)
- [Enabling or disabling bypass mode for AMC bridge modules](#)

Configuring AMC modules

By default, FortiGate units automatically recognize the AMC modules installed in their AMC slots or automatically recognize that an AMC slot is empty. If the module contains interfaces, FortiOS automatically adds the interfaces to the FortiGate configuration. If the module contains a hard disk, the hard disk is automatically added to the configuration. However, when the FortiGate unit is powered down and the module removed from the slot, when the FortiGate unit restarts it automatically recognizes that the slot is empty and will not retain any configuration settings for the missing module.

This default behavior is usually acceptable in most cases. However, it can be useful when a module is present in a slot to add the name of the module to the FortiGate configuration. Then, if the module fails or if you temporarily remove it from the slot, the FortiGate unit keeps the module's configuration settings so that when the module is replaced you will not have to re-configure it.

If you have added the name of a module to a slot and you are planning or removing the module and replacing it with a different type of module (for example, if you are removing a FortiGate-ASM-S08 and replacing it with a FortiGate-ASM-FX2) you should reset the slot to the default before removing the module. Then after adding the new module you should add its name to the slot.

You configure AMC slot settings from the FortiGate CLI using the `config system amc` command. For information about this command, see the [FortiGate CLI Reference](#).

The following procedure shows how to add a FortiGate-ADM-FB8 to the first double-width AMC slot (dw1) and how to add the name of the module to the slot configuration.

To change the default setting for an AMC slot

- 1 Enter the following CLI command to verify that the slot that you will insert the FortiGate-ADM-FB8 module into is set to the default configuration.

This command lists the AMC slots and the settings for each one. Example command output for a FortiGate-5001A with an empty double-width AMC slot:

```
get system amc
dw1          : auto
```

- 2 Power down the FortiGate unit.
- 3 Insert the FortiGate-ADM-FB8 module into the double-width AMC slot.

- 4 Power up the FortiGate unit.

As long as the slot that you have inserted the FortiGate-ADM-FB8 module into is set to `auto` the FortiGate unit should automatically find the module when it powers up.

- 5 Add the name of the FortiGate-ADM-FB8 module to the FortiGate configuration.

```
config system amc
  set dw1 adm-fb8
end
```

Auto-bypass and recovery for AMC bridge module

The FortiGate-ASM-CX4 and FortiGate-ASM-FX2 modules provide fail open protection for interface pairs of FortiGate units operating in Transparent mode and that have a single-width AMC slot. The FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module bridges FortiGate interfaces, monitors the interfaces for traffic failures, and operate as pass-through devices if the interfaces or the entire FortiGate unit fails or for some reason cannot pass traffic between the interfaces. If a failure occurs, traffic bypasses the FortiGate unit and passes through the FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module to make sure that the network can continue processing traffic after a FortiGate failure.

This section describes how to configure a FortiGate unit to use a FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module to bridge FortiGate interfaces. The FortiGate unit must operate in Transparent mode and the FortiGate-ASM-CX4 and FortiGate-ASM-FX2 modules are not compatible with FortiGate HA.

The FortiGate-ASM-CX4 and FortiGate-ASM-FX2 modules include a bypass watchdog that continually verifies that traffic is flowing through the bridged FortiGate interfaces. If traffic stops flowing, for example if the FortiGate unit fails, and if the bypass watchdog detects this, the bridge module switches to bypass mode to ensure the flow of traffic on the network.

In bypass mode all traffic flows between interfaces on the FortiGate-ASM-CX4 and FortiGate-ASM-FX2 modules and not through the FortiGate unit. You can configure a recovery watchdog to verify that the bridged FortiGate interfaces cannot process traffic. If you fix the problem or the problem fixes itself, the recovery watchdog automatically detects that traffic can resume and switches the module back to normal operation by turning off bypass mode.

To configure a FortiGate unit to operate with a FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module

- 1 Switch the FortiGate unit to operate in transparent mode.

```
config system settings
  set opmode transparent
  set manageip <management_IPv4> <netmask_ipv4>
  set gateway <gateway_ipv4>
end
```

After a short pause the FortiGate unit is operating in transparent mode.

- 2 Enter the following command to verify that the slot that you will insert the FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module into is set to `auto`.

This command lists the AMC slots and the settings for each one. Example command output for a FortiGate-620B with an empty AMC slot:

```
get system amc
sw1          : auto
```


- 3 Power down the FortiGate unit.
- 4 Insert the FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module into a single-width AMC slot.
- 5 Power up the FortiGate unit.

As long as the slot that you have inserted the module into is set to `auto` the FortiGate unit should automatically find the module when it powers up.

- 6 Add the name of the module to the FortiGate configuration and configure bypass and recovery settings.

The following command configures AMC single width slot 1 (sw1) for a FortiGate-ASM-CX4.

This command also enables the bypass watchdog and increases the bypass timeout from the default value of 10 seconds to 60 seconds. This means that if a failure occurs the bridge module will change to bypass mode 60 seconds after the bypass watchdog detects the failure.

This command also enables watchdog recovery and sets the watchdog recovery period to 30 seconds. This means that if a failure occurs, while the FortiGate-ASM-CX4 module is bridging the connection the AMC bypass watchdog monitors FortiGate processes and will revert to normal operating mode (that is disable the bridging the interfaces with the FortiGate-ASM-CX4 module) if the FortiGate unit recovers from the failure.

```
config system amc
    set sw1 asm-cx4
    set bypass-watchdog enable
    set bypass-timeout 60
    set watchdog-recovery enable
    set watchdog-recovery-period 30
end
```

Enabling or disabling bypass mode for AMC bridge modules

Use the `execute amc bypass` command to switch between normal mode and bypass mode for a FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module installed in an single-width AMC slot in a FortiGate unit. Normally the FortiGate-ASM-CX4 and FortiGate-ASM-FX2 modules operate with bypass mode disabled and traffic passes through the FortiGate interfaces bridged by the FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module. You can use this command manually enable bypass mode and force traffic to bypass the FortiGate interfaces and pass through the FortiGate-ASM-CX4 or FortiGate-ASM-FX2 module.

Also, if bypass mode has been enabled (using this command or because of a failure), you can also use this command to manually disable bypass mode and resume normal operation. This can be useful if the problem that caused the failure has been fixed and normal operation can resume.

To manually enable bypass mode

- 1 Use the following command to manually enable bypass mode:
`execute amc bypass enable`

- 2 Use the following diagnose command to view the status of the AMC modules installed in a FortiGate unit, including whether they are operating in bypass mode.

For example if you have installed a FortiGate-ASM-CX4 module in AMC slot 2 of a FortiGate-3810A and bypass mode is enabled:

```
diagnose sys amc bypass status
ASM-CX4 in slot 2:
    amc-sw2/1 <--> amc-sw2/2: mode=bypass (admin action)
    amc-sw2/3 <--> amc-sw2/4: mode=bypass (admin action)
```

```
Daemon heartbeat status: normal
Last heartbeat received: 0 second(s) ago
```

- 3 Log into the web-based manager and go to *System > Dashboard > Status* and view the *Unit Operation* widget to see the status of the AMC bridge module.

To manually disable bypass mode

- 1 Use the following command to manually disable bypass mode:

```
execute amc bypass disable
```

- 2 Use the following diagnose command to view the status of the AMC modules installed in a FortiGate unit, including whether they are operating in bypass mode.

For example if you have installed a FortiGate-ASM-CX4 module in AMC slot 2 of a FortiGate-3810A and bypass mode is disabled:

```
diagnose sys amc bypass status
ASM-CX4 in slot 2:
    amc-sw2/1 <--> amc-sw2/2: mode=normal
    amc-sw2/3 <--> amc-sw2/4: mode=normal
```

```
Daemon heartbeat status: normal
Last heartbeat received: 1 second(s) ago
```

- 3 Log into the web-based manager and go to *System > Dashboard > Status* and view the *Unit Operation* widget to see the status of the AMC bridge module.



FortiGate hardware accelerated processing

Many FortiGate models can offload some types of network traffic processing from main processing resources to specialized network processors. If your network has a significant volume of traffic that is suitable for offloading, this hardware acceleration can significantly improve your network throughput.

Some FortiGate models incorporate network processors in the main unit, others support the addition of AMC (Advanced Mezzanine Card) modules. The FortiGate-5000 series supports rear transition modules (RTMs) that incorporate network processors.

This chapter contains the following topics:

- [How hardware acceleration alters packet flow](#)
- [Network processors overview](#)
- [Content processors overview](#)
- [Security processing modules overview](#)
- [Configuring overall security priorities](#)
- [Configuring traffic offloading](#)
- [Configuring IPsec VPN offloading](#)
- [Configuring IPS offloading](#)

How hardware acceleration alters packet flow

Hardware acceleration generally alters packet processing flow as follows:

- 1 Packets initiating a session pass to the FortiGate unit's main processing resources.
- 2 The FortiGate unit assesses whether the session matches fast path (offload) requirements.

To be suitable for offloading, traffic must possess only characteristics that can be processed by the fast path. For a list of requirements, see [“Configuring traffic offloading” on page 2940](#).

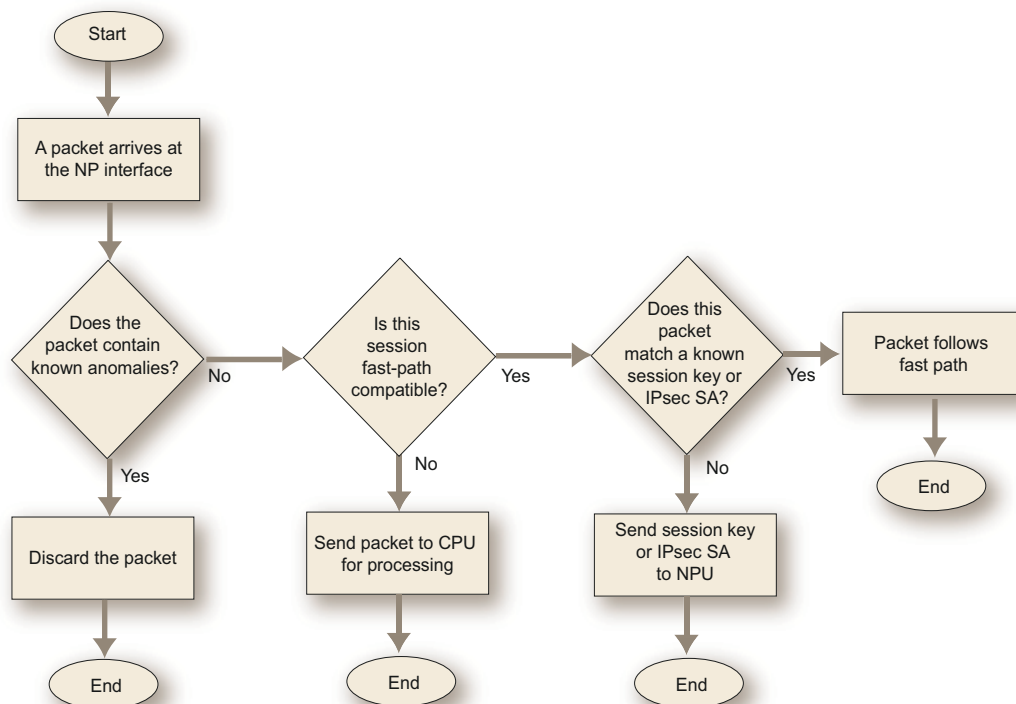
If the traffic is categorized as fast path friendly, the FortiGate unit sends the session key or IPsec security association (SA) and configured processing action to the network processor(s).

- 3 Network processors continuously match packets arriving on their attached ports against the session keys and SAs they have received from the FortiGate unit's main processing resources.
- If a network processor's network interface is configured to perform hardware accelerated anomaly checks, the network processor drops or accepts packets which match the configured anomaly patterns. These checks are separate from and in advance of anomaly checks performed by IPS, which is not compatible with network processor offloading. See ["Configuring pre-IPS anomaly detection" on page 2955](#).
 - The network processor next checks for a matching session key or SA. If a matching session key or SA is found, and if the packet meets packet requirements, the network processor processes the packet according to the configured action and then sends the resulting packet. Packet processing is hardware accelerated.
 - If a matching session key or SA is not found, or if the packet does not meet packet requirements, the traffic cannot be offloaded. The network processor sends the data to the FortiGate unit's main processing resources, which process the packet. Packet processing is similar to normal network interfaces (that is, packet processing is not hardware accelerated by the network processor, and requires main processing resources). Packet forwarding occurs at normal rates.



Network processors do not count offloaded packets, and offloaded packets will not be included in traffic statistics, such as FortiAnalyzer traffic reports.

Figure 367: Deciding the packet flow for accelerated interfaces



Some traffic processing can still be hardware accelerated, even though it does not meet general offloading requirements. For example, some IPsec traffic originates from the FortiGate unit itself and does not follow the offloading requirement of ingress from a network processor's network interface, but FortiGate units can still utilize network processor encryption capabilities. See [“Configuring IPsec VPN offloading” on page 2949](#).

Packet forwarding rates vary by the percentage of offloadable processing and the type of network processing required by your configuration, but are independent of frame size. For optimal traffic types, network throughput can equal wire speed.

Offloading requirements vary slightly by the model of the network processor.

The following types of acceleration hardware are found on FortiGate units:

- network processors: NP1 (formerly known as FA2), NP2, NP4
- content processors: CP4, CP5, CP6
- accelerated interface modules: ASM-FB4, ADM-FB8, ADM-XB2, ADM-XD4, RTM-XD2
- security processor modules: ASM-CE4, ASM-XE2

Network processors overview

Many Fortinet products contain network processors. Some of these products contain NP1 network processors (also known as FortiAccel, or FA2), while others contain NP2 network processors. Some newer models contain an NP4 processor. Network processor features, and therefore offloading requirements, vary by network processor model. Differing offloading requirements are noted in [“Configuring traffic offloading” on page 2940](#) and [“Configuring IPsec VPN offloading” on page 2949](#).

Network processor models

FortiASIC network processors work at the interface level to support IPsec offload and unicast UDP/TCP traffic forwarding. The maximum throughput and number of network interfaces varies by processor model.

NP1: supports FW and VPN acceleration with 2Gbps capacity. It is found on FortiGate units such models 1000A-FA2, 3600A, and 3810A, and also on FortiGate-5000 series 5001FA2 and 5005FA2 blades.

NP2: supports FW and VPN acceleration with 4Gbps capacity. It is found on newer, B-series FortiGate units ranging from models 200B to 3016B, and on most AMC accelerated interface cards.

NP4: supports FW and VPN acceleration with 40 Gbps capacity. It is found on the ADM-XD4 AMC card and on the FortiGate-5000 series RTM-XD2 blade.

Table 151: Network processor models

Processor	Interfaces
NP1	2 x 1 Gb/s
NP2	1 x 10Gb/s, 4 x 1Gb/s
NP4	2 x 10Gb/s



The NP1 network processor does not support frames greater than 1500 bytes. If your network uses jumbo frames, you may need to adjust the MTU (Maximum Transmission Unit) of devices connected to NP1 ports. Maximum frame size for NP2 and NP4 processors is 9000 bytes.



For both NP1 and NP2 network processors, ports attached to a network processor cannot be used for firmware installation by TFTP.

Some Fortinet products contain multiple network processors. Depending on the product, network processors may or may not be directly connected to each other on the circuit board through an EEI (Enhanced Extension Interface).

- Directly connected network processors have an EEI, and can pass traffic between them without involving the FortiGate unit's main processing resources.
- Indirectly connected network processors have no EEI, and **cannot** pass traffic between them without involving the FortiGate unit's main processing resources.

Sessions can only be offloaded if both the source and destination port are connected to the same network processor or directly (EEI) connected network processor pair.

For information about the network processors in any specific FortiGate model, refer to the product brochure.

Determining the network processors installed on your FortiGate unit

To list the network processors on your FortiGate unit, use the following CLI command.

```
get hardware npu <model> list
```

<model> can be np1, np2 or np4.

The output lists the interfaces that have the specified processor. For example,

```
# get hardware npu np1 list
ID          Interface
0           port9 port10
```

This command does not detect Security processing modules.

Content processors overview

The FortiASIC Content Processor (CP) works at the system level. Its main functions are SSL VPN key generation and SSL offloading. Capabilities vary by model.

CP4

- FIPS-compliant DES/3DES/AES encryption and decryption

- SHA-1 and MD5 HMAC
- IPSEC protocol processor
- Random Number generator
- Public Key Crypto Engine
- Content processing engine
- ANSI X9.31 and PKCS#1 certificate support

CP5

- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321/2104/2403/2404 and FIPS180/FIPS198
- IPsec protocol processor
- High performance IPSEC Engine
- Random Number generator compliant with ANSI X9.31
- Public Key Crypto Engine supports high performance IKE and RSA computation
- Script Processor

CP6

- Dual content processors
- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321 and FIPS180
- HMAC in accordance with RFC2104/2403/2404 and FIPS198
- IPsec protocol processor
- High performance IPsec engine
- Random Number generator compliance with ANSI X9.31
- Key exchange processor for high performance IKE and RSA computation
- Script Processor
- SSL/TLS protocol processor for SSL content scanning and SSL acceleration

CP8

- Over 10Gbps throughput IPS content processor for packet content matching with signatures
- High performance VPN bulk data engine
 - IPSEC and SSL/TLS protocol processor
 - DES/3DES/AES in accordance with FIPS46-3/FIPS81/FIPS197
 - ARC4 in compliance with RC4
 - MD5/SHA-1/SHA256 with RFC1321 and FIPS180
 - HMAC in accordance with RFC2104/2403/2404 and FIPS198

- Key Exchange Processor support high performance IKE and RSA computation
 - Public key exponentiation engine with hardware CRT support
 - Primarily checking for RSA key generation
 - Handshake accelerator with automatic key material generation
 - Random Number generator compliance with ANSI X9.31
 - Sub public key engine (PKCE) to support up to 4094 bit operation directly
- Message authentication module offers high performance cryptographic engine for calculating SHA256/SHA1/MD5 of data up to 4G bytes (used by any application like WAN opt.)
- PCI express Gen 2 four lanes interface
- Cascade Interface for chip expansion

Determining the content processor in your FortiGate unit

Use the `get hardware status` CLI command to determine which content processor your FortiGate unit contains. The output looks like this:

```
# get hardware status
Model name: Fortigate-620B
ASIC version: CP6
ASIC SRAM: 64M
CPU: Intel(R) Core(TM)2 Duo CPU      E4300   @ 1.80GHz
RAM: 2020 MB
Compact Flash: 493 MB /dev/sda
Hard disk: 76618 MB /dev/sdb
USB Flash: not available
Network Card chipset: Broadcom 570x Tigon3 Ethernet Adapter
(rev.0x5784100)
```

The ASIC version line lists the content processor model number.

If you have a CP6 processor, you can view the status of SSL acceleration using the command `get vpn status ssl hardware-acceleration`.

Security processing modules overview

FortiGate Security Processing (SP) modules, such as the ASM-CE4 and ADM-XE2, work at both the interface and system level to increase overall system performance by accelerating some security and networking processing on the interfaces they provide. The SP frees the FortiGate unit's processor for other tasks by offloading firewall, application control, and IPS processing, including flow-based antivirus protection. You can configure the SP to favor IPS over firewall processing in hostile high-traffic environments.

The ASM-CE4 and ADM-XE2 are Advanced Mezzanine cards (AMCs) that are the first generation of SP modules. The next generation of SP modules are Fortinet Mezzanine cards (FMCs) found on newer FortiGate models, such as the 3950. FMC modules take advantage of the Integrated Switch Fabric (ISF) backplane, meaning that accelerated performance is available between any two interfaces, not just interfaces on the same FMC module.

Security processor module models

The ADM-XE2 is a dual-width AMC card with two 10 Gb/s interfaces that can be used on FortiGate-3810A and FortiGate-5001A-DW systems.

The ADM-FE8 is a dual-width AMC card with eight 1 Gb/s interfaces that can be used with the FortiGate-3810A.

The ASM-CE4 is a single-width AMC card with four 10/100/1000 Mb/s interfaces that can be used on FortiGate-3016B and FortiGate-3810A units.

Displaying information about security processing modules

You can display information about installed AMC modules using the CLI command

```
diagnose npu spm dos synproxy <sp_id>
```

Variable	Description
<sp_id>	Enter the ID of the security processing device that you want to display information for. The first device is 0, the second is 1, and so on.



Security processing modules are also called network processing units (NPU).


Example

This example shows how to display details about how the module is processing sessions using the syn proxy. This is a partial output of the command:

```
Number of proxied TCP connections : 0
Number of working proxied TCP connections : 0
Number of retired TCP connections : 0
Number of valid TCP connections : 0
Number of attacks, no ACK from client : 0
Number of no SYN-ACK from server : 0
Number of reset by server (service not supported): 0
Number of established session timeout : 0
Client timeout setting : 3 Seconds
Server timeout setting : 3 Seconds
```

Setting switch-mode mapping on the ADM-XD4

The ADM-XD4 SP has four 10 Gb/s ports, but the NP4 processor it contains has only two 10 Gb/s ports. The external ports you use are important to optimize the SP for your application.

Figure 368: ADM-XD4 mapping mode

Ports 1 and 3 share one NP4 processor and ports 2 and 4 share the other. Performance ports sharing the same NP4 processor is far better than when forcing network data to move between NP4 processors by using one port from each, for example ports 1 and 2 or ports 3 and 4.

Configuring overall security priorities

You can set the priority for security processing using the CLI:

```
config system global
  set optimize {antivirus | throughput | session}
end
```

antivirus - Allow all CPU cores to process traffic – typically used with proxy style services (AntiX, content filtering)

throughput - Prevents code synchronisation delays from impacting raw throughput.

session - Allows distributed session set up across all cores for high session per second environments. This option is available on newer FortiGate models such as the 1240B.

Configuring traffic offloading

Offloading traffic to a network processor requires that the FortiGate unit configuration and the traffic itself is suited to hardware acceleration. There are requirements for path the sessions and the individual packets.

Session fast path requirements

Sessions must be fast path ready. Fast path ready session characteristics are:

- Layer 2 type/length must be 0x0800 (IEEE 802.1q VLAN specification is supported); link aggregation between any network interfaces sharing the same network processor(s) may be used (IEEE 802.3ad specification is supported)
- Layer 3 protocol must be IPv4
- Layer 4 protocol must be UDP, TCP or ICMP
- Layer 3 / Layer 4 header or content modification must not require a session helper (for example, SNAT, DNAT, and TTL reduction are supported, but application layer content modification is not supported)

- FortiGate unit firewall policy must not require antivirus or IPS inspection
- origin must not be local host (the FortiGate unit)
- ingress and egress network interfaces are both attached to the same network processor(s)



If you disable anomaly checks by Intrusion Prevention (IPS), you can still enable hardware accelerated anomaly checks using the `fp-anomaly` field of the `config system interface` CLI command. See “[Configuring pre-IPS anomaly detection](#)” on page 2955.



For session offloading to NP1 network processors, the session must not use an aggregated link or require QoS, including rate limits and bandwidth guarantees. Traffic shaping and link aggregation are not supported.

If a session is not fast path ready, the FortiGate unit will not send the session key to the network processor(s). Without the session key, all session key lookup by a network processor for incoming packets of that session fails, causing all session packets to be sent to the FortiGate unit’s main processing resources, and processed at normal speeds.

If a session is fast path ready, the FortiGate unit will send the session key to the network processor(s). Session key lookup then succeeds for subsequent packets from the known session.

Packet fast path requirements

Packets within the session must then also meet packet requirements.

- Incoming packets must not be fragmented.
- Outgoing packets must not require fragmentation to a size less than 385 bytes. Because of this requirement, the configured MTU (Maximum Transmission Unit) for network processors’ network interfaces must also meet or exceed the network processors’ supported minimum MTU of 385 bytes.

If packet requirements are not met, an individual packet will use FortiGate unit main processing resources, regardless of whether other packets in the session are offloaded to the specialized network processor(s).

In some cases, due to these requirements, a protocol’s session(s) may receive a mixture of offloaded and non-offloaded processing.

For example, FTP uses two connections: a control connection and a data connection. The control connection requires a session helper, and cannot be offloaded, but the data connection does not require a session helper, and can be offloaded. Within the offloadable data session, fragmented packets will not be offloaded, but other packets will be offloaded.

Some traffic types differ from general offloading requirements, but still utilize some of the network processors’ encryption and other capabilities. Exceptions include IPsec traffic and active-active high availability (HA) load balanced traffic.

Fast path connections for specific FortiGate models

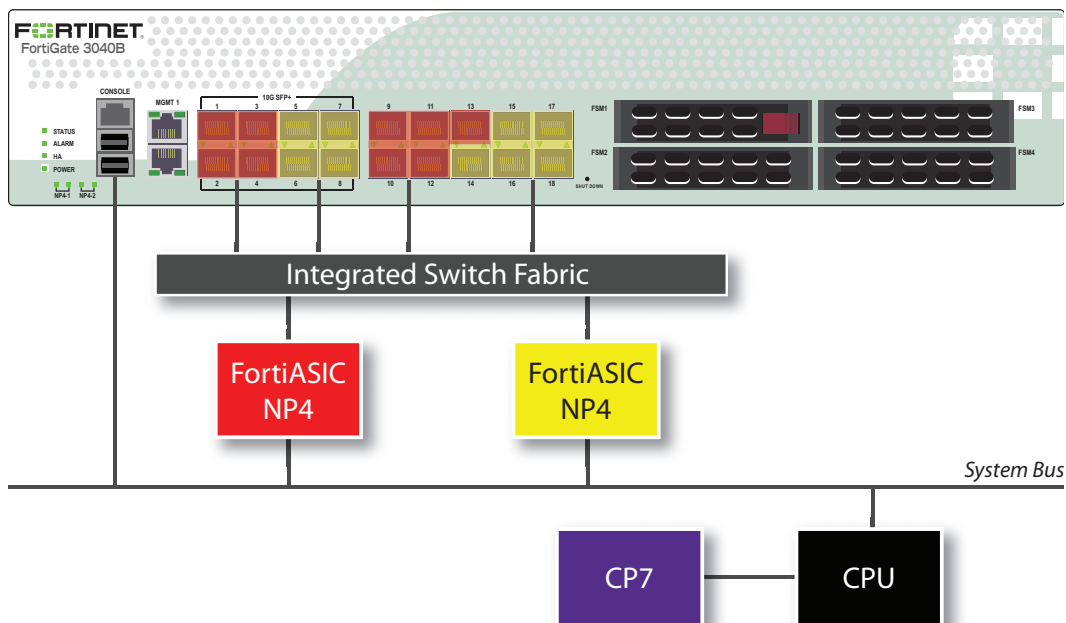
A number of FortiGate models contain multiple NP4 processors that require care when connecting network cables. For accelerated processing, traffic must enter and exit interfaces connected to the same NP4 processor.

FortiGate-3040B

The FortiGate-3040B features two NP4 processors to accelerate network traffic to wire speeds. Traffic between interfaces that use the same processor experience the highest acceleration.

- The 10 Gb interfaces, port1, port2, port3, port4, and the 1 Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10 Gb interfaces, port5, port6, port7, port8, and the 1 Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.

Figure 369: The FortiGate-3040B

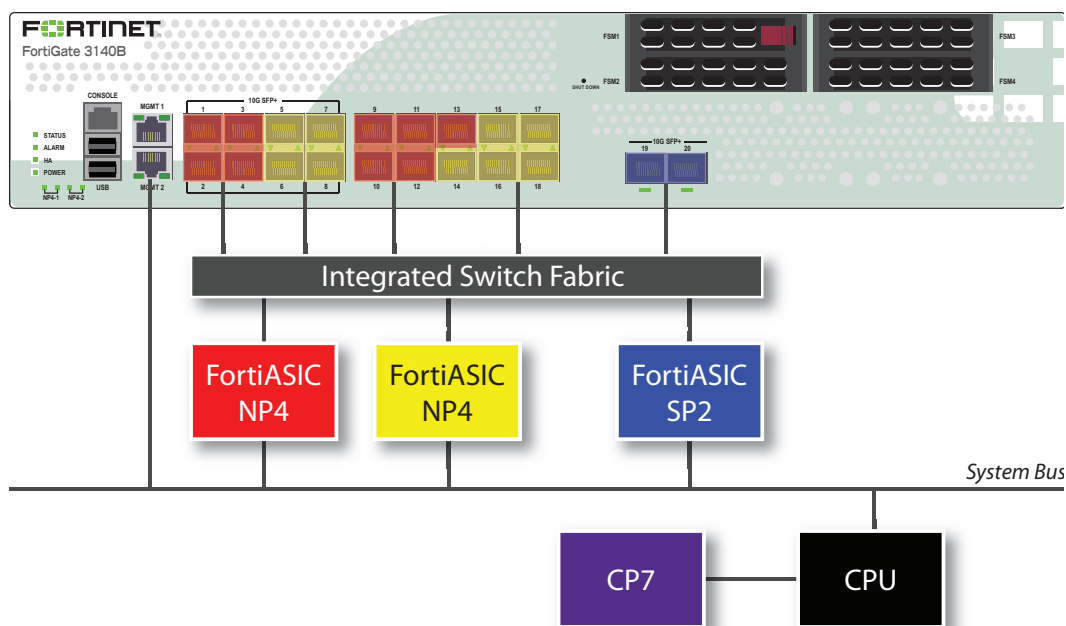


For example, for maximum NP4 acceleration of traffic received on port1, the traffic must exit the FortiGate-3040B unit on port2, port3, or port4 if the bandwidth exceeds 1 Gb. If the traffic bandwidth does not exceed 1 Gb, the traffic can also exit on port9, port10, port11, port12, or port13. Also, for maximum NP4 acceleration of traffic received on port5, the traffic must exit the FortiGate-3040B unit on port6, port7, or port8 if the bandwidth exceeds 1 Gb. If the traffic bandwidth does not exceed 1 Gb, the traffic can also exit on port14, port15, port16, port17, or port18.

FortiGate-3140B

The FortiGate-3140B features two NP4 processors and one SP2 processor to accelerate network traffic to wire speeds. Traffic between interfaces that use the same processor experience the highest acceleration.

- The 10 Gb interfaces, port1, port2, port3, port4, and the 1 Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10 Gb interfaces, port5, port6, port7, port8, and the 1 Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.
- The 10 Gb interfaces, port19 and port20, share connections to the SP2 processor.

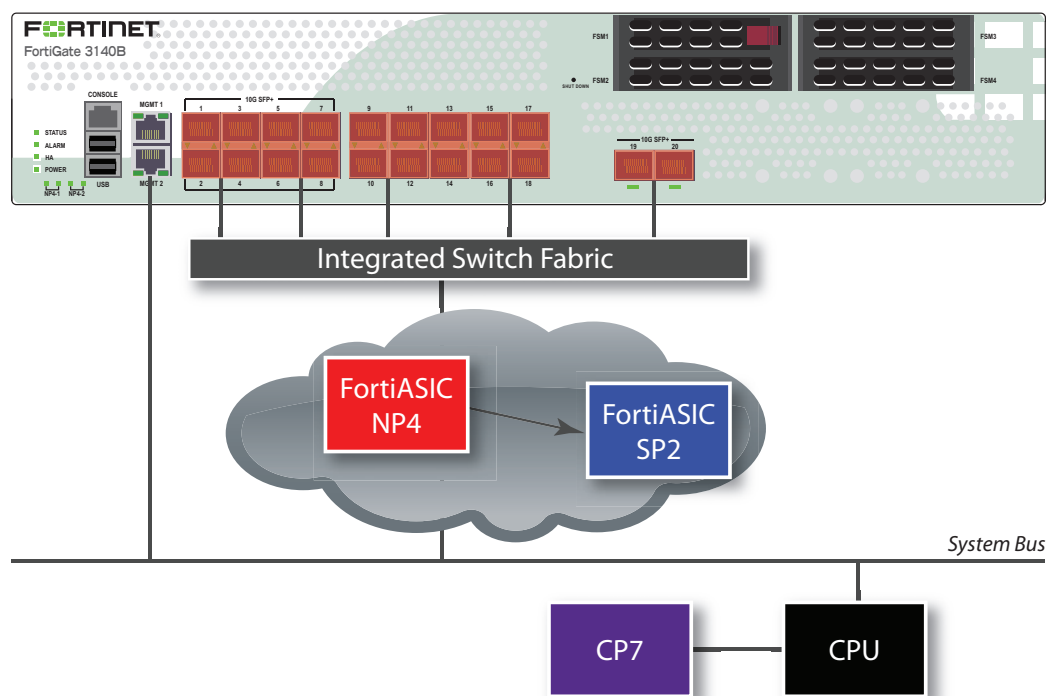
Figure 370: The FortiGate-3140B

For example, for maximum NP4 acceleration of traffic received on port1, the traffic must exit the FortiGate-3140B unit on port2, port3, or port4 if the bandwidth exceeds 1 Gb. If the traffic bandwidth does not exceed 1 Gb, the traffic can also exit on port9, port10, port11, port12, or port13. For maximum NP4 acceleration of traffic received on port5, the traffic must exit the FortiGate-3140B unit on port6, port7, or port8 if the bandwidth exceeds 1 Gb. If the traffic bandwidth does not exceed 1 Gb, the traffic can also exit on port14, port15, port16, port17, or port18. Also, for maximum SP2 acceleration of traffic received on port 19, the traffic must exit the FortiGate-3140B unit on port20.

FortiGate-3140B — load balance mode

The FortiGate-3140B load balance mode allows you increased flexibility in how you use the interfaces on the FortiGate unit. When enabled, traffic between any two interfaces (excluding management and console) is accelerated. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.

Figure 371: The FortiGate-3140B in load balance mode

To enable this feature, issue this CLI command.

```
config system global
  set sp-load-balance enable
end
```

The FortiGate unit will then restart.

To return to the default mode, issue this CLI command.

```
config system global
  set sp-load-balance disable
end
```

FortiGate-3950B and FortiGate-3951B

The FortiGate-3950B features one NP4 processor to accelerate network traffic to wire speeds.

- The 1 Gb SPF interfaces, port1, port2, port3, port4, and the 10 Gb SPF+ interfaces, port5, port6, share connections to one NP4 processor.



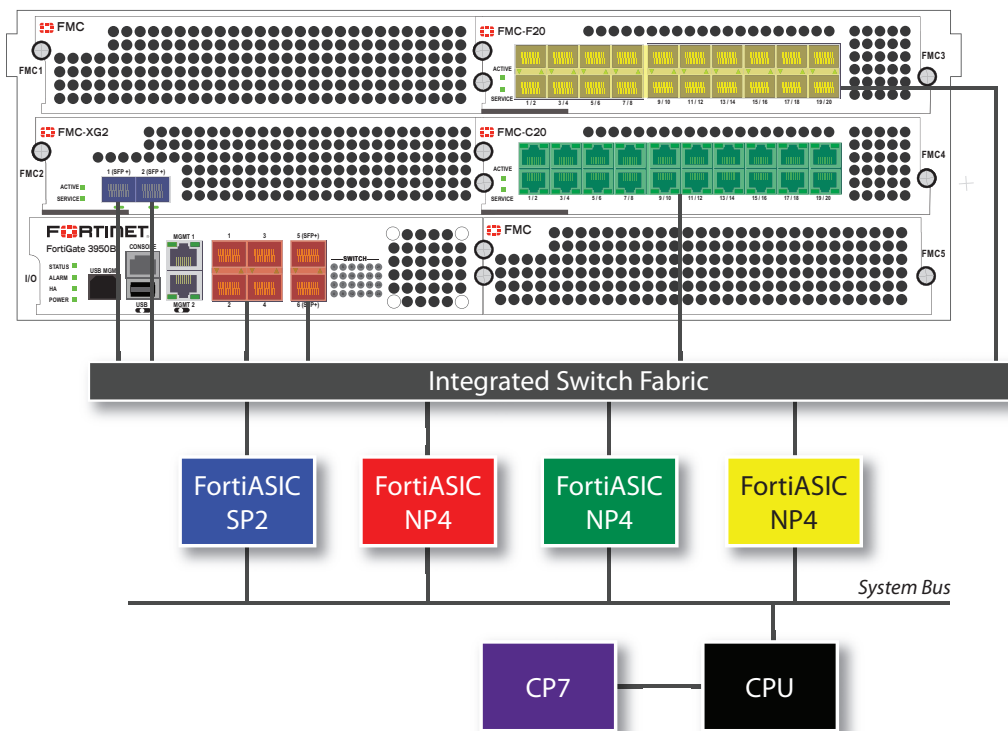
The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

You can add additional FMC interface modules. If they support network traffic acceleration, it is only between interfaces on the module. Figure 372 shows a FortiGate-3950B with three modules installed: an FMC-XG2, an FMC-F20, and an FMC-C20.

- The FMC-XG2 has one NP4 processor and one SP2 processor. The 10 Gb SPF+ interfaces, port1 and port2, share connections to both processors.

- The FMC-F20 has one NP4 processor and the twenty 1 Gb SPF interfaces, port1 through port20, share connections to the NP4 processor.
- The FMC-C20 has one NP4 processor and the twenty 10/100/1000 interfaces, port1 through port20, share connections to the NP4 processor.

Figure 372: The FortiGate-3950B with an FMC-XG2, an FMC-F20, and an FMC-C20



The processor in the FortiGate-3950B and on each FMC module can accelerate only the network traffic entering and leaving its own interfaces. For example, for maximum NP4 acceleration of traffic that exceeds 1 Gb bandwidth must enter and leave the FortiGate-3950B on its own port5 and port6, or the FMC-XG2 port1 and port2. If the traffic bandwidth does not exceed 1 Gb, the traffic can enter and exit using port1 through port4, or port1 through port20 on either the FMC-F20 or the FMC-C20.

Also, for maximum SP2 acceleration of traffic received on port1 of the FMC-XG2, the traffic must exit port2 of the FMC-XG2.

Traffic can enter an interface on one module and leaving an interface on another module, but it will not take advantage of any network or security acceleration.

FortiGate-3950B and FortiGate-3951B — load balance mode

Adding one or more FMC-XG2 modules to your FortiGate-3950B allows you to enable load balance mode. This feature allows you increased flexibility in how you use the interfaces on the FortiGate unit.

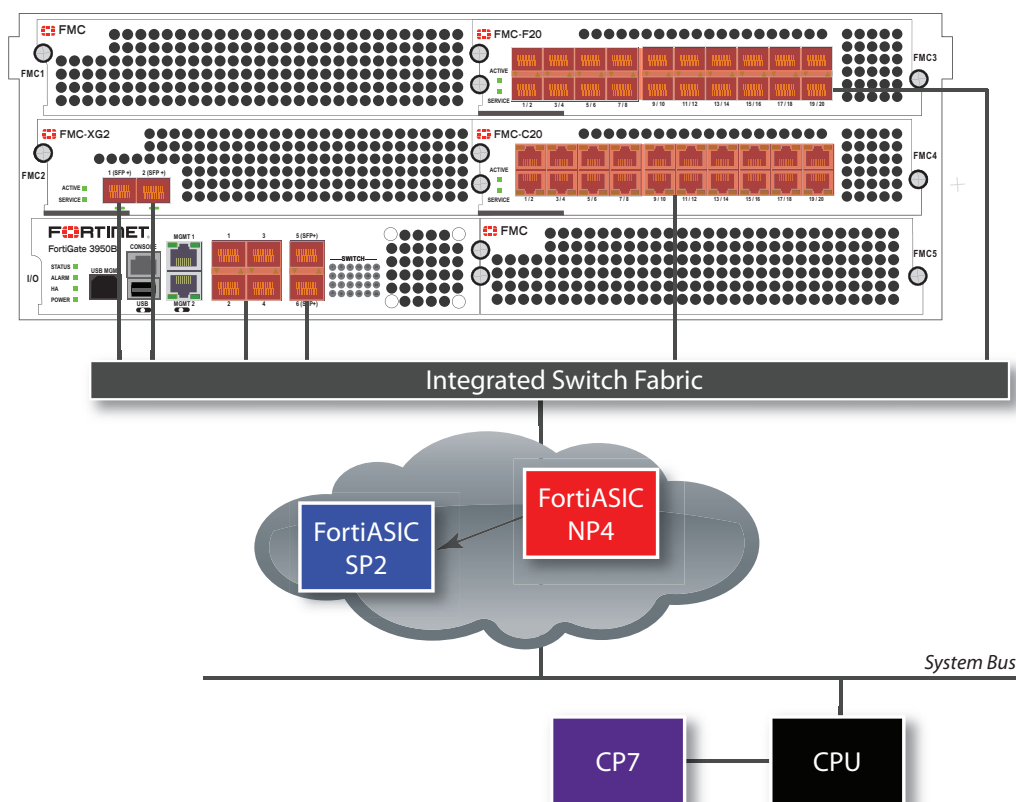


The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

When enabled, traffic between any two interfaces (excluding management and console) is accelerated whether they are the six interfaces on the FortiGate-3950B itself, or on any installed FMC modules. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.

Figure 373: The FortiGate-3950B in load balance mode



To enable this feature, issue this CLI command.

```
config system global
    set sp-load-balance enable
end
```

The FortiGate unit will then restart.

To return to the default mode, issue this CLI command.

```
config system global
    set sp-load-balance disable
end
```

Session offloading in HA active-active configuration

Fortinet's specialized network processors can improve network performance in active-active (load balancing) high availability (HA) configurations, even though traffic deviates from general offloading patterns, involving more than one network processor, each in a separate FortiGate unit. No additional offloading requirements apply.

Once the primary FortiGate unit's main processing resources send a session key to its network processor(s), network processor(s) on the primary unit can redirect any subsequent session traffic to other cluster members, reducing traffic redirection load on the primary unit's main processing resources.

As subordinate units receive redirected traffic, each network processor in the cluster assesses and processes session offloading independently from the primary unit. Session key states of each network processor are not part of synchronization traffic between HA members.

Configuring traffic shaping offloading

Accelerated Traffic shaping is supported with some limitations on NP2 and NP4 interfaces. Security processor modules do not perform any traffic shaping. Any traffic on which traffic shaping is enabled is handled by the FortiGate unit's main processing resources.

For traffic shaping and QoS through accelerated NP2 and NP4 ports,

- Accelerated ports support policy-based traffic policing. However, fast path traffic and traffic handled by the FortiGate CPU (slow path) are controlled separately, which means the policy setting on fast path does not consider the traffic on the slow path.
- The port based traffic policing as defined by the inbandwidth and outbandwidth CLI commands is not supported on the NP2 processor or the NP4 processor.
- NP2 and NP4 ports support DSCP configurations.
- Per-IP traffic shaping is not supported with NP2 interfaces due to hardware limitations.
- QoS in general is not supported by NP2 and NP4.

You can also use the traffic shaping features of the FortiGate unit's main processing resources by disabling the acceleration features of the NP2 and NP4 ports. See [“Disabling offloading” on page 2948](#).

Network processing unit (`npu`) settings configure offloading for traffic shaping. Configured behavior applies to all network processors contained by the FortiGate unit itself or any installed AMC modules.

```
config system npu
    set traffic-shaping-mode {bidirection | unidirection}
```

end

Variables	Description	Default
traffic-shaping-mode {bidirection unidirection}	<p>Select the offloaded traffic shaping bandwidth calculation method.</p> <ul style="list-style-type: none"> unidirection: The bandwidth limit applies per direction. For example, a unidirectional limit of 10 KBps would result in an overall limit of 20 KBps — 10 KBps per direction. bidirection: The bandwidth limit applies to both directions overall. For example, a bidirectional limit of 10 KBps would result in an overall limit of 10 KBps — 5 KBps per direction. <p>This option applies only if the FortiGate unit itself or any installed AMC modules contain a network processor that supports offloading of traffic shaping.</p>	Varies by model.

Example

You could configure the traffic shaping limit to be applied as a bidirectional total limit during hardware accelerated sessions.

```
config system npu
  set traffic-shaping-mode bidirection
end

config system interface
  edit <interface_name>
    set outbandwidth <real outbandwidth>
  end
```

Checking that traffic is offloaded

You can determine whether traffic is offloaded by using the CLI command:

```
diagnose sys session list
```

The output provides detailed information about each session. Look for the “state=” line. If “npu npr” appears on that line, the session was offloaded to a network processor.

You can also use the diagnose command:

```
diagnose sniffer packet <interface_name>
```

Disabling offloading

If you want to completely disable offloading for test purposes or other reasons, you can do so by security policy.

```
config firewall policy
  edit <policy_id_int>
    set auto-asic-offload disable
  end
```

Multicast offloading / acceleration

Only security processor modules such as the CE4, CE8, or XE2 can offload multicast traffic from the FortiGate unit's CPU-based resources. To make use of this capability, the multicast traffic must enter and exit the FortiGate unit on network interfaces on the same SPM card. Also, the session fast path requirements must be met. These are the same requirements that apply to unicast traffic. See [“Session fast path requirements” on page 2940](#).

Like any other traffic between interfaces, multicast traffic requires a firewall policy, in this case a multicast firewall policy. These policies, for example, permit multicast traffic between the first port and each of the other ports on an ASM-CE4 card:

```
config firewall multicast-policy
  edit 1
    set srcintf amc-sw1/1
    set dstintf amc-sw11/2
    set action accept
  next
  edit 2
    set srcintf amc-sw1/1
    set dstintf amc-sw11/3
    set action accept
  next
  edit 3
    set srcintf amc-sw1/1
    set dstintf amc-sw11/4
    set action accept
end
```

Note that simple forwarding of multicast packets is not accelerated. Also, if the FortiGate unit or VDOM is in Transparent mode, multicast is not accelerated.

Use `diagnose ip multicast npu-session list` to verify the NPU session is established

Configuring IPsec VPN offloading

Fortinet's specialized network processors contain features to improve IPsec tunnel performance. For example, network processors can encrypt and decrypt packets, reducing cryptographic load on the FortiGate unit's main processing resources.

IPsec offloading requirements

Requirements for hardware accelerated IPsec encryption or decryption are a modification of general offloading requirements. Differing characteristics are:

- origin can be local host (the FortiGate unit)
- in Phase I configuration, Local Gateway IP must be specified as an IP address of a network interface for a port attached to a network processor
- SA must have been received by the network processor

- in Phase II configuration:
 - encryption algorithm must be DES, 3DES, AES-128, AES-192, AES-256, or null
 - authentication must be MD5, SHA1, or null
 - if encryption is null, authentication must not also be null
 - if replay detection is enabled, `enc-offload-antireplay` must also be `enable` in the CLI



If replay detection is enabled in the Phase II configuration, you can enable or disable IPsec encryption and decryption offloading from the CLI. Performance varies by those CLI options and the percentage of packets requiring encryption or decryption. For details, see [“Configuring VPN encryption/decryption offloading” on page 2950](#).



For session offloading to NP1 network processors, in Phase II configuration, the encryption algorithm must be 3DES and authentication must be MD5. Other encryption and authentication algorithms are not supported.

To apply hardware accelerated encryption and decryption, the FortiGate unit's main processing resources must first perform Phase I negotiations to establish the security association (SA). The SA includes cryptographic processing instructions required by the network processor, such as which encryption algorithms must be applied to the tunnel. After ISAKMP negotiations, the FortiGate unit's main processing resources send the SA to the network processor, enabling the network processor to apply the negotiated hardware accelerated encryption or decryption to tunnel traffic.

