ArubaOS-Switch Management and Configuration Guide for KA/KB.16.04

a Hewlett Packard Enterprise company

Part Number: 5200-3123
Published: July 2017
Edition: 1

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Chapter 1 About this guide.......... 24
Applicable products.......... 24
Switch prompts used in this guide.......... 24

Chapter 2 Time synchronization.......... 25
NTP.......... 25
NTP related commands.......... 25
Elements of time synchronization.......... 38
Time synchronization protocols.......... 39
timesync.......... 39
Setting a time protocol on the switch.......... 39
The SNTP protocol.......... 40
Selecting and configuring SNTP.......... 41
Viewing SNTP parameters.......... 44
Enabling SNTP client authentication.......... 45
SNTP poll interval.......... 46
SNTP server priority.......... 47
SNTP software version.......... 48
SNTP server address.......... 48
SNTP authentication trusted keys.......... 49
Configuration files and the `include-credentials` command 49
Configuring the key-identifier, authentication mode, and key-value.......... 50
SNTP messages in the event log.......... 56
Storing security information in the running-config file.......... 57
The TimeP Protocol.......... 57
Enabling TimeP mode.......... 57
Viewing, enabling, and modifying the TimeP protocol(Menu).......... 59
TimeP operation in manual mode.......... 60
Viewing, enabling, and modifying the TimeP protocol (Menu).......... 61
Current TimeP configuration.......... 62
Change from one TimeP server to another 63
TimeP poll interval.......... 63
Disable time synchronization protocols 63
Disabling TimeP in manual mode.......... 63
Disabling time synchronization.......... 63
Disabling the TimeP mode.......... 64
Disabling time synchronization without changing the SNTP configuration.......... 65
Disabling SNTP mode.......... 65
Disabling SNTP Mode.......... 65
Disabling time synchronization in DHCP mode by disabling the TimeP mode parameter.......... 66
Other time protocol commands.......... 67
Show management command.......... 67
Show SNTP command.......... 68
Show TimeP command.......... 69

Chapter 3 Resource usage.......... 72
Viewing current resource usage.......... 72
showquos.......... 72

Viewing information on resource usage....73
When insufficient resources are available....74
Policy enforcement engine....74
Usage notes for show resources output....75

Chapter 4 Hardware components....76
Services....76
Show services....76
No parameters....76
Show services locator....77
Show services device....78
Requesting a reboot....79
Services in Operator/Manager/Configure context....79
Show services set locator module....83
Reloading services module....83
Connection to the application via a serial port....84
Shutdown the services module....84
Transceiver status....84
Operating notes....84
show interfaces transceivers....85
Configuring the type of a module....85
module type....85
Clearing the module configuration....86
Configuring transceivers and modules that have not been inserted....86
Transceivers....86
Modules....86
Clearing the module configuration....86
Power consumption....87
show system power-supply....87
Fans....90
show system....90
`show system fans`....92
show system power-supply....94
Fan failures and SNMP traps....98
System boot diagnostics....98
`show system post`....98
`show system post member`....99
`show system post vsf member`....100

Chapter 5 Port status and configuration....102
Viewing port status and configuration....102
`show interfaces`....102
Viewing transceiver information....104
The port VLAN tagged status....105
Dynamically updating the show interfaces command....106
command name....106
Customizing the show interfaces command....107
show interfaces custom....107
show interface smartrate....108
show interface port utilization....109
Enabling or disabling ports and configuring port mode....109
interface....109
Basic USB port commands....110
`usb-port`....111