Possible accelerated cryptographic paths are:

- IPsec decryption offload
 - Ingress ESP packet > Offloaded decryption > Decrypted packet egress (fast path)
 - Ingress ESP packet > Offloaded decryption > Decrypted packet to FortiGate unit's main processing resources
- IPsec encryption offload
 - Ingress packet > Offloaded encryption > Encrypted (ESP) packet egress (fast path)
 - Packet from FortiGate unit's main processing resources > Offloaded encryption > Encrypted (ESP) packet egress

Configuring HMAC check offloading

Hash-based Message Authentication Code (HMAC) checks can be offloaded to network processors. To enable HMAC check offloading, enter

```
configure system global
  set ipsec-hmac-offload (enable|disable)
end
```

Configuring VPN encryption/decryption offloading

Network processing unit (npu) settings configure offloading behavior for IPsec VPN. Configured behavior applies to all network processors contained by the FortiGate unit itself or any installed AMC modules.

```
config system npu
  set enc-offload-antireplay {enable | disable}
  set dec-offload-antireplay {enable | disable}
  set offload-ipsec-host {enable | disable}
```

end

Variables	Description	Default
enc-offload-antireplay {enable disable}	Enable or disable offloading of IPsec encryption. This option is used only when replay detection is enabled in Phase II configuration. If replay detection is disabled, encryption is always offloaded.	disable
dec-offload-antireplay {enable disable}	Enable or disable offloading of IPsec decryption. This option is used only when replay detection is enabled in Phase II configuration. If replay detection is disabled, decryption is always offloaded.	enable
offload-ipsec-host {enable disable}	Enable or disable offloading of IPsec encryption of traffic from local host (FortiGate unit). Note: For this option to take effect, the FortiGate unit must have previously sent the security association (SA) to the network processor. For details on SA offloading, see “Configuring IPsec VPN offloading” on page 2949 .	disable

Example

You could configure the offloading of encryption and decryption for an IPsec SA that was sent to the network processor.


```
config system npu
  set enc-offload-antireplay enable
  set dec-offload-antireplay enable
  set offload-ipsec-host enable
end
```

Examples of ASM-FB4 accelerated VPNs

This section contains example IPsec configurations whose IPsec encryption and decryption processing is hardware accelerated by FortiGate-ASM-FB4 modules. [Figure 374](#) illustrates the example network topology. [Table 152](#) lists the example network interfaces and IP addresses.



Hardware accelerated IPsec does not require both tunnel endpoints to have the same network processor model. However, if hardware is not symmetrical, the packet forwarding rate is limited by the slower side.

Figure 374: Example network topology for offloaded IPsec processing**Table 152: Example ports and IP addresses for offloaded IPsec processing**

	FortiGate_1		FortiGate_2	
	Port	IP	Port	IP
IPsec tunnel	FortiGate-ASM-FB4 port 2	3.3.3.1/24	FortiGate-ASM-FB4 port 2	3.3.3.2/24
Protected network	FortiGate-ASM-FB4 port 1	1.1.1.0/24	FortiGate-ASM-FB4 port 1	2.2.2.0/24

Tunnel mode IPsec VPN example

The following steps create a hardware accelerated tunnel mode IPsec tunnel between two FortiGate units, each containing a FortiGate-ASM-FB4 module.

To configure hardware accelerated tunnel mode IPsec

- 1 On FortiGate_1, go to *VPN > IPsec*.
- 2 Configure Phase I.
For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.
Select Advanced. In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2's FortiGate-ASM-FB4 module port 2.
- 3 Configure Phase II.
If you enable the check box "Enable replay detection," set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see ["Configuring VPN encryption/decryption offloading" on page 2950](#).
- 4 Go to *Firewall > Policy*.
- 5 Configure one policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.
- 6 Go to *Router > Static*.

- 7 Configure a static route to route traffic destined for FortiGate_2's protected network to VPN IP address of FortiGate_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "AMC-SW1/2"
    set dst 2.2.2.0 255.255.255.0
    set gateway 3.3.3.2
  end
```

- 8 On FortiGate_2, go to *VPN > IPsec*.

- 9 Configure Phase I.

For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.

Select Advanced. In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1's FortiGate-ASM-FB4 module port 2.

- 10 Configure Phase II.

If you enable the check box "Enable replay detection," set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see ["Configuring VPN encryption/decryption offloading" on page 2950](#)

- 11 Go to *Firewall > Policy*.

- 12 Configure one policy to apply the Phase 1 IPsec tunnel you configured in step 9 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.

- 13 Go to *Router > Static*.

- 14 Configure a static route to route traffic destined for FortiGate_1's protected network to VPN IP address of FortiGate_1's VPN gateway, 3.3.3.1, through the FortiGate-ASM-FB4 module's port 2 (device).

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "AMC-SW1/2"
    set dst 1.1.1.0 255.255.255.0
    set gateway 3.3.3.1
  end
```

- 15 Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > IPSEC > Monitor*.

Interface mode IPsec VPN example

The following steps create a hardware accelerated interface mode IPsec tunnel between two FortiGate units, each containing a FortiGate-ASM-FB4 module.

To configure hardware accelerated interface mode IPsec

1 On FortiGate_1, go to *VPN > IPsec*.

2 Configure Phase I.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Select Advanced.
- Enable the check box “Enable IPsec Interface Mode.”
- In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2’s FortiGate-ASM-FB4 module port 2.

3 Configure Phase II.

If you enable the check box “Enable replay detection,” set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see [“Configuring VPN encryption/decryption offloading” on page 2950](#)

4 Go to *Firewall > Policy*.

5 Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.

6 Go to *Router > Static*.

7 Configure a static route to route traffic destined for FortiGate_2’s protected network to the Phase 1 IPsec device, `FGT_1_IPsec`.

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "FGT_1_IPsec"
    set dst 2.2.2.0 255.255.255.0
  end
```

8 On FortiGate_2, go to *VPN > IPsec*.

9 Configure Phase I.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Enable the check box “Enable IPsec Interface Mode.”
- In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1’s FortiGate-ASM-FB4 module port 2.

10 Configure Phase II.

If you enable the check box “Enable replay detection,” set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see [“Configuring VPN encryption/decryption offloading” on page 2950](#)

11 Go to *Firewall > Policy*.

12 Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 9 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.

13 Go to *Router > Static*.

- 14** Configure a static route to route traffic destined for FortiGate_1's protected network to the Phase 1 IPsec device, FGT_2_IPsec.

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "FGT_2_IPsec"
    set dst 1.1.1.0 255.255.255.0
  next
end
```

- 15** Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > IPSEC > Monitor*.

Configuring IPS offloading

Security modules offload IPS. Requirements are:

- Source port is on CE4
- Destination port is on the same CE4
- UTM configuration must enable only IPS, not AV or content archive.
- Packet protocol is ICMP, UDP or TCP.

IPS offloading functions with policy-based IPS and interface-based IPS, as well as sniffer-mode.

Configuring pre-IPS anomaly detection

Network interfaces associated with a port attached to a network processor can be configured to use hardware acceleration to drop or allow certain anomaly types, separately from and in advance of any anomaly checks specified by Intrusion Prevention (IPS). Configured behavior applies separately to each of these network interfaces.

```
config system interface
  edit <name_str>
    set fp-anomaly
      {drop_icmpland | pass_icmpland}
      {drop_ipland | pass_ipland}
      {drop_iplsrr | pass_iplsrr}
      {drop_iprr | pass_iprr}
      {drop_ipsecurity | pass_ipsecurity}
      {drop_ipssrr | pass_ipssrr}
      {drop_ipstream | pass_ipstream}
      {drop_iptimestamp | pass_iptimestamp}
      {drop_ipunknown_option | pass_ipunknown_option}
      {drop_unknown_prot | pass_ipunknown_prot}
      {drop_tcpland | pass_tcpland}
      {drop_udpland | pass_udpland}
      {drop_winnuke | pass_winnuke}
  end
```

where:

icmpland	ICMP land
ipland	IP land
iplsrr	IP with loose source record route
iprr	IP with record route option
ipsecurity	IP with security option
ipssrr	IP with strict source record route option
ipstream	IP with stream option
iptimestamp	IP with timestamp option
ipunknown_option	IP with unknown option
ipunknown_prot	IP with unknown protocol
tcpland	TCP land
udpland	UDP land
winnuke	TCP WinNuke

Example

You might configure a FortiGate-ASM-FB4 module to drop packets with TCP WinNuke or unknown IP protocol anomalies, but to pass packets with an IP time stamp, using hardware acceleration provided by the network processor.

```
config system interface
  edit AMC-SW1/1
    set fp-anomaly drop_winnuke drop_ipunknown_prot
    pass_iptimestamp
  end
```

Configuring policy-based IPS on SP modules

In the firewall policy, enable UTM, then enable IPS and select the desired IPS profile.

Configuring interface-based IPS on SP modules

- 1 Define the IPS sensor. This step is the same with current policy-based IPS. For system predefined sensor, this step can be ignored.
- 2 Define on which interface IPS should be enabled and what sensor you want to use to scan traffic. Both physical interface and VLAN interface are valid interface choices.

The followed is an example to enable IPS sensor “all_default” on physical port AMC-SW1/2.

```
config ips interface
  edit AMC-SW1/2
    set ips-sensor all_default
  end
```

This command will enable IPS on all traffic ingress and egress through AMC-SW1/2.

Do not enable policy-based IPS when either the source or destination port has interface IPS enabled. Doing so provides no additional security and results in reduced performance.


Examples

Hardware accelerated IPsec processing, involving either partial or full offloading, can be achieved in either tunnel or interface mode IPsec configurations.

To achieve offloading for both encryption and decryption:

- In Phase I configuration’s Advanced section, Local Gateway IP must be specified as an IP address of a network interface associated with a port attached to a network processor. (In other words, if Phase 1’s Local Gateway IP is Main Interface IP, or is specified as an IP address that is not associated with a network interface associated with a port attached to a network processor, IPsec network processing is not offloaded.)
- In Phase II configuration’s P2 Proposal section, if the checkbox “Enable replay detection” is enabled, `enc-offload-antireplay` and `dec-offload-antireplay` must be set to enable in the CLI.
- `offload-ipsec-host` must be set to enable in the CLI.

This section contains example IPsec configurations whose IPsec encryption and decryption processing is hardware accelerated by FortiGate-ASM-FB4 modules. [Figure 374](#) illustrates the example network topology. [Table 152](#) lists the example network interfaces and IP addresses.



Hardware accelerated IPsec does not require both tunnel endpoints to have the same network processor model. However, if hardware is not symmetrical, the packet forwarding rate is limited by the slower side.

Figure 375: Example network topology for offloaded IPsec processing




Table 153: Example ports and IP addresses for offloaded IPsec processing

	FortiGate_1		FortiGate_2	
	Port	IP	Port	IP
IPsec tunnel	FortiGate-ASM-FB4 port 2	3.3.3.1/24	FortiGate-ASM-FB4 port 2	3.3.3.2/24

Table 153: Example ports and IP addresses for offloaded IPsec processing

Protected network	FortiGate-ASM-FB4 port 1	1.1.1.0/24	FortiGate-ASM-FB4 port 1	2.2.2.0/24
--------------------------	--------------------------	------------	--------------------------	------------

This section includes the following topics:

- [Tunnel mode IPsec VPN example](#)
- [Configuring traffic offloading](#)

Accelerated tunnel mode IPsec

The following steps create a hardware accelerated tunnel mode IPsec tunnel between two FortiGate units, each containing a FortiGate-ASM-FB4 module.

To configure hardware accelerated tunnel mode IPsec

- 1 On FortiGate_1, go to *VPN > IPsec*.
- 2 Configure Phase I.
For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.
Select Advanced. In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2's FortiGate-ASM-FB4 module port 2.
- 3 Configure Phase II.
If you enable the checkbox "Enable replay detection," set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see "[Configuring VPN encryption/decryption offloading](#)" on page 2950
- 4 Go to *Firewall > Policy*.
- 5 Configure one policy to apply the Phase 1 IPsec tunnel you configured in step 2 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.
- 6 Go to *Router > Static*.
- 7 Configure a static route to route traffic destined for FortiGate_2's protected network to VPN IP address of FortiGate_2's VPN gateway, 3.3.3.2, through the FortiGate-ASM-FB4 module's port 2 (device).

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "AMC-SW1/2"
    set dst 2.2.2.0 255.255.255.0
    set gateway 3.3.3.2
  end
```

- 8 On FortiGate_2, go to *VPN > IPsec*.
- 9 Configure Phase I.
For tunnel mode IPsec and for hardware acceleration, specifying the Local Gateway IP is required.
Select Advanced. In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1's FortiGate-ASM-FB4 module port 2.

10 Configure Phase II.

If you enable the checkbox “Enable replay detection,” set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see [“Configuring VPN encryption/decryption offloading” on page 2950](#)

11 Go to *Firewall > Policy*.**12** Configure one policy to apply the Phase 1 IPsec tunnel you configured in step 9 to traffic between FortiGate-ASM-FB4 module ports 1 and 2.**13** Go to *Router > Static*.**14** Configure a static route to route traffic destined for FortiGate_1’s protected network to VPN IP address of FortiGate_1’s VPN gateway, 3.3.3.1, through the FortiGate-ASM-FB4 module’s port 2 (device).

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "AMC-SW1/2"
    set dst 1.1.1.0 255.255.255.0
    set gateway 3.3.3.1
  end
```

15 Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > IPSEC > Monitor*.

Accelerated interface mode IPsec

The following steps create a hardware accelerated interface mode IPsec tunnel between two FortiGate units, each containing a FortiGate-ASM-FB4 module.

To configure hardware accelerated interface mode IPsec

1 On FortiGate_1, go to *VPN > IPsec*.**2** Configure Phase I.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Select Advanced.
- Enable the checkbox “Enable IPsec Interface Mode.”
- In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.2, which is the IP address of FortiGate_2’s FortiGate-ASM-FB4 module port 2.

3 Configure Phase II.

If you enable the checkbox “Enable replay detection,” set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see [“Configuring VPN encryption/decryption offloading” on page 2950](#)

4 Go to *Firewall > Policy*.**5** Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 2 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.**6** Go to *Router > Static*.

- 7** Configure a static route to route traffic destined for FortiGate_2's protected network to the Phase 1 IPsec device, FGT_1_IPsec.

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "FGT_1_IPsec"
    set dst 2.2.2.0 255.255.255.0
  end
```

- 8** On FortiGate_2, go to *VPN > IPsec*.

- 9** Configure Phase I.

For interface mode IPsec and for hardware acceleration, the following settings are required.

- Enable the checkbox "Enable IPsec Interface Mode."
- In the Local Gateway IP section, select Specify and type the VPN IP address 3.3.3.1, which is the IP address of FortiGate_1's FortiGate-ASM-FB4 module port 2.

- 10** Configure Phase II.

If you enable the checkbox "Enable replay detection," set `enc-offload-antireplay` to `enable` in the CLI. For details on encryption and decryption offloading options available in the CLI, see ["Configuring VPN encryption/decryption offloading" on page 2950](#)

- 11** Go to *Firewall > Policy*.

- 12** Configure two policies (one for each direction) to apply the Phase 1 IPsec configuration you configured in step 9 to traffic leaving from or arriving on FortiGate-ASM-FB4 module port 1.

- 13** Go to *Router > Static*.

- 14** Configure a static route to route traffic destined for FortiGate_1's protected network to the Phase 1 IPsec device, FGT_2_IPsec.

You can also configure the static route using the following CLI commands:

```
config router static
  edit 2
    set device "FGT_2_IPsec"
    set dst 1.1.1.0 255.255.255.0
  next
end
```

- 15** Activate the IPsec tunnel by sending traffic between the two protected networks.

To verify tunnel activation, go to *VPN > IPSEC > Monitor*.



Configuring RAID

This section describes how to configure RAID on a FortiGate unit with multiple disk support. RAID arrays can provide faster disk access, redundancy in case of partial failure, or both depending on the RAID level you select.

The following topics are included in this section:

- [RAID levels](#)
- [Configuring a RAID array](#)
- [Checking the status of a RAID array](#)
- [Rebuilding a RAID array](#)

RAID levels

Some FortiGate models have two or more hard disks configured in a RAID array to store log messages locally on the FortiGate unit. A RAID array can provide faster disk access, redundancy in case of partial failure, or both depending on the RAID level you select.

When changing the RAID level, the available levels depend on the number of working disks that are actually present in the unit. For example, RAID-5 is not available on units with fewer than three disks. When a disk fails, becomes corrupt, or is removed you must rebuild the RAID array. For more information, see [“Rebuilding a RAID array” on page 2964](#).

If the FortiGate unit has only one disk installed, the RAID monitor widget will not be displayed as it is not possible to configure a RAID array with only one disk.

Available RAID levels include:

- [RAID-0](#)
- [RAID-1](#)
- [RAID-5](#)

RAID-0

A RAID-0 array is also referred to as striping. The FortiGate unit writes information evenly across all hard disks. The total space available is that of all the disks in the RAID array. There is no redundancy available. If any single drive fails, the data on the array is lost and cannot be recovered. Because of this lack of redundancy, a RAID-0 array will never report a degraded condition. This RAID level is beneficial because it provides better performance, since the FortiGate unit can distribute disk writing across multiple disks.

For example if your FortiGate unit has three disks each with a 1 terabyte (TB) capacity, your RAID-0 array will have a 3TB capacity.

RAID-1

A RAID-1 array is also referred to as mirroring. The FortiGate unit writes information to one hard disk, and writes a copy (a mirror image) of all information to all other hard disks. The total disk space available is that of only one hard disk, as the others are solely used for mirroring. This provides redundant data storage. Should any of the hard disks fail, there one or more backup hard disks available. For example, if one disk fails, the unit can still access three other hard disks and continue functioning.

RAID-5

A RAID-5 array employs striping with a parity check. Similar to RAID-0, the FortiGate unit writes information evenly across all drives but additional parity blocks are written on the same stripes. The parity block is staggered for each stripe. RAID-5 requires three or more hard disks. The total disk space is the total number of disks in the array, minus the capacity of one disk for parity storage. For example, with four hard disks, the total capacity available is the capacity of three hard disks. RAID-5 performance is typically better with reading than with writing, although performance is degraded when one disk has failed. With RAID-5, one disk can fail without the loss of data. If a drive fails, it can be replaced and the FortiGate unit will restore the data on the new disk by using reference information from the parity volume.

Configuring a RAID array



Do not remove a disk while the RAID array is synchronizing — you may lose stored information. This will also degrade the array, requiring a rebuild.

A RAID array provides no redundancy in a degraded state. Any disk failure while the RAID is in a degraded state will cause data loss.

When switching RAID levels, you may see the message “RAID status is OK and RAID is doing background synchronization.” Synchronization of the disks in the array will take considerable time — it will take longer for larger arrays and for disks with more storage capacity.

To configure a RAID array

- 1 Go to *System > Dashboard > Status* where the RAID Monitor widget is located, and then select *Configure* in the widget title bar area.



Changing the RAID level will erase any stored log information on the array, and reboot the FortiGate unit. The unit will remain offline while it reconfigures the RAID array. When it reboots, the array will need to synchronize before being fully operational.

- 2 Confirm that the FortiGate unit recognizes the installed hard disks. Each slot in which you have installed a hard disk displays a green check mark for *Member* and OK for *Status*. The Capacity figure for each hard disk simply lists its size.

The available space on the array will depend on the size of the member drives, but it may not be equal to the total size of the member drives. Further, the hard disks in a RAID array need to have the same capacity. If you use disks with differing capacities, the member hard disks will be treated as if they all have the capacity of the smallest drive in the array. The RAID level determines how the size of the RAID array relates to the size of the member hard disks. For example, an array of three 1TB hard disks will result in 3TB of usable space with RAID-0, 2TB of usable space with RAID-5, and 1TB of space with RAID-1.

3 Select the RAID level.

RAID-0 (Striping)	Better performance than a single disk, but no redundancy. If either disk fails, all data is lost.
RAID-1 (mirroring)	Performance comparable to a single disk, and data is protected by redundancy. One disk can fail with no data loss.
RAID-5 (striping with parity)	Performance is mixed with disk writes slower than a single disk and disk reads faster. Data is protected by redundancy. One disk can fail with no data loss.

For more information on RAID levels, see [“RAID levels” on page 2961](#).

4 Select *Apply*.

The FortiGate unit reboots and reconfigures the RAID array. You may log in again when it is complete.

Checking the status of a RAID array

Once a RAID array is configured, it requires no regular maintenance. Attention is required only when a member hard drive fails. The RAID widget reports the RAID array condition and disk space utilization.

To check the status of a RAID array

- 1 Go to *System > Dashboard > Status* where the RAID Monitor widget is located.
- 2 The widget shows three pieces of status information about the RAID array.

Array Status	<p>Displays the RAID level and status of the RAID array. The hard disks installed in the FortiGate unit are also displayed, with indicators to show which are part of the RAID array and the status of each disk.</p> <p>The status can be:</p> <p>OK — standard status, everything is normal</p> <p>OK (Background-Synchronizing) (%) — synchronizing the disks after changing RAID level, Synchronizing progress bar shows percent complete</p> <p>Degraded — One or more of the disks in the array has failed, been removed, or is not working properly. A warning is displayed about the lack of redundancy in this state. Also, a degraded array is slower than a healthy array. Select <i>Rebuild RAID</i> to fix the array after replacing the defective or missing disk.</p> <p>Degraded (Background-Rebuilding) (%) — The same as degraded, but the RAID array is being rebuilt in the background. The array continues to be in a fragile state until the rebuilding is completed.</p>
---------------------	---

Disk Space Usage	Shows a bar graph of the used space as well as text listing the used space, free space, and total disk space available in the array.
Synchronize Status	Shows that the array is synchronized or reports the synchronization progress, as well as any information about the current synchronization status.

Rebuilding a RAID array

A RAID array has multiple disks with writing to the disks being spread out so that if one disk in the array fails, the array can still provide all the stored information. Some forms of RAID do not provide redundancy, however most do.

When a disk fails, or the RAID array becomes degraded

The Alert Message Console widget, located in *System > Dashboard > Status*, displays any messages about events or activities that need urgent attention, such as a failed hard disk. This widget provides detailed messages that contain the date and time of the event or activity, as well as an explanation about what happened.

Why rebuild a RAID array?

When the RAID array has redundancy and one disk in the array fails, becomes corrupted, or is removed the array becomes degraded. In a degraded state the array can still function, but there are some changes. The two main changes are that there is no longer redundancy and accessing the array takes longer than before.

There is no redundancy because with one disk removed from the array, the information that was stored on that disk can be retrieved using the other disks in the array. However, removing another disk from the array would remove information that has no backup or parity data. This second disk's removal would result in data loss and the array will fail. This delicate state of the RAID array is displayed in the warning message on the dashboard RAID monitor when the status is degraded in the form of a warning.

The array takes longer to access data because instead of the data being retrieved in the format and order it is expected, the array has to jump around to find it and at times recreate the missing data from the parity information. This all takes longer than just the usual straight read operation and will continue until the RAID array has been rebuilt.

The reasons you rebuild a RAID array include:

- a disk has failed
- the array has become corrupted
- a disk has been removed

How to rebuild the RAID array

When the RAID array is in its normal OK state, there is no option to rebuild the array because there is no need for it. You only need to rebuild the array when it is in a degraded state and in danger of losing data.

Before you rebuild the RAID array, you should have a replacement disk for the one that failed if that is the cause of the degraded array. You cannot rebuild an array that is missing a disk. A replacement disk should be the same storage capacity as the disk it is replacing.

Also before rebuilding the array, you should backup the data if possible. As soon as the RAID array becomes degraded you should backup the array if possible to prevent data loss.

To rebuild the RAID array

- 1 Go to *System > Dashboard > Status*, and then in the RAID Monitor widget, select *[Configure]*.
- 2 Verify the status of the RAID array is degraded, and the Rebuild button is not greyed out.
- 3 Remove the failed disk from the FortiGate unit.
 - Ensure you have the correct disk.
 - Press the green button to unlock the disk.
 - Gently push the lever to the left as far as it will go to disconnect the disk.
 - Remove the disk from the FortiGate unit by pulling on the lever.
- 4 Insert the new disk into the FortiGate unit that is replacing the failed disk.
 - Insert the disk carefully into the FortiGate unit.
 - Push the front panel of the disk to make the connection—the lever will start to move to the right. Ensure that both sides of the disk are in line with the other disks.
 - When in place push the bar fully to the right, until the green button clicks.
- 5 Refresh your display to ensure the new disk is installed properly. If it is not recognized, repeat steps 3 and 4 with the new disk to ensure it is properly installed.
- 6 On the configure screen, select *Rebuild RAID*.

Rebuilding the RAID array will normally take several hours. You can follow its progress on the RAID Monitor display on the dashboard.
- 7 When the rebuild is complete, the status of the RAID array will change to OK.



FortiBridge installation and operation

This chapter describes a typical transparent mode FortiGate network and how to add a FortiBridge unit to this network to provide fail open protection. This chapter also contains detailed information about how FortiBridge units operate and concludes with descriptions of adding a FortiBridge unit to an HA cluster and connecting a FortiBridge unit other FortiGate interfaces.

This chapter contains the following sections:


- [Example FortiBridge application](#)
- [Normal mode operation](#)
- [Bypass mode operation](#)
- [FortiBridge power failure](#)
- [Example FortiGate HA cluster FortiBridge application](#)
- [Example configuration with other FortiGate interfaces](#)

Example FortiBridge application

A typical application of a FortiGate unit operating in transparent mode is to insert the FortiGate unit into an internal network, between the network and the router that connects the network to the Internet. In this configuration, the FortiGate unit can provide security services for all traffic passing between the internal network and the internet. These security services can include:

- applying firewall policies and IPS attack prevention to all traffic,
- applying virus scanning to HTTP, FTP, POP3, SMTP, and IMAP traffic,
- applying web filtering to HTTP traffic,
- applying Spam filtering to POP3, SMTP, and IMAP traffic.

The internal network is connected to the FortiGate unit internal interface. The router is connected to the FortiGate unit external interface. The FortiGate unit can be added to the network without changing the configuration of the network (except to add the FortiGate management IP address).

Figure 376: Example transparent mode network

To allow users on the internal network to connect to resources on the Internet, add Internal → External firewall policies to the FortiGate unit. Add protection profiles to the firewall policies to apply security services such as virus scanning, web filtering, spam filtering and IPS to the traffic that passes through the FortiGate unit.

The FortiGate unit acts as an extra layer of protection for your internal network. While it is operating, the FortiGate unit protects the internal network from threats originating on the Internet. All users on the internal network connect through the FortiGate unit to the Internet. This also means that if a failure or other interruption caused the FortiGate unit to stop functioning, users on the internal network would not be able to connect to the Internet.

You can install a FortiBridge unit to maintain internet connectivity for the internal network if the FortiGate unit stops functioning. The FortiBridge unit provides fail open protection for your network by bypassing the FortiGate unit if a failure occurs.

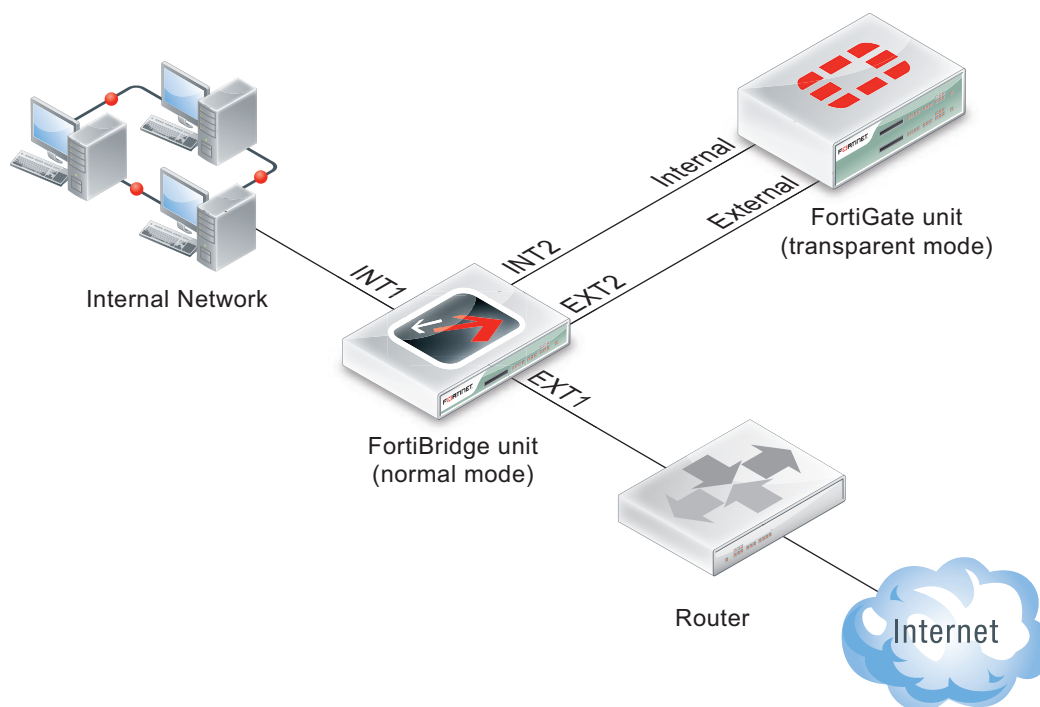
Connecting the FortiBridge unit

Operating in normal mode, the FortiBridge unit functions like a layer-2 bridge, passing all traffic to the FortiGate unit. The FortiGate unit processes the traffic, which passes through the FortiBridge unit again and then to its final destination.

In most cases, you do not have to make changes to the FortiGate unit configuration or to the network to add a FortiBridge unit. The only network requirement for FortiBridge is the availability of a single management IP address for the FortiBridge unit. The FortiBridge management IP address is required in addition to the FortiGate management IP address.

The connection procedure is different depending on whether the FortiBridge unit uses copper gigabit ethernet network connections or fiber gigabit ethernet network connections. This section includes the following connection procedures:

- [Connecting the FortiBridge-2002 \(copper gigabit ethernet\)](#)
- [Connecting the FortiBridge-2002F \(fiber gigabit ethernet\)](#)

Figure 377: FortiBridge unit providing fail open protection

Connecting the FortiBridge-2002 (copper gigabit ethernet)

The FortiBridge-2002 unit contains 4 auto-sensing 10/100/1000 Ethernet interfaces that connect to the internal and external networks and to the FortiGate interfaces that were connected to these networks. Use the following steps to connect a FortiBridge-2002 unit to the network as shown in [Figure 377](#).



Normally, you would use straight-through ethernet cables to connect the FortiBridge-2002 unit to the FortiGate unit and to your networks. However, for some connections you may need a crossover ethernet cable (for example, for compatibility with network devices that do not support Auto MDI/MDIX).

- 1 Connect the FortiBridge-2002 INT2 interface to the FortiGate internal interface.
- 2 Connect the FortiGate external interface to the FortiBridge-2002 EXT2 interface.
- 3 Connect the internal network to the FortiBridge-2002 INT1 interface.
- 4 Connect the FortiBridge-2002 EXT1 interface to the router.

Connecting the FortiBridge-2002F (fiber gigabit ethernet)

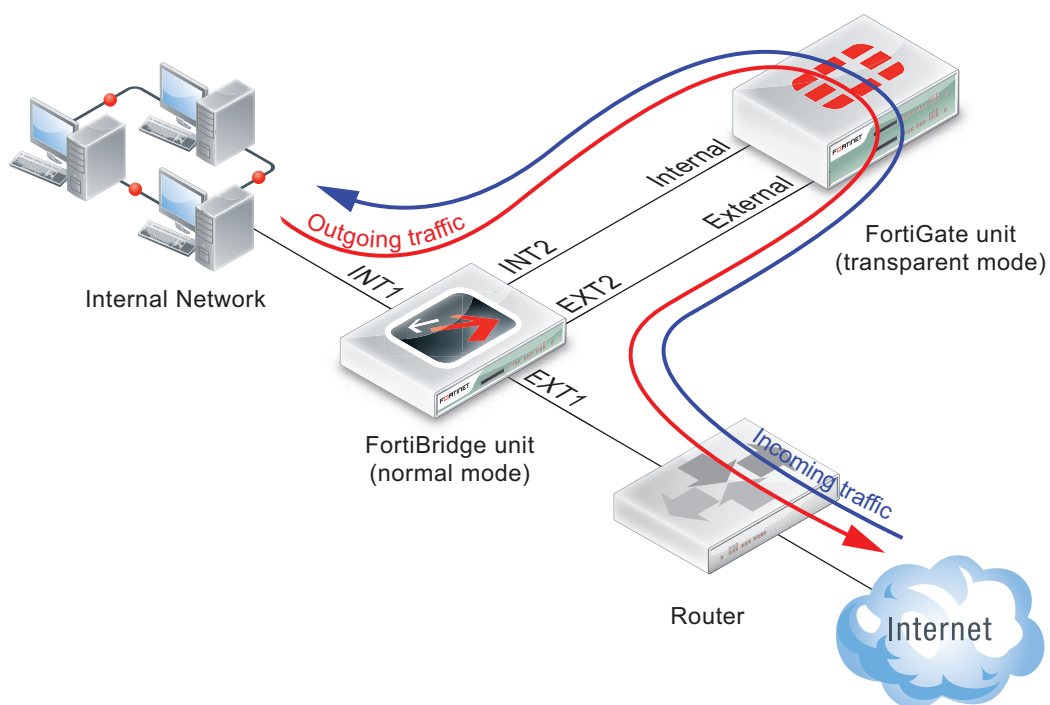
The FortiBridge-2002F unit contains 4 multimode fiber optic gigabit interfaces that connect to the internal and external networks and to the FortiGate interfaces that were connected to these networks. Use the following steps to connect a FortiBridge-2002F unit to the network as shown in [Figure 377](#).

- 1 Connect the FortiBridge-2002F INT2 interface to the FortiGate internal interface.
- 2 Connect the FortiGate external interface to the FortiBridge-2002F EXT2 interface.
- 3 Connect the internal network to the FortiBridge-2002F INT1 interface.
- 4 Connect the FortiBridge-2002F EXT1 interface to the router.

Normal mode operation

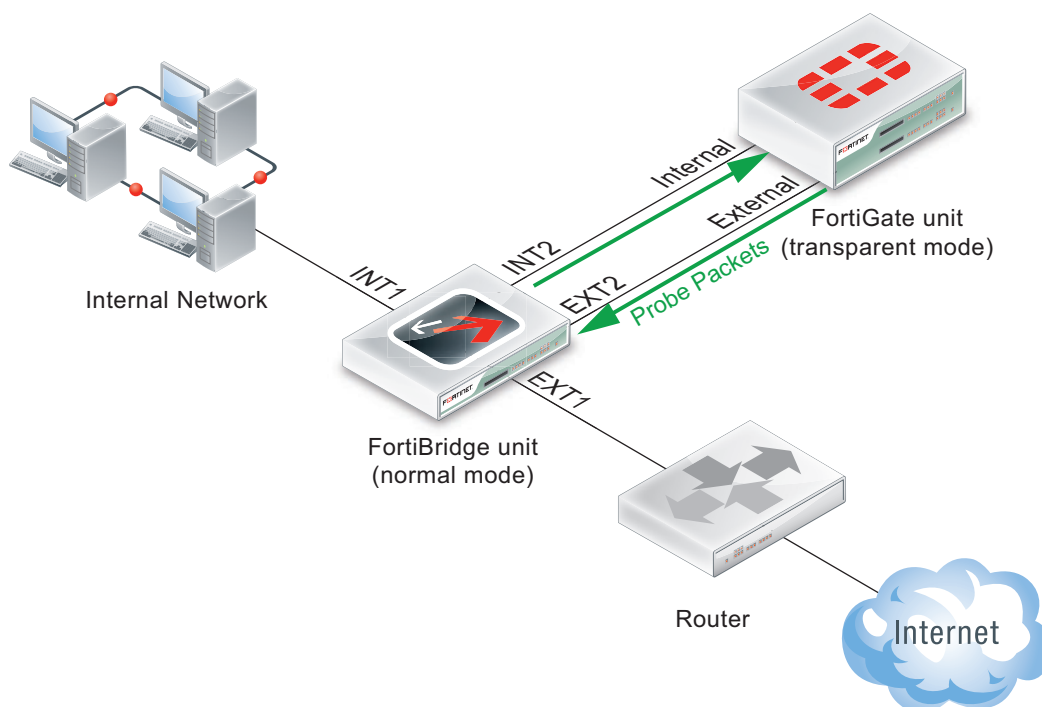
If the FortiGate unit is processing traffic normally, the FortiBridge unit operates in Normal mode. Traffic from the internal network enters the FortiBridge INT1 interface then exits the INT2 interface to the FortiGate unit. The traffic from the FortiBridge INT2 interface enters the FortiGate internal interface. Firewall policies and protection profiles are applied to the traffic by the FortiGate unit. Accepted traffic exits the FortiGate External interface and enters the FortiBridge EXT2 interface. The traffic then exits the FortiBridge EXT1 interface and goes to the external network. Traffic from the external network follows this sequence in the opposite direction.

Figure 378: Normal mode traffic flow



How the FortiBridge unit monitors the FortiGate unit

To monitor the FortiGate unit for failure, you must enable probes on the FortiBridge unit. When you enable a probe, the FortiBridge unit sends packets from the FortiBridge INT2 interface, through the FortiGate unit to the FortiBridge EXT2 interface. If the EXT2 interface receives the probe packets, the FortiGate unit is operating normally. If the EXT2 interface does not receive probe packets the FortiBridge unit assumes that the FortiGate unit has failed.

Figure 379: FortiBridge unit operating in normal mode sending probe packets

You can enable ICMP (ping), HTTP, FTP, POP3, SMTP, and IMAP probes to test connectivity through the FortiGate unit for each of these protocols. The FortiBridge unit simultaneously tests connectivity through the FortiGate unit for each probe that is enabled.

The first probe that registers a failure causes the FortiBridge unit to stop sending all probe packets. The FortiBridge unit responds to the failure according to the action on failure that you configure. The action on failure can include fail open, send alert email, send a syslog message, and send an SNMP trap. You can enable any combination of these actions on failure. Fail open switches the FortiBridge unit to bypass mode. Other actions on failure alert system administrators that the FortiBridge has determined that a failure occurred.

Probes and FortiGate firewall policies

Probe packets are accepted and passed through the FortiGate unit by firewall policies added to the FortiGate unit. When enabling probes, you must make sure that the firewall policies added to the FortiGate unit can accept probe packets. For example, if your FortiGate unit does not accept FTP packets, you should not enable the FTP probe. [Table 154](#) describes FortiGate firewall policy requirements for each FortiBridge probe.

Table 154: FortiBridge probes and FortiGate firewall policy requirements

Probe	Description	FortiGate Firewall policy	
		Direction	Service
Ping	ICMP packets are sent from the INT2 interface to the EXT2 interface. The EXT2 interface responds to the ping.	Internal -> External	ICMP or ANY

Table 154: FortiBridge probes and FortiGate firewall policy requirements

Probe	Description	FortiGate Firewall policy	
		Direction	Service
HTTP	HTTP requests are sent from an HTTP client at the INT2 interface to a web server at the EXT2 interface. The web server sends a response from the EXT2 interface to the INT2 interface.	Internal -> External	HTTP or ANY
SMTP	SMTP packets are sent from an SMTP server at the INT2 interface to an SMTP server at the EXT2 interface. The SMTP server sends a response from the EXT2 interface to the INT2 interface.	Internal -> External	SMTP or ANY
POP3	POP3 packets are sent from a POP3 client at the INT2 interface to a POP3 server at the EXT2 interface. The POP3 server sends a response from the EXT2 interface to the INT2 interface.	Internal -> External	POP3 or ANY
IMAP	IMAP packets are sent from an IMAP client at the INT2 interface to an IMAP server at the EXT2 interface. The IMAP server sends a response from the EXT2 interface to the INT2 interface.	Internal -> External	IMAP or ANY
FTP	FTP requests are sent from an FTP client at the INT2 interface to an FTP server at the EXT2 interface. The FTP server sends a response from the EXT2 interface to the INT2 interface.	Internal -> External	FTP or ANY
mm1	MM1 packets are sent from the INT2 interface to the EXT2 interface, through the FortiGate unit. When the packet is received, an MM1 response is sent back from the EXT2 interface to the INT2 interface.	Internal -> External	custom or ANY
mm3	MM3 packets are sent from the INT2 interface to the EXT2 interface, through the FortiGate unit. When the packet is received, an MM3 response is sent back from the EXT2 interface to the INT2 interface.	Internal -> External	custom or ANY
mm4	MM4 packets are sent from the INT2 interface to the EXT2 interface, through the FortiGate unit. When the packet is received, an MM4 response is sent back from the EXT2 interface to the INT2 interface.	Internal -> External	custom or ANY

Table 154: FortiBridge probes and FortiGate firewall policy requirements

Probe	Description	FortiGate Firewall policy	
		Direction	Service
mm7	MM7 packets are sent from the INT2 interface to the EXT2 interface, through the FortiGate unit. When the packet is received, an MM7 response is sent back from the EXT2 interface to the INT2 interface.	Internal -> External	custom or ANY

*No predefined service selections are offered for the MMS protocols. To allow the probes for these protocols, you can select the ANY service or create custom services for TCP packets with the destination ports listed in *Probe > Settings*.

Enabling probes to detect FortiGate hardware failure

A FortiGate unit can stop processing network traffic because of a hardware failure such as the failure of a hardware component, a loss of power, or a loss of connectivity if a network cable is unplugged.

If a hardware failure occurs, the FortiGate unit stops processing all traffic. You can enable any FortiBridge probe for the FortiBridge unit to detect a FortiGate hardware failure.

Enabling probes to detect FortiGate software failure

A FortiGate unit can also stop processing network traffic because of a software failure. For example, a firmware issue could cause a specific software process to crash. Also, network traffic could increase to a point where the FortiGate unit cannot process all traffic. As a result, the FortiGate unit could stop processing some or all traffic without a hardware failure occurring.

To detect a FortiGate software failure, you can enable probes for FortiGate services that you want to provide fail open protection for. For example, if SMTP email services are a high priority for your network, you should enable the SMTP probe. If the SMTP probe detects a failure of SMTP traffic through the FortiGate unit, the FortiBridge unit switches to bypass mode to maintain SMTP traffic flow.

If you do not consider FTP traffic a high priority, you can leave the FTP probe disabled. In this configuration, if only FTP traffic fails, the FortiBridge does not switch to bypass mode.

Probe interval and probe threshold

For each probe, you set a probe interval and a probe threshold. The probe interval defines how often to test the connection. The probe threshold defines how many consecutive failed probes can occur before the FortiBridge considers the connection to have failed.

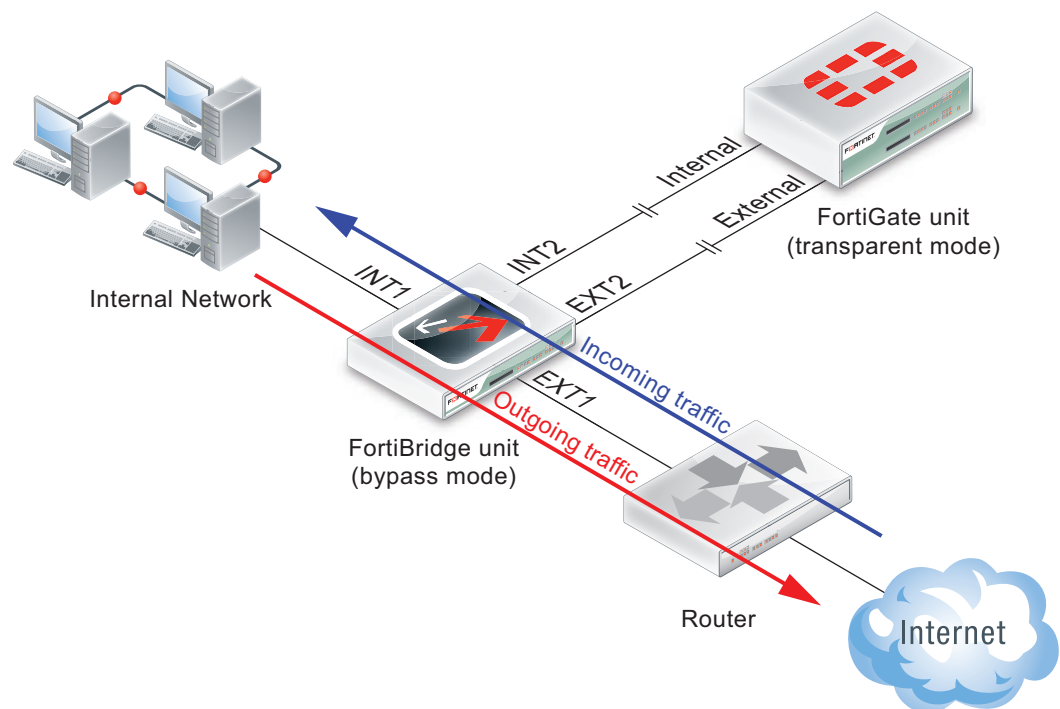
Bypass mode operation

When the FortiBridge unit operates in bypass mode, the FortiBridge INT1 and EXT1 interfaces are directly connected. All traffic between the internal and external network segments flows, whether or not the FortiGate unit is operating normally.

Because the INT1 and EXT1 interfaces are directly connected, you cannot use Telnet or SSH to connect to the FortiBridge CLI. Instead you must use a console connection.

The FortiBridge unit remains in bypass mode even if the FortiGate unit recovers. To restore the FortiGate unit, you must manually switch the FortiBridge unit back to normal mode. You can switch the FortiBridge unit to normal mode by pressing the mode switch on the FortiBridge front panel or by using a console connection to the CLI and entering the command `execute switch-mode`. You can also use the mode switch and the `execute switch-mode` command to manually switch the FortiBridge unit from normal mode to bypass mode.

Figure 380: FortiBridge unit operating in bypass mode



When the FortiBridge unit is operating in bypass mode you can still connect to the FortiBridge CLI and manage the FortiBridge unit (for example to switch the FortiBridge unit to normal mode). When the FortiBridge unit operates in bypass mode, you cannot connect to the FortiGate interfaces that are connected to the FortiBridge unit.

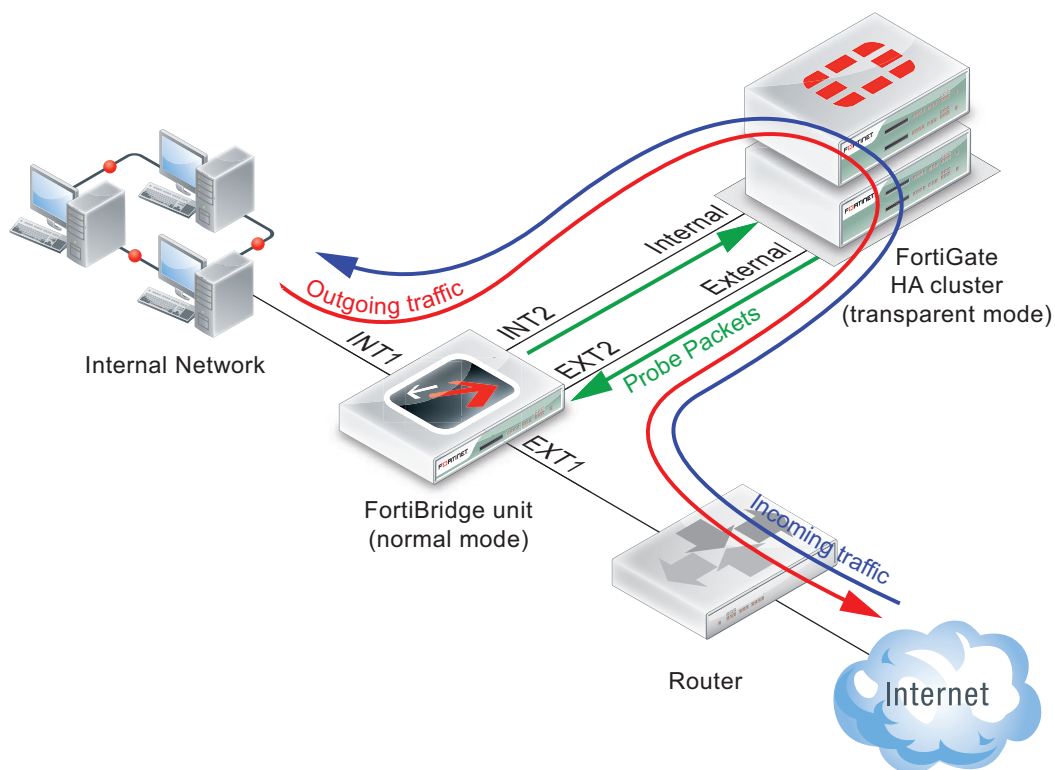
FortiBridge power failure

If a power failure occurs and the FortiBridge unit loses power, zero power fail-open technology causes FortiBridge unit to fail open. The FortiBridge unit bypasses the FortiGate unit and all traffic passes between the FortiBridge INT1 and EXT1 interfaces. If power is restored to the FortiBridge unit, it starts up in bypass mode and then switches to normal mode when its start up sequence is complete, reconnecting the FortiGate unit to the network.

Example FortiGate HA cluster FortiBridge application

A FortiBridge unit can provide fail open protection for a FortiGate HA cluster operating in transparent mode in much the same way as for a standalone FortiGate unit. To provide fail open protection for an HA cluster, connect the FortiBridge unit to the switches that connect the internal and external interfaces of the cluster. Use the following steps to connect a FortiBridge unit to the HA cluster, as shown in [Figure 381](#):

Figure 381: FortiBridge unit providing fail open protection for a FortiGate HA cluster



The network configuration and FortiBridge configuration are the same for a cluster and for a standalone FortiGate unit. In normal mode, packets pass through the FortiBridge unit and through the FortiGate HA cluster and back through the FortiBridge unit. For the cluster to process this traffic, you must add Internal -> External firewall policies to the cluster configuration. If a failure occurs and the cluster no longer processes traffic, the FortiBridge unit switches to bypass mode, bypassing the cluster.

The connection procedure is different depending on whether the FortiBridge unit uses copper gigabit ethernet network connections or fiber gigabit ethernet network connections. This section includes the following connection procedures:

- [Connecting the FortiBridge-2002 \(copper gigabit ethernet\)](#)
- [Connecting the FortiBridge-2002F \(fiber gigabit ethernet\)](#)

Connecting the FortiBridge-2002 (copper gigabit ethernet)

The FortiBridge-2002 unit contains 4 auto-sensing 10/100/1000 Ethernet interfaces that connect to the internal and external networks and to the cluster interfaces that were connected to these networks. Use the following steps to connect a FortiBridge-2002 unit to the network as shown in [Figure 381](#).



Normally, you would use straight-through ethernet cables to connect the FortiBridge-2002 unit to the FortiGate unit and to your networks. However, for some connections you may need a crossover ethernet cable (for example, for compatibility with network devices that do not support Auto MDI/MDIX).

- 1 Connect the FortiBridge-2002 INT2 interface to the switch connected to the HA cluster internal interface.
- 2 Connect the switch connected to the HA cluster external interface to the FortiBridge-2002 EXT2 interface.
- 3 Connect the internal network to the FortiBridge-2002 INT1 interface.
- 4 Connect the FortiBridge-2002 EXT1 interface to the router.

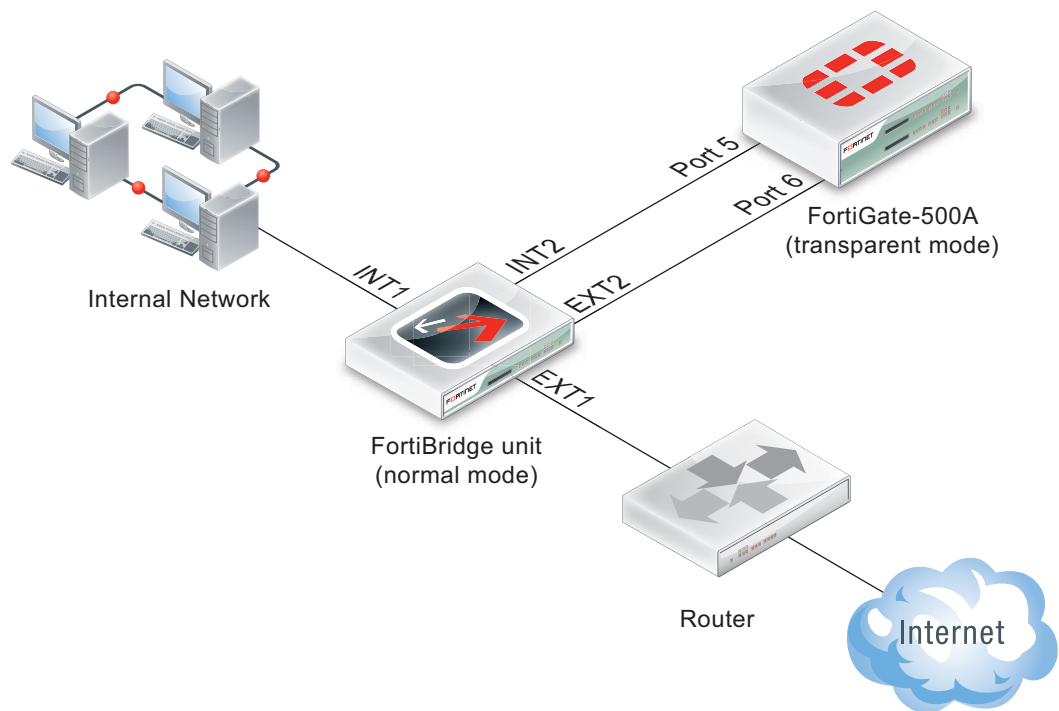
Connecting the FortiBridge-2002F (fiber gigabit ethernet)

The FortiBridge-2002F unit contains 4 multimode fiber optic gigabit interfaces that connect to the internal and external networks and to the FortiGate cluster interfaces that were connected to these networks. Use the following steps to connect a FortiBridge-2002F unit to the network as shown in [Figure 377](#).

- 1 Connect the FortiBridge-2002F INT2 interface to the switch connected to the HA cluster internal interface.
- 2 Connect the switch connected to the HA cluster external interface to the FortiBridge-2002F EXT2 interface.
- 3 Connect the internal network to the FortiBridge-2002F INT1 interface.
- 4 Connect the FortiBridge-2002F EXT1 interface to the router.

Example configuration with other FortiGate interfaces

All of the examples in this chapter describe using the FortiBridge unit to provide fail open protection for traffic passing between the FortiGate unit internal and external interfaces. You can actually use a FortiBridge unit to provide fail open protection for any two FortiGate unit interfaces. No limitation is implied by naming the FortiBridge interfaces INT and EXT. These names are used to simplify installation procedures. [Figure 382](#) shows a FortiBridge unit providing fail open protection for network traffic between ports 5 and 6 of a FortiGate-500A unit.

Figure 382: FortiBridge unit providing fail open protection for a single FortiGate unit

To connect a FortiBridge unit to the network shown in [Figure 382](#):

- 1 Connect the FortiBridge INT2 interface to the FortiGate-500A port 5 interface.
 - 2 Connect the FortiGate-500A port 6 interface to the FortiBridge EXT2 interface.
 - 3 Connect the internal network to the FortiBridge INT1 interface.
 - 4 Connect the FortiBridge EXT1 interface to the router.
- You must add port 5 -> port 6 firewall policies to the FortiGate-500A unit configuration.

Completing the basic FortiBridge configuration

Now that you have connected the FortiBridge unit to your network and connected to the FortiBridge CLI, use the following procedures to complete the basic configuration of the FortiBridge unit.



Not all of the following procedures are required to complete the basic FortiBridge unit configuration. Choose the procedures that apply to your installation.

- [Adding an administrator password](#)
- [Changing the management IP address](#)
- [Changing DNS server IP addresses](#)
- [Changing the default gateway and adding static routes](#)
- [Allowing management access to the EXT1 interface](#)
- [Changing the system time and date](#)
- [Adding administrator accounts](#)

When you complete the procedures in this chapter, the FortiBridge unit will be operating and connected to your network and to your FortiGate unit. See [“Example network configuration” on page 2987](#) to configure the FortiBridge unit to monitor the status of the FortiGate unit and to fail open if the FortiBridge unit detects that the FortiGate unit has failed.

Adding an administrator password

Add an administrator password to the default admin administrator account to prevent unauthorized users from connecting to and managing the FortiBridge unit.

To add an administrator password — Web-based manager

- 1 Go to *System > Status*.
- 2 In the Administrators section of the dashboard, select the *Edit* icon of the admin user.
- 3 Select the *Change Password* link.
- 4 Enter the new password.
- 5 Enter the new password again in the second field.
- 6 Select *OK*.

To add an administrator password — CLI

```
config system admin
  edit admin
    set password <password_str>
  end
```

Changing the management IP address

Change the FortiBridge unit management IP address so that you can connect to the FortiBridge CLI from your network (instead of being required to use a direct console connection). The management IP should be a valid IP address for your network.

To change the management IP address — Web-based manager

- 1 Go to *System > Status*.
- 2 Select the *Change* link in the Management Port section of the dashboard.
- 3 Enter the new management IP address and netmask in the IP/Netmask field.
- 4 Select *OK*.

To change the management IP address — CLI

```
config system manageip
    set ip <management_ipv4mask>
end
```

Changing DNS server IP addresses

Change the FortiBridge DNS server IP addresses to the IP addresses of your DNS servers. The correct DNS server configuration is required for alert email.

To change DNS server IP addresses — Web-based manager

- 1 Go to *System > Status*.
- 2 Select the *Change* link in the Management Port section of the dashboard.
- 3 Enter the primary DNS IP address in the *Primary DNS Server* field.
- 4 Enter the secondary DNS IP address in the *Secondary DNS Server* field.
- 5 Select *OK*.

To change DNS server IP addresses — CLI

```
config system dns
    set primary <primary_ipv4>
    set secondary <secondary_ipv4>
end
```

Changing the default gateway and adding static routes

Add static routes if you need to route packets from the FortiBridge unit through a router to another network. For example, if alert email sends email messages from the internal network to an email server on the Internet, you should add a route to the Internet.

The web-based manager allows you to enter only the default gateway. If you require additional static routes, use the CLI to enter them.

To change the default gateway — Web-based manager

- 1 Go to *System > Status*.
- 2 Select the *Change* link in the Management Port section of the dashboard.
- 3 Enter the default gateway IP address in the *Default Gateway* field.

To change the default gateway — CLI

```
config system route
    edit <sequence_int>
        set gateway <gateway_ipv4>
    end
```

To add additional static routes — CLI

```
config system route
```

```

edit <sequence_int>
  set gateway <gateway_ipv4>
  set dst <destination_ipv4mask>
end

```

Allowing management access to the EXT1 interface

By default no management access is configured for the EXT1 interface. Use the following procedure to add management access to this interface if required.

Configuring the EXT1 interface to allow management access is possible only using the CLI.

To allow management access to the EXT1 interface — CLI

```

config system interface external
  set allowaccess ssh
end

```

Changing the system time and date

Use the following procedure to change the system time and date.

To change the system time and date — Web-based manager

- 1 Go to *System > Status*.
- 2 Select the *Change* link beside *System Time* in the System Information section of the dashboard.
- 3 Enter the time, date, and timezone as required.
- 4 Select *OK*.

To change the system time and date — CLI

```

execute time <hh:mm:ss>
execute date <mm/dd/yyyy>
config system global
  set timezone <timezone_int>
end

```

Enter the number corresponding to your time zone. Type ? to list time zones and their numbers.

For example, to set the time zone to Central time (time zone number 8), enter:

```

config system global
  set timezone 8
end

```

For information about configuring other global settings, see “system global” in the [FortiBridge CLI Reference](#).

Adding administrator accounts

The factory default FortiBridge configuration includes the admin administrator account. Use this procedure to add more administrator accounts.

To add administrator accounts — Web-based manager

- 1 Go to *System > Status*.
- 2 In the Administrators section of the dashboard, select *Create New*.

- 3 Enter the administrator account name.
- 4 Enter the administrator account password.
- 5 Enter the password again in the second field.
- 6 Select *OK*.

To add administrator accounts — CLI

```
config system admin
  edit <admin_name_str>
    set password <password_str>
    set accprofile prof_admin
  end
```

For more information about configuring administrators see “system admin” in the [FortiBridge CLI Reference](#).

Resetting to the factory default configuration

Use the following procedure to reset the FortiBridge unit to the factory default configuration. You might want to reset the FortiBridge to the factory default condition if the FortiBridge unit is not functioning as expected and you would like to re-start the configuration process. Resetting to the factory default configuration resets all configuration changes that you have made, including the management IP address.

To reset to factory default configuration from the FortiBridge front panel

- 1 Use a pen or other pointed object to press the Factory reset button.
After a few seconds the FortiBridge unit restarts; reset to the factory default configuration. You can now re-configure the FortiBridge unit.

To reset to factory defaults — CLI

```
execute factoryreset
```

A few seconds after confirming your command, the FortiBridge unit restarts, reset to the factory default configuration. You can now re-configure the FortiBridge unit.

Installing FortiBridge unit firmware

Before beginning any of the procedures in this section, you must have the FortiBridge firmware image file that you are going to install on the FortiBridge unit. During these procedures you are required to enter the name of the firmware image file.

Changing firmware versions

You can use these procedure to upgrade to a newer version of the FortiBridge firmware, re-install the current version, or revert to an older version of the firmware.

The CLI-based procedure requires that you have a TFTP server you can connect to from the FortiBridge unit.

Changing firmware versions — Web-based manager

- 1 Go to *System > Status*.
- 2 In the System Information section, the Firmware Version displays the currently installed firmware version.

- 3 Select *Update* to install another version of the firmware.
- 4 Select *Browse* to choose the firmware file on your computer.
- 5 Select *OK* to install the firmware file.
- 6 If you are installing an older version of the firmware, you must confirm your selection before the installation can proceed.
- 7 The FortiBridge installs the firmware and restarts. This process takes a few minutes.
- 8 To confirm that the firmware you selected is installed, log into the web-based manager, go to *System > Status*, and confirm that the firmware version is correct.

Changing firmware versions — CLI

- 1 Make sure that the TFTP server is running.
- 2 Copy the new firmware image file to the root directory of your TFTP server.
- 3 Log into the CLI as an administrator with `syssshutdowngrp` access.
Normally this would be the admin administrator. But you can use access profiles to control administrative access. See “system accprofile” in the [FortiBridge CLI Reference](#) for more information.
- 4 Make sure the FortiBridge unit can connect to the TFTP server.
You can use the following command to ping the computer running the TFTP server. For example, if the TFTP server IP address is 192.168.1.168:

```
execute ping 192.168.1.168
```

- 5 Enter the following command to copy the firmware image from the TFTP server to the FortiBridge unit:

```
execute restore image <name_str> <tftp_ip>
```

Where `<name_str>` is the name of the firmware image file on the TFTP server and `<tftp_ip>` is the IP address of the TFTP server. For example, if the firmware image file name is `FBG_2002-v30-build010-FORTINET.out` and the IP address of the TFTP server is 192.168.1.23, enter:

```
execute restore image FBG_2002-v30-build010-FORTINET.out
192.168.1.168
```

- 6 If you are downgrading to an older firmware version, a message is displayed:
Get image from tftp server OK.
This operation will downgrade the current firmware version!
Do you want to continue? (y/n)
If you are certain that you want to downgrade to the older firmware version, press `Y`.
- 7 The FortiBridge installs the firmware and restarts. This process takes a few minutes.
- 8 Reconnect to the CLI.
- 9 To confirm that the new firmware image has been loaded, enter:

```
get system status
```

Installing firmware from a system reboot

This procedure installs a specified firmware image and resets the FortiBridge unit to default settings. You can use this procedure to upgrade to a new firmware version, revert to an older firmware version, or to re-install the current firmware.

To use this procedure you:

- access the CLI by connecting to the FortiBridge console port,

- install a TFTP server that you can connect to from the FortiBridge EXT2 interface. The TFTP server should be on the same network as the EXT2 interface. The FortiBridge unit cannot access the TFTP server if its behind a router.

During this procedure you will be asked to enter a local IP address for the FortiBridge unit. This is a temporary address used for downloading the firmware image.

This procedure reverts your FortiBridge unit to its factory default configuration. Before running this procedure you can backup the FortiBridge unit configuration using the command `execute backup config`.

To install firmware from a system reboot

- 1 Connect to the CLI using the FortiBridge console port.
- 2 Make sure the TFTP server is running.
- 3 Copy the new firmware image file to the root directory of the TFTP server.
- 4 Make sure the EXT2 interface of the FortiBridge unit can connect to the TFTP server.
- 5 Enter the following command to restart the FortiBridge unit:

```
execute reboot
```

As the FortiBridge unit starts, a series of system startup messages are displayed. When the following messages appears:

```
Hit any key to stop autoboot:
```

- 6 Immediately press any key to interrupt the system startup.



You only have 3 seconds to press any key. If you do not press any key soon enough, the FortiBridge unit reboots and you must log in and repeat the `execute reboot` command.

When you successfully interrupt the startup process, the `=>` prompt appears:

- 7 Type `upgrade` and press Enter to get the new firmware image from the TFTP server. The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

- 8 Type the address of the TFTP server and press Enter.

The following message appears:

```
Enter local address [192.168.1.188]:
```

- 9 Type an IP address that the FortiBridge unit can use to connect to the TFTP server press Enter.



The local IP address is a temporary address used to download the firmware image. The local IP address should be on the same subnet as the TFTP server IP address.

The following message appears:

```
Enter firmware image file [image.out]:
```

- 10 Type the firmware image file name and press Enter.

The TFTP server uploads the firmware image file to the FortiBridge unit and the FortiBridge unit installs the new firmware image, resets the configuration to factory defaults, and restarts. This process takes a few minutes.

- 11 Reconnect to the CLI.

12 To confirm that the firmware image has been loaded, enter:

```
get system status
```


Example network configuration

This chapter describes how to configure a FortiBridge unit to provide fail open protection for a FortiGate unit operating in transparent mode. This chapter also describes some commonly required FortiBridge operating procedures such as recovering from a fail open event, manually switching between FortiBridge operating modes and backing up and restoring the FortiBridge configuration.

The procedures in this chapter assume that you have connected the FortiBridge unit to your network and completed its basic configuration as described in [“Completing the basic FortiBridge configuration” on page 2979](#).



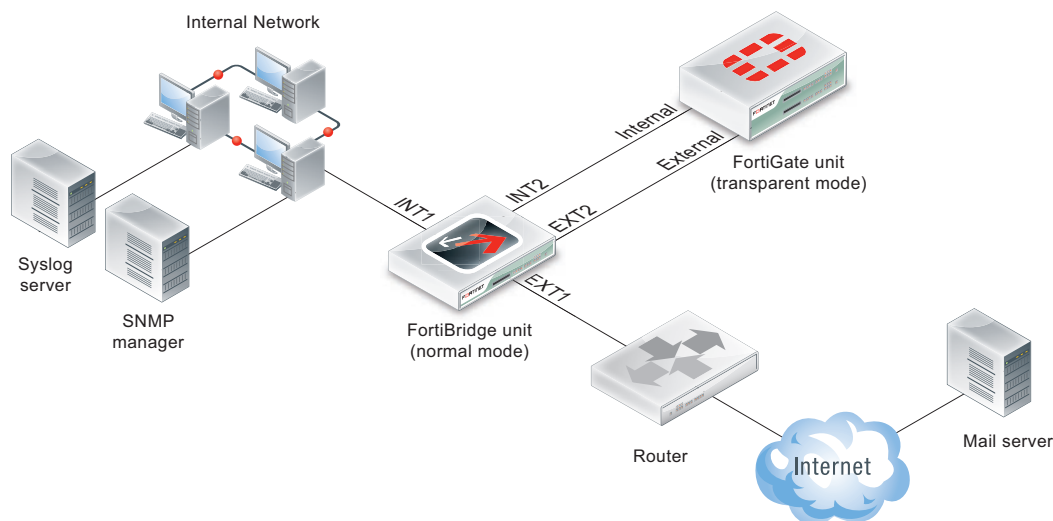
The information in this chapter can be applied to any standalone FortiGate transparent mode network configuration. These procedures can also be applied to a FortiBridge unit providing fail open protection for a FortiGate HA cluster operating in transparent mode.

The descriptions and procedures in this section assume that the FortiGate unit is installed between an internal network and the router that connects the internal network to the Internet as show in [Figure 383](#). The FortiGate unit can provide the following security services for all traffic passing between the internal network and the internet:

- Internal -> External firewall policies for HTTP, FTP, POP3, SMTP, and IMAP connections from Internal network to the Internet.
- Virus scanning of HTTP, FTP, POP3, SMTP, and IMAP traffic,
- Web filtering of HTTP traffic,
- Spam filtering of POP3, SMTP, and IMAP traffic.

In addition to the above security services, a FortiCarrier unit can process MM1, MM3, MM4, and MM7 traffic.

Figure 383: Example FortiBridge application



[Table 155](#) lists the internal network configuration.

Table 155: Internal network configuration

FortiGate management IP address	172.20.120.10/24
Internal network subnet IP address	172.20.120.0/24

Table 155: Internal network configuration (Continued)

Router internal IP address	172.20.120.1/24
Internal network default route	172.20.120.1
Primary DNS server	172.20.120.2
Secondary DNS server	172.20.120.3
Syslog Server IP address	172.20.120.11
SNMP Manager IP address	172.20.120.12
Mail Server Name	mail.myorg.com

[Table 156](#) lists the basic FortiBridge unit configuration settings.

Table 156: Basic FortiBridge unit configurations settings

Administrator password	passWORD
Management IP address	172.20.120.20/24
Default route	172.20.120.1
Primary DNS server	172.20.120.2
Secondary DNS server	172.20.120.3

Configuring FortiBridge probes

To monitor a FortiGate unit for failure, you configure the FortiBridge unit to send probe packets through the FortiGate unit. Using probe packets, the FortiBridge unit can confirm that the FortiGate unit can process ICMP (ping), HTTP, FTP, POP3, SMTP, IMAP, MM1, MM3, MM4, and MM7 traffic. Until you configure probes, the FortiBridge unit cannot detect if the FortiGate unit has failed.

This section describes:

- [Probe settings](#)
- [Enabling probes](#)
- [Verifying that probes are functioning](#)
- [Tuning the failure threshold and probe interval](#)

Probe settings

Configure probe settings to control the response when a FortiBridge probe detects that the FortiGate unit has failed. Probe settings consist of:

Table 157: Probe settings

Probe Setting	Description	Default
Action on failure	Set the FortiBridge unit response when a probe detects that the FortiGate unit has failed. The FortiBridge unit can, <ul style="list-style-type: none"> • Send alertmail • Fail open • Send an SNMP trap • Send a message to a syslog server You can add up to four actions on failure. All of the configured actions on failure occur when the FortiBridge unit detects a failure.	fail open
Dynamic IP pattern	Configure the INT2 and EXT2 interfaces with dynamic probe IP addresses. The dynamic probe IP addresses should not conflict with IP addresses on the network that the FortiGate unit is connected to. These IP addresses are not visible from the outside network, but they should not conflict with IP addresses in packets passing through the FortiBridge unit. You cannot change the dynamic IP pattern if any probes are enabled.	(none)
FortiGate unit serial number	The serial number of the FortiGate unit that the FortiBridge unit is connected to. The serial number appears in FortiBridge alert mail, and syslog messages to identify the FortiGate unit.	(none)

To configure probe settings

This procedure shows how to configure the following probe settings:

- The FortiBridge unit responds to a FortiGate unit failure by failing open and by sending an alert email, a syslog message, and an SNMP trap
- The dynamic IP pattern is 2.2.2.*
- The FortiGate unit serial number is FGT8002803923050



The FortiBridge unit does not have to fail open if the FortiGate unit fails. The FortiBridge unit can be configured just to send alerts if the FortiGate unit fails.

Configure probe settings — Web-based manager

- 1 Go to *Probe > Settings*.
- 2 Enter the IP pattern in the *Probe IP Address Pattern* field.
- 3 Select *Apply*.
- 4 Go to *Probe > Notifications*.

5 Select the notification types you require.

6 Select *Apply*.

You cannot set the failopen or failcutoff action, nor the FortiGate serial number using the web-based manager.

Configure probe settings — CLI

```
config probe setting
  set action_on_failure alertmail failopen snmp syslog
  set dynamic_ip_pattern 2.2.2.*
  set fgt_serial FGT8002803923050
end
```

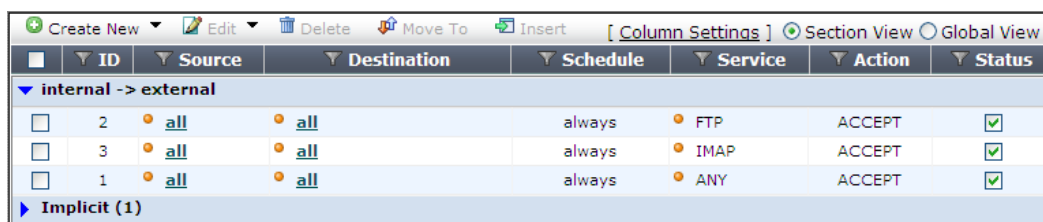
Enabling probes

Enable probes to control the protocols that the FortiBridge unit uses to confirm that the FortiGate unit is functioning normally. You can configure probes for ping (ICMP), HTTP, FTP, POP3, SMTP, IMAP, MM1, MM3, MM4, and MM7 protocols. For all probes you can configure the probe interval (the time between consecutive probe packets) and the probe threshold (the number of probe packets lost before the FortiBridge unit registers a failure). For HTTP, FTP, POP3, SMTP, and IMAP probes you can also change the probe port. You would change the probe port for a protocol if the FortiGate unit uses a non-standard port for that protocol.

The FortiBridge unit simultaneously tests connectivity through the FortiGate unit for each probe that you have enabled. The first probe that registers a failure causes all probes to stop and the configured action on failure to occur.

Before you configure probes, the FortiGate unit must be configured to pass the probe traffic. A single Internal->External firewall policy that allows all traffic also allows all probe packets. You can also configure individual policies for each protocol. For example, you could add the firewall policies shown in [Figure 384](#) to the FortiGate unit.

Figure 384: Sample firewall policies



ID	Source	Destination	Schedule	Service	Action	Status
2	all	all	always	FTP	ACCEPT	✓
3	all	all	always	IMAP	ACCEPT	✓
1	all	all	always	ANY	ACCEPT	✓

Policy 1 processes any network traffic. Policy 2 processes all FTP traffic. Policy 2 is above Policy 1 in the policy list, so FTP traffic is matched by policy 2. In the same way, Policy 3 processes all IMAP traffic.

FTP and IMAP probes would be processed by policies 2 and 3 respectively. All other probes would be processed by policy 1. This would include pings, SMTP traffic and so on.

To enable and configure FortiBridge probes — Web-based manager

The following steps show examples for configuring ping, HTTP, FTP, POP3, SMTP, and IMAP probes. For a complete description of FortiBridge probes see “probe probe_list {ping | http | ftp | pop3 | smtp | imap | mm1 | mm3 | mm4 | mm7}” in the [FortiBridge CLI Reference](#).

1 Go to *Probe > Settings*.

- 2 For the ping protocol, select *Enable*.
This enables ping probes with the default settings.
- 3 For the FTP protocol, select *Enable*, enter 5 for the *Interval*, and enter 8 for the *Failure-Threshold*.
These settings have the FortiBridge unit send an FTP probe every 5 seconds and fail open if 8 consecutive FTP probe packets are not received.
- 4 For the IMAP protocol, select *Enable*.
This enables IMAP probes with the default settings.
- 5 For the SMTP protocol, select *Enable* and enter 26 for the *Port Number*.
This enables SMTP probes on port 26.

To enable and configure FortiBridge probes — CLI

- 1 Enable the ping probe using the default ping probe parameters. Enter:


```
config probe probe_list ping
  set status enable
end
```
- 2 Display ping probe settings, enter:


```
get probe probe_list ping
  name           : ping
  failure_threshold : 3
  probe_interval  : 1
  status          : enable
```
- 3 Enable the FTP probe. Increase the failure threshold to 5 and the probe interval to 8.


```
config probe probe_list ftp
  set status enable
  set failure_threshold 8
  set probe_interval 5
end
```

The FortiBridge unit sends an FTP probe every 5 seconds and fails open if 8 consecutive FTP probe packets are not received.
- 4 Display FTP probe settings. Enter:


```
get probe probe_list ftp
  name           : ftp
  failure_threshold : 8
  probe_interval  : 5
  status          : enable
  test_port       : 21
```
- 5 Enable the IMAP probe. Enter:


```
config probe probe_list IMAP
  set status enable
end
```
- 6 Enable the SMTP probe and change the port used by the probe from 25 to 26. Enter:


```
config probe probe_list SMTP
  set status enable
  set test_port 26
end
```

Verifying that probes are functioning

You verify that the probes are functioning by viewing the sessions being processed by the FortiGate unit.

To verify that probes are functioning

- 1 Go to *System > Dashboard > Status*.
- 2 In the Top Sessions widget, select *Details* at the bottom of the widget.
The current sessions list appears. Optionally select *Detach* to detach and expand the browser window to see the entire list.
- 3 View the sessions on the *Session* list.

Figure 385: FortiGate Session list showing FortiBridge probes

<div> 1 / 1 Total: 4 Clear All Filters Return </div>									
#	Protocol	Source Address	Source Port	Destination Address	Destination Port	Policy ID	Expiry (sec)	Duration (sec)	
1	tcp	2.2.2.13	1053	2.2.2.14	143	3	3352	4486	
2	tcp	2.2.2.13	1054	2.2.2.14	21	2	3600	4486	
3	icmp	2.2.2.13		2.2.2.14		1	3598	28	
4	tcp	2.2.2.13	1052	2.2.2.14	26	1	3598	4486	

This session list shows the following:

- The FortiBridge dynamic probe IP addresses are 2.2.2.213 and 2.2.2.214.
- IMAP probe packets (port 143) are processed by firewall policy 3.
- FTP probe packets (port 21) are processed by firewall policy 2.
- ping probe packets are processed by firewall policy 1.
- SMTP packets using port 26 are processed by firewall policy 1.

Tuning the failure threshold and probe interval

If you find the FortiBridge unit failing open when the FortiGate unit has not failed or if the FortiGate unit fails and there is an unacceptably long delay before the FortiBridge unit fails open, you should adjust the failure threshold and probe interval.

Failing open when the FortiGate unit has not failed indicates that you should increase the time the FortiBridge unit waits to fail open. During startup, if the FortiBridge unit begins sending probe packets before the FortiGate unit has completed its start up sequence the FortiBridge unit may detect a failure and switch to bypass mode. Also, if the FortiGate unit is processing high traffic volumes, a fail open could occur if the FortiGate unit delays FortiBridge probe packets. You can increase the fail open delay by increasing the failure threshold and probe interval.

An unacceptable delay before failing open means network traffic can be interrupted for the time period between when the FortiGate unit fails and the FortiBridge unit fails open. You can minimize the delay by reducing the failure threshold and probe interval.

Configuring FortiBridge alerts

Configure FortiBridge alerts so that the `alertemail`, `syslog`, and `snmp` actions on failure cause the FortiBridge unit to notify system administrators that the FortiGate unit has failed. Until you configure alert email, syslog, and SNMP alerts, the FortiBridge cannot notify system administrators of a FortiGate failure.

You can configure the following FortiBridge alerts:

- [FortiBridge alert email](#)

- [FortiBridge syslog](#)
- [FortiBridge SNMP](#)

FortiBridge alert email

If you set the probe action on failure to `alertmail`, you can configure alert email so that the FortiBridge unit sends an email message to up to three email addresses if the FortiBridge unit detects a failure. The alert email informs the recipient that a FortiGate unit has failed, includes the protocol for which the failure was detected, and includes the serial number of the FortiGate unit that failed.

Only the first probe to detect a failure triggers the actions on failure. So, even if multiple probes are configured, when a failure is detected, the FortiBridge unit sends one alert email.

Figure 386: Sample FortiBridge alert email message

```
FortiBridge detect FortiGate failure

Time: Tue Feb  1 19:58:46 2010
failed protocol: http
failed FortiGate serial number: FGT8002803923050
```

To configure alert email — Web-based manager

Configuring FortiBridge alert email is similar to configuring FortiGate alert email.

- 1 Go to *Probe > Notifications*.
- 2 Enable *Email*.
- 3 Enter your email server name in the *SMTP Server* field.
- 4 Enter the email addresses to which the alert email messages are sent in the *Email to* fields. Three fields are provided for up to three addresses.
- 5 If your email server requires authentication to send messages, select *Authentication* and enter your SMTP user name and password.
- 6 Select *Apply*.

To configure alert email — CLI

```
config alertemail setting
    set server mail.myorg.com
    set username user@company.com
    set password PassWORD
    set mailto1 user@company.com
    set mailto1 user2@company.co.uk
    set mailto1 user3@company.com
end
```

FortiBridge syslog

If you set the probe action on failure to `syslog`, you can configure the FortiBridge unit to send a syslog message to one syslog server if the FortiBridge unit detects a failure. The message informs the recipient that a FortiGate unit has failed, includes the protocol for which the failure was detected, and includes the serial number of the FortiGate unit that failed.

Only the first probe to detect a failure triggers the actions on failure. So, even if multiple probes are configured, when a failure is detected, the FortiBridge unit sends one message.

Figure 387: Sample FortiBridge syslog messages

```
02-01-2010 18:22:50 Local7.Alert 172.20.120.13 date=2010-02-01
time=15:28:22 device_id= log_id=0100020001 type=event
subtype=system pri=alert msg="FortiBridge detect FortiGate
failure: [failed time: Tue Feb 1 15:28:22 2010][failed protocol:
http] [failed FortiGate serial number: FGT8002803923050]"
02-01-2010 8:21:27 Local7.Alert 172.20.120.13 date=2010-02-01
time=15:26:59 device_id= log_id=0100020001 type=event
subtype=system pri=alert msg="FortiBridge detect FortiGate
failure: [failed time: Tue Feb 1 15:26:59 2010][failed protocol:
ftp] [failed FortiGate serial number: FGT8002803923050]"
02-01-2010 18:17:17 Local7.Alert 172.20.120.13 date=2010-02-01
time=15:22:49 device_id= log_id=0100020001 type=event
subtype=system pri=alert msg="FortiBridge detect FortiGate
failure: [failed time: Tue Feb 1 15:22:49 2010][failed protocol:
ping] [failed FortiGate serial number: FGT8002803923050]"
02-01-2010 8:13:43 Local7.Alert 172.20.120.13 date=2010-02-01
time=15:19:15 device_id= log_id=0100020001 type=event
subtype=system pri=alert msg="FortiBridge detect FortiGate
failure: [failed time: Tue Feb 1 15:19:15 2010][failed protocol:
smtp] [failed FortiGate serial number: FGT8002803923050]"
```

To configure FortiBridge syslog — Web-based manager

In most cases you should only need to configure the IP address of the syslog server to receive FortiBridge syslog messages. See “log syslogd setting” in the [FortiBridge CLI Reference](#) for more FortiBridge syslog options.

- 1 Go to *Probe > Notifications*.
- 2 Enable *Syslog*.
- 3 In the IP address field, enter the syslog server address
- 4 If required, configure the port, facility, and format.
- 5 Select *Apply*.

To configure FortiBridge syslog — CLI

```
config log syslogd setting
    set server 172.20.120.11
end
```

FortiBridge SNMP

If you set the probe action on failure to `snmp`, you can configure FortiBridge SNMP settings so that the FortiBridge unit sends SNMP v1 and v2c compliant traps to SNMP v1 and v2c compliant SNMP managers if the FortiBridge unit detects a failure. The traps inform the recipient that a FortiGate unit has failed and include the protocol for which the failure was detected.

Only the first probe to detect a failure triggers the actions on failure. So, even if multiple probes are configured, when a failure is detected, the FortiBridge unit sends one v1 SNMP trap and one v2c SNMP trap.

Configure FortiBridge SNMP by adding and configuring an SNMP community. An SNMP community is a grouping of equipment for network administration purposes. You can add up to three SNMP communities. Each community can have a different configuration for SNMP traps. You can add the IP addresses of up to 8 SNMP managers to each community.

To add and enable an SNMP community — Web-based manager

- 1 Go to *Probe > Notifications*.
- 2 Enable *SNMP*.
- 3 In the *Community Name* field, enter an SNMP community name.
- 4 Enter the address of an SNMP manager in the *IP Address* field.
You may enter up to eight SNMP manager addresses.
- 5 Select *Apply*.

To add and enable an SNMP community — CLI

```
config system snmp community
  edit 1
    set name snmp_1
  end
```

The new SNMP community, named snmp1, is enabled by default. SNMP v1 and v2 traps are also enable by default. You can disable traps and change ports. See “[system snmp community](#)” in the *FortiBridge CLI Reference* for more information.

Add the IP addresses of two SNMP managers that can receive traps.

```
config system snmp community
  edit 1
    config hosts
      edit 1
        set ip 172.20.120.12
      next
      edit 2
        set ip 192.168.20.102
      end
    end
  end
```

Recovering from a FortiGate failure

After a FortiBridge probe detects a FortiGate failure, the FortiBridge unit stops sending probes. To restart probes you can restart the FortiBridge unit, connect to the FortiBridge CLI and enter the `execute switch-mode` command, or press the mode button on the FortiBridge unit front panel.

Normally, an action on failure causes the FortiBridge unit to fail open. When the FortiBridge unit fails open, it begins operating in Bypass mode. In bypass mode the INT1 and EXT1 interfaces are directly connected and you cannot use Telnet or SSH to connect to the FortiBridge CLI. Use the following procedure to recover from bypass mode after a FortiGate failure and resume normal operation.

To resume normal operation from bypass mode

When the FortiBridge unit is operating in bypass mode, you need to do the following to resume normal operation:

- 1 Review FortiBridge alerts and check the status of your FortiGate unit and network components to determine the source of the failure.

A network component or the FortiGate unit could have experienced a general hardware failure or a specific software failure.

- 2 Make the required changes to fix the problem.

Depending on the cause, this could mean re-connecting and restarting the FortiGate unit, or diagnosing a problem with the FortiGate unit or other network component.

If all network and FortiGate unit hardware and software is functioning normally, you may have to adjust FortiBridge probe settings. See [“Tuning the failure threshold and probe interval” on page 2992](#).

- 3 Manually switch the FortiBridge unit from bypass to normal mode.

Connect to the FortiBridge CLI using the console connection and enter the command:

```
execute switch-mode
```

Or press the Mode button on the FortiBridge unit front panel.

Or restart the FortiBridge unit by cycling the power or from the console using the `execute reboot` command. The FortiBridge unit always restarts in normal mode.

Manually switching between FortiBridge operating modes

You can manually switch between FortiBridge operating modes from the FortiBridge CLI or by pressing the Mode button on the FortiBridge front panel. To switch operating modes from the CLI enter:

```
execute switch-mode
```

Backing up and restoring the FortiBridge configuration

Use the following procedures to backup and restore your FortiBridge configuration. For both of these procedures, you must have a TFTP server that you can connect to from any FortiBridge unit interface. The FortiBridge unit must be operating in normal mode.

To back up the FortiBridge configuration — Web-based manager

- 1 Go to *System > Status*.
- 2 In the System Configuration section, select the *Configuration Backup* link.
- 3 Your browser prompts you for the file name and location of the configuration file.

Using the web-based manager, the configuration backup file is saved to the computer you are using.

To back up the FortiBridge configuration — CLI

- 1 Make sure that the TFTP server is running.
- 2 Log into the FortiBridge CLI.
- 3 Backup the system configuration to a text file on the TFTP server. Enter:

```
execute backup config <filename_str> <tftp-server_ipv4>
```

The config file is copied to the TFTP server and saved with the specified file name.

To restore the FortiBridge configuration — Web-based manager


- 1 Go to *System > Status*.
- 2 In the System Configuration section, select the *Configuration Restore* link.
- 3 Select *Browse* and find the configuration backup file you want to restore.
- 4 Select *OK* to begin the restore procedure.
- 5 The FortiBridge unit reboots after loading the configuration file. While the FortiBridge unit is rebooting, all network traffic passes directly from INT1 and EXT1 bypassing the FortiGate unit.

To restore the FortiBridge configuration — CLI

- 1 Make sure that the TFTP server is running.
- 2 Log into the FortiBridge CLI.
- 3 Restore the system configuration from a text file on the TFTP server. Enter:

```
execute restore config <filename_str> <tftp-server_ipv4>
```

The config file is copied from the TFTP server to the FortiBridge unit. The FortiBridge unit reboots after loading the new configuration. While the FortiBridge unit is rebooting, all network traffic passes directly from INT1 and EXT1 bypassing the FortiGate unit.



Chapter 20 Certifications and Compliances

This FortiOS Handbook chapter contains the following sections:

[FIPS-CC operation of FortiGate units](#) describes how to install and use special FortiOS firmware builds certified to either Federal Information Processing Standards (FIPS) or Common Criteria (CC) requirements.

[Configuring FortiGate units for PCI DSS compliance](#) explains the Payment Card Industry Data Security Standard (PCI DSS). It provides information about configuring your network and FortiGate unit to help you comply with PCI DSS requirements.



FIPS-CC operation of FortiGate units

Fortinet produces special FortiOS firmware builds that are compliant with U.S. Federal Information Processing Standards (FIPS), Common Criteria (CC) security requirements, or both. These are enhanced security options for some FortiGate Unified Threat Management System models.

This section describes how to install these special builds on a FortiGate Unified Threat Management System and how to operate the unit in the FIPS-CC compliant mode. It provides information about features that differ from the standard firmware for your FortiGate unit. Installation of FIPS-CC certified firmware is required only if the unit was not ordered with this firmware pre-installed.

At the publication date of this document, the latest FIPS certified systems use firmware based on FortiOS version 4.0 and the latest Common Criteria certified systems use firmware based on FortiOS version 3.0 MR4.

This document is intended to be used by a system administrator.

This chapter contains the following sections:

- [Introduction to FIPS-CC](#)
- [Overview of Common Criteria compliant operation](#)
- [Initial configuration of the FortiGate unit](#)
- [Administration](#)
- [Firewall](#)
- [Logging](#)
- [Alarms](#)
- [Error modes](#)
- [Disabling FIPS-CC mode](#)

Introduction to FIPS-CC

Security level summary

Fortinet performs Common Criteria certifications on specific FortiOS versions in combination with specific FortiGate models. Fortinet performs FIPS 140-2 certifications on specific FortiOS versions in combination with specific FortiGate models (FIPS 140-2 Level 2 certification) and on FortiOS independent of the FortiGate hardware (FIPS 140-2 Level 1 certification). Information on Common Criteria certification is found in the relevant Security Target. Information on FIPS 140-2 certification is found in the relevant Security Policy. These documents are available on the Fortinet Support web site from the same directory where you download the firmware.

Documentation

The documentation for FortiGate units operated in FIPS-CC mode consists of this FortiOS Handbook chapter and the standard FortiGate unit documentation set for the version of FortiOS that the FIPS-CC build is based on. This documentation is available from the Fortinet Technical Documentation web site at <http://docs.forticare.com>.

Overview of Common Criteria compliant operation

Common Criteria compliant operation requires both that you use the FortiGate Unified Threat Management System in its FIPS-CC mode and that you follow secure procedures for installation and operation of the FortiGate unit. You must ensure that:

- The FortiGate unit is installed in a secure physical location.
- Physical access to the FortiGate unit is restricted to authorized operators.
- Strong password policies are enabled by default.
- Administration of the FortiGate unit is permitted using only certified administrative methods. These are:
 - console connection
 - web-based manager via HTTPS
 - command line interface (CLI) access via SSH
- The FortiGate unit can be used in either of its two operation modes: NAT/Route or Transparent. NAT/Route mode applies security features between two or more different networks (for example, between a private network and the Internet). Transparent mode applies security features at any point in a network. The current operation mode is displayed on the web-based manager Status page and in the output of the `get system status` CLI command. Also, on -equipped units, Transparent mode is indicated by “FIPS-CC-TP” and NAT/Route by “FIPS-CC-NAT” on the display.

Use of non-FIPS-CC compliant features

FIPS-CC mode does not prevent you from using non-FIPS-CC compliant features that are not permanently disabled. If you use these features, however, you are not operating the FortiGate unit in strict FIPS-CC compliance according to the Security Target or Security Policy.

Effects of FIPS-CC compliant mode

The following list describes, not necessarily in order, the effects of enabling FIPS-CC mode with respect to the normal mode of operation.

Interfaces

- Immediately after switching to FIPS-CC mode, all network interfaces are down and have no IP address assigned. Configure interfaces as needed.
- By default, admin access (except for ping access) is disabled and must be enabled on a per-interface basis.
- Network interfaces cannot be configured for HTTP or Telnet administrative access.
- NPU support is disabled by default, but can be re enabled.
- Some FortiGate models have grouped interfaces that by default operate as a switch with a single IP address. Optionally, these interfaces can be configured as individual interfaces, each with its own IP address. FIPS-CC supports both configurations.

Administration

- Administrative access via HTTPS or SSH requires strong cryptography: AES or 3DES encryption with SHA1 digest. DES encryption and MD5 digest are not available.
- By default, after three failed attempts to log on to an administrator account, the account is locked out for one hour. You can change the number of attempts permitted and the length of the lockout. See [“Administrator account lockout settings” on page 3010](#).
- Optionally, you can limit administrator access to scheduled times. See [“Scheduled administrator access” on page 3011](#).
- On a CLI session, when an administrator logs out or the session times out, the FortiGate unit sends 300 carriage return characters to clear the screen. Note: if your terminal buffer is large, not all information from the session is cleared.
- The USB auto install options are disabled.
- The control panel keys cannot be used to modify the FortiGate unit configuration.
- The FortiGate unit front panel displays “FIPS-CC-” followed by the operation mode, “NAT” or “TP”. You might have to press a panel key to deactivate the screen saver and view this display.
- The `get system status` CLI command display includes “FIPS-CC mode: enable”.
- By default, all administrators must accept a disclaimer statement at logon. This disclaimer can be disabled or modified. See [“Disclaimer access banner” on page 3010](#).
- When configuring passwords or keys, the FortiGate unit requires you to enter the password or key a second time as confirmation.
- Configuration backups use 3DES encryption with a HMAC-SHA1 checksum and a user-defined password. FIPS-CC mode backup files are not valid in non-FIPS-CC mode and vice-versa.
- The FortiGate unit performs self-tests at startup, when cryptographic keys are generated, and on a recurring basis. If any of these tests fail, the unit goes into FIPS Error mode and shuts down. The self-tests cannot be disabled, except by leaving FIPS-CC mode. Also, the administrator can run self-tests at any time. See [“Running self-tests manually” on page 3008](#).
- TFTP communication is insecure and is disabled by default. In non-FIPS-CC operation TFTP can be used to back up or restore the configuration remotely. In FIPS-CC mode, you should use a USB drive for this purpose. TFTP can be re-enabled using the `tftp` keyword in the `config system global` CLI command, but this is not FIPS-CC compliant operation.
- Remote access clients must meet security requirements. See [“Remote access requirements” on page 3009](#).
- There is an alarm capability. See [“Alarms” on page 3018](#).
- USB auto-install options are disabled.
- The `fnsysctl` command, which provides some access to the underlying operating system, is not available.
- Virus attack reporting to FortiGuard Distribution Service (FDS) is disabled.

HA

- In HA mode, HA heartbeat data is exchanged using AES encryption and SHA1 authentication. The key is automatically generated, but can be overridden by setting the `key` field to a new 16-byte hexadecimal value, like this:

```
config system ha
    set key fdoa0803e0157d4e
end
```

FortiOS will require the key value to be entered again to confirm it.

Routing

- Immediately after switching to FIPS-CC mode, no DNS addresses are configured.
- Immediately after switching to FIPS-CC mode, no default route is configured.

Logging

- Logging is enabled by default for:
 - new security policies
 - interfaces where administrative access is enabled
 - attempts to gain administration access on network interfaces where administrative access is not enabled
 - failed connection attempts to the FortiGate unit using TCP/IP ports other than 22 (ssh), 23 (telnet), 80 (HTTP), and 443 (HTTPS).
 - all configuration changes
 - configuration failures
 - remote IP lockout due to reaching maximum number of failed login attempts
 - log viewing
 - interface going up or down
 - other traffic: dropped ICMP packets, dropped invalid IP packets, session start and session deletion
- Logging is enabled for all event types at debug severity level.
- Memory logging is enabled on units that do not contain a hard disk. Logging includes traffic logging and all event types.
- Traffic logging to memory is available only in FIPS-CC mode.
- Reaching 95% of the log storage capacity results in the FortiGate unit entering an error mode that shuts down all of the interfaces until the administrator intervenes.

Firewall

- Immediately after switching to FIPS-CC mode, all security policies are removed.
- Newly-created security policies have Log allowed traffic enabled by default.
- Newly-created security policies are disabled and must be explicitly enabled.
- Blocking of spoofed TCP RST packets is enabled by default.

VPN

- The DES and MD5 algorithms are not available.
- Diffie-Hellman groups 14 through 18 are available to VPN configurations and group 15 is the default. DH groups 15 through 18 use 3072 to 8192-bit keys. You should use these groups for FIPS-CC compliant VPNs between FortiGate units. Current versions of the FortiClient Host Security application support only DH groups 1, 2 and 5.
- ANSI X9.31 RSA signature is an optional authentication method for IPsec VPNs. This method is supported by default on FortiGate units in FIPS-CC mode.
- By default, 2048-bit RSA certificates are configured, but 1024-bit certificates can be configured.

Initial configuration of the FortiGate unit

This section describes how to configure your FortiGate unit in the FIPS-CC mode of operation. Proceed as follows:

- Install the unit following the procedures in the documentation.
- Register your FortiGate unit with Fortinet.
- If you are upgrading an existing FortiGate unit to FIPS-CC firmware, download the appropriate firmware from Fortinet and install it on your unit.
- Verify the firmware version of your FortiGate unit.
- Enable FIPS-CC mode.

Installing the unit

Both the [Quick Start Guide](#) and the Getting Started section of the [Installation Guide](#) for your FortiGate unit provide instructions on the physical installation and initial configuration of your unit. When you have completed these procedures you will be able to access both the web-based manager and Command Line Interface (CLI).

Configuration of units with AMC/FMC modules

To use AMC/FMC modules, you must insert and configure them before enabling FIPS-CC mode. Modules inserted during FIPS-CC mode operation cause intermittent failures of integrity self-tests.

For more information about using AMC/FMC modules, refer to the documentation provided with your FortiGate unit.

Downloading and installing FIPS-CC compliant firmware

Unless you purchased a FortiGate unit with FIPS-CC firmware pre-installed, you need to download and install the appropriate firmware for your FortiGate unit. The Support web site provides FIPS-certified and CC-certified versions of FortiOS firmware for specific FortiGate models. Refer to the relevant CC Security Target or FIPS Security Policy document to determine which specific build you need.

To download the firmware

- 1 With your web browser, go to <https://support.fortinet.com/> and log in using the name and password you received when you registered with Fortinet Support.

- 2 Navigate to the download page for the appropriate version of FortiGate firmware and select the FIPS-CC-Certified folder. Download the FIPS-CC compliant firmware build you need. Save the file on the management computer or on your network where it is accessible from the FortiGate unit.

Installing the FIPS-CC firmware

You can install the FIPS-CC compliant firmware as an upgrade from the standard firmware.

To install the FIPS-CC firmware

- 1 Using the management computer, connect to the unit's web-based manager. See the [Quick Start Guide](#) or the [Installation Guide](#) for information.
- 2 Type `admin` in the name field. If you have assigned a password, type it in the *Password* field. Select *Login*.
- 3 Go to *Dashboard > Status*.
- 4 Under *System Information > Firmware* version, select *Update*.
- 5 Type the path and filename of the firmware image file, or select *Browse* and locate the file.
- 6 Select *OK*.

The unit uploads the firmware image file, upgrades to the new firmware version, restarts, and displays the Login page. This process takes a few minutes.

Verifying the firmware version of the unit

Execute the following command from the command line:

```
get system status
```

The version line of the status display shows the FortiGate model number, firmware version, build number and date. For example:

```
Version: Fortigate-500A 4.00,build6204,100928
```

Verify in the relevant security target or security policy document that your firmware version, build number and date are appropriate.

A note about non FIPS-CC functionality

Even when operated in non-FIPS-CC mode, the FIPS-CC firmware functionality differs in some ways from the standard FortiGate firmware on which it is based.

Access Profiles

Log & Report access is split into Log & Report Configuration and Log & Report Data. In the web-based manager, the Log & Report access control item expands to show these two access controls. In the CLI, you can independently control administrator access to logging configuration and data as follows:

```
config system accprofile
  edit <profile_name>
    set loggrp custom
    config loggrp-permission
      set config {none | read | read-write}
      set data-access {none | read | read-write}
    end
  end
```

Memory log size

Maximum memory log size is configurable from the CLI:

```
config log memory global-setting
  set max-lines <value>
end
```

Log full limit thresholds are configurable from CLI

```
config log memory global-setting
  set full-first-warning-threshold [1-100 default=75]
  set full-second-warning-threshold [1-100 default=90]
  set full-final-warning-threshold [1-100 default=95]
end
```

These threshold values are a percentage of the max-lines limit. When the first threshold is reached, an informational event is logged. When the second and final thresholds are reached, warning events are logged. The log message in each case is in the form, "Memory [log-type] log is [percentage]% full." After the final threshold is reached, new log messages overwrite the oldest log message.

Enabling FIPS-CC mode

If you have verified the firmware version, you are ready to enable FIPS-CC mode. As part of enabling FIPS-CC mode, you must define the administrator password. You must use a console connection to enable FIPS-CC mode. If you try to use another type of connection, a "check permission failed" error occurs.



When you enable FIPS-CC mode, all of the existing configuration is lost.

To enable FIPS-CC mode

- 1 Log in to the CLI using default admin account or another account with super_admin access profile. Enter the following commands:

```
config system fips-cc
  set status enable
end
```



If the FortiGate unit is currently in multi-VDOM mode, you need to precede the above commands with the command `config global`.

- 2 In response to the following prompt, enter the password for the administrator:

```
Please enter administrator password:
```

- 3 When prompted, re-enter the administrator password.

The CLI displays the following message:

```
Warning: most configuration will be lost,
do you want to continue? (y/n)
```

- 4 Enter `y`.

The FortiGate unit restarts and runs in FIPS-CC compliant mode.

Configuring interfaces

When FIPS-CC mode is initially enabled, all network interfaces are down and have no IP addresses assigned. This example shows how to configure port1 with an IP address of 192.168.0.99 and administrative access to permit use of the web-based manager.