show usb-port ... 111
Enabling or disabling flow control ... 111
interface flow-control ... 112
Configuring auto-MDIX ... 113
interface mdix-mode ... 113
show interfaces config ... 113
show interfaces brief ... 114
Viewing port configuration (Menu) ... 114
Configuring ports (Menu) ... 115
Configuring friendly port names ... 116
interface name ... 116
Configuring a single port name ... 116
Configuring the same name for multiple ports ... 116
Viewing friendly port names with other port data ... 117
Listing all ports or selected ports with their friendly port names ... 117
show name ... 117
Including friendly port names in per-port statistics listings ... 118
Searching the configuration for ports with friendly port names ... 120
Configuring uni-directional link detection ... 120
interface link-keepalive ... 121
Enabling UDLD ... 121
Changing the keepalive interval ... 122
Changing the keepalive retries ... 122
Configuring UDLD for tagged ports ... 122
Viewing UDLD information ... 122
clear link-keepalive ... 123
Viewing summary information on all UDLD-enabled ports ... 123
Viewing detailed UDLD information for specific ports ... 123
Port status and Port parameters ... 124
Connecting transceivers to fixed-configuration devices ... 124
Error messages associated with the show interfaces command ... 126
Auto-MDIX configurations ... 127
Manual override ... 127
About using friendly port names ... 127
Configuring and operating rules for friendly port names ... 128
Uni-directional link detection (UDLD) ... 128
Configuring UDLD ... 129
Uplink failure detection ... 129
Configuration Guidelines for UFD ... 130
UFD enable/disable ... 131
UFD track data configuration ... 131
UFD minimum uplink threshold configuration ... 132
show uplink-failure-detection ... 132
UFD operating notes ... 133
Error log ... 133
Invalid port error messages ... 133
Port Shutdown with Broadcast Storm ... 134
Configuration Commands ... 134
Viewing broadcast-storm configuration ... 135
Broadcast-storm event logs ... 138

Chapter 6 Power over ethernet (PoE/PoE+) operation ... 139
PoE ... 139
PoE terminology ... 139
Planning and implementing a PoE configuration ... 139

Power requirements....139
Assigning PoE ports to VLANs....140
Applying security features to PoE configurations....140
Assigning priority policies to PoE traffic....140
PoE operation....140
PoE configuration options....141
PD support....141
PoE power priority....142
About configuring PoE....142
Disabling or re-enabling PoE port operation....144
interface....144
Enabling support for pre-standard devices....144
power-over-ethernet....144
Configuring the PoE port priority....145
interface....145
Controlling PoE allocation....145
int....145
Manually configuring PoE power levels....146
Detection status: fault....147
Configuring PoE redundancy (chassis switches only)....147
power-over-ethernet redundancy....148
Changing the threshold for generating a power notice....148
power-over-ethernet slot....148
Enabling or disabling ports for allocating power using LLDP....148
int poe-lldp-detect....149
Enabling PoE detection via LLDP TLV advertisement....149
lldp config....149
Negotiating power using the DLL....149
int poe-lldp-detect....150
Initiating advertisement of PoE+ TLVs....151
lldp config....151
Temporary PoE+ power drop....152
Viewing PoE when using LLDP information....152
show lldp config....152
Viewing the global PoE power status of the switch....154
show power-over-ethernet....154
Viewing PoE status on all ports....155
show power-over-ethernet....155
Viewing the PoE status on specific ports....157
show power-over-ethernet....157
Configuring thresholds for generating a power notice....160
PoE/PoE+ allocation using LLDP....160
LLDP with PoE....160
LLDP with PoE+....160
Operation note....161

Chapter 7 Port trunking....162
Port trunking overview....162
Port trunk connections and configuration....162
Viewing and configuring port trunk groups....163
Viewing static trunk type and group for all ports or for selected ports....163
Viewing static LACP and dynamic LACP trunk data....164
Configuring a static trunk or static LACP trunk group....164
Removing ports from a static trunk group....165
Enabling dynamic LACP trunk groups....165

Remove ports from a dynamic LACP trunk group....166
Set the LACP key....166
Viewing and configuring a static trunk group (Menu)....167
Enable L4-based trunk load balancing....168
trunk-load-balance....168
Viewing trunk load balancing....169
show trunks....169
Operating notes....170
Distributed trunking....170
Configure ISC ports....170
Configuring distributed trunking ports....171
Configuring peer-keepalive links....171
Viewing distributed trunking information....172
Viewing peer-keepalive configuration....174
Viewing switch interconnect....174
Port trunk operations....174
Fault tolerance 174
Trunk configuration methods....174
Dynamic LACP trunk....174
Dynamic LACP Standby Links....175
Viewing LACP Local Information....175
Viewing LACP Peer Information....175
Viewing LACP Counters....176
Using keys to control dynamic LACP trunk configuration....176
Static trunk....176
Static or dynamic trunk group overview....179
Enabling a dynamic LACP trunk group....180
Dynamic LACP standby links....180
Viewing LACP local information....181
Viewing LACP peer information....181
Viewing LACP counters....181
Trunk group operation using LACP....182
Default port operation....184
LACP operating notes and restrictions....185
Trunk group operation using the "trunk" option....188
Viewing trunk data on the switch....188
Outbound traffic distribution across trunked links....189
Trunk load balancing using Layer 4 ports....190
Distributed trunking overview....190
Distributed trunking interconnect protocol....192
Configuring distributed trunking....192
Configuring peer-keepalive links....193
Maximum DT trunks and links supported....194
Forwarding traffic with distributed trunking and spanning tree....194
IP routing and distributed trunking....197
Distributed trunking restrictions....199
Updating software versions with DT....200