```
config system interface
edit port1
set ip 192.168.0.99 255.255.255.0
set allowaccess https
set status up
end
```

For detailed information about configuring network interfaces, refer to the FortiGate documentation supplied with your unit.

Re-enabling NPU support

Support for NPU accelerated interfaces is disabled by default in FIPS-CC mode. The following CLI command will re enable NPU support:

```
config system npu
set offload-ipsec-host enable
end
```

FIPS-CC mode status indicators

There are two status indicators that show when the FortiGate unit is running in the FIPS-CC mode of operation:

Table 158: FIPS-CC mode status indicators

Location	Indication
Front panel (some models) (press a button to deactivate screen saver)	FIPS-CC-NAT (NAT/Route mode) FIPS-CC-TP (Transparent mode)
Output of <code>get system status</code> command	FIPS-CC mode: enable

Self-test settings

The default self-test period is every 1440 minutes. The following CLI command can change this to any period from 1 to 1440 minutes, inclusive.

```
config system fips-cc
set self-test-period <minutes_int>
end
```

Self-tests on key generation (SSL, SSH, and IPsec) are disabled by default. The following command will re-enable the tests:

```
config vpn fips-cc
set key-generation-self-test enable
end
```

Running self-tests manually

The administrator can run self-tests manually at any time. To run all of the tests, enter the following CLI command:

```
execute fips kat all
```

To run an individual test, enter `execute fips kat <test_name>`. To see the list of valid test names, enter `execute fips kat ?`

Administration

When you invoke FIPS-CC mode for the first time, the FortiGate unit prompts you for a password to assign to the administrator account. After the initial configuration of administrators when you enable FIPS-CC mode, you can create additional administrator accounts as needed.

User guidance

It is the administrator's responsibility to ensure users know how to use the user authentication functions of the FortiGate unit, as described in the Authentication chapter of the FortiOS Handbook.

Remote access requirements

In FIPS-CC mode, remote administration is not allowed via HTTP or Telnet, which are not secure. SSH and HTTPS access are permitted but must meet certain security requirements.

Setting minimum DH primes size

By default, in FIPS-CC mode the FortiGate unit requires values at least 3072 bits long to be used in the Diffie-Hellman key exchange when an SSL or HTTPS session begins. Using the CLI, you can set this minimum to any of the safe standard values specified in RFC 3526: 1024, 1536, 2048, 3072, 4096, 6144 or 8192 bits. For example, to use commercially available browsers, you might need to set the key size to 1024, like this:

```
config system global
    set dh-params 1024
end
```

Enabling administrative access

In FIPS-CC mode, the network interfaces by default do not allow administrative access, preventing you from using the web-based manager. You can re-enable use of the web-based manager using CLI commands on the console. This example enables HTTPS administrative access on the port1 interface to allow use of the web-based manager and SSH clients:

```
config system interface
    edit port1
        set allowaccess https ssh
    end
```

For detailed information about accessing the web-based manager, see "Connecting to the web-based manager" in the *Installation Guide* for your unit.

SSH client requirements

To access the CLI through network interfaces in FIPS-CC mode, your SSH client must support the following:

Authentication:

- RSA X9.31 or HMAC SHA-1

Encryption:

- AES128, AES192, AES256 or 3DES

Web browser requirements

To use the web-based manager in FIPS-CC mode, your web browser application must meet the following requirements:

- Authentication algorithm: RSA X9.31, PKCS1 RSA or DSS (in descending order of preference)
- Connection security: TLS 1.0

Disclaimer access banner

By default, in FIPS-CC mode, each time you log on as an administrator, you see a warning statement that usage is monitored and that unauthorized usage can result in disciplinary or legal action. You must accept the statement to continue. If you decline the statement, you are immediately logged out. Logs record response to the disclaimer at each logon.

You can disable the disclaimer in the CLI as follows:

```
config system global
  set access-banner disable
end
```

Similarly, you can enable the disclaimer, like this:

```
config system global
  set access-banner enable
end
```

Modifying the disclaimer text

You can modify the disclaimer to meet the requirements of your organization. The disclaimer is an editable replacement message. Using the CLI, enter the disclaimer text in the buffer field. To put a carriage return in the message, use Shift-Enter. For example,

```
config system replacemsg admin admin-disclaimer-text
  set buffer "Warning! Warning! Warning!
  This system is monitored at all times.
  Unauthorized use may be prosecuted.
  "
end
```

To restore the default message, enter

```
config system replacemsg admin admin-disclaimer-text
  unset buffer
end
```

Administrator account lockout settings

By default, after three failed attempts to log on to an administrator account, the account is locked out for one hour. The lockout applies only to the IP address from which the failed attempts were made. The login name is logged. You can change the number of logon attempts permitted and the length of the lockout using the following CLI commands:

```
config system global
  set admin-lockout-threshold <tries>
  set admin-lockout-duration <seconds>
end
```

where <tries> is permitted number of attempts, range 1 to 10 (default 3) and <seconds> is the lockout duration in seconds, range 1 to 4,294,967,295 (default 60).

The Security Administrator can clear a lockout with the following CLI command:

```
execute clear system login-lockout <index>
```

Use a ? as the index to see the list of locked-out accounts.

Scheduled administrator access

For additional security, you can limit administrator access to certain times, business days for example. To do this, you need to create a firewall schedule and then assign the schedule to the administrator.

You can create a firewall schedule in the web-based manager or the CLI. For more information, refer to the documentation provided with your FortiGate unit.

To assign a schedule to an administrator, enter the following CLI commands:

```
config system admin
  edit <admin-name>
    set schedule <schedule-name>
  end
```

where <admin-name> is the name of the administrator account and <schedule-name> is the name of the firewall schedule.

Using custom administrator access keys (certificates)

You, as cryptographic administrator, can upload VPN certificates to use as custom RSA keys to authenticate administrators. To do this, you must upload the signed public certificate to the FortiGate unit. If the private key was not generated on the FortiGate unit, it also must be uploaded. Certificates must have a modulus of at least 2048 bits.

Importing the custom RSA key

In FIPS-CC mode, you cannot import certificates using TFTP. You must use a USB storage device instead. Put the files you want to upload on the device and connect the device to the FortiGate unit. Use one of the following commands to import the files:

If you have a PKCS12 format key-certificate file,

```
execute vpn certificate local import usb pkcs12 <file_name>
  <password>
```

If you have separate certificate and key files,

```
execute vpn certificate local import usb cert <cert_file_name>
  <keyfile_name> <password_for_keyfile>
```

Enabling the custom RSA key

To enable the custom RSA key you imported, use the following CLI command:

```
config system global
  set admin-server-cert <certificate-name>
end
```

This applies to both HTTPS and SSH connections.

Configuration backup

Configuration backup files created in FIPS-CC mode are not compatible with backup files created in non-FIPS-CC mode. A FIPS-CC mode configuration backup cannot be restored in non-FIPS-CC mode and vice-versa.

You can create FIPS-CC configuration backup files to use for disaster recovery. They are valid on a replacement FortiGate unit or to restore configuration after you exit and then re-enter FIPS-CC mode.

For detailed information about creating configuration backup files, refer to the documentation provided with your FortiGate unit.



Configuration backup or restoration using TFTP is not permitted in FIPS-CC mode.

Firewall

FIPS-CC mode has additional requirements for security policies and firewall authentication, compared to the standard firmware.

Security policies

When you create a security policy in FIPS-CC mode, by default the policy is not enabled. You must explicitly enable it. In the web-based manager, after creating the policy, select the checkbox at the beginning of the policy entry on the **Policy > Policy** page. In the CLI, enable a policy by setting its status to enable. You can do this when you create the policy or later:

```
config firewall policy
  edit 2
    set status enable
  end
```

Policies are identified by policy ID. In the preceding example, the ID was 2.

Firewall authentication

In FIPS-CC mode, user passwords must be 8 characters or more. FTP and Telnet mechanisms for Proxy User Authentication are not allowed, and SSL redirection must be enabled for the HTTP mechanism.

User account lockout settings

Optionally, you can lock out a user's account for a period of time after a number of unsuccessful attempts to authenticate. You can configure this in the CLI using the following commands:

```
config system global
  set auth-lockout-threshold <tries>
  set auth-lockout-duration <seconds>
end
```

where <tries> is permitted number of attempts, range 1 to 10 (default 3) and <seconds> is the lockout duration in seconds, range 1 to 4,294,967,295, or 0 to disable lockout. The default is 0.

This lockout applies to end users only. Administrator lockout is configured separately. See [“Administrator account lockout settings” on page 3010](#).

Logging

The Common Criteria protection profile requires logging of all traffic and logging of system events, including startup and shutdown of functional components. The severity threshold for logging is set to the lowest level: debug. This ensures that the maximum amount of information is logged.

Logs are written to the FortiGate unit hard disk on all models except model 5001, which contains a flash memory drive and models 50A and 100A that log to system memory.



Traffic logging to system memory is available only in FIPS-CC mode.

The FortiGate unit generates warning log entries when the space allocated for logging is filled to 75%, then 90% and finally 95% of capacity. For information about setting the log size, see [“Memory log size” on page 3007](#). When logs exceed 95% of capacity, the default action is to block further traffic and switch to Error mode. See [“CC Error mode” on page 3022](#) for more information.

Logging to external devices

Logging to external devices is disabled due to the security requirements of FIPS-CC operation, except for logging to a FortiAnalyzer unit through a secure tunnel. By default, the secure tunnel uses the SHA-256 HMAC algorithm. You can also select SHA-1 using the following CLI command:

```
config log fortianalyzer setting
    set hmac-algorithm sha1
end
```

Downloading of logs to the management computer is also permitted. See [“Backing up log messages” on page 3017](#).

Required logging settings

[Table 160](#) and [Table 161](#) list the logging settings required for FIPS-CC mode. The `config log memory setting` command settings apply to models 50A and 100A. The `config log disk setting` command settings apply to all other models. If you change these options from the default, the operation of your FortiGate unit is no longer compliant with the FIPS-CC Security Target.

Table 159: config log disk filter command keywords and variables

Keywords and variables	Description	Default
admin {disable enable}	Enable or disable logging all administrative events, such as user logins, resets, and configuration updates in the event log. This is available only if <code>event</code> is set to <code>enable</code> .	enable
allowed {disable enable}	Enable or disable logging all traffic that is allowed according to the firewall policy settings in the traffic log. This is available only if <code>traffic</code> is set to <code>enable</code> .	enable

Table 159: config log disk filter command keywords and variables (Continued)

Keywords and variables	Description	Default
auth {disable enable}	Enable or disable logging all firewall-related events, such as user authentication in the event log. This is available only if <code>event</code> is set to <code>enable</code> .	enable
event {disable enable}	Enable or disable the event log.	enable
severity {alert critical debug emergency error information notification warning}	Select the logging severity level. The FortiGate unit logs all messages at and above the logging severity level you select. For example, if you select <code>error</code> , the unit logs <code>error</code> , <code>critical</code> , <code>alert</code> and <code>emergency</code> level messages. emergency - The system is unusable. alert - Immediate action is required. critical - Functionality is affected. error - An erroneous condition exists and functionality is probably affected. warning - Functionality might be affected. notification - Information about normal events. information - General information about system operations. debug - Information used for diagnosing or debugging the FortiGate unit.	debug
system {disable enable}	Enable or disable logging of all system-related events, such as ping server failure and gateway status, in the event log. This is available only if <code>event</code> is set to <code>enable</code> .	enable
traffic {disable enable}	Enable or disable the traffic log.	enable
violation {disable enable}	Enable or disable logging of all traffic that violates the firewall policy settings in the traffic log. This is available only if <code>traffic</code> is set to <code>enable</code> .	enable

Table 160: config log disk setting command keywords and variables

Keywords and variables	Description	Default
diskfull {blocktraffic nolog overwrite}	Enter the action to take when the log disk is full.	blocktraffic

Table 161: config log memory setting command keywords and variables

Keywords and variables	Description	Default
diskfull {blocktraffic nolog overwrite}	Enter the action to take when the log memory is full.	blocktraffic
status {disable enable}	Enter <i>enable</i> to enable logging to the FortiGate system memory.	enable

Excluding specific logs (selective audit)

Use the exclude-list option of the log filtering command to define log entries that will not be recorded:

```
config log disk filter
  config exclude-list
    edit <number>
      set category <category>
      config fields
        edit <field_name>
          set args <argvalue>
          set negate {enable | disable}
        end
      end
    end
  end
end
```

Table 162: log filter exclude-list command keywords and variables

Keywords and variables	Description	Default
category <category>	category is one of: attack, content, event, im, spam, traffic, virus, webfilter	No default.
edit <field_name>	Enter the name of the field on which to base exclusion. Available <i>field_name</i> values depend on the category setting. If you enter an invalid field name, valid field names are listed.	No default.
args <argvalue>	Enter the field value to match.	No default.

Table 162: log filter exclude-list command keywords and variables (Continued)

Keywords and variables	Description	Default
negate {enable disable}	Enable to exclude logs where the value of the <field_name> field does not match <argvalue>. Disable to exclude logs where the value of the <field_name> field matches <argvalue>.	disable

Viewing log messages from the web-based manager

To view log messages from the web-based manager, go to *Log & Report > Log Access*. For detailed instructions about viewing the logs, consult the online Help system or see the Logging & Reporting chapter of the *FortiOS Handbook*.

Viewing log messages from the CLI

You can view and clear log messages from the CLI. Before viewing logs, you must set filter options to select the logs that you want to view. You can view one log category on one device at a time. Optionally, you can filter the listing to show only specified date ranges or severities of log messages. For traffic logs, you can filter log messages by source or destination IP address.

Setting filtering for log messages

Use `execute log filter` commands to select which logs to display with the `execute log display` command. Commands are cumulative. Enter `execute log filter list` to see the current settings. For more information about log filtering, see the *FortiGate CLI Reference*.

The command syntax is:

```
execute log filter <keyword> <variable>
```

Table 163: execute log filter command keywords and variables

Keywords and variables	Description	Default
category {event ids spam traffic virus webfilter list }	Type of log, except <code>list</code> , which displays the current setting.	event
device {disk memory list}	Device where the logs are stored, except <code>list</code> , which displays the current setting.	disk
field <field_name>	Filter on log field. Use ? as <code>field_name</code> to see a list of valid fields.	No default.
lines-per-view <number>	Set lines per view. Range: 5 to 1000	10
list	Display current filter settings.	No default.
reset	Reset filter settings.	No default.
rolled-number <number>	Select logs from rolled log files. 0 selects current log file.	0

Table 163: execute log filter command keywords and variables (Continued)

Keywords and variables	Description	Default
sortby	Set display order. See “Sorting log messages” on page 3017 .	No default.
start_line <integer>	Select first line of logs to display.	1

Use as many `execute log filter` commands as you need to define the log messages that you want to view. For example, to select the memory event logs from 10-14 July 2006, you use the following commands:

```
execute log filter category event
execute log filter device memory
execute log filter field date 2006-07-10 2006-07-14
```

Sorting log messages

In addition to selecting logs to display, the `execute log filter` command can sort logs by field.

```
execute log filter sortby <field_name>
```

Enter the command without a field name to see a list of valid field names.

Viewing log messages

After you have selected the log messages that you want to view using the `execute log filter` command, you can display them with the following command:

```
execute log display
```

The console displays the first 10 log messages. To view more messages, run the command again. You can do this until you have seen all of the selected log messages. To restart viewing the list from the beginning, use the commands

```
execute log filter start_line 1
execute log display
```

Resetting log filters

You can restore the log filters to their default values using the command

```
execute log reset
```

Backing up log messages

You can back up log messages to your Administrative computer or other computer on the network.

Backing up log messages using the web-based manager

The FortiGate unit downloads log files to the Administrative computer using HTTPS.

- 1 Go to *Log & Report > Log Access*.
- 2 Select either the *Disk* or *Memory* tab as appropriate.
- 3 From the *Log Type* list, select the type of log you want to back up.
- 4 Select the download icon for the log file you want to back up.
- 5 Select either *Download file in the normal format* or *Download file in CSV format*, as appropriate.
- 6 Follow your browser's procedure for saving the downloaded file.

Viewing log file information

You can view the list of current and rolled log files on the console. The list shows the file name, size and timestamp. The CLI command is as follows:

```
execute log list <category>
```

<category> must be one of: event, ids, spam, traffic, virus or webfilter.

The output looks like this:

```
elog                8704      Fri Jan 28 14:24:35 2005
elog.1              1536      Thu Jan 27 18:02:51 2005
elog.2              35840     Wed Jan 26 22:22:47 2005
```

At the end of the list the total number of files in the category is displayed. For example:

```
501 event log file(s) found.
```

Deleting filtered log messages

You can select log messages with the `execute log filter` command and then delete them with the `execute log delete-filtered` command. On units that provide only memory logging, be sure to specify memory as the log device.

For example, to delete all the traffic logs from memory, enter the following commands:

```
execute log filter category traffic
execute log filter device memory
execute log delete-filtered
```

For information about the `execute log filter` command, see [“Setting filtering for log messages” on page 3016](#).

Deleting rolled log files

You can delete rolled log files using the `execute log delete-rolled` command:

```
execute log delete-rolled <category> <start> [<end>]
```

<category> must be one of: event, ids, spam, traffic, virus or webfilter. The <start> and <end> values represent the range of log files to delete. If <end> is not specified, only the log number specified by <start> is deleted.

For example, to delete all of the rolled traffic log files, enter the following command:

```
execute log delete-rolled traffic 1 9999
```

Alarms

In FIPS-CC mode, the FortiGate unit can raise alarms for the following types of events:

- failed administrator authentication
- packet replay attempts (IPS)
- bootup self-test failure
- cryptographic failure
- security policy violation (blocked sessions)

Alarms for all of these events are based on logs that report these events. For firewall events, traffic violation logs are used.

Configuring alarms

An alarm consists of one or more trigger events that occur a specified number of times in a particular time period. For example, you could configure the FortiGate unit to raise an alarm if there are three unsuccessful administrative login attempts in the same minute.

You can configure alarms only in the CLI. Each alarm is defined as an “alarm group”. There are separate alarm groups for each virtual domain (VDOM). You can select whether the alarms in each VDOM are audible. Within each alarm group, you specify:

- the threshold for each triggering event, 0 for events that will not trigger the alarm
- the period over which the number of triggering events is counted, 0 to count from startup

If you include more than one trigger event, the threshold for all the trigger events must be met to trigger the alarm. Security policy violations are configured together.

Alarm notification messages appear on the web-based manager, in SSH administrator sessions and on the console. The messages repeat until the administrator acknowledges the alarm.

Alarm CLI configuration

The `system alarm` command syntax is as follows:

```
config system alarm
  set status {enable | disable}
  set audible {enable | disable}
  config groups
    edit <group_id>
      set admin-auth-failure-threshold <integer>
      set decryption-failure-threshold <integer>
      set encryption-failure-threshold <integer>
      set log-full-warning-threshold <integer>
      set period <integer>
      set replay-attempt-threshold <integer>
      set self-test-failure-threshold <integer>
      set user-auth-failure-threshold <integer>
    config fw-policy-violations
      edit <violation_id>
        dst-port <dport_number>
        dst-ip <dst_ip>
        src-port <sport_number>
        src-ip <src_ip>
        threshold <integer>
      end
    end
  end
```

The keywords and variables are:

Table 164: system alarm keywords and variables

Keywords and variables	Description	Default
audible {enable disable}	If enabled, the console beeps when the alarm notification appears.	disable
status {enable disable}	Enable or disable all alarms.	disable

Table 164: system alarm keywords and variables

Keywords and variables	Description	Default
Alarm group keywords and variables		
edit <group_id>	<group_id> is the group identifier. Use 0 to automatically assign the next available number.	No default.
admin-auth-failure-threshold <integer>	Enter threshold for administrator authentication failures. Use 0 to disregard in this alarm group.	0
decryption-failure-threshold <integer>	Enter threshold for cryptographic failure in decryption. Use 0 to disregard in this alarm group.	0
encryption-failure-threshold <integer>	Enter threshold for cryptographic failure in encryption. Use 0 to disregard in this alarm group.	0
log-full-warning-threshold <integer>	Enter the threshold for log full warnings. Use 0 to disregard in this alarm group.	0
period <integer>	Enter the period over which triggering events are counted. Use 0 for no limit (events are counted from startup).	0
replay-attempt-threshold <integer>	Enter threshold for packet replay attempts. Use 0 to disregard in this alarm group.	0
self-test-failure-threshold <integer>	Enter threshold for failure of startup integrity tests. Use 0 to disregard in this alarm group. A self-test failure alarm is visible only after you recover from Error mode.	0
user-auth-failure-threshold <integer>	Enter threshold for user authentication failures. Use 0 to disregard in this alarm group.	0
fw-policy-violations keywords and variables		
edit <violation_id>	<violation_id> is the identifier for this trigger. Use 0 to automatically assign the next available number.	No default.
dst-port <dport_number>	Enter the destination port number to match in the traffic violation log.	0
dst-ip <dst_ip>	Enter the destination IP or subnet address to match in the traffic violation log.	0
src-port <sport_number>	Enter the source port number to match in the traffic violation log.	0
src-ip <src_ip>	Enter the source IP or subnet address to match in the traffic violation log.	0

Alarm notifications

Alarm notifications appear on both the CLI console and the web-based manager. On the CLI console alarm notifications look like this:

```
***** !!! A L A R M !!! *****
* ID: 1                               Time: Tue Sep 5 09:39:55 2006
* Group ID: 1                         VD: root
* Type: Authentication failures
* Message: Alarm is triggered
*****
```

On the web-based manager, alarm notifications appear in a separate browser window and look like this:

Figure 388: Alarm notification - web-based manager

Alarms				
ID	Time	Group ID	Type	Message
2	2006-08-01 04:40:23	1	Authentication failures	Alarm is triggered
Acknowledgements				
ID	Time	Ack by	Ack from	Message
No new acknowledgements				
OK				

The notification clearly shows the time, virtual domain, alarm group and type. Alarm notifications repeat until you acknowledge them. On the CLI console, the notification repeats every time you use the Enter key. In the web-based manager, you can close the alarm notification window, but the alarm will reappear in a few seconds.

Acknowledging alarms

To acknowledge an alarm in the web-based manager, you simply select OK in the alarm notification window. On the CLI console, you acknowledge alarms using one of the following commands:

To acknowledge a single alarm

```
execute ack-alarm <alarm-ID>
```

To acknowledge all alarms

```
execute ack-alarm all
```

Alarm polling

A terminal connected to the console connector can display alarm messages periodically if no one is logged in to the console. You can set how often alarm messages are reported on the console.

```
config system global
  set alarm-poll-interval <second>
end
```

You can set the polling interval to a value from 1 to 60 seconds. The default is 5 seconds.

Error modes

There are two error modes in FIPS-CC mode: FIPS Error and CC Error.

FIPS Error mode

When one or more of the self-tests fail, the FortiGate unit switches to FIPS Error mode. The FortiGate unit shuts down all interfaces including the console and blocks traffic.

To resume normal FIPS-CC mode operation, switch the unit off and then on again. If the self-tests pass after the reboot, the unit will resume normal FIPS-CC compliant operation. If a self-test continues to fail after rebooting, there is likely a serious firmware or hardware problem and the unit should be removed from the network until the problem is solved.

If the self-test failure persists across reboots, you can attempt to reload the firmware after resetting the unit to the factory default configuration. If the self-test failure persists after reloading the firmware and re-enabling the FIPS-CC mode of operation, contact Fortinet technical support.

CC Error mode

When current logs and rolled log files consume more than 95% of log capacity, the FortiGate unit switches to CC Error mode, shuts down network interfaces and blocks traffic.

The FortiGate unit indicates Error mode as follows:

- The console displays “FIPS-CC-ERR”. You might have to press a panel key to see this display.
- “CC-ERR” is prepended to the CLI prompt, `CC-ERR FortiGate-500A$`, for example.

To resume normal FIPS-CC mode operation, you first must reduce logs to less than 95% of device capacity and exit error mode.

To reduce logs

From the console, do any of the following:

- Delete selected logs. See [“Deleting filtered log messages” on page 3018](#). Ideally, you should reduce logs to 50% or less of device capacity.
- Delete rolled log files using the command
`execute log delete-rolled.`
- Delete all current log entries using the command
`execute log delete-all.`

To exit error mode

From the console, enter the following CLI command:

```
execute error-mode exit
```

The FortiGate unit resumes normal FIPS-CC compliant operation unless there is still too little free space on the log device.

Disabling FIPS-CC mode

The only way that you can return the FortiGate unit to the normal mode of operation is to restore the factory default configuration. Enter the following CLI command:

```
execute factoryreset
```

Disabling FIPS-CC mode erases the current configuration, including VPN certificates and encryption keys for SSH and HTTPS.



Configuring FortiGate units for PCI DSS compliance

This chapter provides information about configuring your network and FortiGate unit to help you comply with PCI DSS requirements. The following topics are included in this section:

- [Introduction to PCI DSS](#)
- [Network topology](#)
- [Security policies for the CDE network](#)
- [Wireless network security](#)
- [Protecting stored cardholder data](#)
- [Protecting communicated cardholder data](#)
- [Protecting the CDE network from viruses](#)
- [Monitoring the network for vulnerabilities](#)
- [Restricting access to cardholder data](#)
- [Controlling access to the CDE network](#)

Introduction to PCI DSS

The primary source of information for your PCI DSS compliance program is the [Payment Card Industry \(PCI\) Data Security Standard](#) itself. Version 1.2.1 of the standard was published in July 2009. The following is only a brief summary of PCI DSS.

What is PCI DSS?

The Payment Card Industry Data Security Standard (PCI DSS) sets data handling requirements for organizations that hold, process, or exchange cardholder information.

What is the Customer Data Environment

Throughout the PCI DSS requirements, there are references to the Customer Data Environment (CDE). The CDE is the computer environment wherein cardholder data is transferred, processed, or stored, and any networks or devices directly connected to that environment.

PCI DSS objectives and requirements

PCI DSS consists of 7 control objectives and 12 requirements.

Table 165: PCI DSS Control Objectives and Requirements

Control Objective	Requirement	Fortinet Solution
Build and Maintain a Secure Network	1) Install and maintain a firewall configuration to protect cardholder data	FortiGate firewall functionality. See “Security policies for the CDE network” on page 3028.
	2) Do not use vendor-supplied defaults for system passwords and other security parameters	FortiDB vulnerability assessment and auditing FortiScan OS vulnerability management FortiWeb web application password checking See “Password complexity and change requirements” on page 3034.
Protect Cardholder Data	3) Protect stored cardholder data	FortiDB vulnerability assessment and monitoring FortiWeb web application firewall See “Protecting stored cardholder data” on page 3031.
	4) Encrypt transmission of cardholder data across open, public networks	FortiGate IPsec VPN. See “Protecting communicated cardholder data” on page 3031.

Table 165: PCI DSS Control Objectives and Requirements (Continued)

Control Objective	Requirement	Fortinet Solution
Maintain a Vulnerability Management Program	5) Use and regularly update anti-virus software	FortiGate integrated AV FortiClient integrated AV FortiMobile integrated AV FortiMail integrated AV FortiGuard automated AV updates See “Protecting the CDE network from viruses” on page 3032.
	6) Develop and maintain secure systems and applications	FortiDB vulnerability assessment, auditing and monitoring FortiWeb web application security FortiScan OS vulnerability management FortiAnalyzer network vulnerability scanning See “Monitoring the network for vulnerabilities” on page 3033.
Implement Strong Access Control Measures	7) Restrict access to cardholder data by business need-to-know	FortiDB vulnerability assessment, auditing and monitoring. See “Restricting access to cardholder data” on page 3034.
	8) Assign a unique ID to each person with computer access	FortiGate integrated database or hooks to Active Directory. See “Controlling access to the CDE network” on page 3034.
	9) Restrict physical access to cardholder data	Fortinet professional services in partnership with partner solutions

Table 165: PCI DSS Control Objectives and Requirements (Continued)

Control Objective	Requirement	Fortinet Solution
Regularly Monitor and Test Networks	10) Track and monitor all access to network resources and cardholder data	FortiDB auditing and monitoring FortiAnalyzer event reporting, vulnerability scanning. See “Monitoring the network for vulnerabilities” on page 3033 .
	11) Regularly test security systems and processes	FortiDB vulnerability assessment FortiScan OS vulnerability management. See “Monitoring the network for vulnerabilities” on page 3033 .
Maintain an Information Security Policy	12) Maintain a policy that addresses information security	FortiManager security policy management appliance

This chapter describes how the FortiGate unit’s features can help your organization to be compliant with PCI DSS. Requirements that the FortiGate cannot enforce need to be met through organization policies with some means determined for auditing compliance.

Be sure to read the section, [“Wireless guidelines”](#), below. Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that wireless networking has not been introduced into the CDE.

Wireless guidelines

While wired networks usually connect fixed known workstations, wireless networks are more dynamic, introducing a different set of security concerns.

Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that unauthorized wireless networking has not been introduced into the CDE. Wireless networking could be introduced quite casually by adding a wireless device to a PC on the CDE network.

For all PCI DSS networks, whether they use wireless technology or not, the following requirement applies:

- Test for the presence of wireless access points by using a wireless analyzer at least quarterly or deploying a wireless IDS/IPS to identify all wireless devices in use. (11.1)

If your organization uses wireless networking outside the CDE network and the firewall prevents communication with the CDE network, the wireless network is outside the PCI DSS scope, but the firewall configuration must meet PCI DSS requirements.

If your organization uses wireless networking inside the CDE network, the wireless network is within the PCI DSS scope. For information about wireless network requirements, see [“Wireless network security” on page 3029](#).


Network topology

The cardholder data environment must be protected against unauthorized access from the Internet and from other networks in your organization. FortiGate unit firewall functionality provides tight control over the traffic that can pass between the following network interfaces:

- Internet
- CDE wired LAN
- CDE wireless LAN
- Other internal networks

Figure 389 shows how the Customer Data Environment can be delineated in a typical network.

Figure 389: Enterprise network with a customer data environment



Internet

The FortiGate unit has at least one network interface connected to the Internet. If your organization uses more than one Internet service provider, there could be additional network interfaces that function as a route to the Internet.

The CDE wired LAN

The CDE network typically contains point-of-sale (POS) terminals, databases, and servers. The only security policies between the CDE network and the Internet should be for encrypted connections. For remote point-of-sale terminals or off-site databases, VPN connections are required and they should use strong encryption. For a web server that handles online purchases, only HTTPS (SSL or TLS) connections can be permitted. The security policies that enable these connections should have the narrowest possible definitions for source address, destination address and service.

PCI DSS does not require the CDE network to be isolated from the rest of your corporate LAN. But isolating the CDE network reduces the scope of required data protection measures and may reduce the scope of PCI DSS assessments that are periodically required.

The CDE wireless LAN

Wireless networking is a special issue. Even if you do not use wireless technology you must monitor to ensure that unauthorized wireless access has not been added to the CDE network. For this purpose, [Figure 389](#) shows a FortiAP device in the CDE. The FortiAP device can provide dedicated wireless monitoring, an access point, or both.

A small retail outlet could reduce costs by using a FortiWiFi unit, a FortiGate unit with integrated wireless networking. The FortiWiFi unit would have to be located where it could provide sufficient wireless monitoring (or access point) coverage for the entire premises.

Other internal networks

Other internal networks such as your office LAN, unless they provide connection to the CDE, are not subject to PCI DSS requirements.

Security policies for the CDE network

The FortiGate unit's firewall functionality is ideally suited to PCI DSS requirement 1.2.1, "Restrict inbound and outbound traffic to that which is necessary for the cardholder data environment." Security policies control the source, destination, and type of traffic passing between networks.

The PCI DSS standard includes requirements to document your network topology and configuration. As part of that requirement, and to assist the auditing of your network, make use of the *Comment* field available in FortiGate security policies. Describe the purpose of each policy.

Controlling the source and destination of traffic

The source and destination are the first parameters you specify in a security policy. (Go to *Policy > Policy > Policy* and select *Create New*.)

Source Interface/Zone	port4	
Source Address	Branch_Office	Multiple
Destination Interface/Zone	port2	
Destination Address	Server1	Multiple

The *Interface/Zone* settings depend on network topology. The Address settings define the IP addresses to which the policy applies. These should be as narrow as possible, so that only the appropriate hosts are included. For example, if the destination is a server with a single IP address, the named Destination Address should be defined as that single address, not the entire subnet on which the server resides.

Addresses are defined in *Firewall Objects > Address > Address*. You can also define a new address by selecting [Create New...] from either the *Source Address* or *Destination Address* drop-down lists in a security policy. Some addresses will be used in several security policies, so it is best to plan ahead and define the addresses first.

Controlling the types of traffic in the CDE

The *Service* setting in each security policy determines which types of traffic can pass based on protocol.

Service ANY Multiple

You can select a single protocol from the *Service* drop-down list or select *Multiple* and create a list of services to permit in this policy. If several security policies will need the same list of services, consider creating a named service group. (Go to *Firewall Objects > Service > Group*.) In the security policy, service groups are available at the bottom of the *Service* drop-down list.

The default deny policy

All traffic not specifically allowed by a security policy that you create is blocked by the Implicit policy listed at the bottom of the *Policy > Policy > Policy* page.

Implicit (1)						
<input type="checkbox"/>	all	all	always	ANY	DENY	Implicit

You cannot delete this policy and you can edit the policy only to enable or disable logging of the traffic that it handles.

Wireless network security

Scanning for rogue access points is the minimum requirement for wireless security. Even if your organization does not use wireless networking, PCI DSS requires you to verify periodically that wireless networking has not been introduced into the CDE.

If you use wireless networking, the wireless network is only within the PCI DSS scope if it can connect to the CDE.

Scanning for rogue access points

A FortiGate unit with a connected FortiAP unit can perform wireless scanning. Each of the FortiAP radios can act as a dedicated monitor or can perform scanning in the background while acting as a wireless access point.

▼ Radio 1

Mode ☐ Disable ☐ Access Point ☒ Dedicated Monitor

▼ Radio 2

Mode ☐ Disable ☒ Access Point ☐ Dedicated Monitor

Background Scan ☒

Band 802.11n 5G

To configure rogue AP scanning

- 1 Go to *Wifi Controller > Configuration > AP Profile*.
- 2 Select an existing *AP Profile* and edit it, or select *Create New*.
- 3 For each radio, select either *Access Point* or *Dedicated Monitor*, as required.
- 4 If you selected *Access Point*, enable *Background Scan*.
- 5 If needed, modify other settings.
- 6 Select *OK*.

Radio 1 operates in the 2.4GHz band and Radio 2 operates in the 5GHz band. Both bands should be monitored. The FortiAP unit(s) used for scanning must be located within the coverage area that would result if an access point were added to the CDE.

Automatic detection of rogue APs

Some FortiGate units include an “on-wire” detection technique that correlates the SSID MAC addresses of the unknown access points with MAC addresses detected on your wired networks. This helps to differentiate unrelated neighboring APs from security-compromising unauthorized APs connected to your network.

Viewing the results of rogue AP scanning

Go to *Wifi Controller > Monitor > Rogue AP* to view information about detected wireless access points.

Logging the results of rogue AP scanning

To ensure that detection of rogue access points is logged, go to *Log&Report > Log Config > Log Setting* and enable logging for *Wifi activity event*.

In the logs, the *Type* is event and the *Sub Type* is wireless.

Securing a CDE network WAP

If your wireless network is within PCI DSS scope, it must meet the following requirements:

- Default settings such as SSID and passphrases must be changed.
- Use WPA security, not WEP.
- Log wireless activity.

Setting wireless security

On FortiGate units, go to *Wifi Controller > WiFi Network > SSID* to configure wireless security settings for either a new or existing virtual access point.

WiFi Settings

SSID: fortinet

Enable DHCP: ☐

Security Mode: WPA/WPA2-Enterprise

Data Encryption: ☒ AES ☐ TKIP

Authentication: ☐ RADIUS Server ☒ Usergroup

Guest-group: Guest-group

The default SSID for the FortiAP is “fortinet”. You must change this.

The *Security Mode* **must** be set to one of the WPA/WPA2 modes. Both WPA or WPA2 clients can be served. In the CLI, you can optionally select exclusively WPA or WPA2 operation.

AES is stronger *Data Encryption* than TKIP.

WPA/WPA2-Enterprise *Authentication* uses separate logon credentials for each user. Either FortiGate user group security or an external RADIUS server performs the authentication. Optionally, certificate-based security can also be applied. WPA/WPA2-Personal authentication requires a single pre-shared key that is used by all clients and is thus less secure.

For detailed information about wireless access points, see the *Deploying Wireless Networks* chapter of this Handbook.

Logging wireless network activity

To ensure that wireless network activity is logged, go to *Log&Report > Log Config > Log Setting* and enable logging for *WiFi activity event*. In the logs, the *Type* is event and the *Sub Type* is wireless.

Protecting stored cardholder data

The Fortinet FortiDB and FortiWeb products can provide security for your sensitive cardholder data.

The Fortinet Database Security (FortiDB) device provides vulnerability assessment, database activity monitoring, auditing and monitoring. For more information about this product, see the Fortinet web site, www.fortinet.com.

The Fortinet FortiWeb Web Application Firewall deployed in front of public-facing web applications protects Web applications, databases, and the information exchanged between them. In particular, it addresses the PCI DSS requirements 6.5 and 6.6 regarding web application vulnerabilities such as cross-site scripting, SQL injection, and information leakage. For more information about this product, see the Fortinet web site, www.fortinet.com.

Protecting communicated cardholder data

If cardholder data must be communicated over an untrusted network, such as the Internet, use the FortiGate unit's IPsec VPN capability to exchange the data securely. If you support customer on-line transactions, use HTTPS (SSL or TLS encryption) for security. The relevant PCI DSS requirement is:

- Use strong cryptography and security protocols such as SSL/TLS or IPsec to safeguard sensitive cardholder data during transmission over open, public networks. (4.1)



This does not prescribe particular cryptography, but it can be interpreted as a requirement to follow industry best practices.

Configuring IPsec VPN security

The security considerations for IPsec VPNs are encryption and authentication.

Encryption

Go to *VPN > IPsec > Auto Key (IKE)* to configure an IPsec VPN. In both Phase 1 and Phase 2 parts of the configuration, you select the encryption to use.

1 - Encryption Authentication
 2 - Encryption Authentication  

These are advanced settings, overriding defaults that are not necessarily the strongest algorithms. VPNs negotiate over standards, so you can list multiple proposed algorithms. The VPN will use the strongest encryption that both ends support.

Choose strong encryption. The available encryption algorithms in descending order of strength are AES256, AES192, AES128, 3DES, DES. DES encryption is the weakest with only a 64-bit key and does not meet the 80-bit key length minimum that PCI DSS requires. NULL means no encryption and must not be used.

The message digest (authentication) algorithms in descending order of strength are SHA256, SHA1 and MD5. MD5 is particularly weak and should be avoided. NULL means no message digest and must not be used.

Authentication

VPN peers authenticate each other before establishing a tunnel. FortiGate units support two different authentication methods: pre-shared key and RSA signature (certificate). Certificates provide the best security. PCI DSS does not prohibit pre-shared keys, but you should limit access to the keys to the personnel who are responsible for the FortiGate units or other equipment at either end of the VPN.

Configuring SSL VPN security

The SSL VPN configuration includes a choice of encryption algorithm. Go to *VPN > SSL > Config*. The *Default* selection, *RC4 (128 bits)* is acceptable, but the *High* option, *AES (128/256 bits)* and *3DES* is more secure. The *Low* option, *RC4 (64 bits)*, *DES* and *higher* does not meet PCI DSS requirements.

Protecting the CDE network from viruses

PCI DSS requires the use of regularly updated antivirus protection. The antivirus functionality of the FortiGate unit protects both the FortiGate unit and the networks it manages. Workstations on these networks can be protected using FortiClient Endpoint Security. Both FortiGate and FortiClient antivirus protection can receive updates from Fortinet's FortiGuard service. Workstations can also use third-party antivirus applications with update services.

The FortiGate unit can enforce the use of antivirus software, denying unprotected workstations access to the network.

Enabling FortiGate antivirus protection

To create the antivirus profile

- 1 Go to *UTM Profiles > Antivirus > Profile*.
- 2 Edit the *default* predefined profile or select *Create New*.

	Web		Email						File Transfer		
	HTTP	HTTPS	SMTP	SMTPS	POP3	POP3S	IMAP	IMAPS	FTP	FTPS	IM
Virus Scan and Removal	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

- 3 Ensure that all check boxes are selected.
- 4 Select *OK*.

To select the antivirus database

- 1 Go to *UTM Profiles > Antivirus > Virus Database*.
- 2 Select the *Extended Virus Database*.
- 3 Select *Enable Grayware Detection*.
- 4 Select *Apply*.

For detailed information about the Antivirus feature, see the *UTM* chapter of this Handbook.

Configuring antivirus updates

On the system dashboard, check the *License Information* widget. The *Support Contract* section should show Valid Contract and the contract expiry date. If your FortiGate unit is not registered, you need to visit the Fortinet Support web page (<http://support.fortinet.com/>) to register. Go to Product Registration and follow the instructions.

In the *FortiGuard Services* section, check the *Antivirus* field. If the service is unreachable, see the online Help for information about troubleshooting your connectivity to FortiGuard Services.

Enforcing firewall use on endpoint PCs

PCI DSS requires you to “install personal firewall software on any mobile and/or employee-owned computers with direct connectivity to the Internet (for example, laptops used by employees), which are used to access the organization’s network. (1.4)” Consider using the Endpoint Control feature of the FortiGate unit to enforce use of this software.

Monitoring the network for vulnerabilities

There are several tools that can assist you in monitoring your network for vulnerabilities and provide evidence to the PCI DSS auditor of such monitoring.

Using the FortiOS Network Vulnerability Scan feature

As part of its UTM features, FortiGate units provide a Network Vulnerability Scan. You define assets to monitor, such as servers, workstations, or point-of-sale terminals. Then, the FortiGate unit scans those devices on a regular schedule. The scan checks TCP and UDP ports against a list of known vulnerabilities provided by FortiGuard Services. Scan settings determine how many of the ports are checked. Optionally, all ports are scanned. To view scan logs, go to *Log&Report > Log & Archive Access > Vulnerability Scan Log*.

FortiGate units can be configured to send logs to FortiAnalyzer unit. In a larger network, this enables you to collect log information, including vulnerability scan information, in a central location from several FortiGate units.

For more information, see “[Vulnerability Scan](#)” on page 1113.

Monitoring with other Fortinet products

In addition to your FortiGate unit and its FortiOS firmware, there are several other Fortinet products that can assist your organization to comply with PCI DSS requirements.

FortiAnalyzer network vulnerability scan

FortiAnalyzer units provide a Network Vulnerability Scan similar to the FortiGate vulnerability scan but with more features. In particular, the FortiAnalyzer scan generates compliance reports specifically tailored to PCI DSS requirements. For more information, see the Vulnerability Management chapter of the [FortiAnalyzer Administration Guide](#).

Fortinet Database Security (FortiDB)

A FortiDB appliance or FortiDB software can provide vulnerability scanning and activity monitoring for your databases. For more information, see the [FortiDB User Guide](#).

FortiScan Vulnerability and Compliance Management platform

The FortiScan Vulnerability and Compliance Management (VCM) platform combines a FortiScan appliance with FortiScan agent software to monitor your network assets such as servers, workstations, or point-of-sale terminals. This system can perform vulnerability scans and apply software patches provided by the software vendors. The scan profiles include a predefined one for PCI DSS. The FortiScan appliance produces compliance reports detailing the results of the vulnerability scan.

For more information, see the [FortiScan Administration Guide](#).

FortiWeb Web Application Security

If your organization engages in e-Commerce, you can use FortiWeb Application Security to protect your web servers against attack. The FortiWeb application protects against HTTP and XML-based attacks, guards against attempts to deface your web sites, and scans web servers for vulnerabilities. For more information, see the [FortiWeb Web Application Security Administration Guide](#).

Restricting access to cardholder data

In addition to security policies and authentication governing access to the CDE, you can deploy the Fortinet Database Security (FortiDB) device, which provides vulnerability assessment, database activity monitoring, auditing and monitoring. For more information about this product, see the Fortinet web site, www.fortinet.com.

Controlling access to the CDE network

PCI DSS requires each user to be uniquely identified and authenticated. On the FortiGate unit, this applies to administrators and to users of SSL VPN and IPsec VPNs.

Password complexity and change requirements

By default, the FortiGate unit admin account has no password. Be sure to define a password.

PCI DSS password requirements are

- Require a minimum password length of at least seven characters. (8.5.10)
- Use passwords containing both numeric and alphabetic characters. (8.5.11)
- Change user passwords at least every 90 days. (8.5.9)

To facilitate creation of compliant administrator passwords, you can set a password policy. Go to *System > Admin > Settings*. In the *Password Policy* section, enter the following and then select OK at the bottom of the page.

☒ **Enable Password Policy**

Minimum Length (8-64 characters)

Must Contain ☒

Upper Case Letters Lower Case Letters

Numerical Digits Non-alphanumeric Letters

Apply Password Policy to ☒ Admin Password ☐ IPsec Preshared Key

Enable Password Expiration ☒ (days)

Enable	Select the check box.
Minimum Length	8 or more. (Field does not accept a value less than 8.)
Must Contain	At minimum, set a required number of <i>Numerical Digits</i> and either <i>Upper Case Letters</i> or <i>Lower Case Letters</i> . Also setting a required number of <i>Non-alphabetic Letters</i> is acceptable.
Apply Password Policy to	Select <i>Admin Password</i> .
Enable Password Expiration	Set to 90 days or less. The default is 90 days.

Note that the FortiGate password policy does not apply to user passwords. Both password complexity and password expiry for users would need to be addressed by making them a policy in your organization.

Password non-reuse requirement

PCI DSS requires that passwords are not re-used to satisfy the change requirement:

“Do not allow an individual to submit a new password that is the same as any of the last four passwords he or she has used.” (8.5.12)

FortiGate users don't set their own passwords. The super_admin administrators can and so can admins with appropriate access. There is, however, no FortiGate-based mechanism to enforce non re-use of passwords.

Administrator lockout requirement

PCI DSS requires a user account lockout for administrators to guard against unauthorized access attempts:

- Limit repeated access attempts by locking out the user ID after not more than six attempts. (8.5.13),
- Set the lockout duration to a minimum of 30 minutes or until administrator enables the user ID. (8.5.14)

You can meet these requirements with the following CLI commands:

```
config system global
    set admin-lockout-threshold 6
    set admin-lockout-duration 1800
end
```

The threshold can be less than 6 and the lockout duration can be more than 1800.

Administrator timeout requirement

PCI DSS requires:

- If a session has been idle for more than 15 minutes, require the user to re-enter the password to reactivate the terminal. (8.5.15)

By default, the idle timeout is five minutes. You can go to *System > Admin > Settings* and change the *Idle Timeout* timeout to any value up to the permitted value of 15 minutes.

Administrator access security

To accommodate the requirement for unique identification of each user, the generic admin account should either be assigned to only one administrator or not used at all. You can create an administrator account for each administrator in *System > Admin > Administrators*.

If an administrator always works from the same workstation, consider using the Trusted Host feature. The administrator will be able to log in only from a trusted IP address. You can define up to three trusted IP addresses per administrator.

Administrative access must also be enabled per network interface. Go to *System > Network > Interface* to edit the interface settings. Enable administrative access only on interfaces where you would expect the administrator to connect. Allow only secure connection protocols, HTTPS for web-based access, SSH for CLI access.

Remote access security

For remote access, PCI DSS requires two-factor authentication: a password and some other authentication, such as a smart token or certificate. This applies to employees, administrators, and third parties.

SSL VPN users

For SSL VPN users, implement two-factor authentication by requiring users to have a certificate in addition to the correct password. Go to *VPN > SSL > Config*, enable *Require Client Certificate*, and then select *Apply*. For more information, see the SSL VPN chapter of this Handbook.

IPsec VPN users

If users access your network using an IPsec VPN, you can implement two-factor authentication by enabling extended authentication (XAUTH). This requires the user to enter a password in addition to the VPN authentication provided by the certificate or pre-shared key. As PCI DSS requires each user to have a unique identifier, you should already have user accounts and user groups defined.

To configure XAUTH on your VPN

- 1 Go to *VPN > IPsec > Auto Key (IKE)* and edit your Phase 1 configuration.
- 2 Select *Advanced*.
- 3 In *XAUTH*, select *Enable as Server*.
Enable as Server is available only if *Remote Gateway* is *Dialup User*.
- 4 Set *Server Type* to *PAP*, *CHAP*, or *AUTO* as appropriate.
- 5 Select the *User Group* to which the VPN users belong.
- 6 Select *OK*.



Appendix

Document conventions

Fortinet technical documentation uses the conventions described below.

IPv4 IP addresses

To avoid publication of public IPv4 IP addresses that belong to Fortinet or any other organization, the IP addresses used in Fortinet technical documentation are fictional and follow documentation guidelines specific to Fortinet. The addresses used are from the private IP address ranges defined in RFC 1918: Address Allocation for Private Internets, available at <http://ietf.org/rfc/rfc1918.txt?number-1918>.

Most of the examples in this document use the following IP addressing:

IP addresses are made up of A.B.C.D:

- A - can be one of 192, 172, or 10 - the private addresses covered in RFC 1918.
- B - 168, or the branch / device / virtual device number.
 - Branch number can be 0xx, 1xx, 2xx - 0 is Head office, 1 is remote, 2 is other.
 - Device or virtual device - allows multiple FortiGate units in this address space (VDMs).
 - Devices can be from x01 to x99.
- C - interface - FortiGate units can have up to 40 interfaces, potentially more than one on the same subnet
 - 001 - 099- physical address ports, and non -virtual interfaces
 - 100-255 - VLANs, tunnels, aggregate links, redundant links, vdom-links, etc.
- D - usage based addresses, this part is determined by what the device is doing. The following gives 16 reserved, 140 users, and 100 servers in the subnet.
 - 001 - 009 - reserved for networking hardware, like routers, gateways, etc.
 - 010 - 099 - DHCP range - users
 - 100 - 109 - FortiGate devices - typically only use 100
 - 110 - 199 - servers in general (see later for details)
 - 200 - 249 - static range - users
 - 250 - 255 - reserved (255 is broadcast, 000 not used)
 - The D segment servers can be farther broken down into:
 - 110 - 119 - Email servers
 - 120 - 129 - Web servers
 - 130 - 139 - Syslog servers
 - 140 - 149 - Authentication (RADIUS, LDAP, TACACS+, FSAE, etc)
 - 150 - 159 - VoIP / SIP servers / managers
 - 160 - 169 - FortiAnalyzers
 - 170 - 179 - FortiManagers
 - 180 - 189 - Other Fortinet products (FortiScan, FortiDB, etc.)
 - 190 - 199 - Other non-Fortinet servers (NAS, SQL, DNS, DDNS, etc.)
 - Fortinet products, non-FortiGate, are found from 160 - 189.

Example Network

Variations on network shown in [Figure 390](#) are used for many of the examples in this document. In this example, the 172.20.120.0 network is equivalent to the Internet. The network consists of a head office and two branch offices.

Figure 390: Example network

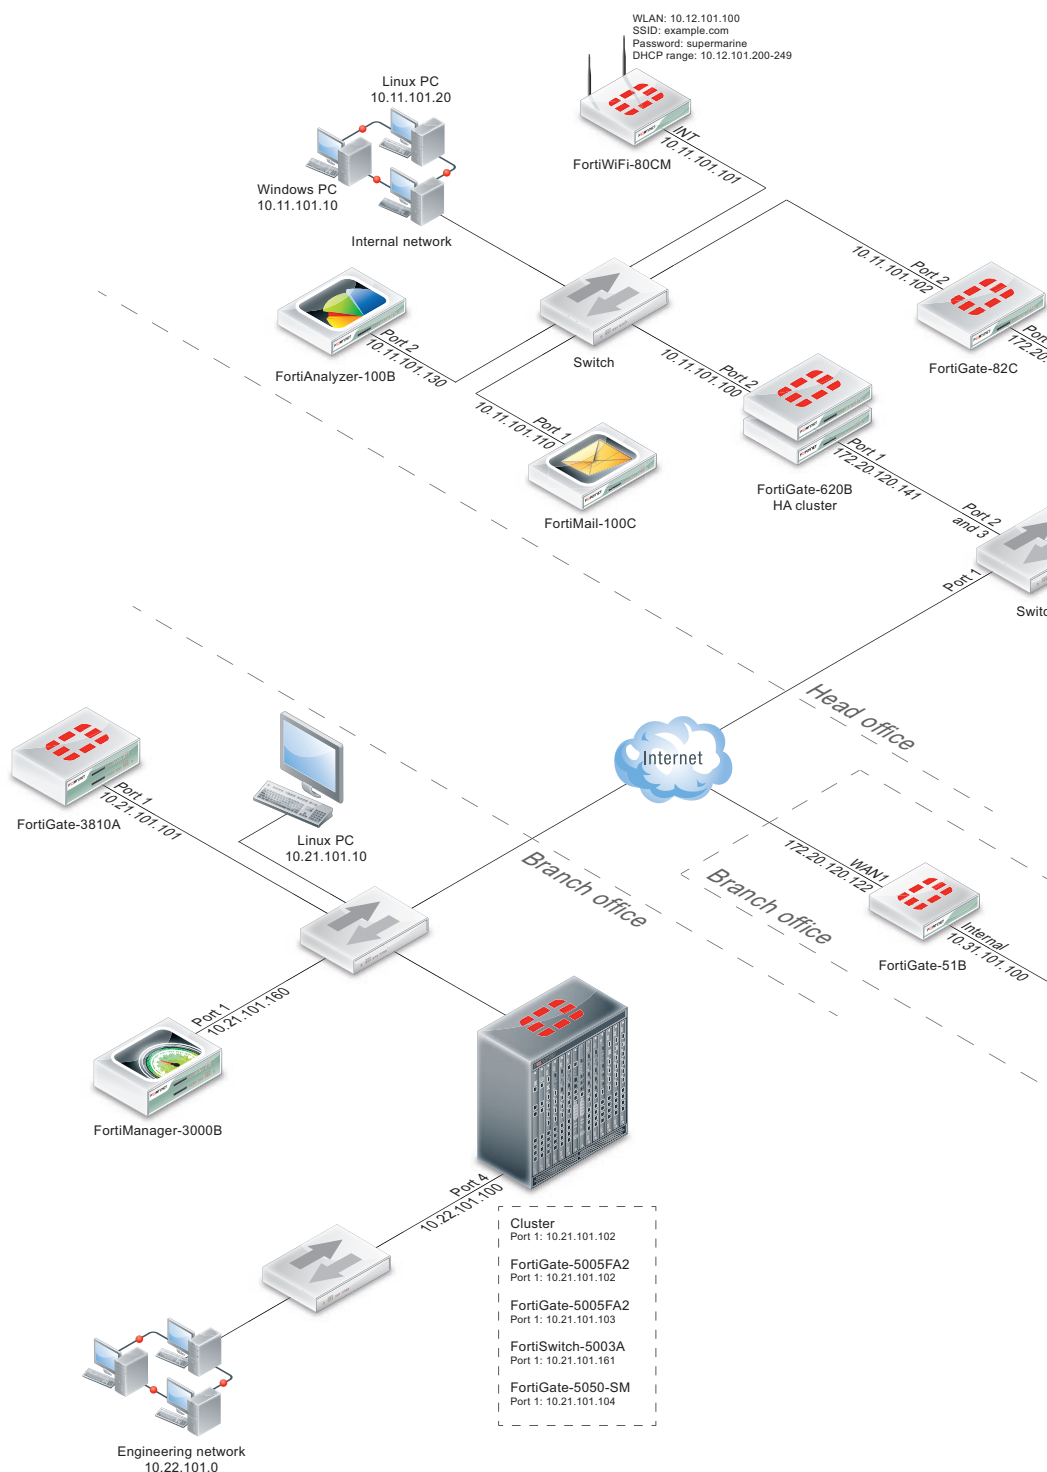


Table 166: Example IPv4 IP addresses

Location and device	Internal	Dmz	External
Head Office, one FortiGate	10.11.101.100	10.11.201.100	172.20.120.191
Head Office, second FortiGate	10.12.101.100	10.12.201.100	172.20.120.192
Branch Office, one FortiGate	10.21.101.100	10.21.201.100	172.20.120.193
Office 7, one FortiGate with 9 VDOMs	10.79.101.100	10.79.101.100	172.20.120.194
Office 3, one FortiGate, web server	n/a	10.31.201.110	n/a
Bob in accounting on the corporate user network (DHCP) at Head Office, one FortiGate	10.0.11.101.200	n/a	n/a
Router outside the FortiGate	n/a	n/a	172.20.120.195

Tips, must reads, and troubleshooting



A Tip provides shortcuts, alternative approaches, or background information about the task at hand. Ignoring a tip should have no negative consequences, but you might miss out on a trick that makes your life easier.



A Must Read item details things that should not be missed such as reminders to back up your configuration, configuration items that must be set, or information about safe handling of hardware. Ignoring a must read item may cause physical injury, component damage, data loss, irritation or frustration.



A Troubleshooting tip provides information to help you track down why your configuration is not working.

Typographical conventions

Table 167: Typographical conventions in Fortinet technical documentation

Convention	Example
Button, menu, text box, field, or check box label	From <i>Minimum log level</i> , select <i>Notification</i> .
CLI input	<pre>config system dns set primary <address_ipv4> end</pre>
CLI output	<pre>FGT-602803030703 # get system settings comments : (null) opmode : nat</pre>

Table 167: Typographical conventions in Fortinet technical documentation

Emphasis	HTTP connections are <i>not</i> secure and can be intercepted by a third party.
File content	<HTML><HEAD><TITLE>Firewall Authentication</TITLE></HEAD><BODY><H4>You must authenticate to use this service.</H4>
Hyperlink	Visit the Fortinet Technical Support web site, https://support.fortinet.com .
Keyboard entry	Type a name for the remote VPN peer or client, such as Central_Office_1.
Navigation	Go to VPN > IPSEC > Auto Key (IKE).
Publication	For details, see the <i>FortiOS Handbook</i> .

Registering your Fortinet product

Access to Fortinet customer services, such as firmware updates, support, and FortiGuard services, requires product registration. You can register your Fortinet product at <http://support.fortinet.com>.

Training Services

Fortinet Training Services offers courses that orient you quickly to your new equipment, and certifications to verify your knowledge level. Fortinet training programs serve the needs of Fortinet customers and partners world-wide.

Visit Fortinet Training Services at <http://campus.training.fortinet.com>, or email training@fortinet.com.

Technical Documentation

Visit the Fortinet Technical Documentation web site, <http://docs.fortinet.com>, for the most up-to-date technical documentation.

The Fortinet Knowledge Base provides troubleshooting, how-to articles, examples, FAQs, technical notes, and more. Visit the Fortinet Knowledge Base at <http://kb.fortinet.com>.

Comments on Fortinet technical documentation

Send information about any errors or omissions in this or any Fortinet technical document to techdoc@fortinet.com.

Customer service and support

Fortinet is committed to your complete satisfaction. Through our regional Technical Assistance Centers and partners worldwide, Fortinet provides remedial support during the operation phase of your Fortinet product's development life cycle. Our Certified Support Partners provide first level technical assistance to Fortinet customers, while the regional TACs solve complex technical issues that our partners are unable to resolve.

Visit Customer Service and Support at <http://support.fortinet.com>.

Fortinet products End User License Agreement

See the [Fortinet products End User License Agreement](#).



Index

Symbols

_email, 322
_fqdn, 322
_index, 322
_int, 322
_ipv4, 322
_ipv4/mask, 322
_ipv4mask, 322
_ipv4range, 322
_ipv6, 322
_ipv6mask, 322
_name, 322
_pattern, 322
_str, 322
_v4mask, 322
_v6mask, 322

Numerics

3DES, 318, 2950
3DES-Triple-DES, 1419
3G/4G modem list, 111
802.11 wireless protocols, 2435
802.1Q, 503, 507, 511
802.3ad, 577
 aggregate interface, 2057
802.3ad aggregate interface
 full mesh HA, 2113
 HA MAC addresses, 2058
 port monitoring, 2058

A

A, 1248
a-a
 load balance schedule, 2221
abort, 325
accelerated interfaces, 782, 1906, 1980
accept, 218
accept any peer, 2697
Accept peer certificate, 1414
Accept this peer certificate
 group only, 1414
access control list (ACL), 1679
access controls, 326
access point
 adding, 2460
 enabling, 2461

Access Point Number (APN), 2418
access profiles
 VDOM and global privileges, 169
access profiles, example
 VDOM and global privileges, 169
accounting system
 RADIUS, 1330
action on failure, FortiBridge
 fail open, 2989
 probe, 2989
 send alertmail, 2989
 SNMP trap, 2989
 syslog, 2989
Active Directory - see Directory Service
Active Directory (AD), 1167, 1216, 1283
 access mode, 1300
active mode
 real server, 2874
active sessions
 HA statistics, 2150
active-active
 best practice, 2013
 device failover, 2218
 IPsec VPN, 2217
 link failover, 2218
 load balancing, 1999, 2217
 network processor accelerated interfaces, 2220
 operation mode, 1998
 redundant interfaces, 2072
 session failover, 2218
 SSL VPN, 2217
 traffic processed by primary unit, 2217
 UTM sessions continue after a failover, 2209
active-active HA, 2941, 2946
active-passive
 device failover, 1998
 failover, 2167
 LACP, 2058
 link failover, 1998
 operating mode, 1998
 WAN optimization rules, 2705
active-passive mode
 redundant interfaces, 2072
adding
 default route, 341
 DHCP relay agent, 572
 SNMP community, 432

- adding configuring defining, 644
 - adding charts, reports, 687
 - alert email message, 669
 - attack logging, 659
 - connecting using automatic discovery, FortiAnalyzer, 652
 - datasets for charts, reports, 681
 - deny security policy, 224
 - explanation of log messages, 627
 - FortiGuard license expiry alert email, 671
 - how firewall components create a FortiGate firewall, 184
 - how packets flow, 186
 - how to apply VLANs and zones to security policies, 195
 - how to arrange security policies, 221
 - how to create a basic security policy for Internet access, 227
 - importing images, reports, 688
 - interfaces and zones, 195
 - IPS packet logging, 662
 - ipv4 tunneling, 237
 - ipv6, dual stack routing, 236
 - log messages, 627
 - logging of events, 660
 - logging practices, 625
 - logging within a firewall policy, 660
 - multiple FortiAnalyzer units, 647
 - multiple syslog servers, 648
 - nac quarantine, 663
 - overview, sql, 633
 - remotely connecting to an IPv6 over the Internet, 237
 - styles, reports, 683
 - syslog server, 645
 - system memory, 642
 - testing FortiAnalyzer configuration, 651
 - themes, reports, 682
 - viewing FortiOS reports, 688
 - webtrends server, 646
- adding or configuring
 - authenticated access, 1243
 - carrier end point IP filter, 1345, 1354
 - carrier end point MMS filter, 1342
 - carrier end point MMS filter list, 1343, 1352
 - dynamic profile, 1332
 - local users, 1224
 - log messages, FortiOS Carrier, 1350
 - logging of events, 1315
 - peer user groups, 1240
 - peer users, 1227
 - RADIUS, 1330
 - user groups, 1237
- adding tags to firewall policies and address lists, 151
- adding tags to predefined signatures and application, 152

- adding, configuring defining
 - administrator password, 349
 - administrator settings, 359
 - advanced RIP options, 1851
 - anomalies, 989
 - antivirus profile, 910
 - authentication settings, 1176
 - backing up configuration, 300
 - BFD, 1867
 - BGP, 1862
 - browser cookie-based FortiGuard web filtering overrides, 1017
 - carrier end point MMS filter list, 1353
 - central NAT table, 266
 - changing administrator's password, 303
 - changing gateway for default route (static route), 1844
 - custom AP profile, 2512
 - custom signature, 995
 - custom signatures, 995
 - dashboards, 295
 - dead gateway detection, 1847
 - DHCP interface settings, 384
 - DHCP server, 570
 - Directory Service user groups, 1180
 - DLP document fingerprinting, 1054
 - DLP file filter, 1056, 1057
 - DLP sensor, 1050
 - document sources, 1055
 - DoS firewall policy, 1086
 - DoS sensor, 988
 - ECMP load balancing method, 1847
 - email address black/white list, 940
 - email filter banned word, 934
 - email filter black/white IP address list, 937
 - email filter profile, 930
 - endpoint profile, 1105
 - endpoint vulnerability result, 1126
 - endpoint, asset definition, 1124
 - endpoint, client installers, 1110
 - endpoint, scan schedule, 1125
 - firewall user groups, 1179
 - firmware version, 299
 - formatting USB disks, 301
 - FortiToken, 1189
 - FSSO, 1192
 - general system settings, 359
 - HA, 2018
 - health check monitor, 2876
 - interface, 379
 - IPS filters, 984
 - IPS sensor, 981
 - IPSec VPN phase 1, 1394
 - IPSec VPN phase 1 advanced options, 1396
 - IPSec VPN phase 2, 1399
 - IPSec VPN phase 2 advanced options, 1399
 - LDAP authentication for administrators, 354
 - LDAP server, 1184
 - local wifi radio, 2510
 - managed FortiAP, 2510, 2511
 - manually updating FortiGuard definitions, 304
 - multicast, 1865
 - OSPF areas, 1857
 - OSPF basic settings, 1854
 - OSPF interface, operating parameters, 1859
 - OSPF networks, 1858
 - OSPF settings, advanced, 1856
 - password authentication, 349
 - password, administrator, 349
 - per-IP traffic shaping, 2265
 - PKI authentication, administrators, 355
 - policy route, 1845
 - port monitoring (HA), 2020
 - pre-defined overrides and custom overrides, 985
 - profile group, 1154
 - protocol options, 1148
 - RADIUS authentication, administrators, 353
 - RAID disk, 310
 - replacement message images, 587
 - replacement messages, 588
 - restoring configuration, 302
 - RIP, 1850
 - RIP enabled interface, 1853
 - RIP settings, advanced, 1851
 - RIP-enabled interface, 1853
 - rogue AP settings, 2509
 - secondary IP address, 388
 - shared traffic shapers, 2263
 - sniffer firewall policy, 1134
 - SSID, 2506
 - SSL VPN user groups, 1180
 - static route to routing table, 1845
 - static route, adding to routing table, 1845
 - static routes, 1843
 - synchronizing with NTP server, 299
 - system configuration backup and restore, FortiManager, 301
 - system time, 299
 - TACACS+ authentication, 354
 - TACACS+ server, 1187
 - tags, 255
 - tags, application control, 1077
 - text strings (names), 293
 - uploading scripts, 620
 - URL filter, 1018
 - user authentication settings, 1176
 - user groups, 1176
 - web filter local ratings, 1021
 - web filter profile, 1012
 - widgets, 295
- adding, configuring or defining
 - anti-overbilling protection, 2349
 - APN filtering, 2340
 - carrier duplicate message, 2333
 - carrier end point MMS filter list, 2357
 - carrier endpoint filter list, 2334
 - carrier GTP profile, advanced filtering rule, 2345
 - carrier utm, message flood, 2332
 - carrier utm, notification list, 2329
 - GTP log settings, 2349
 - GTP message type filtering, 2340
 - GTP profile, 2336
 - MMS profile, 2316
 - policy route, 1694
 - server load balance port forwarding virtual IP, 2905
 - server load balance virtual IP, 2899

- address
 - CIDR format, 196
 - FDQN, 204
 - geography-based, 201
 - groups, 204
 - IP pool, 198
 - IP range, 197
 - IPv6, 235, 393
 - matching, IP pool, 199
- Address Resolution Protocol (ARP), 532, 1922
- address spoofing, 2402
- Address Translation, 1347
- address translation
 - MMS, 2295
- addresses
 - ipv6, 235
- addresses, firewall, 196
- ADM, 2220
- admin, 3013
 - administrator account, 312
 - password, 348
 - password length, 350
- administration
 - schools, 579
- administrative access, 386
 - changing, 313
- administrative distance, 1681, 1682
- administrative interface. **See** web-based manager
- Administrative Status, 770
- administrator
 - account, 312
 - adding a FortiBridge password, 2979
 - lockout, 351
 - password, 312
- administrator accounts, FortiBridge
 - adding, 2981
- administrator profiles
 - global, 356
 - vdom, 356
- administrator settings, 359
- administrators
 - LDAP authentication, 354
 - management access, 350
 - monitoring *See also* widgets, 302
 - viewing list, 349
- ADM-XD4
 - security processing module, 2939
- advanced mode (FSSO), nested groups, 1291
- advertisement message
 - adjusting the interval, 2242
- advertisement messages
 - VRRP, 2237
- AES-128, 2950
- AES128,192 ,256, 1419
- AES-192, 2950
- AES-256, 2950
- AFS3, advanced file security encrypted file
 - AFS3, 206
- age
 - age difference margin, 2003
 - changing the age difference margin, 2003
 - displaying cluster unit age, 2004
 - primary unit selection, 2003
 - reset the cluster age, 2005
 - resetting cluster unit age, 2005
- age difference margin, 2003
 - changing, 2003
- Agent, sFlow, 422
- aggregate interface
 - best practice, 2013
 - HA MAC addresses, 2058
 - interface monitoring, 2058
- aggregate interfaces, 577
- aggregated subnets
 - for hub-and-spoke VPN, 1454
- aggregation, link, 2940, 2941
- aggressive mode, 1407
- AH, predefined service, 206
- air flow, 401
- alert email, 1903
 - configuring the FortiBridge, 2993
 - HA, 2155, 2156
 - sample FortiBridge message, 2993
- alert message console
 - viewing, 305
- alert notification, 2367, 2381
- alertmail, FortiBridge
 - action on failure, 2989
- alerts
 - configuring the FortiBridge, 2992
- ALG
 - changing the port numbers that the SIP ALG listens on, 2545
- all policies, 2262, 2263
- allow
 - pattern, 1042
- allow access, 386
- Allow inbound, encryption policy, 1433
- Allow outbound, encryption policy, 1433
- allowed, 3013
- always revalidate, 2753
- ambient temperature, 401
- ambiguous routing
 - resolving in FortiGate dialup-client configuration, 1502
 - VPN routing, 1444
- AMC
 - bridge module, 2930
 - configuring AMC modules, 2929
 - hard disk, 2013
- AMC (Advanced Mezzanine Card), 864, 2933
- AMC module
 - configuring, 2929
- anomalies
 - IPS, 990
- anomaly
 - checks, 2955
 - hardware checks, 2955
 - IPS checks, 2955

- anomaly protection
 - DoS, 880
- antenna, 2436
- anti-overbilling, 2349
- antireplay, 1598, 2950, 2951, 2952, 2953, 2954, 2957, 2958, 2959, 2960
- antispam, **see** email filtering **and** FortiGuard, AntiSpam
- antispam. **See also** Email filter, 928
- anti-spoofing, 780, 1681, 1979, 2430
- antivirus, 895, 2941
 - and PCI DSS, 3032
 - archive scan depth, 902
 - change default database, 900
 - concepts, 895
 - databases, 898
 - enabling scanning, 900
 - example, 907
 - explicit FTP proxy, 2837
 - explicit web proxy, 2817
 - file filtering, 879
 - flow-based scanning, 896
 - FortiAnalyzer, 879
 - HTTPS, IMAPS, POP3S, SMTPS, 1142
 - maximum file size, 903
 - override default database, 901
 - proxy-based scanning, 895
 - scan buffer size, 902
 - scanning order, 896
 - SymbOS/Commwar.A!worm, 2354
 - SymbOS/Commwar.B!wm, 2353
 - SymbOS/Commwarrie.C-wm, 2354
 - virus database, 911
 - virus list, 979
- antivirus monitor, 1156
- antivirus profile, 910
- antivirus profiles, archive inspection, 105
- antivirus quarantine
 - HTTPS, IMAPS, POP3S, SMTPS, 1142
- antivirus scanning, 1942
- antivirus updates, 410
 - manual, 304
- ANY
 - service, 206
- AOL
 - service, 206
- AP profile
 - creating, 2447
 - described, 2445
- APN filtering, 2340
- application
 - database, viewing, 1097
 - detection, 1093
- application control, 880, 1073, 2266
 - explicit FTP proxy, 2837
 - explicit web proxy, 2817
 - monitor, 1068
 - packet logging, 1070
- application database
 - endpoint, 1107
- application layer, 2940
- application monitor, 1068, 1158
 - Application Monitor submenu, 146
 - application sensor, 107
 - Archive & Data Leak Monitor submenu, 145
 - archive and data leak monitor, 1158
 - archive antivirus scan depth, 902
 - archive inspection for antivirus profiles, 105
 - archiving
 - DLP, 1046
 - area, 1794
 - area border router (ABR), 1857
 - ARP, 2872
 - cache, 743
 - duplicate packets, 779
 - gratuitous, 2178
 - proxy ARP, 2872
 - request, 1924
 - resolution, 782, 1980
 - arp table, 2198, 2234
 - arp-reply
 - load balance virtual server, 2872
 - arps
 - CLI command, 2178
 - gratuitous, 2178
 - arps-interval
 - CLI command, 2178
 - arranging security policies, 221
 - AS
 - multihomed, 1710
 - number (ASN), 1710
 - stub, 1710
 - ASCII, 329, 1215
 - ASM-CX4, 2930
 - ASM-cx4, 2930
 - ASM-FX2, 2930
 - asset definition, 121
 - assets
 - adding manually, 1114
 - discovering, 1113
 - selecting to scan, 1113
 - asymmetric routing, 535, 780, 1979, 2430
 - attached network equipment
 - failover, 2210
 - attack updates
 - manual, 304
 - scheduling, 410
 - attributes, RADIUS, 1202
 - auth, 3014
 - authenticating
 - based on peer IDs, 1414
 - FortiGate unit pre-shared key, 1411
 - IPsec VPN peers and clients, 1412
 - L2TP clients, 548
 - PPTP clients, 539
 - through IPsec certificate, 1409
 - through XAuth settings, 1422
 - authenticating users
 - FortiGate, local, 1223
 - with LDAP servers, 1224
 - with RADIUS servers, 1224
 - with TACACS+ servers, 1224

- authentication, 2438, 2706
 - authentication method, 2700
 - certificate-based, 1246
 - Citrix, 2816
 - client certificates, 1621
 - configuring access, 1243
 - defining settings, 1176
 - disclaimer, 1251
 - explicit web proxy, 2815, 2816
 - firewall policy, 1246, 1250
 - guest users, 1315
 - heartbeat, 2177
 - HTTP, 2816
 - IP Based, 1890
 - IPsec VPN, 1258
 - L2TP, 1169, 1261
 - MD5, 1858
 - NAT device, 2816
 - overview, 1264
 - peer, 2698
 - PKI certificate, administrators, 355
 - PPTP, 1169, 1260
 - protocols, 1246
 - proxy, 2816
 - RADIUS for administrators, 353
 - replacement messages, 1248
 - RIP, 1853
 - SCP, 363
 - server certificate and SSL VPN, 1621
 - SHA-1, 256, 384, 512, 1419
 - SSL VPN, 1258
 - SSL VPN timeout, 1258
 - timeout, 1243
 - two-factor, 1189
 - VPN, 1257
 - VPN client-based, 1170
 - WAN optimization peer authentication, 2697
 - web proxy, 2816
 - web-based user, 1169
 - Windows Terminal Server, 2816
 - XAuth, 1259
- Authentication Algorithm
 - Manual Key, 1552
- authentication group
 - authentication method, 2700
 - certificate, 2700
 - password, 2700
 - pre-shared key, 2700
- Authentication Key, Manual Key, 1553
- authentication protocols, 1247
 - ASCII, 1215
 - CHAP, 1215
 - MS-CHAP, 1215
 - PAP, 1215
 - setting, 1246
 - TACACS+ servers, 1215
- authentication realm
 - explicit web proxy, 2812
- authentication server, external
 - for L2TP, 548
 - for PPTP, 539
- authentication server, XAuth, 1422

- authentication servers
 - Directory Service, 1216
 - LDAP, 1207
 - RADIUS, 1201
 - TACACS+, 1215
- authentication timeout setting, 1619
- authorization, LDAP, 357
- authorize-manager-only, CLI, 159
- auto-install, 368
- Autokey
 - IPSec VPN, 1393
 - keep alive, 1427
 - Keep Alive, IPsec interface mode, 1429
- automatic radio resource provisioning ARRP, 113
- autonomous system (AS), 1860
- AV Monitor submenu, 145
- av-failopen, 783

B

- back to HA monitor
 - HA statistics, 2150
- backing up
 - FortiBridge configuration, 2996
- backing up and restoring, per VDOM, 153
- backing up config file, 100
- backing up configuration
 - See** widgets, system information
- backup
 - cluster configuration, 2155
- backup and restore configuration, central management, 301
- backup configuration
 - SCP, 361
 - USB, 372
- backup password, 1212
- backup unit, 1992
 - See Also** subordinate unit, 1992
- backup VPN, 1541
- band
 - radio bands for wireless LANs, 2435
- bandwidth, 2255, 2441
 - calculation method, 2948
 - guaranteed, 2253, 2262, 2264, 2265
 - limitation, 2948
 - maximum, 2262, 2264, 2265
 - zero, 2262
- bandwidth guarantees, 2941
- banned user list
 - quarantining attackers, 986
- banned users
 - cause or rule, 1235
- banned word (spam filter)
 - list, 936
- basic FortiBridge configuration, 2979
- basic FortiBridge settings, 2988
- baud rate, 332
- Berkeley Packet Filtering (BPF), 738
- best practice, 2014
- BFD
 - disabling, 1868
 - disabling, interface, 1868

BGP

- AS, 1860
- configuring, 1862
- graceful restart, 2191
- RFC 1771, 1860
- RFC 2385, 1860
- service, 206

bgp

- attribute
 - AS_PATH, 1757
 - ATOMIC_AGGREGATE, 1759
 - COMMUNITY, 1758
 - MULTI_EXIT_DESC, 1758
 - NEXT_HOP, 1759
- BGP-4+, 1751
- clearing routes, 1753, 1763
- control plane, 1765
- flap, 1765
- graceful restart, 1765
- MED, 1758
- neighbors, 1753
- password, MD5, 1753
- RFC 1997, 1758
- route reflectors (RR), 1755
- stabilizing the network, 1765

- BGP, IPv6, 251

- bidirection, 2948

- Bi-directional Forwarding Detection (BFD), 1766

- billing, 1347

- binding

- LDAP servers, 1208

- bits per second (bps), 316

- black list, 880

- blackhole route, 1680

- block, 2590

- pattern, 1042

- block traffic, 2262

- blocking

- http access by ip, 273

- port 25, 270

- blocking of users

- Endpoint Control, 1090

- Blowfish, 318

- body

- SIP message, 2520

- boot interrupt, 315

- Boot Strap Router (BSR), 471

- Bootstrap issues, 871

- border gateway protocol (BGP). See routing, BGP

- BPDU

- message exchange, 2235

- branch, 2596

- brctl,netlink, 779

- bridge mode, 2930

- bridge module

- AMC, 2930

- bridge protocol data unit, 2235

- bridge, Transparent mode, 779

- broadcast

- domains, 503

- storm, 532

- broken cluster unit

- replacing, 2056

- browser cookie-based overrides

- FortiGuard web filtering, 1016

- buffer size

- IPS, 966

- busy

- load balance, 2223

- bypass mode, FortiBridge, 2973

- connecting to a FortiBridge CLI, 2974

- resuming normal mode, 2996

- switching to normal mode, 2974

- byte cache, 2673

- changing the relative amount of disk space, 2694

C

- CA certificate, 1140

- cache

- exempting from web caching, 2752

- cache cleaner, 1616

- cache engine

- WCCP, 2845

- cache expired objects, 2755

- call-keepalive, 2547, 2548

- captive portal, 2437

- carrier

- duplicate message, configuring, 2333

- GTP profile, 2336

- HA, 2403

- message flood, 2332

- MMS profile, 2316

- notification list, 2329

- carrier end point, 1327, 1341

- block, 1342, 1352

- blocking, 1342, 1352

- configuring the filter list, 1343, 1353

- filter list, 2357

- filtering, 1342, 1352, 2295, 2389

- IP filter, 1345, 1354

- IP filtering, 1344, 1354, 2358

- MMS filtering, 1342, 1352, 2318

- pattern, Perl, 1345, 2334

- pattern,regular expression, 1345, 2334

- pattern,wildcard, 1345, 2334

- patterns, 1342, 1343, 1352, 1353

- user context list, 1330

- wildcard pattern, 1352

- carrier end point MMS filter list

- adding to an MMS protection profile, 1342, 1352

- adding to MMS protection profile, 1342, 1352

- blocking MMS messages, 1342, 1352

- carrier end point pattern

- Perl regular expression, 1342, 1352

- wildcard, 1342

- carrier endpoint

- filter list, 2334

- case sensitivity

- Perl regular expressions, 333

- CB2, 2220

- CDE

- defined, 3023

- central management
 - backup and restore configuration, 301
- central NAT, 265
- central NAT table
 - configuring, 266
- certificate
 - authentication group, 2700
 - key size, 1141
 - SSL, 1140
- certificate authority (CA), 1264
- Certificate Name, Phase 1, 1410
- certificate request, 1267, 1277
 - generating, 1267, 1277
 - key size, 1268
- certificate revocation list (CRL)
 - importing, 1270
- certificate signing request (CSR), 1267
- certificate, IPsec
 - group, 1413
 - Local ID setting, 1414
 - using DN to establish access, 1412
 - viewing local DN, 1413
- certificate, security, 336
- certificate, server, 1621
- certificates
 - import, 1274
 - importing CRL, 1270
 - OSCP, 1265
 - root CA, installing, 1270
 - self-signed, 1264
 - signed server, installing, 1269
 - Single Sign On (SSO), 1263
- certification, 3040
- Challenge-Handshake Authentication Protocol (CHAP), 1215, 1423
- changing unit's host name, 297
- channels
 - for 802.11a, 2503
 - for 802.11b, 2503
 - for 802.11n 5GHz, 2503
 - radio channels for wireless LANs, 2435
- CHAP, 537, 1260
- Charging Data Function (CDF), 2310
- Charging Data Record (CDR), 2311
- chart display, 143
- chat message log, MSNP21, 135
- checking windows version, 1643
- CIDR, 322, 2710
- CIFS
 - protocol optimization, 2689
- Cisco
 - router configuration, 515, 530
 - switch configuration, 515, 521, 530
- Cisco switch configuration, 1937
- Cisco VPN, 1579
- Citrix
 - authentication, 2816
- Classless Inter-Domain Routing (CIDR), 238, 1700
- CLI, 289, 2957
 - authorize-manager-only, 159
 - connecting, 315
 - connecting to a FortiBridge unit in bypass mode, 2974
 - connecting to from the web-based manager, 313
 - connecting to the, 315
 - Console widget, 317
 - custom maximum invalid firewall auth attempts, 162
 - customize number of invalid firewall auth attempts, 162
 - DDNS commands, 159
 - deleting all local logs, archives and user-configured report templates, 137
 - example for uploading logs to FTP server (text format), 136
 - get commands, traffic shaper and per-IP shaper, 161
 - management checksum config for FortiManager, 161
 - MTU config support (non-IPsec tunnels), 162
 - netscan asset auth, 121
 - NTLM authentication, 120
 - OSPFv3 NSSA extension, 122
 - real-time session, traffic shaper bandwidth and CP6 statistics, 160
 - resetting a FortiBridge unit to factory defaults, 2982
 - session-pickup, 2171, 2205
 - SSL connection encrypt level option, 135
 - upgrading the firmware, 367
 - uploading logs to FTP server (text format), 136
 - WiFi, user group authentication, 114
- CLI command, 2024, 2029, 2036, 2041, 2048, 2051, 2061, 2066, 2074, 2079, 2096, 2101, 2116, 2121, 2181
- CLI console, 308
- client
 - downloading, 1652
 - using FortiWiFi unit as a WiFi client, 2501
 - WCCP, 2845
- client certificates, 1621
- client comforting, 2359
 - amount, 2360
- client IP
 - assigning with RADIUS, 1487
- client mode, 2501
 - using FortiWiFi unit as a WiFi client, 2501
- Client to FortiGate
 - SSL offloading, 2890
- Client to FortiGate to Server
 - SSL offloading, 2890
- CLOSE_WAIT, 845
- cluster
 - adding a new FortiGate unit, 2055
 - configuring in transparent mode, 2034
 - connecting an HA cluster, 1996
 - converting a standalone FortiGate unit, 2054
 - definition, 2015
 - FortiBridge application, 2975
 - operating, 2127
 - replacing a failed cluster unit, 2056
 - virtual cluster, 2089
- cluster configuration
 - backup, 2155
 - restore, 2155
 - troubleshoot, 2027, 2032, 2039, 2045, 2049, 2053, 2062, 2068, 2075, 2081, 2085, 2117, 2123

- cluster member, 2147, 2947
 - cluster members list, 2148
 - priority, 2148
 - role, 2148
- cluster name, 1999
- cluster unit
 - connect to a cluster, 2164
 - definition, 2015
 - disconnect from a cluster, 2164
 - getting information using SNMP, 2145, 2146
 - getting cluster unit serial numbers, 2146
 - getting serial numbers using SNMP, 2146
 - SNMP get, 2145, 2146
- cluster units, 1992
- CNAME, 575
- collector agent
 - Ignore User list, 1300
 - LDAP access, 1299
 - logs, 1319
 - settings, 1296
 - specifying, 1310
 - TCP ports, 1302
- collector agent, sFlow, 422
- collision domain, 779
- column settings
 - configuring, 291
- column settings, security policies, 218
- command, 320
 - abbreviation, 328
 - completion, 327
 - help, 327
 - multi-line, 327
- comments
 - firewall policy, 255
- common name, LDAP servers, 1208
- community
 - adding to a FortiBridge unit, 2995
 - SNMP on a FortiBridge unit, 2995
- concentrator, 1458
 - IPSec tunnel mode, 1404
 - IPSec VPN, policy-based, 1404
- concepts
 - antivirus, 895
 - web filtering, 880
- config file, backing up, 100
- configuration
 - backing up and restoring a FortiBridge unit, 2996
 - backup, 2155
 - basic FortiBridge configuration, 2979
 - FortiExplorer, 337
 - restore, 2155
 - synchronization, 2184
- configuration example, FortiBridge
 - HA cluster, 2975
 - other FortiGate interfaces, 2976
 - standalone FortiGate unit, 2967
- configuration lock, 398
- configuration revisions, 364
- configuration synchronization, 2169
 - disabling, 2184
- configure
 - DNS, 340, 343
 - FortiGuard, 347, 2607
 - interfaces, 338
 - restore, 363
- configuring
 - alert email message, 669
 - collector agent, 1296
 - dynamic DNS VPN, 1472
 - firewall policy authentication, 1250
 - FortiClient dialup-client VPN, 1488
 - FortiClient in dialup-client VPN, 1493
 - FortiGate dialup-client VPN, 1504
 - FortiGate firewall policies, 1312
 - FortiGate in dialup-client IPsec VPN, 1506
 - gateway-to-gateway IPsec VPN, 1439
 - hub-and-spoke IPsec VPN, 1453
 - IPsec VPN authentication, 1258
 - L2TP VPN authentication, 1261
 - LDAP server, FortiGate unit, 1308
 - local users, 1224
 - manual keys, 1552
 - multiple FortiAnalyzer units, 647
 - multiple syslog servers, 648
 - peer user groups, 1240
 - peer users, 1227
 - PPTP VPN authentication, 1260
 - SSL VPN authentication, 1258
 - transparent mode IPsec VPN, 1547
 - WAN optimization peer, 2699
 - XAuth authentication, 1259
- configuring a FortiGate unit for HA operation, 1995
- configuring anomalies, 989
- connect
 - FortiBridge unit, 2968
- connected monitored interfaces
 - primary unit selection, 2002
- connecting
 - to the CLI using SSH, 318
 - to the CLI using Telnet, 319
 - to the console, 316
 - web-based manager, 335
- connecting a FortiGate HA cluster, 1996
- connecting using automatic discovery, FortiAnalyzer, 652
- connectionless, 700
- conservation mode, 440
- conserve mode, 305, 783, 1137
 - no SYN-ACK, 727
- console, 316
- console messages
 - synchronization fails, 2186
- contact-fixup, 2575
- content archive
 - metadata, 2371
 - MMS protection profile, 2371
 - summary, 2371
- content archiving
 - DLP archiving, 1060
- content blocking, 2295
- content provider (CP), 2325
- content scanning
 - SSL, 1139

- Content-Length
 - SIP header, 2590
- control plane, 1765
- control plane (GTP-C), 2409
- controlled upgrade, 374
- controlling connection between FortiGate and FortiManager, CLI, 162
- conventions, 319, 3037
- convergence, 1699, 1767
- cookie, 1349, 2322
 - persistence, 2886
- cookie persistence
 - HTTP host-based load balancing, 2889
- country code, 2323
- coverage, 2440
- cp1252, 330
- CP6 statistics, diag, CLI, 160
- CPU load, 1892, 1942
- CPU usage, 772
 - HA statistics, 2150
 - weight, 172, 2224
- cpu usage
 - weighted load balancing, 172, 2223
- cross site scripting (XSS), 1256
- Cross-Site Scripting
 - protection from, 293
- cryptographic load, 1597, 2949
- custom AP profile
 - configuring, 2512
- custom maximum invalid firewall auth attempts, CLI, 162
- custom services, 206
- custom signature
 - adding, 949
- customer communication session, 1327, 1341
- customer identifying information, 1327
- customer service, 3040
- customize number of invalid firewall auth attempts, CLI, 162
- CVSPSERVER, concurrent versions system proxy server, 207
- cx4, 2930

D

- dampening, 1765
 - reachability half-life, 1765
- dashboards
 - adding, 295
- data leak prevention (DLP), 1050
- data leak prevention (DLP), **see** DLP
- data leak protection, 1049
- date, 720, 744
 - changing the FortiBridge system date, 2981
 - quarantine files list, 666
- date and time, 346
- DB-9, 316
- DC
 - quarantine files list, 667
- DC Agent mode, 1292, 1293, 1304
 - dcagent.dll, 1322
- DCE-RPC, 555
 - firewall service, 207
- dcerps
 - session helper, 555
- Dead entry timeout
 - collector agent configuration, 1298
- dead gateway detection, 800, 1451, 1728, 1831, 1847, 2200
- Dead Peer Detection (DPD), 1421, 1451
 - Phase 1, 1420, 1421
- dead peer detection (dpd), 800
- debug
 - diagnose, 2189
- debug flow, 782
- decryption, 2951, 2952, 2953, 2954, 2957, 2958, 2959, 2960
- dedicated monitoring
 - interface, 2129
- deep scan, 1142
- deep scanning
 - firewall protocol, HTTPS, 1150
- deep SIP message inspection, 2586
- default
 - probe settings on a FortiBridge unit, 2989
 - resetting a FortiBridge unit to factory defaults, 2982
 - VoIP profile, 2544
- default gateway, 1844
- default password, 1306
- default port
 - RADIUS servers, 1205
 - TACACS+ servers, 1215
- default route, 341, 514, 1844
 - NAT/Route example, 1917
 - VDOM example, 1919
 - VLAN, 514
- default TTL
 - web cache, 2754
- definitions, 320
- delaying session pickup, 2206
- delete
 - local users from FortiGate configuration, 1226
 - user group from FortiGate configuration, 1241
- delete, shell command, 324
- Denial of Service (DoS), 704
- dense mode, 472, 1864
- deny, 218
- deny policy, 224
- deployment, 2440
- DES, 2950
- designated router (DR), OSPF, 1591
- Designated Routers (DRs), 471
- destination
 - firewall policy, 252
- destination NAT
 - SIP, 2571
- details, security policies, 218
- device
 - failure, 2168

- device failover, 2167, 2169
 - active-active, 2218
 - active-passive, 1998
 - configuration synchronization, 2169
 - definition, 2015
 - HA heartbeat, 2169
 - IPsec SA synchronization, 2169
 - route synchronization, 2169
 - virtual MAC address, 2169
- device priority, 2152
 - primary unit selection, 2001, 2006
 - subordinate unit, 2152
- DH Group
 - IPsec interface mode, 1429
 - Phase 1, 1417, 1420
 - Phase 2, 1426
- DH key size, FIPS-CC, 1417
- DHCP, 2012
 - for WiFi clients, 2450
 - IP reservation, 573
 - relay, 2012
 - server, 2012
 - servers and relays, 570
 - service, 572
- DHCP (Dynamic Host Configuration Protocol)
 - configuring on an interface, 384
 - service, 207
- DHCP interface, 384
- DHCP ipv6, 123
- DHCP Monitor submenu, 143
- DHCP relay
 - in FortiClient dialup-client configuration, 1491
 - in FortiGate dialup client configuration, 1503
- DHCP server, 165, 1427
 - in FortiClient dialup-client configuration, 1491
 - IP Reservation, 165
 - ipv6, 123
- DHCP6
 - service, 207
- DHCP-IPsec
 - IPsec interface mode, 1429
 - IPSec VPN, phase 2, 1401
 - phase 2, 1427
- diag netlink brctl, 767
- diag system icap profile list , ICAP, 110
- diag system icap server list , ICAP, 110
- diagnose
 - firewall vip realserver, 2880
 - firewall vip virtual-server, 2880
 - sys ha reset-uptime, 2005
 - sys ha showcsum, 2189
- diagnose commands, 788, 803
 - diag debug, 782
 - diag netlink, 779
- diagnose debug, 2189
- diagnose hardware deviceinfo nic, 2024, 2029, 2036, 2041, 2048, 2051, 2061, 2066, 2074, 2079, 2096, 2101, 2116, 2121
 - CLI command, 2181
- diagnose sys ha dump, 2004
- diagnose sys sip, 2541, 2546
- diagnose sys sip debug-mask, 2541
- diagnose sys sip dialog, 2541
- diagnose sys sip mapping list, 2541
- diagnose sys sip status, 2541
- diagnose sys sip-proxy calls, 2545
- diagnose sys sip-proxy filter, 2545
- diagnose sys sip-proxy log-filter, 2545
- diagnose sys sip-proxy meters, 2545
- diagnose sys sip-proxy stats, 2545
- diagnose test application sip, 2545
- diagnostics
 - debug the packet flow, 1980
 - packet sniffing, 1979
 - traceroute, 1920
 - tracert, 1920
- diagnostics, tracert, 522
- dialog
 - SIP, 2519, 2525
- dialup users
 - configuring authentication for, 1258