Chapter 8 Port traffic controls....202

Rate-limiting....202
Configuring rate-limiting on all traffic....202
Viewing the current rate-limit configuration....203
Configuring ICMP rate-limiting....205
Viewing the current ICMP rate-limit configuration....206
Resetting the ICMP trap function of the port....207

Configuring an egress/outbound broadcast limit on the switch....208
Egress per-queue rate-limiting....211
Guaranteed Minimum Bandwidth (GMB) for outbound traffic....213
int bandwidth-min output....214
Non-default GMB settings....214
Viewing the current GMB configuration....216
Validation rules....219
Event log....219
All traffic rate-limiting....220
Operating notes for rate-limiting....221
ICMP rate-limiting....222
ICMP rate-limiting....223
Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface....224
Operating notes for ICMP rate-limiting....224
Guaranteed minimum bandwidth (GMB)....226
GMB operations....226
Jumbo frames....228
Operating rules for jumbo frames....228
Jumbo frame maximum size....230
Jumbo IP MTU....230
Configuring jumbo frame operation....230
Troubleshooting Jumbo frames....234

Chapter 9 Fault-Finder port-level link-flap....235
Overview....235
fault-finder link-flap....235
Show fault-finder link-flap....237
Event Log....238
Restrictions....238

Chapter 10 Configuring for Network Management Applications....239
Configuring the switch to filter untagged traffic....239
ignore-untagged-mac....239
Viewing configuration file change information....239
show running-config....239
Minimal interval for successive data change notifications....241
setmib....241
Viewing the current port speed and duplex configuration on a switch port....241
show interfaces....242
Viewing the configuration....243
show running-config....243
RMON advanced management....244
rmon alarm....244
Configuring UDLD verify before forwarding....247
UDLD time delay....247
UDLD configuration commands....247
show link-keepalive....248
RMON generated when user changes UDLD mode....249
MAC configurations....249
Configuring the MAC address count option....249
Configuring the MAC address table change option....249
Per-port MAC change options for mac-notify....250
Viewing the mac-count-notify option....251
Viewing mac-notify traps configuration....252

Configuring sFlow 253
- sflow 254
- sFlow Configuring multiple instances 255
- Viewing sFlow Configuration and Status 255
- show snmpv3 user 257

Configuring SNMP 258
- Network security notifications 258
- Group access levels 259
- SNMPv2c informs 260
- SNMP notifications 261
- SNMPv3 users 267
- SNMP tools for switch management 268
- Enabling SNMPv3 270
- Configuring users in SNMPv3 270
- Switch access from SNMPv3 agents 271
- Restrict access from SNMPv3 agents 271
- Restrict non-SNMPv3 agents to read-only access 271
- Operating status of SNMPv3 272
- Non-SNMPv3 message reception status 272
- Non-SNMPv3 write message status 272
- Viewing and configuring non-version-3 SNMP communities (Menu) 272
- SNMP trap receiver configuration 273
- SNMPv2c inform option 274
- Configuring SNMPv3 notifications 275
- SNMPv3 community mapping 278
- Running configuration changes and SNMP traps 279
- Startup configuration changes and SNMP traps 279
- Source IP address for SNMP notifications 280
- SNMP replies and traps configuration 282
- SNMP notification configuration 282
- Assign users to groups 283
- snmp-server community 284
- Community names and values 285
- Notification/traps for network security failures and other security events 286
- Current network security notification configuration 288
- Link-Change Traps 288
- Listening mode 289