- dialup-client IPsec configuration
 - DHCP server and relay, FortiClient VIP, 1491
 - dialup server for FortiClient dialup clients, 1488
 - dialup server for FortiGate dialup clients, 1504
 - FortiGate client configuration, 1506
 - FortiGate dialup client configuration, 1504
 - requirements for FortiClient access, 1487
 - requirements for FortiGate client access, 1503
- dictionary
 - RADIUS attributes, 1204
- differentiated services, 2269
 - mapping, 2275
- Diffie-Hellman algorithm, 1417, 1426
- Digital Encryption Standard, 1385
- Dijkstra's algorithm, 1796
- directory
 - LDAP servers, 1207
- Directory Service
 - servers, 1216
 - user groups, 1240
- Directory service, 1216
- disabled mode
 - real server, 2874
- disabling, 554
- disclaimer
 - customized, 1252
 - default, 1252
- disconnecting a unit from a cluster
 - override, 2012
- disk space
 - byte cache, 2694
 - web cache, 2694
- disk status, viewing, 619
- diskfull keyword, config log memory command, 3015
- distance vector protocols, 1701
- distinguished names
 - elements, 1209
 - LDAP servers, 1208
 - list of, 1211
 - max size, 1210
- distributed ARP, 113
- Distributed Computing Environment Remote Procedure Call

- (DCE-RPC), 555
- DLP, 1035
 - archiving, 1046, 1060, 2295
 - content archiving, 1060
 - creating rules, 1045, 1046
 - default rules, 1044
 - document fingerprinting, 1054
 - document sources, document fingerprinting, 1055
 - explicit FTP proxy, 2837
 - explicit web proxy, 2817
 - file filter, 1056
 - MMS, 2379
 - sensor, 1050
- DLP archive
 - displaying on dashboard, 1144
 - HTTPS, IMAPS, POP3S, SMTPS, 1143
 - viewing, 307
- DLP archives, sending to multiple FortiAnalyzer units, 133
- DLP archiving, 1060
- DLP document fingerprinting, 108
- DLP. See data leak protection
- DNAT, 1946, 2940
 - multicast service reflection, 1866
- DNS, 555, 574
 - CNAME, 575
 - external servers, 574
 - local domains, 574
 - override, 339
 - public, 576
 - recursive, 576
 - server
 - server, DNS, 575
 - service, 207
 - shadow, 576
 - slave, 575
 - split, 576
 - TTL, 204
- DNS lookups, 1903
- DNS master, 575
- DNS server, 165
 - changing DNS IP addresses on a FortiBridge unit, 2980
- DNS server, dynamic DNS configuration, 1472
- dns-tcp, session helper, 555
- dns-udp, session helper, 555
- document fingerprinting, 1054
- documentation
 - conventions, 3037
 - Fortinet, 3040
- domain component, LDAP servers, 1208
- Domain Name Identifier (DNI), 1309
- domain name server, 574
 - configure, 343
- domain name server (DNS), 773
 - domain name server, configure, 340
 - domain name, dynamic DNS configuration, 1471, 1473
 - DoS
 - anomaly protection, 880
 - DoS policy
 - viewing, 1085
 - DoS sensor
 - configuring, 988
 - SCCP, 2592
 - SIP, 2592
 - dotted decimal, 322
 - dotted-decimal notation, 1857
 - double NAT example, 279
 - downloading
 - quarantine logs, 667
 - tunnel client, 1652
 - downloading firmware, 365
 - downloading log messages in Log Access, 131
 - dual internet connection, 561
 - dual stack routing, ipv6, 236
 - dual WAN
 - link redundancy, 561
 - load sharing, 564
 - Duplicate ARP packet, 779
 - duplicate MAC, 533
 - Dynamic DNS (DDNS), 1469
 - configuration steps, 1472
 - domain name configuration, 1473
 - overview, 1469
 - remote VPN peer configuration, 1478
 - dynamic DNS, CLI, 159
 - dynamic IP address
 - for remote host, 1483
 - FortiGate DDNS peer, 1471
 - FortiGate dialup client, 1501
 - dynamic IP pattern
 - FortiBridge probe setting, 2989
 - dynamic IP pool, 253
 - SIP, 2572
 - dynamic profile, 118, 1325, 1332
 - assigning profile groups, 1164
 - carrier end point, 1341
 - configuring the accounting system, 1330
 - event log messages, 1350
 - log message period, 1351
 - log settings, 1350
 - profile group, 1329
 - RADIUS, 1164
 - RADIUS Stop record, 1330
 - timeout options, 1349
 - user context list, 1330
 - Users Only, 1331
 - VDOMs, 1332
 - WAP traffic, 1327

- dynamic routing, 1849
 - access list, 1868
 - BGP, configuring, 1862
 - Bi-directional forwarding detection (BFD), 1867
 - disabling BFD for an interface, 1868
 - failover, 2191
 - multicast, configuring, 1865
 - OSPF, 1854
 - OSPF networks, 1858
 - OSPF, advanced, 1856
 - OSPF, defining areas, 1857
 - OSPF, operating parameters for interface, 1859
 - PIM, 1863
 - RIP, 1850
 - RIP enabled interface, 1853
 - RIP, advanced, 1851
- dynamic VPN address
 - mode-cfg, 1512
- dynamic-gateway, 1469

E

- E Reload, 2755
- earthing, 403
- ECMP, 1680
- ECMP load balancing, 1847
- eDirectory, 1216, 1284
 - servers, 1307
- eDirectory - see Directory Service
- edit, shell command, 324
- EI (Enhanced Extension Interface), 2936
- EICAR, 906
- elements, distinguished names, 1209
- email
 - alert, FortiBridge, 2993
- email address black/white list, 940
- email filter, 928, 940
 - banned word, 934
 - order of, 929
 - IMAP/POP3/IMAPS/POP3S, 930
 - SMTP/SMTPS, 929
 - profile, 930
- email filtering
 - IMAPS, POP3S, SMTPS, 1143
- email filtering, **see also** FortiGuard, AntiSpam, 880
- email monitor, 1157
- Email Monitor submenu, 145
- email token, 1228
- Enable perfect forward secrecy (PFS)
 - IPsec interface mode, 1429
 - Phase 2, 1426
- Enable replay detection
 - IPsec interface mode, 1429
- Enable replay detection, Phase 2, 1426
- enable session pickup, 2205
- Encapsulation, 1381
- encryption, 2951, 2952, 2953, 2954, 2957, 2958, 2959, 2960
 - heartbeat, 2177
- Encryption Algorithm
 - IPSec VPN, manual key, 1404
- Encryption Algorithm, Manual Key, 1552
- Encryption Key, Manual Key, 1552

- encryption policy
 - allow outbound and inbound, 1433
 - defining IP addresses, 1431
 - defining IPsec, 1434
 - defining multiple for same IPsec tunnel, 1435
 - evaluating multiple, 1435
 - outbound and inbound NAT, 1433
 - traffic direction, 1433
- encryption types, 2436
- end
 - command in an edit shell, 325
- end point, 1327
- End User Address, 2349
- end, shell command, 324
- endpoint
 - application database, 1107
 - asset definition, 1124
 - client installers, 1110
 - configuring a profile, 1105
 - scan schedule, 1125
 - vulnerability result, 1126
- Endpoint Control
 - blocked users, 1090
 - modifying download portal, 1100
 - modifying recommendation portal, 1100
 - modifying replacement pages, 1100
 - monitoring endpoints, 1099
- endpoint control, 120
- Endpoint Mapper (EPM), 555
- endpoint security
 - endpoint control, 120
- endpoints
 - monitoring, 1099
- engine algorithm
 - IPS, 965
- engine count
 - IPS, 965
- enhanced packet-matching, 1711
- entering text strings (names), 293
- entry
 - IPS signature, 948
- environment variables, 328
- Equal Cost Multipath (ECMP), 1680
- equal cost multipath (ECMP), 1683
- equal-cost multi-path (ECMP), 1686
- equipment
 - FortiAP unit, 2440
 - FortiWiFi unit, 2438
 - wireless, 2438
- escape sequence, 328
- ESP, 2950
 - service, 207
- event, 3014
- event log
 - dynamic profile, 1350
- event logs, 134
- event-system logs, 135
- example
 - Endpoint Control configuration, 1101
 - inter-VDOM, 1953
 - NAT/Route VDOM, 1911
 - VDOM, 1911

- example configuration, FortiBridge, 2967
 - HA cluster, 2975
 - other FortiGate interfaces, 2976
 - example for two-factor authentication, 116
 - example for VDOM and global privileges for access profiles, 169
 - example IPSec configurations, 2951, 2957
 - example of source IP address, FortiGate-originating traffic, 176
 - examples
 - configuring multiple FortiAnalyzer units, 649
 - hub-and-spoke VPN, 1463
 - report styles, 684
 - reports, 689
 - sql statements, 634
 - execute
 - ha synchronize all, 2184
 - execute shutdown, 403, 2927
 - exempt
 - web cache, 2752
 - expired objects
 - cache, 2755
 - explicit, 1890
 - explicit FTP proxy, 2831
 - antivirus, 2837
 - application control, 2837
 - DLP, 2837
 - incoming IP address, 2837
 - intrusion protection, 2837
 - outgoing IP address, 2837
 - protocol options, 2837
 - replacement message, 124, 2833, 2842
 - reverse, 2681
 - UTM, 2835
 - explicit HTTP proxy
 - incoming IP address, 2830
 - outgoing IP address, 2830
 - explicit mode
 - WAN optimization, 2706, 2712
 - explicit proxy, 1228, 1890
 - authentication cookie for session-based authentication, 125
 - form-based user authentication, 125
 - forwarding servers, 124
 - FTP, 123
 - explicit web proxy, 2208, 2807, 2818
 - antivirus, 2817
 - application control, 2817
 - authentication, 2815, 2816
 - authentication realm, 2812
 - DLP, 2817
 - flow-based scanning, 2818
 - FortiGuard web filtering, 2817
 - FTP, 2807, 2827
 - HTTPS, 2807, 2827, 2831
 - intrusion protection, 2817
 - IPS, 2817, 2837
 - load balancing, 2217
 - PAC, 2807, 2827
 - protocol options, 2817
 - proxy auto-config, 2807, 2827
 - proxy chaining, 124, 2813
 - realm, 2812
 - SOCKS, 2807, 2827
 - SSL content scanning and inspection, 2819
 - unknown HTTP version, 2812
 - UTM, 2810, 2817
 - web filtering, 2817
 - EXT1
 - FortiBridge management access, 2981
 - Extended Authentication (XAuth), 1220, 1409, 1422, 1494
 - extended authentication (XAuth), 1410
 - Exterior Gateway Protocol (EGP), 1751
 - exterior gateway protocol (EGP), 1701
- ## F
- FA2, 2220
 - FA2 (NP1) processor, 2935
 - factory default
 - resetting a FortiBridge unit, 2982
 - fail open, FortiBridge, 2989
 - recovering from, 2995
 - failed authentication attempts, 1245
 - failed cluster unit
 - replacing, 2056
 - fail-open
 - IPS, 965
 - failover, 2000
 - active-passive, 2167
 - and attached network equipment, 2210
 - attached network equipment, 2235
 - definition, 2015
 - delayed, 2235
 - device, 2167, 2169
 - dynamic routing, 2191
 - enabling session failover, 2205
 - GTP and HA session failover, 2209
 - HA, 1992
 - heartbeat, 2016
 - issues with layer-3 switches, 2234
 - link, 2016, 2167, 2194
 - monitoring cluster units, 2155
 - session, 2168, 2205
 - subsecond, 175, 2200
 - failover protection, 2089
 - active-passive operating mode, 1998
 - virtual clustering, 2089

- failure
 - definition, 2015
 - device, 2168
 - link, 2168, 2194
 - multiple link failures, 2199
- failure threshold
 - tuning a FortiBridge unit, 2992
- failure, FortiBridge
 - recovering from, 2995
- FAMS, 137
- fast path
 - required session characteristics, 1597, 2940
- fast path requirements, 2940
- fast roaming, 2443
- FB4, 2220
- FB8, 2220
- FDN, 2135
- FDQN, 204
- FGCP
 - definition, 2015
 - description, 1987
- FGT_ha_admin
 - HA administrator account, 2164
- field, 320
 - SIP, 2531
- file
 - quarantine, 2146
- file block
 - default list of patterns, 1056
 - default patterns, 2364
- file filtering, 1042
 - antivirus, 879
 - general configuration steps, 1042
- file name
 - quarantine files list, 666
- file pattern, 896, 1042
 - creating, 1043
- file quarantine
 - configuring, 905
 - general configuration steps, 905
- file sharing, 1949
- file size, 896
- File transfer protocol (FTP), 556
- file type, 896, 1042
 - creating, 1043
- filter
 - APN, 2340
 - carrier end point, 1342, 1352
 - carrier end point MMS filtering, 1342, 1352
 - filtering information on web-based manager lists, 290
 - GTP message type, 2340
 - IPS, 946
 - list, carrier end point, 2357
 - message type, GTP, 2340
 - quarantine files list, 666
 - web-based manager lists, 290
- filter list
 - carrier end point, 1353
 - IP, 1330
- filtering
 - using filter settings, 1135
- filtering and customizing log messages, 668
- filtering lists, 147
- filtering log messages, 131
- FIN packet, 845
- final
 - SIP response message, 2529
- FINGER
 - service, 207
- fingerprint
 - DLP, 1040
 - document fingerprint, 1040
- FireFox, 1256
- firewall
 - applying VLANs and zones to security policies, 195
 - central NAT table, 266
 - configuring user groups, 1237
 - creating user groups, 1237
 - DoS policy, 1086
 - dynamic IP pool, 253
 - how firewall components create a FortiGate firewall, 184
 - identity based policy, 1339
 - interfaces and zones, 195
 - IPsec VPN dialup user access, 1236
 - ipv6, 235
 - load balancing, 2208
 - option, any, 251
 - per-IP traffic shaping, 2265
 - policy authentication, 1246, 1250
 - predefined services, 206
 - protection profile, 1931
 - protocol options, 177
 - protocol options, configuring, 1148
 - schedule, 1930
 - service group, 1909
 - shared traffic shapers, 2263
 - sniffer policy, 1134
 - source interface, 251
 - user authentication timeout, 1243
 - user groups, 1236
 - virtual IP port forwarding, 178
 - virtual IP source address filter, 178
 - what is it, 183
 - wildcard addresses, 202
- firewall address, 196, 1915, 1930, 1934
 - geography-based filtering, 107
 - NAT/Route VDOM example, 1915
 - simple VDOM NAT/Route example, 1917
 - VDOM NAT/Route example, 1917
- firewall address lists
 - adding tags, 151
- firewall addresses
 - ipv6, 235
 - wildcard address, 202
- firewall configuration object icons, 153
- firewall IP addresses
 - defining, 1431
- firewall IP addresses, defining L2TP, 548
- firewall load balancing, 2706
- firewall objects, 139

- firewall policies, 2455
 - adding NAT policies to transparent mode, 274
 - adding tags, 151
 - and Endpoint Control, 1098
 - local-in, 177
 - see security policies, 341
- firewall policy, 902, 1915
 - and FortiBridge probes, 2971
 - comments, 255
 - defining for policy-based VPN, 1433
 - defining for route-based VPN, 1436
 - destination, 252
 - guaranteed bandwidth, 2264
 - hub to spoke, 1460
 - inter-VDOM, 1942
 - log traffic, 252
 - maximum bandwidth, 2264, 2265
 - policy-based, for FortiGate dialup client, 1507
 - policy-based, for gateway-to-gateway, 1444
 - policy-based, for hub-and-spoke, 1457
 - route-based, for FortiGate dialup client, 1507
 - route-based, for gateway-to-gateway, 1442
 - route-based, for hub-and-spoke, 1457
 - schedule, 252
 - See security policy, 139
 - service, 252
 - spoke to spoke, 1462
 - traffic shaping, 254
 - user groups, 1179
 - using as route-based "concentrator", 1459
 - VDOM, 1909, 1910
 - VDOM example, 1915, 1918, 1935
 - VLAN Transparent, 1926
- firewall policy and strong authentication, 1276

- firewall service
 - AFS3, 206
 - AH, 206
 - ANY, 206
 - AOL, 206
 - BGP, 206
 - CVSPSERVER, 207
 - DCE-RPC, 207
 - DHCP, 207
 - DHCP6, 207
 - DNS, 207
 - ESP, 207
 - FINGER, 207
 - FTP, 207
 - FTP_GET, 207
 - FTP_PUT, 207
 - GOPHER, 207
 - GRE, 207
 - H323, 208
 - HTTP, 208
 - HTTPS, 208
 - ICMP_ANY, 208
 - IKE, 208
 - IMAP, 208
 - INFO_ADDRESS, 208
 - INFO_REQUEST, 208
 - Internet-Locator-Service, 208
 - IRC, 208
 - L2TP, 208
 - LDAP, 208
 - MGCP, 208
 - MS-SQL, 209
 - MYSQL, 209
 - NetMeeting, 209
 - NFS, 209
 - NNTP, 209
 - NTP, 209
 - ONC-RPC, 209
 - OSPF, 209
 - PC-Anywhere, 209
 - PING, 209
 - PING6, 209
 - POP3, 209
 - PPTP, 210
 - QUAKE, 210
 - RAUDIO, 210
 - REXEC, 210
 - RIP, 210
 - RLOGIN, 210
 - RSH, 210
 - RTSP, 210
 - SAMBA, 210
 - SCCP, 210
 - SIP, 211
 - SIP-MSNmessenger, 211
 - SMTP, 211
 - SNMP, 211
 - SOCKS, 211
 - SQUID, 211
 - SSH, 211
 - SYSLOG, 211
 - TALK, 211
 - TCP, 211
 - TELNET, 211
 - TFTP, 212
 - TIMESTAMP, 212
 - UDP, 212
 - UUCP, 212
 - VDOLIVE, 212
 - viewing predefined list, 206
 - VNC, 212
 - WAIS, 212
 - WINFRAME, 212
 - WINS, 212
 - X-WINDOWS, 212
- firewall session setup rate, 732
- firewall traffic shaper monitor, 2276
- firewall vip realserver
 - diagnose, 2880
- firewall vip virtual-server
 - diagnose, 2880
- firmware
 - backup and restore from USB, 372
 - download, 365
 - from system reboot, 370
 - install on a FortiBridge unit from a system reboot, 2983
 - installing, 370
 - installing on a FortiBridge unit, 2982
 - revert from CLI, 369
 - reverting with web-based manager, 365
 - testing before use, 372
 - testing new firmware, 372
 - upgrade from CLI, 367
 - upgrade with web-based manager, 365
 - upgrading a FortiBridge unit to a new version, 2982
 - upgrading using the CLI, 367
- firmware install, 2936
- firmware practices, 99
- firmware upgrade
 - HA, 2152
- firmware, updating
 - FortiAP unit, 2463
- first alive
 - load balancing, 2873
- fixed ports, IP pools, 199
- flood, 2339, 2373
- flow control, 316
- flow inspection, 701, 702
- flow, reverse shaping, 2267
- flow-based DLP, 104
- flow-based scanning, 2818
- flow-based web filtering, 103
- forget
 - administrator password, 312
- forgot
 - administrator password, 312
- format
 - hard disk, 2693
- formatting multiple disk partitions, 163
- formatting USB disks, 301
- form-based user authentication, explicit proxy, 125
- FortiAccel (NP1) processor, 2935
- FortiAnalyzer, 2136
 - antivirus, 879
 - quarantine, 905
- FortiAnalyzer traffic reports, 2934

- FortiAP
 - configuring access point, 2510, 2511
- FortiAP unit, 2440
 - connecting to CLI, 2466
 - updating firmware, 2463
- FortiASIC, 782, 1418, 1419
 - NP2, 1597
- FortiBridge-2002
 - connecting, 2969
- FortiBridge-2002F
 - connecting, 2969
- FortiClient, 1655, 1658
 - download location, 1091
 - endpoint, 1110
 - required version, 1091
- FortiClient Connect, 158
- FortiClient dialup client configuration
 - example, 1495
- FortiClient dialup-client configuration
 - configuration steps, 1488
 - FortiClient configuration, 1493
 - overview, 1483
 - VIP address assignment, 1486
- FortiClient Endpoint Security peer, 1168
- FortiClient peer, 2687
- Forticlient VPN, 1402
- FortiExplorer, 140, 336
 - configuration, 337
 - updates, 337
- FortiFone
 - configuring, 2615
 - SIP phone, 2520
- FortiGate
 - authenticating users, 1223
 - authenticating with XAuth, 1259
 - configuring to use LDAP, 1209
 - configuring to use RADIUS, 1204
 - dynamic profile, 118
 - IPsec VPN, 1258
 - source IP address, 176
- FortiGate Cluster Protocol
 - description, 1987
- FortiGate dialup client IPsec configuration
 - FortiGate acting as client, 1501
 - policy-based firewall policy, 1507
 - route-based firewall policy, 1507
 - using DHCP relay in, 1503
- FortiGate firewall
 - creating, 184
- FortiGate HA cluster
 - FortiBridge application, 2975
- FortiGate LDAP configuration, 1307
- FortiGate setup wizard, 140
- FortiGate unit
 - adding to a cluster, 2055
 - converting from standalone to a cluster, 2054
 - replacing a failed, 2056
- FortiGate unit serial number
 - primary unit selection, 2007
- FortiGate-ASM-CX4, 2930
- FortiGate-ASM-FB4, 2951, 2957
- FortiGate-ASM-FX2, 2930
- FortiGate-ASM-S08 module, 905
- FortiGate-ASM-SAS module, 905
- FortiGuard, 347, 2135, 2607
 - AntiSpam, 880
 - Antivirus, 899, 905, 3040
 - as source of antivirus signatures, 1091
 - as source of application signatures, 1091
 - as source of FortiClient installer, 1091
 - manually configuring definition updates, 304
 - push update, 409, 410, 411
 - Web Filtering, 880, 1143
 - HTTPS, 1143
- FortiGuard Analysis and Management Service, 2371, 2372
- FortiGuard Analysis server, 644
- FortiGuard Antispam, 2135
- FortiGuard Antivirus, 2135
- FortiGuard Center, 899
- FortiGuard definitions
 - manually updating, 304
- FortiGuard Distribution Network, 2135
- FortiGuard Distribution System (FDS), 746
 - servers, 748
- FortiGuard Intrusion Protection, 2135
- FortiGuard quota, monitoring, 1159
- FortiGuard service, 1903
- FortiGuard Services
 - analysis service options, 408
 - licenses, 303
 - management and analysis service options, 408
 - support contract, 408
 - web filtering and antispam options, 413
- FortiGuard services, 304
- FortiGuard Web Filter quota, 1027
- FortiGuard Web Filtering, 2135
- FortiGuard web filtering
 - explicit web proxy, 2817
- FortiGuard web filtering overrides
 - browser cookie based, 1016
- FortiGuard, 3G/4G modem list, 111
- FortiGuard, backup and restore configuration, 301
- FortiGuard, Distribution Network, 899
- FortiManager
 - remote backup and restore options, 301
- Fortinet
 - Technical Documentation, conventions, 3037
 - Technical Support, 3040
 - Technical Support, registering with, 3040
 - Training Services, 3040
- Fortinet customer service, 3040
- Fortinet documentation, 3040
- Fortinet MIB, 434, 441
- Fortinet Server Authentication Extension (FSAE), 1167, 1250, 1283
- Fortinet Single Sign On (FSSO), 1283
- Fortinet Single Sign On Agent (FSSO), 1192
- FortiOS
 - ipv6, 236
- FortiOS 3.0 MR7, 1303
- FortiOS Carrier, 1341
 - URL extraction, 180
- FortiOS UTM default report, 678

- FortiToken, 115, 1189, 1230
- FortiWiFi unit, 2438
 - configuring as an AP unit, 2466
- forward delay
 - spanning tree parameter, 2234
- forward domain, 779
- forwarding
 - MAC forwarding table, 2198, 2234
- forwarding servers, 124
- fragmented packets, 1598, 2941
- frame size, 2935
- frame size, maximum, 2936
- framed-ip-addr
 - RADIUS field, 1336
- fresh factor
 - web cache, 2753
- front panel
 - resetting FortiBridge unit to factory defaults, 2982
- FSAE 4.0 MR1, 1288
- FSAE version 3.5.27, 1303
- FSM
 - hard disk, 2013
- FSSO
 - components, 1290
 - DC Agent mode, 1285
 - DNI, 1309
 - DNI field, 1320
 - guest, 1308, 1315
 - identity-based policy, 1312
 - LDAP, 1311
 - logoff detection, 1302
 - Mac OS, 1286
 - nested groups, 1291
 - Polling Mode, 1287
 - Polling mode, 1285
 - ports, 874
 - supported operating systems, 1284, 1290
 - TCP port 8000, 1310
 - trust relation, 1288
- FSSO agent
 - FortiGate configuration, 1310
- FSSO_Guest_Users, 1321
- FSSO_Guest_Users group, 1315
- FTP, 2941
 - configuring FortiBridge probe, 2991
 - explicit web proxy, 2807, 2827
 - FortiBridge probe, 2972
 - protocol optimization, 2689
 - service, 207
- FTP proxy, 2831
 - antivirus, 2837
 - change the prompt, 124, 2833, 2842
 - DLP, 2837
 - protocol options, 2837
 - UTM, 2835
- FTP_GET
 - service, 207
- FTP_PUT
 - service, 207
- FTPS, 177

- full mesh
 - HA, 2111
 - redundant HA heartbeat interfaces, 2112
- full mesh HA, 1993, 2111
 - configuration example, 2113
 - definition, 2015
 - troubleshooting, 2125
- full mode
 - SSL offloading, 2890
- fully qualified domain name (FQDN), 322
- fuzzing protection
 - SIP, 2586
- FX2, 2930

G

- GARP, 2178
- gateway, 341
- Gateway Function (CGF), 2310
- gateway-to-gateway IPsec configuration
 - configuration steps, 1439
 - overview, 1437
 - policy-based firewall policy, 1444
 - route-based firewall policy, 1442
- GB2312, 330
- general configuration steps
 - file filtering, 1042
 - file quarantine, 905
- General Packet Radio Service (GPRS), 2309, 2411
 - System Node (GSN), 2313
- generating
 - IPsec phase 1 keys, 1417
 - IPsec phase 2 keys, 1426
- Generic Access Network (GAN), 2418
- Generic Routing Encapsulation (GRE), 538
- geography-based addressing, 201
- geography-based filtering, firewall address, 107
- get
 - edit shell command, 325
 - shell command, 324
 - test vs, 2880
- get hardware nic, 2024, 2029, 2036, 2041, 2048, 2051, 2061, 2066, 2074, 2079, 2086, 2096, 2101, 2116, 2121, 2181
- get system performance
 - status, 772
 - top, 772
- get system performance status, 2142
- get test app, 767
- get test ipldb, 2879
- get test sip, 2545
- gigabit interfaces, SNMP, 432
- global, 733
- Global System for Mobile Communications (GSM), 2309, 2411
- GOPHER
 - service, 207
- Gr interface, 2311
- grace period
 - age difference margin, 2003
 - changing, 2003

- graceful restart, 1765
 - BGP, 2191
 - OSPF, 2192
- graphical user interface. **See** web-based manager
- gratuitous ARP packets, 2178
- gratuitous arps, 2178
- grayware, 896, 899
 - scanning, 906
- GRE, 1853
 - service, 207
- GRE-over-IPsec VPN, 1579
- grounding, 402, 2920
- group filters
 - FortiGate, on collector agent, 1301
- group ID
 - changing, 2182
 - HA configuration option, 2000
 - virtual MAC address, 2182
- group name
 - HA cluster name, 1999
 - HA configuration option, 1999
- group-id
 - CLI command, 2182
- groups
 - user, 1176
 - Windows AD, viewing on FortiGate, 1311
- groups, addressing, 204
- GSM EDGE Radio Access Network (GERAN), 2418
- GTP
 - Access Point Name (APN), 2340
 - control plane, 2409
 - create pdp request, 2340
 - HA session failover, 2209
 - Information Element (IE), 2340
 - International Mobile Station Identity (IMSI), 2341
 - log messages, 2428
 - mobile country code (MCC), 2341
 - mobile network code (MNC), 2341
 - mobile subscriber identification number (MSIN), 2341
 - path management, 2409
 - Routing Area Code (RAC), 2345
 - runtime statistics, 2426
 - Sequence Number, 2338
 - Serial Number (SNR), 2345
 - Software Version Number (SVN), 2345
 - tunnel, 2306
 - tunnel management, 2409
 - Type Allocation Code (TAC), 2345
- GTP billing, 2410
- GTP profile, 2336
 - advanced filtering, 2343, 2345
 - anti-overbilling, 2349
 - APN filtering, 2340
 - encapsulated IP traffic filtering, 2346
 - encapsulated non-IP end user traffic filtering, 2347
 - general settings, 2338
 - IMSI filtering, 2341
 - information element removal policy, 2346
 - logging, 2349
 - message type filtering, 2340
 - specifying logging types, 2351
- GTP protocol anomaly prevention, 2348
 - missing mandatory information elements, 2348
 - out of state, 2348
- GTP UDP
 - session failover, 2209
- GTP⁺ (GTP prime), 2410
- GTP-in-GTP, 2403, 2426
- GTP-U (GTP user data tunnelling), 2426
- GTPv1, 2291
- guaranteed bandwidth, 2255, 2262, 2264, 2265
 - firewall policy, 2264
 - traffic shaping, 2264
- guest account, 1308, 1315
- guest network, 2437
- GUI. **See** web-based manager
- gui-voip-profile, 2543
- Gx interface, 2311
- Gz interface, 2311

H

- H.245, 556
- h245l
 - session helper, 556
- H323
 - service, 208
- H323, session helper, 556

- HA, 1763, 2148
 - alert email, 2155, 2156
 - changing firmware upgrade, 2153
 - cluster member, 2148
 - cluster members list, 2147
 - configure weighted-round-robin weights, 170, 2222
 - configuring virtual clustering, 2093, 2095, 2100
 - connect a cluster unit, 2164
 - definition, 1987
 - disconnect a cluster unit, 2164
 - event log message, 2137
 - FGT_ha_admin administrator account, 2164
 - firmware upgrade, 2152
 - full mesh and 802.3ad aggregate interfaces, 2113
 - full mesh and redundant heartbeat interfaces, 2112
 - full mesh HA configuration example, 2113
 - GTP session failover, 2209
 - GTP tunnels, 2403
 - hello state, 2137
 - host name, 2148
 - IPS processing, 964
 - link failover scenarios, 2199
 - log message, 2137
 - manage individual cluster units, 2163
 - manage logs for individual cluster units, 2136
 - monitor cluster units for a failover, 2155
 - router monitor, 1674
 - routes, 1674
 - SIP session failover, 2594
 - SNMP and reserved management interface, 2130
 - standby state, 2137
 - states, 2137
 - subordinate unit device priority, 2152
 - subordinate unit host name, 2152
 - VDOM partitioning, 2020
 - viewing HA statistics, 2150
 - virtual cluster, 2089
 - virtual clustering, 2018
 - virtual domains, 2089
 - work state, 2137
- HA cluster
 - FortiBridge application, 2975
- HA dynamic weighted load balancing
 - dynamic weighted load balancing, 172
 - weighted-round-robin weights, 170
- HA group ID
 - changing, 2182
- HA group name, 1999
- HA heartbeat, 2169
 - definition, 2015
- HA session offloading, 2946
- HA statistics
 - active sessions, 2150
 - back to HA monitor, 2150
 - CPU usage, 2150
 - intrusion detected, 2151
 - memory usage, 2150
 - monitor, 2150
 - network utilization, 2151
 - refresh every, 2150
 - serial no, 2150
 - status, 2150
 - total bytes, 2151
 - total packets, 2150
 - up time, 2150
 - virus detected, 2150
- HA virtual MAC address
 - definition, 2015
- ha_daemon
 - HA user interface, 2141
- HA, virtual cluster, 1951
- ha-eth-type
 - CLI command, 2175, 2236
- half mode
 - SSL offloading, 2890
- hard disk
 - AMC, 2013
 - byte cache storage, 2693
 - formatting, 2693
 - FSM, 2013
 - Wan optimization storage, 2693
- hardware
 - get hardware nic command, 2024, 2029, 2036, 2041, 2048, 2051, 2061, 2066, 2074, 2079, 2096, 2101, 2116, 2121, 2181
- hardware acceleration
 - RTP, 2527
- hash map, 2173
- Hash-based Message Authentication Code (HMAC), 1418
- hb-interval, 2176
- hb-lost-threshold, 2175
- hc-eth-type
 - CLI command, 2175, 2236
- header
 - SIP, 2531
 - SIP messages, 2520
- health check
 - ping, 2882
- health check monitor
 - configuring, 2876
 - matched content, 2877
 - real server, 2876
- health monitor
 - proxy forwarding, 2814
 - real server, 2876
- heartbeat, 2169
 - authentication, 2177
 - changing the heartbeat interval, 2176
 - changing the hello state hold-down time, 2176
 - changing the lost heartbeat threshold, 2175
 - definition, 2015
 - encryption, 2177
 - modifying heartbeat timing, 2175

- heartbeat device
 - definition, 2015
- heartbeat failover
 - definition, 2016
- heartbeat interface, 2171
 - best practice, 2014
 - configuring, 2171
 - priority, 2171
 - selection, 2172
 - switch interfaces, 2172
 - virtual clustering, 2090
- heartbeat interfaces, 2090
- hello state
 - changing the time to wait, 2176
 - definition, 2016
- hello state hold-down time
 - changing, 2176
- helo-holddown, 2176
- help
 - navigating using keyboard shortcuts, 293
 - searching the online help, 292
 - using FortiGate online help, 291
- heuristics, 896, 899
- hierarchy
 - LDAP servers, 1208
- high availability
 - definition, 1987, 2016
- high availability (HA), 2946
 - active-active, 2941
 - load balancing, 2941
- High Speed Packet Access (HSPA), 2419
- , 1322
- HMAC check offloading, 2950
- HMAC-MD5, 1418
- HMAC-SHA-1, 1418
- HNT, 2580
- hnt-restrict-source-ip, 2585
- Home Location Register (HLR), 2311
- Home Network Identity (HNI), 2417
- host check, 1640
 - custom, 1642
 - introduction, 1616
 - OS, 1646
- host ID
 - peer, 2686
- host name, 297, 2152
 - best practice, 2013
- host OS, patch check, 1646
- host-based
 - load balancing, 2888
- hosted NAT traversal
 - See HNT, 2580
- hosted-nat-traversal, 2583
- hostname
 - cluster members list, 2148
- how to allow DNS queries to only one DNS server, 224
- how to apply VLANs and zones to security policies, 195
- how to arrange policies, 221
- how to create basic security policy for Internet access, 227
- how to test basic security, 227
- how to use match-vip, 205
- how to use UTM profiles to monitor and protect your network, 214
- HTTP, 2754
 - authentication, 2816
 - FortiBridge probe, 2972
 - persistence, 2886
 - protocol optimization, 2689
 - service, 208
 - unknown HTTP sessions, 2714
 - WCCP service ID, 2846
- http
 - blocking, 273
- HTTP 1.1 conditionals, 2754
- HTTP cookie
 - persistence, 2886
- HTTP header, 1327, 1341, 1348
- HTTP Header Field, 1347, 2322
- HTTP host
 - cooke persistence load balancing, 2889
 - load balancing, 2874, 2888
- HTTP multiplexing, 2208
 - load balancing, 2217
- HTTP port
 - web cache, 2736
- HTTP proxy, 842
- HTTP rule
 - non-HTTP sessions, 2714
- HTTPS, 289, 350
 - antivirus, 1142
 - antivirus quarantine, 1142
 - data leak prevention, 1143
 - DLP archive, 1143
 - explicit web proxy, 2807, 2827, 2831
 - FortiGuard Web Filtering, 1143
 - load balancing, 2217
 - persistence, 2886
 - protocol recognition, 1142
 - service, 208
 - web filtering, 1142
- HTTPS deep scanning
 - explicit web proxy.explicit web proxy
 - HTTPS deep scanning, 2819
- HTTP-User-Agent, 1256
- hub
 - HA schedule, 2218
- hub-and-spoke
 - spoke subnet addressing, 1454
- hub-and-spoke IPsec configuration
 - concentrator, defining, 1458
 - configuration example, 1463
 - hub configuration, 1455
 - infrastructure requirements, 1454
 - overview, 1453
 - policy-based concentrator, 1458
 - policy-based firewall policy, 1457
 - route-based firewall policy, 1457
 - route-based inter-spoke communication, 1458
 - spoke configuration, 1460
- humidity, 401

I

- ICAP, 109, 702
 - example of ICAP, 110
 - troubleshooting, 110
- icap, 1150
- icap profile, 1150
- icap server, 1150
- ICMP land, 2956
- ICMP processing, 274
- ICMP_ANY
 - service, 208
- ID tag, 504, 507
- identify-based policies, 706
- Identity based policies (IBP), 1166, 1339
- Identity based policy (IBP), 1255
- identity-based policy, 222, 1312
 - Local ID, 1168
 - NTLM guest, 1255
 - NTLM user agent strings, 1256
 - position, 285
- identity-based security policies, 2706
- Idle timeout
 - VPN connection, 1170
- idle timeout
 - changing for the web-based manager, 313
- idle timeout setting, 1647
- IDS
 - one-armed IDS, 880
- IEEE 1394 (FireWire), 239
- IEEE 802.1, 1923
- IEEE 802.11a, channels, 2503
- IEEE 802.11b, channels, 2504
- IEEE 802.1Q, 503, 507
- IEEE 802.1q, 2940
- IEEE 802.3ad, 2940
- ifHighSpeed, 432
- IF-MIB.ifSpeed, 432
- if-modified-since, 2754
- IGMP
 - RFC 1112, 472, 1864
 - RFC 2236, 472, 1864
 - RFC 3376, 472, 1864
- ignore
 - web cache setting, 2754
- Ignore User List, 1300
- IKE
 - service, 208
- IKE Configuration Method, 1512
- IKE encryption key, 1420
- IKE negotiation, 179
 - parameters, 1417
- IKEv2, 1409
- IM, 880
 - load balancing, 2217
- IMAP
 - FortiBridge probe, 2972
 - service, 208
- IMAPS
 - antivirus, 1142
 - antivirus quarantine, 1142
 - data leak prevention, 1143
 - DLP archive, 1143
 - email filtering, 1143
 - predefined firewall services, 1141
 - protocol recognition, 1142
- inactivity timeout
 - SIP session, 2547
- Inbound NAT, encryption policy, 1433
- incoming-ip
 - explicit FTP proxy, 2837
 - explicit HTTP proxy, 2830
- incremental
 - synchronization, 2185
- indentation, 321
- independent VDOM configuration, 1948
- index, 2173
- index number, 322
- INFO_ADDRESS
 - service, 208
- INFO_REQUEST
 - service, 208
- Information Elements (IE), 2348, 2418
- informational
 - SIP response message, 2529
- Initial Disc Timeout, 339
- initiator, 800, 1450
- insert policy before
 - security policy, 2708
- inspection
 - flow, 701, 702
 - proxy, 702
 - security layers, 703
 - SSL, 1139
 - stateful, 699
- inspection without address translation
 - SIP, 2523, 2543
- installation on Vista, 1615
- installing
 - FortiBridge unit firmware, 2982
- installing firmware on a partition without a reboot, 166
- Instant Messaging (IM), 1949

- interface
 - 802.1Q trunk, 511, 521
 - accelerated NP2, 782, 1980
 - configuring, 338
 - dedicated monitoring, 2129
 - external, VLAN NAT example, 516
 - external, VLAN NAT/Route example, 516
 - failover, 2194
 - Gr, 2311
 - GRE, 1853
 - Gx, 2311
 - Gz, 2311
 - HA heartbeat, 2171
 - heartbeat, 2171
 - link status, 770
 - load balance virtual server, 2872
 - loopback, 1680
 - maximum number, 503, 535, 1924
 - monitor, 2194
 - pairs, 782
 - physical, 1942, 1946
 - point-to-point, 1944
 - proxy ARP, 2872
 - reserved management interface, 2129
 - software switch, 388
 - VDOM link, 1944
 - virtual interface, 1942
 - VLAN subinterface, 511, 512, 515, 517, 521
- interface index
 - hash map order, 2173
- interface mode, 2953, 2959
- interface mode IPSec, 2957
- interface monitoring, 2014
 - aggregate interfaces, 2058
 - definition, 2016
 - redundant interfaces, 2071
- interfaces, 195
 - aggregate, 577
 - AMC card, 376
 - ANY, ANY interface option, 219
 - DHCP, 384
 - disabling BFD, 1868
 - loopback, 383
 - MTU packet size, 387
 - physical, 375
 - PPPoE, 385
 - redundant, 383
 - secondary IP address, 388
 - source, firewall policy, 251
 - switch mode, 382
 - virtual domains, 389
 - virtual LANs, 390
 - wireless, 387
 - zones, 392
- interior gateway protocol (IGP), 1701
- International characters, 329
- International Mobile Equipment Identity (IMEI), 2419
- International Mobile Subscriber Identity (IMSI), 2417
- Internet Assigned Numbers Authority (IANA), 1710
- Internet Control Message Protocol (ICMP), 774
- internet gateway protocol (IGP), 1909
- Internet Service Provider (ISP), 1250
- Internet Traffic Management Practices (ITMP), 789
- Internet-browsing
 - configuring FortiClient, 1518
- Internet-browsing firewall policy
 - VPN server, 1516
- Internet-browsing IPsec configuration
 - FortiClient dialup-client configuration, 1517
 - gateway-to-gateway configuration, 1516
 - infrastructure requirements, 1515
 - overview, 1515
- Internet-Locator-Service
 - service, 208
- interval
 - changing the heartbeat interval, 2176
 - FortiBridge probe, 2973
 - log message, 1351
- interval, comfort clients, 2359
 - protection profile, 2360
- inter-VDOM
 - benefits, 1941
 - firewall policy, 1950
 - independent configuration, 1948
 - management configuration, 1942
 - management VDOM, 1949
 - meshed configuration, 1942, 1950
 - physical interface, 1941
 - stand alone configuration, 1942, 1947
 - virtual interface, 1942
- introduction
 - Fortinet documentation, 3040
- intrusion detected
 - HA statistics, 2151
- intrusion detection system, **see** IDS
- intrusion monitor, 1156
- Intrusion Monitor submenu, 145
- Intrusion Prevention, 2955
- Intrusion Prevention System (IPS), 2941, 2955
- intrusion prevention system, **see** IPS
- intrusion protection
 - explicit FTP proxy, 2837
 - explicit web proxy, 2817
 - signatures, 991
- intrusion protection system, **see** IPS
- IP, 2219
 - load balance virtual server, 2872
 - load balancing, 2897
- IP address
 - email filter black/white IP address list, 937
 - multicasting, 473
 - overlapping, 512
 - peer, 2686
 - private network, 3037
 - tunnel mode range, 1618
- IP address conservation
 - NAT, 2573
- IP address range
 - setting for L2TP VPN, 1261
 - setting for PPTP VPN, 1260
 - setting for SSL VPN, 1258
- IP Based authentication, 1890
- IP filter
 - carrier end point, 1344, 1345, 1354
- IP filter list, 1330

- IP filter, carrier end point, 2358
- IP header
 - differentiated services, 2269
- IP land, 2956
- IP monitoring
 - remote, 2200
- IP pool, 198
 - address matching, 199
 - policies and fixed ports, 199
 - proxy ARP, 2872
 - SIP, 2572
- IP port
 - HA schedule, 2219
- IP protocol 108, 1605
- IP range, 197
- IP Reservation, 165
- IP reservation, 573
- IP stack validation, 704
- IP, protocol 89, 1794
- IPcomp, 1605
- IPS, 2596
 - adding custom signatures, 949
 - buffer size, 966
 - concepts, 943
 - creating tags, 985
 - creating tags for predefined signatures, 993
 - custom signature, 995
 - custom signature keywords, 951
 - custom signature syntax, 950
 - disabling for pinholes, 2596
 - engine algorithm, 965
 - engine count, 965
 - explicit web proxy, 2817, 2837
 - fail-open, 965
 - filter, 946
 - filters, 984
 - in an HA cluster, 964
 - overview, 879
 - packet logging, 967
 - pre-defined overrides and custom overrides, 985
 - predefined signature viewer table, 105
 - protocol decoder, 996
 - protocol decoders, 966
 - scanning, 945
 - sensor, 945, 981
 - session count accuracy, 966
 - signature entry, 948
 - SYN
 - threshold, 990
 - SYN proxy, 990
 - understanding anomalies, 990
 - upgrading protocol decoder list, 996
 - viewing predefined signatures, 992
- IPS for GTP-U, 2426
- IPS signature threshold, 105
- IPS, one-armed, 1942
- IPSec, 2933, 2935, 2941, 2950, 2951, 2952, 2953, 2957, 2958, 2959
 - interface mode, 2957
 - tunnel, 2949
 - tunnel mode, 2957
- IPsec, 218, 1598, 1853, 2263
 - IKE negotiation, 179
 - SAs, 2193
 - security associations, 2193
 - server type, 1260
 - tunnel, 1597
- IPSec Interface Mode, 2954, 2957, 2959, 2960
- IPsec monitor, 1405
- IPsec Monitor submenu, 146
- IPSec VPN
 - adding manual key, 1403
 - Autokey list, 1393
 - concentrator list, 1404
 - configuring phase 1, 1394
 - configuring phase 1 advanced options, 1396
 - configuring phase 2, 1399
 - configuring phase 2 advanced options, 1399
 - Manual Key list, 1403
- IPsec VPN
 - and PCI DSS, 3031
 - authentication methods, 1412
 - authentication options, 1412
 - backup, 1541
 - certificates, 1412
 - configuring authentication for, 1258
 - DDNS routing, 1469
 - dialup users, access to, 1236
 - dialup users, configuring authentication for, 1258
 - extended authentication (XAuth), 1422
 - firewall IP addresses, defining, 1431
 - firewall IPsec policy, 1433
 - keeping tunnel open, 1427
 - load balancing, 2217
 - logging events, 1607
 - manual key, 1392
 - monitoring, dialup connection, 1603
 - monitoring, static or DDNS connection, 1603
 - peer, 1265
 - peer identification, 1415
 - phase 1 parameters, 1407
 - phase 2 parameters, 1425
 - role of encryption policy, 1434
 - route-based firewall policy, 1436
 - testing, 1604
 - troubleshooting, 1609
- IPsec VPN SA
 - synchronization, 2169
- ips-rtp, 2596
- IPv4, 2940
- ipv4 tunneling configuration, ipv6, 237

- IPv6, 235, 393
 - dual stack, 245
 - dynamic routing, 251
 - interfaces, 249
 - IPsec certificate configuration, 257
 - IPSec configuration, 257
 - IPSec phase 1 configuration, 257
 - IPsec phase 2 configuration, 258
 - IPsec routing configuration, 258
 - IPsec security policy configuration, 258
 - Neighbor Discovery (ND), 244
 - security policies, 251
 - SIP, 2585
 - static routing, 250
 - troubleshooting, 258
 - tunnel provider example, 246
 - tunneling, 245
- ipv6, 235
 - dual stack routing configuration, 236
 - ipv4 tunneling configuration, 237
 - remotely connecting over the Internet, 237
- ipv6 in FortiOS, 236
- IPv6 IPsec configurations
 - certificates, 1555
 - configuration, 1556
 - firewall policies, 1556
 - IPv4-over-IPv6 example, 1560
 - IPv6-over-IPv4 example, 1563
 - IPv6-over-IPv6 example, 1557
 - overview, 1555
 - phase 1, 1556
 - phase 2, 1556
 - routing, 1556
- IPX, layer-2 forwarding, 532
- IRC
 - service, 208
- ISAKMP, 1421, 2950
- ISO 8859-1, 330
- ITU-T E.164, 2418

J

- jumbo frames, 2936

K

- K-12, 579
- keepalive, 1421
- Keepalive Frequency, Phase 1, 1420, 1421
- key, 318
- key size
 - certificate, 1141
- keyboard shortcut
 - online help, 293
- Keylife
 - IPsec interface mode, 1429
- keylife, 1426
- Keylife, Phase 1, 1417, 1420
- Keylife, Phase 2, 1426
- keyword, 2008
- keywords
 - IPS custom signatures, 951

L

- l2ep-eth-type
 - CLI command, 2175, 2236
- L2TP, 1169, 2208
 - port 1701, 1568
 - service, 208
 - VPN, configuring authentication for, 1261
- L2TP Access Concentrator (LAC), 1568
- L2TP VPN
 - authentication method, 548
 - configuration steps, 547
 - enabling, 548
 - firewall IP addresses, defining, 548
 - infrastructure requirements, 547
 - network configuration, 546
 - security policy, defining, 549
 - VIP address range, 548
- L2TP-over-IPsec, 1567
- LACP, 2057, 2058
 - active-passive HA mode, 2058
- LACPDU, 2059
- lacp-ha-slave
 - CLI keyword, 2058, 2236
- LAG, 2057
- language
 - changing the web-based manager language, 312
- Layer 2, 2940
- Layer 2 Tunneling Protocol (L2TP), 1567
- Layer 3, 2940
- Layer 4, 2940
- layer 4, 704
- Layer-2, 779
- layer-2, 504, 507, 511
 - example, 505
 - forwarding, 532
 - frames, 504
- layer-2 bridge, FortiBridge, 2968
- layer-2 loops, 1923
- layer-2 switch, 767
 - troubleshooting, 2233
- Layer-3, 774
- layer-3, 507
 - packets, 504
- layer-3 router, 767
- layer-3 switch
 - failover issues, 2234
- LDAP, 745, 1306, 2134, 2135
 - access, collector agent, 1299
 - configuring server, 1184
 - fnbamd, 1214
 - service, 208
 - wildcard admin, 1212
- LDAP authorization, 357
- LDAP server, external
 - for L2TP, 548
 - for PPTP, 539
 - for XAuth, 1422

- LDAP servers, 1207, 1307
 - authenticating users with, 1224
 - binding, 1208
 - common name, 1208
 - configuring FortiGate unit to use, 1209
 - directory, 1207
 - Distinguished Name Query list, 1211
 - distinguished names, 1208
 - domain component, 1208
 - hierarchy, 1208
 - protocols, 1208
 - RFC compliance, 1208
- least round trip time
 - load balancing, 2873
- least RTT
 - load balancing, 2873
- least session
 - load balancing, 2873
- Least-Connection
 - HA schedule, 2218
- length, 350
- length, password, 350
- license key, 1893
- licenses
 - viewing, 303
- life of a packet, 185, 699, 1390
 - UDP, 699
- Lightweight Directory Access Protocol (LDAP), 1283
 - FortiGate configuring, 1308
 - XAuth authentication with, 1259
- limited bandwidth, 2253
- limiting
 - number of SIP dialogs, 2593
- line endings, 332
- link
 - failure, 2168
 - multiple link failures, 2199
- link aggregation, 2940, 2941
- Link Aggregation Control Protocol, 2058
- link failover, 2167, 2194
 - active-active, 2218
 - active-passive, 1998
 - aggregate interfaces, 2058
 - definition, 2016
 - not detected by high-end switches, 2198
 - redundant interfaces, 2071
- link failure
 - remote, 2200
- link redundancy, 561
- link status, 770
- link-failed-signal, 2198
 - CLI, 2198
- link-state advertisement (LSA), 1702
- Linux, 775, 777, 1278
- lists
 - using web-based manager, 290
- load balance
 - according to loading, 2223
 - cpu usage, 172, 2223
 - explicit web proxy, 2217
 - first alive, 2873
 - health check monitor, 2876
 - health check monitoring, 2876
 - health monitoring, 2876
 - how busy, 2223
 - HTTP host, 2874
 - HTTP multiplexing, 2217
 - HTTPS, 2217
 - IM, 2217
 - IPsec VPN, 2217
 - least RTT, 2873
 - least session, 2873
 - memory usage, 172, 2223
 - P2P, 2217
 - proxy UTM sessions, 172, 2223
 - round robin, 2873
 - schedule, 2221
 - source IP hash, 2873
 - SSL offloading, 2217
 - SSL VPN, 2217
 - static, 2873
 - virtual server IP, 2872
 - VoIP, 2217
 - WAN optimization, 2217
 - WCCP, 2217
 - weighted, 2873
- Load Balance Monitor submenu, 144
- load balancing, 1992, 2706, 2869, 2941, 2946
 - active-active, 1999, 2217
 - basic example, 2882
 - definition, 2016
 - HTTP host connections, 178
 - HTTP host-based, 2888
 - IP, 2897
 - load-balance-all, 2217
 - monitoring, 2879
 - real servers, 2874
 - SSL, 2889
 - SSL offloading, 2890
 - TCP, 2897
 - traffic not load balanced by active-active HA, 2217
 - UDP, 2897
- load balancing HTTP host connections, 178
- load sharing, 564
- load-balance-all, 2221
 - best practice, 2013
 - enabling, 2217
- Local certificates
 - generating request, 1267, 1277
 - installing signed, 1269
- local console access, 315
- local domain name, 574
- Local Gateway IP, 1597, 2949, 2952, 2953, 2954, 2957, 2958, 2959, 2960
- local host, 2941, 2949, 2951
- Local ID
 - for certificates, 1414
 - to identify FortiGate dialup clients, 1502

- Local Interface
 - IPSec VPN, manual key, 1404
 - Local SPI
 - IPSec VPN, manual key, 1403
 - Local SPI, Manual Key, 1552
 - local user, 1171
 - local users
 - configuring, 1224
 - creating, 1224
 - deleting from FortiGate configuration, 1226
 - removing from FortiGate configuration, 1226
 - local wifi radio
 - configuring, 2510
 - local-in firewall policies, 177
 - local-in policy, 226
 - Location Area Identity (LAI), 2311
 - location server
 - SIP, 2522
 - locking configuration, 398
 - lockout
 - administrator, 351
 - log
 - traffic, firewall policy, 252
 - log device
 - selecting, 639
 - log file, 1306
 - log message
 - grouping, 1351
 - HA, 2137
 - interval, 1351
 - viewing, 664
 - log message, FortiBridge, 2989
 - sample, 2994
 - log message, FortiGate, 625
 - log messages
 - example log message scenarios, 2138
 - HA, 2138
 - primary unit removed from cluster, 2139
 - logging, 1903, 2136
 - alert email did not send, 802
 - cannot log to log device, 802
 - configuring a FortiBridge unit, 2994
 - downloading quarantine logs, 667
 - dynamic profile, 1350
 - enabling reliable syslog, 646
 - enabling SSL VPN events, 1648
 - example log message scenarios, 2138
 - FortiGate stopped recording logs, 802
 - GTP settings, 2349
 - HA log messages, 2138
 - IPS packets, 984
 - management practices, 631
 - MMS profile, 2327
 - setting event-logging parameters, 1647
 - specifying GTP packets, 2351
 - syslog, FortiBridge, 2993
 - viewing quarantine logs, 666
 - logging in, security messages, 1264
 - Logging Monitor submenu, 134
 - logging out
 - web-based manager, 313
 - logging VPN Events, 544, 1607
 - logging VPN events, 1607
 - login
 - restricting unwanted, 351
 - login grace timer for SSH connections, 159
 - logon blackout period, 1245
 - logon events
 - logging to memory, 1300
 - logs, 658
 - antivirus, 658
 - application control, 657
 - archives, DLP, 659
 - attack, IPS, 658
 - chat message support for MSNP21, 135
 - data leak prevention, 657
 - email filter, 659
 - event, 656
 - managing for individual cluster units, 2136
 - nac quarantine, 657
 - network scan, 659
 - other traffic, 656
 - other-traffic, 135
 - packet, 659
 - traffic, 656
 - loopback interface, 1680
 - loopback interfaces, 383
 - loose source record route, 2956
 - lost heartbeat threshold
 - changing, 2175
 - lost password
 - recovering, 312
- ## M
- M3UA, 268
 - MAC
 - MAC forwarding table, 2198, 2234
 - MAC address, 533
 - aggregate interfaces, 2058
 - redundant interfaces, 2071
 - virtual, 2177
 - VRRP virtual, 174, 2238
 - MAC filter, wireless, 2453
 - MAC forwarding tables, 2198, 2234
 - Mac OS, 1278, 1286
 - MAC table, 779, 1924
 - Main Interface IP, 2957
 - main mode, 1407
 - maintenance
 - configuration revision, 366
 - disk, 619
 - malformed-request-line, 2588
 - manage cluster units
 - HA, 2163
 - managed access points
 - FortiAP, 2510, 2511
 - management access, 350
 - management access to the FortiBridge EXT1 interface, 2981
 - management checksum, FortiManager, 161
 - management configuration, 1949
 - Management Information Base (MIB), 429
 - management interface
 - reserved, 2129

- management IP, 343
 - changing the FortiBridge management IP address, 2979
- FortiBridge, 2968
- management IP address
 - changing, 298
- management services, 1894
- management VDOM, 1330, 1876, 1879, 1894, 1896, 1897, 1942
- managing access points
 - local wifi radio, 2510
- Manual Key
 - IPSec VPN, 1403
- manual key IPsec configuration
 - configuration steps, 1552
 - overview, 1551
- MAPI, 2673
 - protocol optimization, 2689
- Martian addresses, 1681
- master DNS server, 575
- master unit, 2947
 - See Also primary unit, 1992
- matched content, 2878
 - HTTP health check monitor, 2877
- matching
 - security policy, 2708
- match-vip, 205
 - how to, 205
- max cache object size
 - web cache, 2753
- Max HTTP message length
 - web cache, 2754
- Max HTTP request length
 - web cache, 2754
- max TTL
 - web cache, 2754
- max-body-length, 2590
- max-dialogs, 2593
- maximum age
 - spanning tree parameter, 2234
- maximum bandwidth, 2255, 2262, 2264, 2265
 - firewall policy, 2264, 2265
- maximum connections
 - real server, 2874
- maximum file size
 - antivirus, 903
- maximum frame size, 2936
- Maximum Transmission Unit (MTU), 238
- MD5, 1598, 2950
 - OSPF authentication, 1858
- memory, 536, 1892
- memory constraints, 2371
- memory usage, 772
 - HA statistics, 2150
 - WAN Optimization, 2690
 - web caching, 2690
 - weight, 172, 2224
 - weighted load balancing, 172, 2223
- merge interfaces, 388, 585
- meshed configuration, 1942, 1950
- meshed VPN, 1437
- message
 - SIP, 2525
- message fingerprint, 2386
- message flood, 2389
 - alert notifications, 2375
 - threshold, 2375, 2377
- message length
 - SIP, 2590
- message request-line
 - SIP, 2530
- message start line
 - SIP, 2530
- message status-line
 - SIP, 2530
- Message Transfer Part 3, 268
- MGCP, 557
 - service, 208
 - session helper, 557
- MIB, 441, 2143
 - FortiGate, 434
 - HA, 2143
 - RFC 1213, 434
 - RFC 2665, 434
 - wan opt, web cache and explicit proxy, 167
- Microsoft Challenge-Handshake Authentication Protocol v1 (MSCHAP), 1215
- Microsoft Point-to-Point Encryption (MPPE), 538
- Microsoft Windows, 1391
- Microsoft Windows VPN, 1567
- middle-man, 702
- min TTL
 - web cache, 2754
- missing MED, 1758
- MM1 message blocking MMS filter list, 1342, 1352
- MM4, 1342, 1352
- MM7, 1342, 1352
- MMS
 - address translation, 2295
 - carrier end point filtering, 2318
 - content checksum, 2328
 - DLP archive, 2379
 - file filtering, 2363
 - flood prevention, 2373
 - notifications, 2295
 - profile, 2315
 - virus scanning, 2354
- MMS Address Translation, 1347
- MMS filter list
 - adding to an MMS protection profile, 1342, 1352
- MMS message
 - blocking using carrier end point MMS filtering, 1342, 1352
- MMS message flood, 2331
- MMS profile
 - DLP archive, 2326
 - logging, 2327
 - MMS address translation, 2322
 - MMS bulk email filtering, 2319
 - MMS notification, 2324
 - scanning, 2317
- MMS protection profile, 2304
 - adding to a security policy, 2305
 - content archive, 2371

- MMS Service Provider Network (MSPN), 2292
- Mobile Country Code (MCC), 2417
- mobile device
 - Symbian Series 60, 2353
- Mobile Network Code (MNC), 2417
- mobile service provider, 1326
- Mobile Station (MS), 2306
- Mobile Station Identification Number (MSIN), 2417
- Mobile Subscriber Integrated Services Digital Network (MSISDN), 2355, 2375, 2418
- Mobile Subscriber Integrated Services Digital Network Number (MSISDN), 2321
- mode
 - real server, 2874
 - switching between FortiBridge modes, 2974
- Mode, Phase 1, 1410, 1411
- modem, 567
 - routing, 570
- modem modes, 568
- Modem Monitor submenu, 144
- modes of operation
 - overview, 1613
 - port forwarding, 1615
 - tunnel mode, 1614
 - web-only mode, 1614
- monitor
 - application control, 1068
 - HA statistics, 2150
 - how a FortiBridge unit monitors a FortiGate unit, 2970
 - interface, 2194
 - load balancing, 2879
 - port, 2194
 - routing, 1869
 - WAN Optimization peers, 146
 - web cache, 146
- Monitor submenus
 - Application Monitor, 146
 - Archive & Data Leak Monitor, 145
 - AV Monitor, 145
 - DHCP Monitor, 143
 - Email Monitor, 145
 - Intrusion Monitor, 145
 - IPsec Monitor, 146
 - Load Balance Monitor, 144
 - Logging Monitor, 134
 - Modem Monitor, 144
 - Policy Monitor, 144
 - Session Monitor, 144
 - SSL-VPN Monitor, 146
 - Traffic Shaper Monitor, 144
 - Web Monitor, 145
- monitor submenus, 143
- monitored interface
 - definition, 2016
 - primary unit selection, 2002
- monitoring
 - administrators, 302
 - antivirus, 1156
 - applications, 1158
 - archives and dlp, 1158
 - attacks, 1156
 - DHCP, 574
 - email activity, 1157
 - endpoints, 1159
 - FortiGuard quota, 1159
 - ips, 1156
 - IPsec sessions, 1405
 - proxy forwarding, 2814
 - RAID, 308
 - rogue APs, 2470, 2515
 - routing table, 1869
 - traffic shapers, 2276
 - WAN Optimization, 2703
 - WAN optimization, 2690
 - web activity, 1157
 - web caching, 2755
 - wireless clients, 2469, 2514
- monitoring ISIS, 176
- more, 332
- moving a WAN optimization rule, 2709
- Mozilla, 1256
- m-retrieve-conf, 2396
- MS RPC, 555
- MS Windows, 776
- MS Windows Active Directory (AD), 1167, 1255
- MS-CHAP, 1215
- m-send-conf, 2396
- m-send-req, 2396
- MSIE, 1256
- MSISDN, 2391
- MSISDN. **See also** carrier end point, 1327
- MS-SQL
 - service, 209
- MTP3 User Adaptation Layer, 268
- MTU (Maximum Transmission Unit), 1598, 2936, 2941
- MTU config support (non-IPsec tunnels), CLI, 162
- MTU packet size, interface, 387
- multicast, 1863
 - configuring, 1865
 - dense mode, 472, 1864
 - DNAT, 1866
 - IGMP, 472, 1864
 - RFC 3973, 471
 - RFC 4601, 471
- multicast destination NAT, 1866
- multicast-enable command, 476
- multicasting
 - debugging example, 484
 - enabling, 476
 - IP addresses, 473
 - RIPv2, 474
 - security policies, 475
- Multi-Exit Discriminator (MED), 1758
- multi-line command, 327
- Multimedia Broadcast and Multicast Services (MBMS), 2410
- Multimedia Message Service Center (MMSC), 2292
- Multipath routing, 1681

- multiple account sessions, limits, 1257
- multiple group enforcement, 118, 1238
- multiple pages, 332
- MYSQL
 - service, 209

N

- Name
 - IPSec VPN, manual key, 1403
- naming rules, 1895
- NAT, 265, 1946, 2706
 - keepalive frequency, 1421
 - multicast, 1866
 - port translation (NAT-PT), 558
 - SDP, 2575
 - SIP ALG IP address conservation, 2573
 - SIP ALG NAT tracing, 2573
 - SIP contact headers, 2574
 - SIP session helper NAT tracing, 2574
 - traversal, 1420, 1610
 - VLAN example, 517
 - with IP address conservation, 2573
- NAT device
 - authentication, 2816
- NAT mode, 297
- NAT/Route mode, 2687
 - general configuration steps, 2022, 2059, 2072, 2094
 - HA network topology, 2022
 - reserved management interface, 2130
 - web-based manager configuration steps, 2023, 2027, 2034, 2039, 2047, 2050
- nat-trace, 2574
- Nat-traversal, Phase 1, 1420
- negative response duration
 - web cache, 2753
- negotiating
 - IPsec phase 1 parameters, 1417
 - IPsec phase 2 parameters, 1426
- Neighbor
 - Advertisement, 245
 - Solicitation, 245
- nested tunnels, 2403
- NetBIOS, for Windows networks, 534
- netlink, 779
- netmask
 - wildcard firewall addresses, 202
- NetMeeting
 - service, 209
- netscan asset auth, CLI, 121
- network
 - topology, 2951, 2957
 - train, 2178
- Network Access Server (NAS), 1203, 1205
- Network Address Translation (NAT), 1420
- network equipment
 - failover time, 2235
- Network Identifier, 2418
- network instability, 533
- network interface card (NIC), 815
- network processing unit (NPU), 2947, 2950
- network processor accelerated interfaces
 - accelerate active-active HA, 2220
- network processors
 - FA2 (NP1), 2935
 - FortiAccel (NP1), 2935
 - NP1, 2935
 - NP2, 2935
 - NP4, 2935
- network protocol usage, 142
- Network Time Protocol (NTP), 720, 745
- Network Time Protocol server (NTP), 299
- network topologies, 2459
- network topology
 - dynamic DNS, 1469
 - FortiClient dialup-client, 1483
 - FortiGate dialup client, 1501
 - fully meshed network, 1437
 - gateway-to-gateway, 1437
 - hub-and-spoke, 1453
 - Internet-browsing, 1515
 - manual key, 1551
 - NAT/Route mode HA, 2022
 - partially meshed network, 1437
 - redundant-tunnel, 1519
 - supported IPsec VPNs, 1391
 - transparent mode VPN, 1543
- network utilization
 - HA statistics, 2151
- network vulnerability scan
 - asset definition, 121, 1124
 - scan schedule, 121, 1125
 - vulnerability result, 121, 1126
- next, 325
- NFS
 - service, 209
- nic
 - get hardware nic, 2024, 2029, 2036, 2041, 2048, 2051, 2061, 2066, 2074, 2079, 2096, 2101, 2116, 2121, 2181
- NNTP
 - service, 209
- no SYN-ACK, 727
- non-conserve mode, 727
- non-dynamic profile sessions, 1331
- none
 - HA schedule, 2218
- non-HTTP sessions
 - HTTP rule, 2714
- normal mode, FortiBridge, 2968, 2970
 - monitoring the FortiGate unit, 2970
 - probe, 2970
 - resuming from bypass mode, 2996
 - switching to, 2974
 - switching to bypass mode, 2974
 - traffic flow, 2970
- no-sdp-fixup, 2575
- notification alerts, 2367, 2381
- Not-so-stubby Area (NSSA), 1857
- not-so-stubby area (NSSA), 1675, 1870
- Novell eDirectory, 1284
- Novell eDirectory - see Directory Service
- NP1, 2220, 2936, 2941, 2950

- NP1 processor, 2935
- NP2, 2220, 2936
- NP2 interface, 782, 1980, 2265
- NP2 interfaces, 1906, 2429
- NP2 network processor, 1597
- NP2 network processors, 1597
- NP2 processor, 2935
- NP4, 2220
- NP4 processor, 2935
- NT LAN Manager (NTLM), 1167, 1255, 1284
- NTLM, 1250
 - authentication, 1292
 - Guest profile access, 1255
 - HTTP-User-Agent, 1256
 - NTLM mode, 1287
 - user agent strings, 1256
 - XSS vulnerability, 1256
- NTLM authentication, CLI, 120
- NTLM enabled browsers, 1256
- NTLM statistics, 1304
- NTP
 - service, 209
- NTP server, 346
- NT-style domain mode implementation, 1286
- null modem, 316, 317

O

- object, 320
- object identifier (OID), 441
- objects
 - viewing reference, 379
- OID, 2143, 2144, 2145, 2146
- ONC-RPC, 555, 557
 - service, 209
- one time passcode (OTP), 1166, 1217, 1230
- one-armed IDS, 880
- one-armed IPS, 1942
- Online Certificate Status Protocol (OCSP), 1265, 1266
- online help
 - content pane, 291
 - keyboard shortcuts, 293
 - navigation pane, 292
 - search, 292
 - using FortiGate online help, 291
- open shortest path first (OSPF). See routing, OSPF
- Open Systems Interconnect (OSI), 504
- OpenLDAP, 1207
- OpenSSL, 1266
- Opera, 1256
- operating a cluster, 2127
- operating mode
 - active-passive, 1998
- operating modes, FortiBridge
 - switching between, 2996
- operating principles, 2967
- operating temperature, 401
- operation mode, 298
 - active-active, 1998
- option, 320
- order of operations for shapers, 2261
- OS
 - host patch check, 1646
- OS patch check, 1643
- OSFP
 - graceful restart, 2192
- OSI
 - Layer-2, 779
 - Layer-3, 774
- OSPF
 - advanced, 1856
 - area ID, 1859
 - areas, 1857
 - AS, 1855
 - authentication, 1858
 - configuring, 1854
 - Dead Interval, 1860
 - dead packets, 1860
 - IPv6, 251
 - network, 1855
 - networks, 1858
 - NSSA, 1857, 1870
 - operating parameters for interface, 1859
 - protecting with IPsec, 1589
 - regular area, 1857
 - service, 209
 - settings, 1854
 - stub, 1857
 - virtual lan, 1859
 - virtual link, 1857
 - with redundant IPsec tunnels, 1595
- ospf
 - adjacent routers, 1795, 1800
 - area, 1794
 - area border router (ABR), 1794
 - Dijkstra's algorithm, 1796
 - e1, 1675
 - e2, 1675
 - Hello packets, 1794
 - Hello protocol, 1795
 - IP datagrams, 1794
 - link-state, 1794
 - neighbor, 1795
 - NSSA, 1675
 - path cost, 1796
 - state of neighbor, 1800
- OSPF AS
 - defining, 1854
- ospf AS, 1790
- OSPFv3 NSSA extension, CLI, 122
- other-traffic logs, 135
- out of band management, 2129
- out of path
 - topology, 2675
- Outbound NAT, encryption policy, 1433
- outgoing-ip
 - explicit FTP proxy, 2837
 - explicit HTTP proxy, 2830
- overlap
 - resolving IP address, 1502
 - resolving through FortiGate DHCP relay, 1502
- overlapping VPN subnets, 1444

- override, 2008
 - and primary unit selection, 2008
 - configuration changes lost, 2011
 - disconnecting a unit from a cluster, 2012
 - primary unit selection, 2006, 2010
- overrides
 - web filtering, 106
- oversize threshold, 2319, 2360

P

- P1 Proposal, Phase 1, 1417, 1419
- P2 Proposal, 2957
 - IPSec VPN, phase 2, 1400
- P2 Proposal, Phase 2, 1426
- P2P, 880
 - load balancing, 2217
- PAC
 - explicit web proxy, 2807, 2827
- packet
 - flow, 703, 782
 - forwarding rate, 2934, 2935, 2951, 2957
 - gratuitous ARP, 2178
 - ICMP, 274
 - life of, 185, 699
 - processing flow, 2933
 - sniffer, 780
- Packet Data Protocol (PDP), 2409
- packet data protocol (PDP), 2306
- packet flow, 186, 2933
- packet header, 422
- packet logging
 - application control, 1070, 1078
 - custom IPS overrides, 986
 - IPS, 967, 984
 - settings, 1144
 - viewing and saving logged packets, 1144
- packet rates, 2255
- packet sniffer, 1979, 2429
 - verbosity level, 1979, 2430
- Packet verification, 704
- packets
 - flow, 186
 - layer-3 routing, 507
 - VLAN-tagged, 512
- PADT timeout, 339
- page controls
 - web-based manager, 290
- paging, 332
- PAP, 537, 1215, 1260
- parity, 316
- partially meshed VPN, 1437
- password, 350
 - adding to a FortiBridge, 2979
 - authentication group, 2700
 - changing, administrator, 303
 - configuring authentication, 349
 - forgot administrator password, 312
 - HA configuration option, 1999
 - recovering lost password, 312
- Password Authentication Protocol (PAP), 1215, 1423

- password policy
 - PCI DSS requirements, 3034
- password, changing, 348
- patch check
 - host OS, 1646
- pattern, 322, 1042
 - allow, 1042
 - block, 1042
 - carrier end point, 1342, 1343, 1345, 1352, 1353, 2334, 2357
 - creating, 1043
 - default file block list, 2364
 - default list of file block patterns, 1056
- PC-Anywhere
 - service, 209
- PCI DSS, 3030
 - defined, 3023
 - example network topology, 3027
 - firewall policy considerations, 3028
 - logging wireless network activity, 3031
 - objectives and requirements, 3024
 - scanning for rogue wireless APs, 3029
 - wireless guidelines, 3026
- peer, 2152
 - accept any peer, 2697
 - host ID, 2686
 - IP address, 2686
 - monitoring WAN optimization, 2703
 - WAN optimization, 2697
 - WAN Optimization monitor, 146
- peer authentication, 2698
 - WAN optimization, 2697
- peer ID
 - assigning to FortiGate unit, 1414
 - enabling, 1415
- Peer option
 - IPSec VPN, phase 1, 1395
- Peer Options, 1409
- peer user groups
 - configuring, 1240
 - creating, 1240
- peer users, 1223, 1227
 - configuring, 1227
 - creating, 1227
- peer-to-peer
 - WAN optimization rules, 2705
- per policy shaper, 2262, 2263
- perfect forward secrecy (PFS), 1448
- perfect forward secrecy, enabling, 1426
- performance
 - improving session pickup performance, 2205
- periodic
 - synchronization, 2185
- per-IP, 2265
 - NP2 interface, 2265
- per-IP shaper, CLI, 161
- per-IP traffic shaping, 2265
- Perl regular expression
 - carrier end point pattern, 1342, 1352
- Perl regular expressions
 - carrier end point pattern, 1345, 2334
- Perl regular expressions, using, 333

- permissions, 326
- persistence, 2873
 - HTTP cookie, 2886
 - HTTP/HTTPS, 2886
- Phase, 1399
- Phase 1, 2950, 2952, 2953, 2954, 2957, 2958, 2959, 2960
- phase 1
 - IPSec VPN, 1394, 1399
- phase 1 advanced options
 - IPSec VPN, 1396
- phase 1 parameters
 - authenticating with certificates, 1409
 - authenticating with preshared keys, 1410
 - authentication method, 1412
 - authentication options, 1412
 - defining, 1407
 - defining the tunnel ends, 1408
 - IKE proposals, 1418
 - main or aggressive mode, 1408
 - negotiating, 1417
 - overview, 1407
 - peer identifiers, 1414
 - user accounts, 1415
- Phase 2, 2951, 2952, 2953, 2954, 2957, 2958, 2959, 2960
- phase 2
 - Autokey keep alive, 1427
 - IPSec VPN, 1399
 - key expires, 1426
 - PFS, 1448
- phase 2 advanced options
 - IPSec VPN, 1399
- phase 2 parameters
 - autokey keep alive, 1427
 - auto-negotiate, 1426
 - configuring, 1428
 - defining, 1425
 - DHCP-IPsec, 1427
 - keylife, 1426
 - negotiating, 1426
 - perfect forward secrecy (PFS), 1426
 - quick mode selectors, 1427
 - replay detection, 1426
- Phase I, 1597, 2949
- Phase II, 1598, 2950
- physical interface, 1941, 1942, 1946
- PIM
 - dense mode, 1863
 - RFC 2362, 1863
 - RFC 3973, 1863
 - sparse mode, 1863
- PING
 - service, 209
- ping, 774
 - enabling FortiBridge ping probes, 2991
 - FortiBridge probe, 2971
 - health check monitor, 2882
- ping server, 562, 569
- PING6
 - firewall service, 209
- pinhole
 - disabling IPS, 2596
 - more secure, 2578
 - RTP, 2520, 2525
 - SIP, 2520, 2525
 - smaller, 2578
 - strict-register, 2578
- PKI, 1194
- PKI authentication - see peer users
- PKI certificate
 - PKCS, 1269
- PKI certificate authentication, 119
- PKI user, 1227
- planning VPN configuration, 1390
- pmap
 - session helper, 557
- PMK caching, 2443
- point-to-point interface, 1944
- Point-to-Point Tunneling Protocol (PPTP), 537
- policies, 218
 - column settings, 218
 - expiry, 213
 - ICMP packets, 274
 - identity-based, 222
 - multicast, 475
 - NAT to transparent mode, 274
 - order, 219
 - timeout, 213
 - viewing, 219
- policy, 2941
 - comments, 255
 - guaranteed bandwidth, 2264
 - insert policy before, 2708
 - local-in, 226
 - log traffic, 252
 - matching, 2708
 - maximum bandwidth, 2264, 2265
 - menu, 139
 - schedule, 252
 - security, 2685
 - service, 252
 - traffic priority, 2707
 - traffic shaping, 254
- policy 0, 225
- Policy Monitor submenu, 144
- policy route
 - moving in list, 1696
- policy server, VPN
 - configuring FortiGate unit as, 1491
- policy-based routing, 1845
- policy-based VPN
 - vs. route-based, 1390
- Polling mode, 1287, 1304
- polling mode
 - event log polling, 1286
 - NetAPI polling, 1287
- POP3
 - probe, FortiBridge, 2972
 - service, 209

POP3S

- antivirus, 1142
- antivirus quarantine, 1142
- data leak prevention, 1143
- DLP archive, 1143
- email filtering, 1143
- predefined firewall services, 1141
- protocol recognition, 1142

port

- blocking port 25, 270
- forwarding, 1615
- number, web-portal connections, 1638
- virtual server, 2872

port 10443, 1236

port 1701, 1568

port 179, 1760

port 21123, 2349

port 212, 2409

port 2152, 2410, 2426

port 445, 1286

port 4500, 1420

port 47, 557

port 500, 1420

port 8000, 1306

port monitor, 2194

- virtual clustering, 2091

port monitoring, 2014

- aggregate interfaces, 2058
- redundant interfaces, 2071

port number

- changing the port numbers that the SIP ALG listens on, 2545
- changing the port numbers that the SIP session helper listens on, 2538

port pair (transparent mode), 164

port, RADIUS servers, 1205

port, session helper, 552

ports

- 8000 and 8002, 874
- port 1024, 749
- port 1025, 749
- port 443, 781
- port 53, 749
- port 8888, 749
- services, 206
- UDP ports 33434-33534, 776

position

- identity-based policy, 285

power

- security consideration, 2437
- WLAN power level, 2436

power failure

- FortiBridge, 2974

power off, 403, 2927

powering on, 871

PPP, 2012

PPPoE, 2012

PPPoE interface, 385

PPTP, 1169, 2208

- external server, 542
- layer-2 forwarding, 532
- service, 210
- session helper, 557

PPTP PPTP VPN

- configuring authentication for, 1260

PPTP VPN

- authentication method, 539
- configuring pass through, 542
- enabling, 540
- FortiGate implementation, 537
- IP address range, 1260
- security policy, defining, 541
- VIP address range, 540

PRACK

- SIP message, 2529

practices, firmware upgrade, 99

pragma-no-cache, 2755

pre-authentication, 2443

predefined firewall services

- IMAPS, POP3S, SMTPS, 1141

predefined services, 206

predefined signature viewer table, 105

predefined signatures

- adding tags, 152

preempt mode

- VRRP, 2242

preserve-override, 2574

pre-shared key

- authenticating FortiGate unit with, 1411
- authentication group, 2700

preshared key, 1168, 1385

Pre-shared Key, Phase 1, 1411

primary cluster unit

- definition, 2016

primary unit, 1992, 2947

- connected monitored interfaces, 2002
- definition, 2016
- getting information using SNMP, 2143
- override keyword, 2008
- recovery after a failover, 2169
- selection, 2000
- SNMP get, 2143

primary unit selection

- age, 2002, 2003
- basic, 2001
- device priority, 2001, 2006
- FortiGate unit serial number, 2007
- interface monitoring, 2002
- monitored interfaces, 2002
- override, 2006, 2008, 2010
- serial number, 2007

priority

- cluster members, 2148
- heartbeat interface, 2171

priority traffic, 2262

probe interval

- tuning a FortiBridge unit, 2992

- probe list
 - FTP, 2991
 - ping, 2991
 - SMTP, 2991
 - probe, FortiBridge, 2970
 - action on failure, 2989
 - and FortiGate firewall policies, 2971
 - configuring, 2988
 - configuring FortiGate unit, 2990
 - configuring probe settings, 2989
 - default FortiBridge settings, 2989
 - enabling, 2990
 - enabling FortiBridge ping probes, 2991
 - enabling probes, 2990, 2991
 - fail open, 2989
 - FortiBridge dynamic IP pattern, 2989
 - FortiGate hardware failure, 2973
 - FortiGate session list, 2992
 - FortiGate software failure, 2973
 - FortiGate unit serial number, 2989
 - FTP, 2972
 - HTTP, 2972
 - IMAP, 2972, 2991
 - interval, 2973
 - ping, 2971
 - POP3, 2972
 - settings, 2989
 - SMTP, 2972, 2991
 - threshold, 2973
 - verifying, 2992
 - viewing probe configuration, 2991
 - problem scope, 714
 - product registration, 3040
 - profile
 - VoIP, 2544
 - profile group, 110, 1154
 - assigning dynamically, 1164
 - dynamic profile, 1164
 - profile group, dynamic profile, 1332
 - profile, dynamic, 1325
 - proposal
 - IPSec VPN, phase 2, 1400
 - protection profile
 - amount, comfort clients, 2360
 - interval, comfort clients, 2359, 2360
 - protocol
 - ospf Hello, 1795
 - service, 206
 - protocol anomaly attacks, 2402
 - protocol decoder, 996
 - protocol decoders, 966
 - Protocol Independent Multicast (PIM), 471, 1863
 - protocol optimization, 2673
 - CIFS, 2689
 - FTP, 2689
 - HTTP, 2689
 - MAPI, 2689
 - TCP, 2689
 - protocol options, 177
 - explicit FTP proxy, 2837
 - explicit web proxy, 2817
 - FTPS support, 177
 - protocol recognition
 - HTTPS, IMAPS, POP3S, SMTPS, 1142
 - protocol, session helper, 552
 - protocols
 - authentication, 1246
 - LDAP servers, 1208
 - provisional
 - SIP response message, 2529
 - provisional response acknowledgement
 - SIP message, 2529
 - provisional-invite-expiry-time, 2548
 - proxy
 - antivirus, 2817, 2837
 - DLP, 2817, 2837
 - explicit web, 2208
 - explicit web proxy authentication, 2816
 - FortiGuard web filtering, 2817
 - protocol options, 2817, 2837
 - web filtering, 2817
 - proxy ARP, 2872
 - FortiGate interface, 2872
 - IP pool, 2872
 - virtual IP, 2872
 - proxy auto-config
 - explicit web proxy, 2807, 2827
 - proxy chaining
 - explicit web proxy, 124, 2813
 - health monitoring, 2814
 - proxy forwarding server
 - explicit web proxy, 124, 2813
 - health checking, 2814
 - proxy FQDN
 - web cache, 2754
 - proxy inspection, 702
 - proxy server
 - SIP, 2521
 - proxy UTM sessions
 - weighted load balancing, 172, 2223
 - PSTN, 267
 - public key cryptography standard (PKCS), 1267
 - public key infrastructure (PKI), 1263
 - Public Land Mobile Network (PLMN), 2311
 - Public Switched Telephone Network
 - See PSTN, 267
 - publis DNS server, 576
 - purge, shell command, 325
 - push update, 409, 410
 - override, 411
- ## Q
- QoS, 2941, 2948
 - QUAKE
 - service, 210
 - quality of service, 2251
 - quarantine, 905, 2387
 - file, 2146

- quarantine files list
 - apply, 666
 - date, 666
 - DC, 667
 - file name, 666
 - filter, 666
 - service, 666
 - sorting, 666
 - status, 666
 - status description, 667
 - TTL, 667
 - upload status, 667
- quarantine logs, 666
- quarantining attackers to banned user list, 986
- Query list
 - LDAP Distinguished Name, 1211
- queuing, 2252
- Quick mode selectors, Phase 2, 1427
- quota
 - FortiGuard Web Filter, 1027

R

- RADIUS, 745, 1512, 2134, 2135
 - accounting system, 1330
 - assigning client IPs with, 1487
 - attributes, 1202
 - authentication servers, 1201
 - dynamic profile, 1164
 - End record, 1325
 - framed-ip-addr field, 1336
 - long key, 1202
 - server port, 1338
 - servers, 1182
 - Start record, 1164, 1325, 1327, 1341
 - Stop record, 1330
 - vendor ID, 1203
 - viewing server list, 1182
 - XAuth authentication with, 1259
- RADIUS server, external
 - for L2TP, 548
 - for PPTP, 539
 - for XAuth, 1422
- RADIUS servers
 - attribute dictionary, 1204
 - authenticating users with, 1224
 - changing default port, 1205
 - configuring FortiGate unit to use, 1204
 - default port, 1205
 - port, 1205
 - VSA, 1203
- RAID, 2961
 - configuring, 2962
 - levels, 2961
 - rebuilding an array, 2964
- random
 - HA schedule, 2219
- RAS, session helper, 556
- rate limit
 - number of SIP dialogs, 2593
- rate limiting
 - SCCP, 2592
 - SIMPLE, 2592
 - SIP, 2592
- rate limits, 2941
- RAUDIO
 - service, 210
- read & write access level
 - administrator account, 300
- read only access level
 - administrator account, 300
- real server
 - active mode, 2874
 - adding, 2875
 - disabled mode, 2874
 - health check monitoring, 2876
 - health monitoring, 2876
 - load balancing, 2874
 - maximum connections, 2874
 - mode, 2874
 - standby mode, 2874
 - weight, 2874
- Real Time Control Protocol, 2549
- Real Time Protocol, 2559
- realm
 - explicit web proxy, 2812
- real-time session, diag, CLI, 160
- reboot
 - installing FortiBridge firmware, 2983
- reboot, upgrade, 374
- record route option, 2956
- recording log messages, 625
- recover
 - from a FortiGate failure, 2995
- recursive DNS, 576
- Redirect message, 245
- redirect server
 - SIP, 2522
- redistributed routes
 - ospf e1/e2, 1675
- redundant interface, 383
 - active-active mode, 2072
 - active-passive mode, 2072
 - best practice, 2013
 - HA, 1993, 2111
 - HA MAC addresses, 2071
 - port monitoring, 2071
- redundant interfaces, 561
- redundant mode, 568
- redundant VPNs
 - configuration, 1520
 - example, fully redundant configuration, 1523
 - example, partially-redundant configuration, 1534
 - overview, 1519
- Ref. column
 - viewing referenced objects, 379
- reference count column, 148
- refresh every
 - HA statistics, 2150
- regex, 1145
- registering
 - with Fortinet Technical Support, 3040

- registrar
 - SIP, 2522
- Registration, Admission, and Status (RAS), 556
- registry
 - key, 1303
 - remote service, 1319
- regular expression, 322
- regular expressions, 1145
- relay
 - DHCP, 570, 2012
- relay, DHCP, 572
- Release Notes, 1284
- reliable syslog, 646
- remote administration, 350
- remote administrator, 1212
- remote client
 - authenticating with certificates, 1409
 - FortiGate dialup-client, 1501
 - in Internet-browsing IPsec configuration, 1515
- remote client, L2TP VPN, 549
- remote FortiManager options, 301
- Remote Gateway
 - IPSec manual key setting, 1403
- remote gateway
 - dialup user, 1427
- Remote Gateway, Phase 1, 1409, 1411
- remote Internet access, 1655
- remote IP monitoring, 2200
- remote link failover
 - best practice, 2014
 - virtual clustering, 2091
- remote link failure, 2200
- remote logging configuration settings, 133
- remote peer
 - authenticating with certificates, 1409
 - dynamic DNS configuration, 1478
 - gateway-to-gateway IPsec configuration, 1439
 - manual key configuration, 1403
 - manual key IPsec configuration, 1551
 - transparent IPsec VPN configuration, 1544
- remote shell, 559
- Remote SPI
 - IPSec VPN, manual key, 1403
- Remote SPI, Manual Key, 1552
- remote user authentication, 1181
- remotely connecting to IPv6 over the Internet, 237
- removing
 - local users from FortiGate configuration, 1226
 - user group from FortiGate configuration, 1241
- rename, shell command, 325
- Rendezvous Points (RPs), 471
- replacement FortiGate unit
 - adding to a cluster, 2056
- replacement message
 - explicit FTP proxy, 124, 2833, 2842
- replacement message group, 618
- replacement message, to customize web portal login page, 1639
- replacement messages
 - administration, 599
 - alert mail, 597
 - archive and FTP proxy, 168
 - authentication, 1248
 - captive portal default, 602
 - endpoint NAC, 605
 - FortiGuard web filtering, 603
 - FTP, 595
 - FTP proxy, 595
 - HTTP, 591
 - IM, P2P, 604
 - images, 168, 587
 - mail, 590
 - MM1, 608
 - MM3, 612
 - MM4, 615
 - MM7, 617, 618
 - modifying, 588
 - NAC quarantine, 606
 - NNTP, 596
 - RSA
 - SecurID, 1248
 - SMIL, 2394
 - SMS, 2295
 - spam, 598
 - SSL VPN, 608
 - successful firewall authentication, 168
 - tags, 588
 - traffic quota control, 607
 - user authentication, 599
 - video chat block, 168
 - viewing, 587
 - web filtering disclaimer, 168
 - web proxy, 594
- replacing a broken cluster unit, 2056
- replacing a failed cluster unit, 2056
- replay detection, 1598, 2950, 2951, 2952, 2954, 2957, 2958, 2959, 2960
- replay detection, enabling, 1426
- reports
 - adding charts to report layout, 687
 - configuring datasets, 681
 - example, FortiOS UTM default report, 679
 - FortiOS UTM default report, 678
 - importing images, 688
 - styles, reports, 683
 - themes, 682
- reports, vulnerability scans
 - creating, 1122
 - viewing, 1123
- request
 - SIP, 2519
- request messages, 2528
- Request-line
 - deep SIP message checking, 2588
- request-line
 - SIP, 2530
- reserved characters, 328
- reserved management interface, 2129
 - NAT/Route mode, 2130
 - transparent mode, 2130
- reserving addresses, 573

- reset
 - factory default FortiBridge configuration, 2982
- reset age
 - command, 2005
- reset uptime
 - command, 2005
- reset-uptime
 - diagnose command, 2005
- response
 - SIP, 2519
- restore, 363
 - cluster configuration, 2155
- restoring
 - FortiBridge configuration, 2996
- restoring configuration **See** widgets
- restricting login attempts, 351
- Return Material Authorization (RMA), 759
- revalidated pragma-no-cache, 2755
- reverse explicit FTP proxy, 2681
- Reverse Path Forwarding (RPF), 780, 1866
- reverse path lookup, 1681
- reverse proxy
 - web cache, 2682, 2790, 2796
- reverse shaping, 2267
- reverting firmware, 365
- revisions, 364
- revocation list, importing, 1270
- REXEC
 - firewall service, 210
- RFC
 - 1213, 429, 434
 - 1215, 438
 - 1349, 1695
 - 1918, 3037
 - 2071, 238
 - 2080, 251
 - 2185, 245
 - 2516, 385
 - 2545, 251
 - 2665, 429, 434
 - 2740, 251
 - 2858, 251
 - 2893, 245
 - 3315, 384
 - 5237, 1694
 - 791, 1695
 - RFC 1519, 1700
 - RFC 1771, 1751
 - RFC 1965, 1756
 - RFC 1966, 1755
 - RFC 1997, 1758
 - RFC 2385, 1753
 - RFC 2453, 1715
 - RFC 3065, 1756
 - RFC 3509, 1797
 - RFC 4271, 1751
 - RFC 4632, 1700
 - SIP, 2520
- RFC 1112, 472, 1864
- RFC 1321, 1858
- RFC 1700, 2348
- RFC 1771, 1860
- RFC 2236, 472, 1864
- RFC 2362, 1863
- RFC 2385, 1860
- RFC 2474, 2269
- RFC 2475, 2269
- RFC 2543, 2596
- RFC 2865, 1181
- RFC 2866, 1181
- RFC 2986, 1267
- RFC 317, 1605
- RFC 3376, 472, 1864
- RFC 3973, 471, 1863
- RFC 4601, 471
- RFC 5237, 1846
- RFC 5280, 1263
- RFC 791, 2268
- RFC compliance
 - LDAP servers, 1208
- RIP
 - advanced options, 1851
 - authentication, 1853
 - configuring, 1850
 - hop count, 1721
 - interface, 1853
 - RFC 1058, 1715
 - RFC 2453, 1715
 - RIP Next Generation (RIPng), 1716
 - service, 210
 - split horizon, 1853
 - version 1, 1715
 - version 2, 1715
- RIP, IPv6, 251
- RIPv2, 474
- RJ-45, 316
- RJ-45-to-DB-9, 316, 317
- RLOGIN
 - service, 210
- roaming subscribers, 2340
- rogue AP settings, 2509
- role
 - cluster members, 2148
- Role Based Access Control (RBAC), 1204, 1236, 1253, 1333
- root certificate, installing, 1270
- round robin
 - load balancing, 2873
- Round Trip Time (RTT), 748
- Round-Robin
 - HA schedule, 2218
- route, 2953, 2954, 2955, 2958, 2959, 2960
 - adding static routes to a FortiBridge unit, 2980
- route flap, 1765
 - HA, 1763
- route hold, 2193
- route reflectors (RR), 1755
- route synchronization, 2169
- route-based VPN
 - firewall policy, 1436
 - vs. policy-based, 1390
- route-hold, 2192

- router
 - settings, 1847
 - WCCP, 2845
- router monitor
 - HA, 1674
- Router Solicitation message, 245
- route-ttl, 2192
- route-wait, 2192
- routing, 1638
 - administrative distance, 1681
 - asymmetric, 535
 - BGP, 514, 1950
 - blackhole, 1680
 - bridge, 779