CDP configuration 289
- CDP mode 289
- CDPv2 for voice transmission 290
- CDP operation on individual ports 292
- CDP Operation 292
- CDP information filter 293
- CDP switch configuration view 293
- CDP neighbors switch table view 293

LLDP configuration 294
- LLDP and CDP data management 294
- LLDP 296
- LLDP-MED 302
- Configuring per-port transmit and receive modes 304
- Remote management address for outbound LLDP advertisements 304
- Port speed and duplex advertisement support 306
- Location data for LLDP-MED devices 306
- LLDP data change notification for SNMP trap receivers 308
- LLDP operation on the switch 308
- LLDP-MED fast start control 308
- Changing the packet transmission interval 309

Changing the time-to-live for transmitted advertisements....309
Delay interval....310
Changing the reinitialization delay interval....310
PVID mismatch log messages....311
Viewing port configuration details....311
Available switch information available outbound advertisements....312
LLDP statistics....313
Global LLDP, port admin, and SNMP notification status....315
LLDP-MED connects and disconnects—topology change notification....316
Device capability, network policy, PoE status and location data....317
Advertisements currently in the neighbors MIB....327
PoE advertisements....328
LLDP Management TLV Transmission disablement....329
Overview....329
Commands....329
Show commands....330
TVL configuration....330
VLAN ID TLV....330
Advertised TLVs....331
TLVs controlled by medTLvEnable....333
Generic header ID in configuration file....333
DHCP auto deployment....334
Add-Ignore-Tag option....334
Configuration commands for the add-ignore-tag option....335
Show logging commands for the add-ignore-tag option....335
Exclusions....335

Chapter 11 DHCPv4 server....336
Overview....336
IP pools....336
DHCP options....336
BootP support....336
Authoritative server and support for DHCP inform packets....336
Authoritative pools....337
Authoritative dummy pools....337
Change in server behavior....337
DHCPv4 configuration commands....338
DHCPv4 server....338
DHCP address pool name....338
DHCP client boot file....340
DHCP client default router....340
DNS IP servers....340
Configure a domain name....341
Configure lease time....341
NetBIOS WINS servers....341
NetBIOS node type....342
Subnet and mask....342
DHCP server options....342
IP address range....343
Static bindings....343
TFTP server domain name....344
Configure the TFTP server address....344
Number of ping packets....344
Save DHCP server automatic bindings....345
DHCP server and SNMP notifications....345

Conflict logging on a DHCP server 346
Enable the DHCP server on a VLAN 346
Clear commands 346
Reset all DHCP server and BOOTP counters 347
Delete an automatic address binding 347
Show commands 347
show dhcp-server 347
Event log 348
Event Log Messages 348

Chapter 12 DHCPv6 server 351
DHCPv6 hardware address 351
DHCPv6 snooping and relay 351
dhcpv6–snooping 351
dhcpv6 snooping trust 352
dhcpv6–snooping authorized-server 353
ddhcpv6–snooping database file 353
dhcpv6–snooping max-bindings 354
dhcpv6–relay option 79 355
snmp-server enable traps dhcpv6-snooping 355
clear dhcpv6–snooping stats 356
debug security dhcpv6–snooping 356
ipv6 source-lockdown ethernet 356
ipv6 source-binding 358
snmp-server enable traps dyn-ipv6-lockdown 359
debug security dynamic-ipv6-lockdown 360
Show commands for DHCPv6–snooping 360
show dhcpv6-snooping 360
show dhcpv6 snooping bindings 360
show dhcpv6 snooping statistics 360
show ipv6 source-lockdown 361
show ipv6 source-lockdown status 361
show snmp-server traps 362
show distributed-trunking consistency-parameters 362
show distributed-trunking consistency-parameters 363
show dhcpv6 relay 364
DHCPv6 event log 365
DHCPv6 event messages 368

Chapter 13 Captive portal for ClearPass 370
Requirements 370
Best Practices 370
Limitations 371
Features 371
High Availability 371
Load balancing and redundancy 371
Captive Portal when disabled 371
Disabling Captive Portal 371
Configuring Captive Portal on CPPM 372
Importing the HP RADIUS dictionary 372
Creating enforcement profiles 372
Creating a ClearPass guest self-registration 373
Configuring the login delay 374
Configuring the switch 374

The URL key.......... 375
Configuring a certificate for Captive Portal usage.......... 375
Displaying the Captive Portal configuration.......... 376
Showing certificate information.......... 376
Troubleshooting.......... 376
Event Timestamp not working.......... 376
Cannot enable Captive Portal.......... 377
Unable to enable feature.......... 377
Authenticated user redirected to login page.......... 377
Unable to configure a URL hash key.......... 378
ClearPass captive portal authentication commands.......... 378
aaa authentication captive-portal.......... 378
Clearpass captive portal show commands.......... 379
Clearpass captive portal debug command.......... 380

Chapter 14 ZTP with AirWave Network Management.......... 381
Requirements.......... 381
Best Practices.......... 381
Limitations.......... 381
Switch configuration options.......... 382
Configure AirWave details in DHCP (preferred method).......... 382
Configure AirWave details in DHCP (alternate method).......... 387
Zero Touch Provisioning.......... 394
Auto-configuration using ZTP.......... 394
Disabling ZTP.......... 395
Image Upgrade.......... 395
CLI switch configuration.......... 395
Stacking and chassis switches.......... 396
Troubleshooting.......... 396
AMP server messages.......... 396
Validation Rules.......... 396
AirWave configuration details.......... 397
amp-server.......... 397
debug ztp.......... 398

Chapter 15 Auto configuration upon Aruba AP detection.......... 399
Auto device detection and configuration.......... 399
Requirements.......... 399
Limitations.......... 399
Feature Interactions.......... 400
Procedure for creating a device identity and associating a device type.......... 400
device-profile name.......... 401
device-profile type.......... 402
Rogue AP Isolation.......... 403
Limitations.......... 403
Feature Interactions.......... 404
Using the Rogue AP Isolation feature.......... 405
rogue-ap-isolation.......... 406
rogue-ap-isolation action.......... 406
rogue-ap-isolation whitelist.......... 407
clear rogue-ap-isolation.......... 407
Troubleshooting.......... 407
Dynamic configuration not displayed when using “show running-config”.......... 407
Switch does not detect the rogue AP TLVs.......... 408

The `show run` command displays non-numerical value for untagged-vlan....408
Show commands....408
Validation Rules....409

Chapter 16 LACP-MAD....412
LACP configuration....412
interface <PORT-LIST> lacp....412
show lacp....412
clear lacp statistics....413
LACP-MAD Operations....413

Chapter 17 File transfers....414
File transfer methods....414
TFTP....414
Prerequisites....414
Downloading switch software....414
Enabling tftp....416
Automatic software download from a TFTP server....417
Downloading to primary flash using TFTP....418
Disabling TFTP and auto-TFTP for enhanced security....419
SCP and SFTP....424
Enabling SCP and SFTP....424
Using SCP and SFTP....424
Xmodem....425
Downloading software using Xmodem....425
Downloading to primary flash using Xmodem (Menu)....427
USB....427
Downloading switch software using USB....427
USB port status....427
Switch to Switch....429
Switch-to-switch download....429
Copying....431
Software images....431
Copying diagnostic data to a remote host, USB device, PC, or UNIX workstation....432
copy core-dump (standby module)....436
Copy diagnostic data to a remote host, USB device, PC or UNIX workstation....437
Transferring....437
Switch configuration transfer....438
ACL command file transfer....442
Switch software download....444
Switch software download rules....444
Single copy command....445
copy source....446
copy crash-files....449
copy crash-files member....449
copy crash-files crash-file-options....449

Chapter 18 Monitoring and Analyzing Switch Operation....451
Switch and network operations....451
Status and counters data....451
Accessing status and counters (Menu)....452
show system....452
Collecting processor data with the task monitor....454