 - configuring, 2807, 2831
 - domain, 1709
 - ECMP, 1680
 - enhanced packet-matching, 1711
 - hop count, 1909
 - loopback interface, 1680
 - modem, 570
 - monitor, 1869
 - multicast, 1951
 - OSPF, 514, 1950
 - RIP, 514, 1951
 - routing table, searching, 1677
 - static, 1841
 - STP, 535
 - viewing information, 1673
- Routing Area Identifier (RAI), 2419
- routing information protocol (RIP). See routing, RIP
- routing policy
 - protocol number, 1694, 1846
- routing table, 705, 1797
 - removing routes, 1753
 - searching, 1871
- routing table updates
 - synchronizing, 2191
- routing, asymmetric, 2430
- routing, default, 514, 1917
- routing, default route
 - VDOM example, 1917, 1919
- routing, transparent VPN IPsec configuration, 1546
- RPF, 1866
- RPF (Reverse Path Forwarding), 1979, 2430
- RSA, 1217
 - X.509, 1168
- RSA SecurID, 1248
- RSH
 - firewall service, 210
- rsh, session helper, 559
- RTCP, 2520, 2549
- RTP, 2520, 2559, 2571
 - hardware acceleration, 2527
 - pinhole, 2520, 2525
- RTSP
 - firewall service, 210
- RTSP, session helper, 559

- rule, 2705
 - active-passive, 2705
 - changing the position in the rule list, 2709
 - move, 2709
 - moving, 2709
 - non-HTTP sessions, 2714
 - peer-to-peer, 2705
 - unknown HTTP sessions, 2714
 - WAN optimization, 2705

S

- SA
 - IPsec, 2193
- SAMBA
 - service, 210
- scan buffer size
 - antivirus, 902
- scan schedule, 121
- scanning order
 - antivirus, 896
- SCCP
 - DoS sensor, 2592
 - firewall service, 210
 - protection profile, 2592
 - rate limiting, 2592
 - VoIP profile, 2543
- schedule
 - antivirus and attack definition updates, 410
 - firewall policy, 252
 - load balance, 2221
 - timeout, 213
- schedules
 - expiry, 213
 - group, 213
 - one time, 213
 - recurring, 213
- schedule-timeout command, 213
- school administration, 579
- SCP
 - authentication, 363
 - backup configuration, 361
 - client application, 362
 - restore configuration, 363
 - SSH access, 362
- screen resolution
 - minimum recommended, 289
- scripts
 - uploading, 620
- SDP, 2520, 2527, 2533
 - NAT, 2575
 - session profile, 2533
- search
 - online help, 292
 - online help wildcard, 292
- searching
 - routing table, 1871
- Secure Certificate Enrollment Protocol (SCEP), 1265
- secure HTTP (HTTPS), 1264
- Secure Shell (SSH)
 - key, 318
- secure tunnelling, 2673

- SecurID, 1217
 - firewall policy, 1219
- security, 2436
 - policy matching, 2708
- security association
 - IPsec, 2193
- Security Association (SA), 1426, 1551
- security association (SA), 800, 1598, 2933, 2950, 2951
- security certificate, 336
- security IP addresses
 - defining L2TP, 548
- security layer
 - stateful inspection, 1390
- security layers, 703
- security option, 2956
- Security Parameter Index (SPI), 1551
- security policies, 225, 341
 - accept, 218
 - column settings, 218
 - deny, 218
 - deny policy, 224
 - how to apply VLANs and zones, 195
 - how to arrange, 221
 - ICMP packets, 274
 - identity-based, 222
 - IPsec, 218
 - multicast, 475
 - policy order, 219
 - ssl-vpn policies, 218
 - viewing, 219
- security policy, 2267, 2685, 2705
 - defining L2TP, 548, 549
 - defining PPTP, 541
 - firewall policy, 139
 - how to allow Internet access, 227
 - identity-based, 2706
 - insert policy before, 2708
 - local-in, 226
 - matching, 2708
 - MMS protection profile, 2305
 - traffic priority, 2707
 - verifying traffic is hitting a policy, 228
 - VLAN, 514
 - VLAN example, 519
 - VLAN transparent mode, 524, 528
 - web cache support, 125
 - web-only mode access, 1631
- security processing modules, 2938
 - configuring, 966
 - displaying information, 2939
 - example configuration, 975
 - models, 2939
 - proxy statistics, 979
- security vulnerabilities
 - monitoring for (PCI DSS related), 3033
- selecting the primary unit, 2000
- self-signed certificate, 1264
- self-validate, 1266
- send alertmail from FortiBridge unit, 2989
- Sender Address Identifier, 1349
- sensor
 - IPS, 945
- serial communications (COM) port, 316
- serial no
 - HA statistics, 2150
- serial number
 - getting using SNMP, 2146
 - primary unit selection, 2007
- serial port parameters, 872
- Series 60, 2353
- server
 - DHCP, 570, 2012
 - WCCP, 2845
- server certificate, 1621
 - installing signed, 1269
 - obtaining, 1269
- server comforting, 2360
- server load balance port forwarding virtual IP
 - adding, 2905
- server load balance virtual IP
 - adding, 2899
- servers
 - configuring XAuth authentication using, 1259

- service
 - AH, 206
 - ANY, 206
 - AOL, 206
 - BGP, 206
 - CVSPSERVER, 207
 - DCE-RPC, 207
 - DHCP, 207
 - DHCP6, 207
 - DNS, 207
 - ESP, 207
 - FINGER, 207
 - firewall policy, 252
 - FTP, 207
 - FTP_GET, 207
 - FTP_PUT, 207
 - GOPHER, 207
 - GRE, 207
 - H323, 208
 - HTTPS, 208
 - ICMP_ANY, 208
 - IKE, 208
 - IMAP, 208
 - INFO_ADDRESS, 208
 - INFO_REQUEST, 208
 - Internet-Locator-Service, 208
 - IRC, 208
 - L2TP, 208
 - LDAP, 208
 - MGCP, 208
 - MS-SQL, 209
 - MYSQL, 209
 - NetMeeting, 209
 - NFS, 209
 - NNTP, 209
 - NTP, 209
 - ONC-RPC, 209
 - OSPF, 209
 - PC-Anywhere, 209
 - PING, 209
 - PING6, 209
 - POP3, 209
 - PPTP, 210
 - predefined, 206
 - QUAKE, 210
 - quarantine files list, 666
 - RAUDIO, 210
 - REXEC, 210
 - RIP, 210
 - RLOGIN, 210
 - RSH, 210
 - RTSP, 210
 - SAMBA, 210
 - SCCP, 210
 - service name, 206
 - SIP, 211
 - SIP-MSNmessenger, 211
 - SMTP, 211
 - SNMP, 211
 - SOCKS, 211
 - SQUID, 211
 - SSH, 211
 - SYSLOG, 211
 - TALK, 211
 - TCP, 211
 - TELNET, 211
 - TFTP, 212
 - TIMESTAMP, 212
 - UDP, 212
 - UUCP, 212
 - VDOLIVE, 212
 - VNC, 212
 - WAIS, 212
 - WINFRAME, 212
 - WINS, 212
 - X-WINDOWS, 212
- service group
 - VDOM Transparent example, 1934
 - WCCP, 2846
- service ID
 - WCCP, 2846
- service number
 - WCCP, 2846
- service, DHCP, 572
- services, 206
 - custom, 206
 - list, 206
- session
 - ACCEPT queue full, 846
 - ephemeral, 844
 - failover, 2168
 - half-closed, 845
 - half-open, 845
 - key, 2933
 - SYN queue full, 846
 - TCP state, 844
 - timewait, 845
- session count accuracy, 966
- Session creation, 704
- session description protocol
 - See SDP, 2533
- session failover, 1998, 2205
 - active-active, 2218
 - definition, 2017
 - enabling, 2205
 - failover
 - session, 2167
 - GTP and HA, 2209
 - SIP, 2207, 2594

- session helper, 551, 554, 555, 556, 557, 559, 560, 706, 2941
 - changing the configuration, 552
 - changing the port numbers that the SIP session helper listens on, 2538
 - dcerpc, 555
 - disabling the SIP session helper, 2537
 - DNS, 555
 - enabling the SIP session helper, 2537
 - H.245, 556
 - h245O, 556
 - h323, 556
 - mgcp, 557
 - pmap, 557
 - port, 552
 - PPTP, 557
 - protocol, 552
 - ras, 556
 - rsh, 559
 - rtsp, 559
 - sip, 560
 - TFTP, 560
 - tns, 560
 - viewing, 551
- session hijacking, 1275
- Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions
 - See SIMPLE, 2543
- Session Initiation Protocol. **See** SIP
- session list
 - showing FortiBridge probes, 2992
- Session Monitor submenu, 144
- session pick-up, 835
 - definition, 2017
- session pickup, 1998
 - best practice, 2013
 - delay, 2206
 - enable, 2205
 - enhancing performance, 2205
 - improving performance, 2205
 - selecting FortiGate interfaces to use, 2206
- session profile
 - SDP, 2533
- session synchronization
 - between two standalone FortiGate units, 2243
 - improving performance, 2205
 - using multiple FortiGate interfaces, 2206
- session tables, 706
- session timeout, 1240
- session-based authenticated user, 1890
- session-helper, 552
- session-pickup, 2142
 - CLI command, 2171, 2205
- session-pickup-delay, 2142, 2205
- session-sync-dev, 2206
- set, 326
- setting
 - authentication protocols, 1246
 - firewall policy authentication, 1246
 - firewall user authentication timeout, 1243
 - SSL VPN authentication timeout, 1243, 1258
- setting administrative access for SSH or Telnet, 317
- settings, 359
 - administrators, 359
 - configuring FortiBridge probe settings, 2989
- setup wizard, 140, 336
- severity, 3014
- sFlow, 422
- SHA1, 1598, 2950
- SHA-256, 1418
- SHA-256, 384, 512, 1419
- shadow DNS server, 576
- shaper
 - all policies, 2262, 2263
 - application control, 2266
 - per policy, 2262, 2263
 - per-IP, 2265
 - processing order, 2261
 - security policy, 2267
 - shared, 2262
- shared secret, 1168
- shared shaper, 2262
- shared traffic shapers, 2263
- sharing
 - WAN optimization tunnels, 2688
- shell command
 - delete, 324
 - edit, 324
 - end, 324
 - get, 324
 - purge, 325
 - rename, 325
 - show, 325
- shielded twisted pair, 402
- Shift-JIS, 330
- Shortest Path First (SPF), 1796
- show, 326
 - shell command, 325
- shut down, 403, 2927
- signature
 - adding custom IPS signatures, 949
- signature entry
 - IPS, 948
- signature-based IPS, 705
- signatures, update, 347
- SIMPLE
 - protection profile, 2592
 - rate limiting, 2592
 - VoIP profile, 2543
- Simple Internet Transition (SIT), 256
- Single instruction, multiple data (SIMD), 814
- Single Sign On (SSO), 1263, 1619
- Single user Sign On (SSO), 1284

SIP

- accepting register response, 2554
- blocking requests, 2590
- changing the port numbers that the SIP ALG listens on, 2545
- changing the port numbers that the SIP session helper listens on, 2538
- contact headers and NAT, 2574
- deep header inspection, 2586
- deep message inspection, 2586
- destination NAT, 2571
- dialog, 2519, 2525
- different source and destination NAT for SIP and RTP, 2573
- disabling the SIP session helper, 2537
- DoS sensor, 2592
- enabling the SIP session helper, 2537
- fields, 2531
- fuzzing protection, 2586
- HA session failover, 2594
- headers, 2531
- inspection without address translation, 2523, 2543
- IP address conservation, 2573
- IPv6, 2585
- location server, 2522
- message request-line, 2530
- message sequence, 2525
- message start line, 2530
- message status-line, 2530
- NAT tracing, 2573, 2574
- NAT with dynamic IP pool, 2572
- NAT with IP address conservation, 2573
- pinhole, 2520, 2525
- protection profile, 2592
- proxy server, 2521
- rate limiting, 2592
- redirect server, 2522
- registrar, 2522
- request, 2519
- request-line
 - SIP, 2530
- response, 2519
- RFCs, 2520
- service, 211
- session failover, 2207
- source NAT, 2571
- start line
 - SIP, 2530
- status-line
 - SIP, 2530
- Transparent mode, 2523, 2543
- user element, 2519
- VoIP profile, 2544

SIP ALG

- changing the port numbers that the SIP ALG listens on, 2545
- NAT tracing, 2573

SIP dialogs

- limiting the number, 2593

SIP message

- body, 2520
- final, 2529
- headers, 2520
- informational, 2529
- PRACK, 2529
- provisional, 2529

SIP phone

- FortiFone, 2520

SIP requests, 2590

SIP session

- inactivity timeout, 2547

SIP session helper

- changing the port numbers that the SIP session helper listens on, 2538
- disabling, 2537
- enabling, 2537
- NAT tracing, 2574

SIP, session helper, 560

SIP-MSNmessenger

- service, 211

sip-nat-trace, 2574

sip-tcp-port, 2545

sip-udp-port, 2545

Skinny Call Control Protocol

- See SCCP, 2543

Skinny Call Control Protocol. **See** SCCP

slave DNS server, 575

slave unit, 1992, 2947

- See Also subordinate unit, 1992

SMIL, 2362

SMS text message, 2295

SMS token, 1229

SMTP

- FortiBridge probe list, 2991

- probe, FortiBridge, 2972

- service, 211

smtp traffic, 270

SMTPS

- antivirus, 1142

- antivirus quarantine, 1142

- data leak prevention, 1143

- DLP archive, 1143

- email filtering, 1143

- predefined firewall services, 1141

- protocol recognition, 1142

SNAT, 1946, 2940

sniffer policies, 880

sniffer policy

- viewing, 1134

sniffer, verbosity level, 781

- SNMP, 1903, 2143
 - adding a community to a FortiBridge unit, 2995
 - configuring community, 432
 - configuring on a FortiBridge unit, 2994
 - FortiBridge unit community, 2995
 - get command, 436
 - gigabit interfaces, 432
 - HA reserved management interface, 2130
 - manager, 429, 432
 - MIB, 441, 2143
 - MIBs, 434
 - queries, 431, 433
 - RFC 12123, 434
 - RFC 1215, 438
 - RFC 2665, 434
 - service, 211
 - trap, 2143
 - trap, FortiBridge, 2989
 - traps, 436
 - v3, 167, 429, 430
- SNMP Agent, 429
- SNMP get
 - any cluster unit, 2145, 2146
 - primary unit, 2143
 - subordinate unit, 2145, 2146
- snmpget, 2133, 2134, 2145
- SOCKS
 - explicit web proxy, 2807, 2827
 - service, 211
- soft switch, 388, 585
- soft-switch, 388
- softswith, 388
- software switch interface, 388
- sorting
 - quarantine files list, 666
- source IP address, FortiGate-originating traffic, 176
- source IP hash
 - load balancing, 2873
- source NAT
 - SIP, 2571
- spam, 2385
- spam filter
 - banned word list, 936
- spam filter, see email filter, 928
- spanning tree
 - forward delay, 2234
 - maximum age, 2234
- spanning tree protocol, 2235
 - settings and HA, 2234
- Spanning Tree Protocol (STP), 532, 535, 1923
- special characters, 328, 329
- Spill-over, 1689
- split brain, 2171
 - heartbeat, 2014
- split DNS, 576
- split tunnel, 1657
- split tunneling, 1658
- SQL, 1202
- sql
 - tables, 634
- sql statement examples, 634
- sql tables, 634
- SQLNET
 - session helper, 560
- SQUID
 - service, 211
- SS7, 268
- SSH, 317, 318, 350
 - key, 318
 - service, 211
- SSH connections
 - login grace timer, 159
- SSID
 - configuring, 2506
 - described, 2445
 - whether to broadcast, 2436
- SSL, 1264
 - antivirus, 1142
 - antivirus quarantine, 1142
 - certificate, 1140
 - content inspection, 1139
 - content scanning, 1139
 - data leak prevention, 1143
 - DLP archive, 1143
 - email filtering, 1143
 - example, 907
 - FortiGuard Web Filtering, 1143
 - HTTPS, 1143
 - inspection, 1139
 - load balancing, 2889
 - predefined firewall services, 1141
 - protocol recognition, 1142
 - service definition, 208, 209
 - settings, all, 1141
 - supported FortiGate models, 1139
 - web filtering, 1142
- SSL Client Certificate Restrictive option, 1276
- SSL connection encrypt level option, CLI, 135
- SSL content scanning
 - explicit web proxy, 2819
- SSL offloading, 2208, 2673
 - certificates.certificate
 - SSL offloading, 2891
 - Client to FortiGate, 2890
 - Client to FortiGate to Server, 2890
 - full mode, 2890
 - half mode, 2890
 - load balancing, 2217, 2890
- SSL renegotiation for SSL offloading, 179

- SSL VPN
 - allow/deny client renegotiation, 1638
 - authentication timeout, 1243, 1258
 - checking client certificates, 1621
 - downloading client, 1652
 - enabling, 1668
 - event logging, 1647
 - FortiClient, 1658
 - host check, 1640
 - host OS patch check, 1646
 - load balancing, 2217
 - portal widgets, 1623
 - bookmarks, 1623
 - connection tool, 1623
 - session information, 1623
 - tunnel mode, 1623
 - specifying server certificate, 1621
 - specifying timeout values, 1621
 - split tunneling, 1658
 - Subsession, 1359, 1657
 - user authentication, 1258
 - user groups, configuring, 1236
 - user groups, IPsec VPN dialup users, 1236
 - user groups, creating, 1236
 - Virtual Desktop, 1651
 - web portal, 1620
- SSL VPN user groups, 1236
- SSL, port forwarding, 179
- ssl.root, 706, 1656, 1659
- ssl-vpn, 218
- SSL-VPN Monitor submenu, 146
- SSO (Single Sign On), 1619
- standalone FortiGate unit
 - adding to a cluster, 2055
 - converting to a cluster, 2054
- standalone mode, 568
- standalone session synchronization, 2243
 - filters, 2243
- standby mode
 - real server, 2874
- standby state
 - definition, 2017
 - HA, 2137
- start line
 - SIP, 2530
- state
 - hello, 2016
 - standby, 2017
 - work, 2018
- state synchronization
 - definition, 2017
- stateful inspection, 699, 780, 1390, 1978, 1979, 2402, 2430
- stateful SIP tracking, 2547
- stateless, 699
- static
 - load balancing, 2873
- static route, 341, 2953, 2954, 2955, 2958, 2959, 2960
 - adding, 1845
 - adding policy, 1694
 - adding static routes to a FortiBridge unit, 2980
 - administrative distance, 1681
 - changing gateway for default route, 1844
 - creating, 1841
 - default gateway, 1844
 - default route, 1844
 - default route and gateway, 1844
 - editing, 1841
 - moving in list, 1696
 - overview, 1841
 - policy, 1845
 - policy list, 1845
 - table priority, 1682, 1685
 - table sequence, 1682, 1685
 - viewing, 1841
- static route enhancement, 175
- static routes
 - distance, 1844
 - priority, 1844
- static weight, 170, 2222
- statistics
 - viewing HA statistics, 2150
- status, 2707
 - HA statistics, 2150
 - quarantine files list, 666
- status description
 - quarantine files list, 667
- status keyword, config log memory command, 3015
- status-line
 - SIP, 2530
- STP, 2235
- STP, forwarding, 535
- stream option, 2956
- strict
 - VoIP profile, 2544
- strict source record route, 2956
- strict-register, 2578
- string, 322
- strong authentication, 1275
 - for administrators, 1275
 - for SSL VPN users, 1275
- stub
 - OSPF area, 1857
- sub second failover, 175, 2200
- sub-command, 320, 323
- subinterface
 - VLAN NAT/Route, 512
- subordinate cluster unit
 - definition, 2017
- subordinate unit, 1992
 - definition, 2017
 - getting information using SNMP, 2145, 2146
 - getting serial numbers using SNMP, 2146
 - SNMP get, 2145, 2146
- sub-second failover, 175, 2200
- subsecond failover, 175, 2200
- supernetting, 1757
- suppressing rogue AP, 113

- switch, 388, 585
 - link failover, 2198
 - switching between FortiBridge modes, 2974
 - troubleshooting layer-2 switches, 2233
 - switch interface
 - heartbeat interface, 2172
 - switch mode interface, 382
 - switching
 - between FortiBridge operating modes, 2996
 - switching vdoms, 313
 - Symbian OS version 6, 2353
 - SYN proxy, 990
 - SYN_SENT, 845
 - synchronization
 - configuration, 2184
 - failure console messages, 2186
 - incremental, 2185
 - IPsec VPN SA, 2169
 - periodic, 2185
 - route, 2169
 - sessions between standalone FortiGate units, 2243
 - TCP sessions between standalone FortiGate units, 2243
 - synchronize all
 - CLI command, 2184
 - Synchronized Multimedia Integration Language (SMIL), 2394
 - synchronizing routing table updates, 2191
 - synchronizing the configuration
 - disabling, 2184
 - syntax, 319
 - IPS custom signatures, 950
 - sys ha showcsum
 - diagnose, 2189
 - SYSLOG, 745
 - service, 211
 - syslog
 - configuring a FortiBridge unit, 2994
 - sample FortiBridge message, 2994
 - syslog message, 2989
 - system, 3014
 - system idle timeout, 350
 - system reboot, installing, 370
 - system resources, 141, 772
 - memory constraints, 2371
 - viewing, 305
 - system time
 - configuring, 299
 - system, session-helper, 552
- T**
- table, 320
 - arp, 2198, 2234
 - MAC forwarding table, 2198, 2234
 - TACACS+, 2134
 - configuring server, 1187
 - TACACS+ server
 - authentication, 354
 - TACACS+ servers, 1215
 - ASCII, 1215
 - authenticating users with, 1224
 - authentication protocols, 1215
 - changing default port, 1215
 - CHAP, 1215
 - default port, 1215
 - max number, 1216
 - MS-CHAP, 1215
 - PAP, 1215
 - port, 1215
 - tag management, 150
 - tags
 - adding tags, 255
 - application control, 1077
 - applying tags, 255
 - IPS, 993
 - IPS filters, 985
 - replacement messages, 588
 - TALK
 - service, 211
 - TCP
 - load balancing, 2897
 - port 111, 552
 - port 135, 555
 - port 1720, 552
 - port 1723, 552, 558
 - port 21, 556
 - port 512, 552
 - port 514, 552
 - protocol optimization, 2689
 - service, 211
 - TCP header flags, 699
 - TCP land, 2956
 - TCP port, 1306
 - WAN optimization tunnels, 2687
 - web cache, 2736
 - TCP port 49, 1215
 - TCP ports
 - for collector agent, 1302
 - TCP session synchronization, 1988, 2243
 - between two standalone FortiGate units, 2243
 - filters, 2243
 - TCP sessions
 - load-balance-all, 2217
 - TCP SYN packets, 705
 - TCP WinNuke, 2956
 - TCP/IP stack, 706
 - technical
 - documentation conventions, 3037
 - support, 3040
 - technical support, 3040
 - Technology Assistance Center (TAC), 716
 - TELNET
 - service, 211
 - Telnet, 317, 319
 - Terminal Access Controller Access-Control System (TACACS+), 1215
 - test upgrade procedure, 100
 - test vs
 - get, 2880

- testing
 - VDOM, 1920
 - VDOM transparent mode, 531
 - VLAN, 522
- testing a basic security policy, 227
- testing FortiAnalyzer configuration, 651
- testing VPN connections, 1604
- text strings (names), 293
- TFTP, 2936
 - service, 212
- TFTP server, 370
- TFTP, session helper, 560
- third-party products, 2233
- threshold
 - oversize, 2319, 2360
- threshold, FortiBridge
 - probe, 2973
- time, 720, 744
 - and date, 346
 - changing the FortiBridge system time, 2981
 - configuring, 299
 - NTP, 346
 - protocol, 346
 - zone, 346
- time to live (TTL), 775, 850
- time to live for routes, 2192
- TIME_WAIT, 845
- timeout
 - dynamic profile, 1349
 - session, 1240
 - user context creation, 1350
 - user context entry, 1350
 - user group, 1238
- timeout enhancement, authentication, 119
- timeout values, 1621
- timer
 - provisional invite, 2548
- TIMESTAMP
 - service, 212
- timestamp option, 2956
- timing
 - modifying heartbeat timing, 2175
- TKIP, 2451
- TNS, 560
- tns
 - session helper, 560
- top session ipv6 support, 122
- top sessions
 - viewing, 308
- topology, 2951, 2957
 - out of path, 2675
- ToS, 2268
 - byte value, 2254
 - mapping, 2275
- total bytes
 - HA statistics, 2151
- total packets
 - HA statistics, 2150
- trace
 - SIP ALG NAT tracing, 2573
 - SIP session helper NAT tracing, 2574
- traceroute, 1920
- tracert, 522, 1920
- tracert (traceroute), 775, 776
- traffic, 3014
 - policing, 2252
 - priority, 2262
 - reverse shaping, 2267
 - reverse shaping only, 2267
 - shaping, 2252
- traffic flow
 - normal FortiBridge mode, 2970
- traffic history, 141
- traffic offloading, 2940
- Traffic Priority, 2707
- traffic priority
 - security policy, 2707
 - traffic shaping, 2707
- traffic shaper bandwidth, diag, CLI, 160
- Traffic Shaper Monitor submenu, 144
- traffic shaper, CLI, 161
- Traffic shaper, FSAE
 - minimum bandwidth, 1286, 1291
- traffic shaping, 789, 2706, 2941, 2948
 - firewall policy, 254
 - guaranteed bandwidth, 2264
 - traffic priority, 2707
- traffic shaping offloading, 2947
- traffic statistics, 2934
- train the network, 2178
- Training Services, 3040
- Transparent
 - advanced example, 1926
 - firewall address, 1930, 1934
 - firewall policy, 1926
 - firewall schedule, 1930
 - VDOM example, 1928, 1937
- Transparent mode, 2687
 - configuring an active-active HA cluster, 2034
 - general configuration steps, 2034
 - SIP, 2523, 2543
 - VLAN subinterface, 1926
- transparent mode, 298, 522
 - adding NAT policies, 274
 - example FortiBridge network, 2968
 - management IP address, 298
 - reserved management interface, 2130
 - security policy, 524, 528
 - VDOM example, 526, 530
 - VLAN example, 525
 - VLAN subinterface, 523
 - WAN optimization, 2706, 2712
- transparent mode VPN configuration
 - configuration steps, 1547
 - infrastructure requirements, 1546
 - overview, 1543
 - prerequisites to configuration, 1547
- transparent mode, port pair, 164
- transport mode
 - setting, 1581
- trap
 - SNMP, 2143
- traps, SNMP, 436

- troubleshoot
 - cluster configuration, 2027, 2032, 2039, 2045, 2049, 2053, 2062, 2068, 2075, 2081, 2085, 2117, 2123
 - VPN, 1606
 - troubleshooting, 763
 - alert email did not send, 802
 - alert email test issues, examples, 674
 - BFD, 1766
 - bgp, 1763
 - cannot log to log device, 802
 - communication sessions lost after a failover, 2192
 - dampening, 1765
 - debug packet flow, 782, 1980
 - diagnose commands, 788, 803
 - firewall session list, 787
 - FortiGate stopped logging, 802
 - full mesh HA, 2125
 - graceful restart, 1765
 - ha log message indicate lost neighbor information, 673
 - holddown timer, 1764
 - layer-2 loops, 1923
 - layer-2 switch, 2233
 - packet sniffing, 780, 1978, 2429
 - ping, 773
 - route flap, 1763
 - routing table, 778, 1684
 - traceroute, 773
 - traffic shaping, 789
 - troubleshooting sql statements, 636
 - troubleshooting VPNs, 1609
 - trunk
 - interface, 511, 521
 - links, 504
 - TTL
 - quarantine files list, 667
 - web cache default, 2754
 - web cache maximum, 2754
 - web cache minimum, 2754
 - TTL reduction, 2940
 - tunnel
 - bi-directional initiation, 1433
 - sharing WAN optimization tunnels, 2688
 - TCP port, 2687
 - WAN optimization, 2687
 - Tunnel Endpoint Identifier (TEID), 2407
 - Tunnel Mode, 1656
 - tunnel mode, 1614, 2952, 2958
 - configuring FortiGate server, 1634
 - IP address range, 1618
 - routing, 1638
 - SSL VPN IP range, 1258
 - tunnel mode IPSec, 2957
 - tunnel provider, IPv6, 246
 - tunnel request, 2698
 - tunneling, IPv6, 245
 - tunnel-non-http, 2714
 - two-factor authentication, 116, 1166, 1189, 1228
 - email, 1228
 - FortiToken, 1230
 - SMS, 1229
 - two-factor authentication for administrators, 117
 - two-factor authentication, example, 116
 - type of service, 2268
 - Type of service (TOS), 1695
 - types of user groups, 1235
 - types of users, 1223
- ## U
- UA, 2519
 - UAC, 2519
 - UAS, 2519
 - UDP, 699
 - GTP session failover, 2209
 - load balancing, 2897
 - port 111, 552
 - port 135, 555
 - port 1719, 556
 - port 2427, 557
 - port 2727, 557
 - UDP land, 2956
 - UDP service, 212
 - UE
 - See UA, 2519
 - UMTS Terrestrial Radio Access Network (UTRAN), 2418
 - understanding firewall addresses, 196
 - unicast reverse path forwarding (uRPF), 1672
 - Unicode, 330
 - unidirection, 2948
 - Unified Threat Management, **see** UTM
 - unit operation
 - viewing, 305
 - Universal Mobile Telecommunications System (UMTS), 2309, 2411
 - universal unique identifier (UUID), 555
 - unknown action, 320
 - unknown HTTP sessions, 2714
 - unknown HTTP version
 - explicit web proxy, 2812
 - unknown option, 2956
 - unknown protocol, 2956
 - unnumbered IP, 339
 - unset, 326
 - unwanted login attempts, 351
 - up time
 - HA statistics, 2150
 - update signatures, 347
 - updates
 - FortiExplorer, 337
 - updating
 - antivirus and IPS, web-based manager, 347
 - updating switch arp tables, 2198
 - upgrade after reboot, 374
 - upgrade practices, 99
 - upgrading
 - FortiBridge firmware, 2982
 - upgrading protocol decoder list, 996
 - upgrading, firmware using the CLI, 367
 - upload status
 - quarantine files list, 667
 - uploading logs to FTP server (text format), CLI, 136
 - uploading scripts, 620
 - URL block
 - web filter, 1017

- URL extraction, FortiOS Carrier, 180
 - URL filtering, 880
 - URL formats, 1020
 - usage-based ECMP, 1689
 - USB
 - auto-install, 368
 - backup, 372
 - USB disks, formatting, 301
 - user accounts, 1617
 - User Agent, 2519
 - User Agent Client, 2519
 - User Agent Server, 2519
 - user authentication
 - IPsec VPN dialup users, 1258
 - L2TP VPN, 1261
 - logon blackout period, 1245
 - PKI, 1194
 - PPTP VPN, 1260
 - protocols, 1246
 - remote, 1181
 - SSL VPN, 1258
 - timeout, 1243
 - XAuth, 1259
 - User Context Creation Timeout, 1332
 - user context list, 1330
 - carrier end point, 1330
 - dynamic profile, 1330
 - timeout for removing entries, 1350
 - user context creation timeout, 1350
 - user context entry timeout, 1350
 - waiting for new entries, 1350
 - user data tunnelling (GTP-U), 2409
 - user element
 - See UA, 2519
 - user group for wireless users, 2454
 - user groups, 1235, 1617
 - configuring, 1176
 - creating, 1237
 - different access permissions, 1661
 - Directory Service, 1180, 1240
 - firewall, 1179, 1236
 - on authentication servers, 1237
 - on FortiGate unit, 1312
 - peer, configuring, 1240
 - peer, creating, 1240
 - SSL VPN, 1180
 - types of, 1235
 - Windows AD, 1295
 - user IP address, 1303
 - User Location Information (ULI), 2419
 - user logoff
 - ports 139 and 445, 1302
 - usergroup timeouts, 1238
 - username overlap, 1236
 - Users, 1247
 - users, 1223
 - banned, 1235
 - local, creating, 1224
 - local, deleting from FortiGate configuration, 1226
 - local, removing from FortiGate configuration, 1226
 - peer, configuring, 1227
 - peer, creating, 1227
 - types of, 1223
 - users, number of concurrent, 1890
 - using test procedure to install firmware, 100
 - using the CLI, 315
 - using UTM profiles to monitor and protect your network, 214
 - UTF-8, 330
 - UTM, 2705
 - explicit FTP proxy, 2835
 - explicit web proxy, 2810, 2817
 - ftp proxy, 2835
 - overview, 879
 - profiles, 214
 - sessions continue after active-active HA failover, 2209
 - VDOM, 1137
 - web proxy, 2810
 - utm
 - icap, 1150
 - UTM profiles, 139, 881
 - UTM proxy
 - weight, 172, 2224
 - UUCP
 - service, 212
- ## V
- value, 320
 - value-added-service (VAS) ID, 2325
 - value-added-service-provider (VASP) ID, 2325
 - vcluster, 1951
 - VDOLIVE
 - service, 212
 - VDOM, 717, 732, 733, 786, 791, 805, 2263
 - backing up and restoring, 153
 - configuration, 1928
 - firewall policy, 1909, 1910
 - independent configuration, 1948
 - limited resources, 536, 1892
 - link, 1941
 - management, 1330
 - management configuration, 1942, 1949
 - management services, 1894
 - management VDOM, 1876, 1879, 1896, 1897
 - maximum interface, 1924
 - maximum interfaces, 503, 535
 - maximum number, 1892
 - meshed configuration, 1942, 1950
 - partitioning (HA), 2018
 - simple VDOM NAT/Route example, 1915
 - stand alone configuration, 1942, 1947
 - status, 1896
 - Transparent mode, 1922
 - transparent mode, 522
 - UTM, 1137
 - VDOM example, 1912, 1917
 - VLAN subinterface, 1904
 - VPN settings, 1910

- VDOM and global privileges for access profiles, 169
- VDOM partitioning
 - HA, 2020
- VDOMs, 2687
- vdoms, switching, 313
- veiwing
 - DLP archive, log and archive statistics widget, 307
- Vendor specific attributes (VSA), 1203
- Verifications of IP options, 704
- verify settings and testing purposes
 - FortiGuard license expiry, 675
 - logs sent to log device, 675
 - network scan was performed, 674
- verifying
 - FortiBridge probes, 2992
- verifying traffic is hitting a policy, 228
- video chat block replacement message, 168
- viewing
 - administrators list, 349
 - Alert Message Console, 305
 - antivirus list, 979
 - banned word list, 936
 - carrier end point IP filter list, 1345, 1354, 2309, 2411, 2417
 - configuration revisions, 366
 - disk status, 619
 - DLP archive, 307
 - firewall predefined service list, 206
 - FortiGuard support contract, 408
 - IPSec VPN auto key list, 1393
 - IPSec VPN concentrator list, 1404
 - IPSec VPN manual key list, 1403
 - LDAP server list, 1184
 - licenses, 303
 - log, log and archive statistics widget, 308
 - predefined signatures, 992
 - quarantine logs, 666
 - RADIUS server list, 1182
 - referenced objects, 379
 - routing information, 1869
 - session history, widget, 308
 - static route, 1841
 - system information, 296
 - system resources, 305
 - TACACS+ server, 1187
 - top sessions, 308
 - unit operation, 305
- viewing FortiOS reports, 688
- viewing log messages in Log Access, 131
- viewing log messages, quarantine files, 664
- viewing reports, 129
 - reports, viewing, 688
- viewing security policies, 219
- violation, 3014
- vip, 205
- VIP address
 - L2TP clients, 548
 - PPTP clients, 540
- VIP address, FortiClient dialup clients, 1486
- VIP addresses, 1427
- vip, grouping, 205
- vip, match-vip, 205
- virtual AP
 - creating, 2448
- virtual cluster, 2089, 2090
 - and virtual domains, 2089
 - configuring, 2093, 2095, 2100
- virtual clustering, 1993
 - definition, 2018
 - port monitoring, 2091
 - remote link failover, 2091
- Virtual Desktop, 1651
- virtual domain, transparent VPN IPsec configuration, 1547
- virtual domains, 389, 2687
- virtual domains (VDOMs), 1330
- virtual interface, 1942
- virtual interfaces, 2263
- virtual IP, 2208, 2872
 - assigning with RADIUS, 1487
 - WAN optimization, 2706
- virtual IP address (VIP), 1444
- virtual ip addresses, 205
- virtual LANs, 390
- virtual MAC address, 2169, 2177
 - definition, 2015
 - group ID, 2182
 - how its determined, 2179
 - VRRP, 174, 2238
- Virtual Private Network (VPN), 1381
- Virtual Private Network, *see* VPN.
- virtual router MAC address
 - VRRP, 174, 2238
- Virtual Router Redundancy Protocol, 2237
- Virtual Router Redundancy Protocol (VRRP), 1926
- virtual server, 2208
 - arp-reply, 2872
 - interface, 2872
 - IP, 2872
 - port, 2872
- virus
 - explicit web proxy, 2817, 2837
 - MMS scanning, 2354
 - name, 2398
 - viral marketing, 2385
- virus database, 900, 911
- virus detected
 - HA statistics, 2150
- virus list, 979
- virus protection. **See** antivirus
- virus scan, 896, 899
- Visitor Location Register (VLR), 2311
- VLAN, 2263, 2940
 - adding to VDOM, 1904
 - application, 504
 - jumbo traffic frames, 387
 - maximum number, 503, 535, 1924
 - OSPF, 1859
 - security policy, 514
 - subinterface, 511, 512, 515, 517, 521
 - tagged packets, 512
 - Transparent mode, 1922
 - transparent mode, 522

- VLAN ID, 507
 - range, 504
 - tag, 504
 - VLAN subinterface
 - Transparent mode, 1926
 - transparent mode, 523
 - VDOM example, 1929, 1933
 - VDOM NAT/Route, 1904
 - VDOM transparent mode example, 526
 - VLAN NAT example, 517
 - VLAN NAT/Route example, 517
 - VNC
 - service, 212
 - VoIP, 557
 - load balancing, 2217
 - profile, 2544
 - VoIP Profile
 - SCCP, 2543
 - SIMPLE, 2543
 - VoIP profile
 - default, 2544
 - strict, 2544
 - VoIP support
 - enabling on the web-based manager, 2543
 - VPN, 2950
 - authentication, 1257
 - backup, 1541
 - client-based authentication, 1170
 - FortiClient automatic settings, 1493
 - FortiClient manual settings, 1493
 - gateway, 2953, 2958, 2959
 - idle timeout, 1170
 - IPsec, 1258
 - L2TP, 1261
 - logging events, 1607
 - monitoring, dialup connection, 1603
 - monitoring, static or DDNS connection, 1603
 - planning configurations, 1390
 - policy-based vs. route-based, 1390
 - PPTP, 1260
 - preparation steps, 1392
 - SSL, 1258
 - testing, 1604
 - troubleshooting, 1609
 - VDOM, 1910
 - vpn
 - error no SA proposal, 800
 - initiator, 800, 1450
 - P1 proposal, 800
 - R U THERE, 800
 - VPN encryption/decryption offloading, 2950
 - VPN policy server
 - configuring FortiClient to use, 1493
 - configuring FortiGate unit as, 1491
 - VPN, configuring L2TP, 547
 - VPNs
 - Forticlient, 1402
 - VRRP, 2237
 - adjusting the advertisement message interval, 2242
 - advertisement messages, 2237
 - Configuring, 2239
 - destination IP address, 2242
 - example, 2239, 2241
 - preempt mode, 2242
 - startup time, 2242
 - virtual MAC address, 174, 2238
 - VSA
 - dictionary, 1204
 - RADIUS servers, 1203
 - vulnerability
 - Cross-Site Scripting, 293
 - XSS, 293
 - vulnerability result, 121
 - vulnerability scan
 - adding assets manually, 1114
 - configuring scans, 1117
 - creating reports, 1122
 - discovering assets, 1113
 - selecting assets to scan, 1113
 - viewing executive summary graphs, 1122
 - viewing reports, 1123
 - viewing results, 1117, 1121
 - viewing scan logs, 1121
- ## W
- WAIS
 - service, 212
 - WAN optimization, 2208, 2209
 - and virtual IPs, 2706
 - explicit mode, 2712
 - load balancing, 2217
 - memory usage, 2690
 - monitoring, 2690
 - peer authentication, 2697
 - peers, 2697
 - storage, 2693
 - transparent mode, 2712
 - WAN optimization peer
 - configuring, 2699
 - monitoring, 2703
 - WAN Optimization rule, 2707
 - changing the position in the rule list, 2709
 - moving, 2709
 - status, 2707
 - WAN optimization traffic log, 2692
 - wanopt-traffic, 2692
 - WAP, 2295
 - WAP traffic, 1327
 - dynamic profile, 1327
 - warning to install FortiClient, 1090

- WCCP, 2208, 2845
 - cache engine, 2845
 - client, 2845
 - load balancing, 2217
 - router, 2845
 - server, 2845
 - service group, 2846
 - service ID, 2846
 - service number, 2846
 - topology, 2673, 2684
 - well known service, 2846
- WCCP service ID
 - HTTP, 2846
- wdgets
 - unit operation, 305
- web cache, 2673, 2753
 - active-passive WAN optimization, 2744
 - adding to passive WAN optimization rule, 2744
 - always revalidate, 2753
 - changing the relative amount of disk space, 2694
 - client/server WAN optimization, 2744
 - default TTL, 2754
 - exempt, 2752
 - fresh factor, 2753
 - HTTP port, 2736
 - max cache object size, 2753
 - max HTTP message length, 2754
 - max HTTP request length, 2754
 - maximum TTL, 2754
 - minimum TTL, 2754
 - monitor, 146
 - monitoring, 2755
 - negative response duration, 2753
 - non-standard ports, 2741
 - peer to peer WAN optimization, 2748
 - proxy FQDN, 2754
 - reverse proxy, 2682, 2790, 2796
 - storage, 2693
 - TCP port, 2736
- Web Cache Communication Protocol
 - See WCCP, 2845
- web cache, security policy, 125
- web caching
 - memory usage, 2690
- web content filtering, 880
- web filter, 658, 928
 - how URL formats are detected, HTTP, 1021
 - how URL formats are detected, HTTPS, 1020
 - local ratings, 1021
 - quota, 1027
 - URL block, 1017
 - URL filter, 1018
- web filter profile, 1012
- web filtering, 880
 - explicit web proxy, 2817
 - HTTPS, 1142
- web filtering overrides, 106
- web filtering service, 590, 2398
- web monitor, 1157
- Web Monitor submenus, 145
- web portal
 - customizing login page, 1639
 - setting login page port number, 1638
 - SSL VPN,SSL VPN web portal
 - customize, 1620
- web proxy, 2208, 2807
 - antivirus, 2817
 - authentication, 2815, 2816
 - DLP, 2817
 - FortiGuard web filtering, 2817
 - HTTPS deep scanning, 2819
 - protocol options, 2817
 - UTM, 2810
 - web filtering, 2817
- web proxy service, 178
- web proxy service group, 178
- web site, content category, 588, 2396
- Web UI. **See** web-based manager
- web-based manager, 289, 335
 - changing the language, 312
 - connecting to the CLI, 313
 - filter settings, 1135
 - idle timeout, 313
 - logging out, 313
 - online help, 291
 - pages, 289
 - screen resolution, 289
 - tag management, 150
 - using web-based manager lists, 290
- web-based manager configuration steps
 - NAT/Route mode, 2023, 2027, 2034, 2039, 2047, 2050
- web-based manager, filtering lists, 147
- web-based manager, lock, 398
- web-based manager, switching vdoms, 313
- web-based user authentication, 1169
- webcache-storage-percentage, 2695
- web-only mode, 1614
 - security policy for, 1631
- weight
 - real server, 2874
 - static, 170, 2222
- weighted
 - load balancing, 2873
- weighted round-robin
 - HA schedule, 2218
- weighted-round-robin
 - configuring weights, 170, 2222
- well known service
 - WCCP, 2846

- widgets, 302
 - adding, 295
 - alert message console, 305
 - CLI console, 308
 - disk storage, 311
 - IM usage, 312
 - licence information, 303
 - log and archive statistics, 306
 - network protocol usage, 142, 312
 - P2P usage, 311
 - per-IP bandwidth usage, 312
 - RAID monitor, 308
 - session history, 308
 - system information, 296
 - system resources, 141, 305
 - top application usage, 311
 - top history, 308
 - top session, 122
 - top sessions, 308
 - traffic history, 141
 - VoIP usage, 312
 - WiFi controller
 - discovery methods, 2464
 - WiFi, CLI, 114
 - wild cards, 322
 - wildcard, 1042
 - carrier end point pattern, 1342, 1345, 1352, 2334
 - firewall addresses, 202
 - online help search, 292
 - wildcard addresses, 202
 - wildcard admin configuration, 1212
 - wildcard pattern matching, 333
 - wildcards, 1145
 - Windows 2008, 1202
 - Windows Active Directory (AD), 1191
 - forest, 1288
 - trust relation, 1288
 - Windows networks
 - enabling NetBIOS, 534
 - Windows Terminal Server
 - authentication, 2816
 - windows version check, 1643
 - Windows VPN, 1567
 - WINFRAME
 - service, 212
 - WINS, 534
 - service, 212
 - wire speed, 2935
 - wireless, 387
 - client mode, 2501
 - configuring SSID, 2506
 - custom AP profile, 2512
 - managed FortiAP, 2510, 2511
 - rogue AP settings, 2509
 - wireless controller
 - suppressing rogue AP, 113
 - Wireless LAN (WLAN), 2418
 - wireless security considerations, 3030
 - Wireshark, 768
 - wizard, 336
 - wizard, setup, 140
 - WLAN
 - firewall policies, 2455
 - WML, 2362
 - word boundary, Perl regular expressions, 333
 - work state
 - definition, 2018
 - HA, 2137
 - Workstation verify interval
 - collector agent configuration, 1298
 - worm-generated messages, 2385
- ## X
- X.509, 1263, 1275
 - managing security certificates, 1266
 - server certificate, 1264
 - X.509 security certificates, 1619
 - XAUTH, 1257
 - configuring authentication with, 1259
 - XAuth (extended authentication)
 - authenticating users with, 1422
 - FortiClient application as client, 1494
 - FortiGate unit as server, 1422
 - XD4, 2220
 - X-Forwarded-For (XFF), 2813
 - x-mms-response-status, 2355
 - x-mms-response-text, 2355
 - XSS vulnerability
 - protection from, 293
 - XSS vulnerability characters, 1256
 - x-up-calling-line-id, 1348
 - X-WINDOWS
 - service, 212
- ## Z
- zero bandwidth, 2262
 - zone
 - using as route-based "concentrator", 1458
 - zones, 195, 392