Accessing system information (Menu)....455
Switch management address information access....455
Component information views....456
Compatibility mode for v2 zl and zl modules....457
allow-v1-modules....457
Port status....458
show interfaces brief....458
Viewing port status (menu)....458
Accessing port and trunk group statistics....458
Trunk bandwidth utilization....459
Reset port counters....462
Accessing port and trunk statistics (Menu)....464
MAC address tables....464
MAC address views and searches....464
MSTP data....468
show spanning-tree....468
IP IGMP status....469
show ip igmp....469
VLAN information....470
show vlan....471
WebAgent status information....473
Configuring local mirroring....473
Local mirroring sessions....474
Traffic-direction criteria....474
ACL criteria for inbound traffic — deprecated....474
Mirror policy for inbound traffic....474
MAC-based criteria to select traffic....475
Remote mirroring destination on a remote switch....476
Remote mirroring destination on a local switch....476
mirror remote ip....476
Local mirroring destination on the local switch....476
mirror port....476
Monitored traffic....476
interface....476
monitor all....477
service-policy....477
Configuring local mirroring (Menu)....477
Destination mirror on a remote switch....479
mirror endpoint....479
Source mirror on the local switch....479
mirror remote ip....479
Traffic-direction criteria....480
Configure ACL criteria to select inbound....480
interface monitor ip access-group....480
Mirror policy for inbound traffic....480
class [ipv4|ipv6]....480
policy mirror....480
Configuring a destination switch in a remote mirroring session....481
Configuring a source switch in a local mirroring session....482
Configuring a source switch in a remote mirroring session....482
Selecting all traffic on a port interface for mirroring according to traffic direction....483
Selecting all traffic on a VLAN interface for mirroring according to traffic direction....484
Configuring a MAC address to filter mirrored traffic on an interface....485
Configuring classifier-based mirroring....486
Applying a mirroring policy on a port or VLAN interface....488
Viewing a classifier-based mirroring configuration....488
Viewing all mirroring sessions configured on the switch....488

Viewing the remote endpoints configured on the switch....490
Viewing the mirroring configuration for a specific session....491
Viewing a remote mirroring session....492
Viewing a MAC-based mirroring session....493
Viewing a local mirroring session....493
Viewing information on a classifier-based mirroring session....494
Viewing information about a classifier-based mirroring configuration....494
Viewing information about a classifier-based mirroring configuration....495
Viewing information about statistics on one or more mirroring policies....495
Viewing resource usage for mirroring policies....496
Viewing the mirroring configurations in the running configuration file....497
Compatibility mode....498
Port and trunk group statistics and flow control status....499
Traffic mirroring overview....499
Mirroring overview....499
Mirroring destinations....500
Mirroring sources and sessions....500
Mirroring sessions....501
Mirrored traffic destinations....502
Monitored traffic sources....502
Criteria for selecting mirrored traffic....502
Mirroring configuration....502
Remote mirroring endpoint and intermediate devices....504
Migration to release K.12.xx....504
Migration to release K.14.01 or greater....504
Using the Menu to configure local mirroring....505
Menu and WebAgent limits....505
Remote mirroring overview....506
Quick reference to remote mirroring setup....506
High-level overview of the mirror configuration process....507
Determine the mirroring session and destination....507
Configure a mirroring destination on a remote switch....507
Configure a mirroring session on the source switch....507
Configure the monitored traffic in a mirror session....508
About selecting all inbound/outbound traffic to mirror....509
Untagged mirrored packets....509
About selecting inbound traffic using an ACL (deprecated)....510
About selecting inbound/outbound traffic using a MAC address....510
About selecting inbound traffic using advanced classifier-based mirroring....512
Classifier-based mirroring configuration....512
Classifier-based mirroring restrictions....514
About applying multiple mirroring sessions to an interface....515
Mirroring configuration examples....516
Maximum supported frame size....520
Enabling jumbo frames to increase the mirroring path MTU....520
Effect of downstream VLAN tagging on untagged, mirrored traffic....521
Operating notes for traffic mirroring....522
Troubleshooting traffic mirroring....523

Chapter 19 Virtual Technician....525
Cisco Discovery Protocol (CDP)....525
show cdp traffic....525
clear cdp counters....525
Enable/Disable debug tracing for MOCANA code....526
debug security 526

User diagnostic crash via Front Panel Security (FPS) button....526
front-panel-security password-clear....526
front-panel-security diagnostic-reset....527
show front-panel-security....528
Diagnostic table....528
Validation rules....529
FPS Error Log....530
User initiated diagnostic crash via the serial console....531
front-panel-security diagnostic-reset serial-console....531
Serial console error messages....532

Chapter 20 Scalability: IP Address, VLAN, and Routing Maximum Values....533

Chapter 21 Job Scheduler....535
Job Scheduler....535
Commands....535
Job at | delay | enable | disable 535
Show job....536
Show job <Name>....537

Chapter 22 Device Profile for custom device types....539
Procedure for creating a device identity and associating a device type....539

Chapter 23 Dynamically detecting LLDP device profiles....540
device-profile....540
device-profile type-device....540
device-profile device-type enable....541
Associating a profile with a device....542
device-profile device-type associate....542
show device-profile status....542
show device-profile config....543
show device-identity....544

Chapter 24 Virtual Switching Framework (VSF)....546
Overview of VSF....546
Benefits of VSF....547
Member roles....547
Commander....547
Standby....547
Management module....547
VSF member ID....548
VSF link....548
vsf member link....548
Physical VSF ports....550
VSF domain ID....550
VSF split....551
VSF merge....551
Member priority....551
Interface naming conventions....551

Running-configuration synchronization 552
VSF deployment methods....... 552
Discovered configuration mode procedure....... 552
Provisioned configuration mode procedure....... 552
Configuration commands....... 553
copy core-dump....... 553
copy crash-data....... 553
copy crash-files....... 553
copy crash-log vsf-member....... 554
copy fdr-log....... 554
erase fdr-log vsf....... 555
power-over-ethernet vsf-member....... 555
redundancy switchover....... 556
snmp-server enable traps vsf....... 556
vsf domain....... 556
vsf enable....... 556
vsf member reboot....... 557
vsf member remove....... 558
vsf member shutdown....... 559
vsf member priority....... 560
vsf member type 560
Show commands....... 561
show vsf....... 561
show vsf link....... 562
show vsf member....... 563
show system information vsf member 564
show system chassislocate....... 565
show boot-history....... 566
show system temperature....... 566
show system fans....... 567
show CPU....... 568
show CPU process slot....... 569
show modules....... 570
show power-over-ethernet....... 572
show system power-supply....... 573
OOBM-MAD commands....... 577
vsf oobm-mad....... 577
oobm vsf member....... 577
oobm vsf member interface speed-duplex....... 578
show oobm vsf member....... 578
show oobm ip....... 579
show oobm discovery....... 580
show running-config oobm....... 580
LLDP-MAD....... 582
VSF split and MAD operation....... 582
MAD readiness check....... 582
vsf lldp-mad ipv4....... 583
show vsf lldp-mad....... 584
VSF re-join after a split....... 585
Mad assist device requirements....... 585
Limitations of MAD....... 585
VSF restrictions....... 585
Updates for a VSF virtual chassis....... 585

Chapter 25 VSF Fast Software Upgrade....... 586

Upgrading the VSF stack software 586
vsf sequenced-reboot 586

Chapter 26 IP Service Level Agreement 588

Overview 588
How IP SLA works 590
Configuration commands 590
[no] ip-sla <ID> 590
ip-sla <ID> clear 591
[no] ip-sla <ID> history-size 591
[no] ip-sla <ID> icmp-echo 591
[no] ip-sla <ID> udp-echo 592
[no] ip-sla <ID> tcp-connect 592
[no] ip-sla <ID> monitor threshold-config 592
[no] ip-sla <ID> monitor packet-loss 592
[no] ip-sla <ID> monitor test-completion 593
[no] ip-sla <ID> schedule 593
[no] ip-sla <ID> tos 593
[no] ip-sla responder 593
Show commands 594
show ip-sla <ID> 594
show ip-sla <ID> history 595
show ip-sla <ID> message-statistics 595
show ip-sla responder 596
show ip-sla responder statistics 596
show tech ip-sla 596
`clear ip-sla responder statistics` 598
Validation rules 599
Event log messages 602
Interoperability 603
IP SLA UDP Jitter and Jitter for VoIP 603
Overview 603
Significance of jitter 604
Solution components 604

Chapter 27 Aruba Central integration 605

Aruba Central Network Management Solution Overview 605
LED Blink feature 606
Configuration commands 606
aruba-central 606
Show commands 607
show aruba-central 607

Chapter 28 Easing Wired/Wireless Deployment feature integration 608

Overview 608
Configuration commands 608
allow-jumbo-frames 609
Default AP Profile 609
device-profile 609
Associating a profile with a device 610
device-profile type-device 610
Configuring the rogue-ap-isolation command 611
rogue-ap-isolation 611

Show commands ... 612
show device-profile ... 612
show command device-profile status ... 613
Show rogue-ap-isolation ... 613

Chapter 29 IPSec for AirWave Connectivity ... 615
Overview ... 615
Applicable products ... 615
AirWave details ... 615
IPsec for Management Traffic ... 615
IPsec Tunnel Establishment ... 615
IPSec Tunnel Failures ... 615
AirWave IP after discovery ... 616
Configuring the Aruba controller ... 616
AirWave Controller IP configuration commands ... 619
aruba-vpn type ... 619
Show commands ... 620
show aruba-vpn ... 620
show ip route ... 620
show interfaces tunnel aruba-vpn ... 621
show ip counters tunnel aruba-vpn ... 622
show crypto-ipsec sa ... 624
show running-configuration ... 624

Chapter 30 Local user roles ... 626
Overview ... 626
Captive-portal commands ... 628
Overview ... 628
[no] aaa authentication captive-portal profile ... 628
Policy commands ... 629
Overview ... 629
policy user ... 630
[no] policy user ... 630
policy resequence ... 630
Commands in the policy-user context ... 630
User role configuration ... 631
aaa authorization user-role ... 631
captive-portal-profile ... 633
policy ... 634
reauth-period ... 634
VLAN commands ... 635
VLAN range commands ... 635
Applying a UDR ... 636
aaa port-access local-mac apply user-role ... 636
Show commands ... 637
show captive-portal profile ... 637
show user-role ... 637
show port-access clients ... 639
Per-user tunneled node ... 640
`tunneled-node-server-redirect` ... 641
`show user-role` ... 642
show tunneled-node-server information ... 642
`show tunneled-node-users` ... 645
Tagged VLAN for user role ... 647

vlan-id-tagged....647
user-role vlan-id....648
Downloadable user-roles....649
aaa authorization user-role enable download....649
radius-server cppm identity....650
downloadable-role-delete....650
show user-role <XYZ>....651
show port-access clients....651
debug usertn....653

Chapter 31 Port QoS Trust Mode....655
Overview....655
Configuration commands....655
qos trust....655
qos dscp-map....656
Show commands....656
show qos trust....656
Validation rules....657

Chapter 32 Tunneled node....659
Overview....659
Operating notes....659
Protocol Application Programming Interface (PAPI)....659
Configuration commands....660
tunneled-node-server....660
interface tunneled-node-server....660
controller-ip....661
keepalive....661
backup-controller-ip....661
fallback-local-switching....662
Show commands....662
show tunneled-node-server....662
clear statistics tunneled-node-server....663
Interaction table....663
Restrictions....664
PAPI security....665
Protocol Application Programming Interface (PAPI)....665
PAPI configurable secret key....665
papi-security....666

Chapter 33 Time Domain Reflectometry....668
Virtual cable testing....668
Test cable-diagnostics....668
show cable-diagnostics....671
clear cable-diagnostics....671
Limitations....672

Chapter 34 Link Layer Discovery Protocol bypass authentication....673
Overview....673
Configuration commands....673
aaa port-access lldp-bypass....673
Show commands....675

show port-access lldp-bypass clients.......675
show port-access lldp-bypass config.......676
Error Log.......677
Debug log.......678

Chapter 35 Net-destination and Net-service.......680

Net-service Overview.......680
netservice [tcp | udp | port].......680
Net-destination overview.......681
net-destination host |position | network.......682
show net-destination.......682

Chapter 36 Static IP Visibility.......683

Static IP Visibility.......683
IP client-tracker.......683

Chapter 37 Network Out-of-Band Management (OOBM).......686

OOBM concepts.......686
OOBM and switch applications.......687
Example.......688
OOBM Configuration.......688
Entering the OOBM configuration context from the general configuration context.......689
Enabling and disabling OOBM.......689
Enabling and disabling the OOBM port.......689
Setting the OOBM port speed.......690
Configuring an OOBM IPv4 address.......690
Configuring an OOBM IPv4 default gateway.......691
Configuring an IPv6 default gateway for OOBM devices.......691
IPv6 default router preferences.......692
OOBM `show` commands.......693
Showing the global OOBM and OOBM port configuration.......694
Showing OOBM IP configuration.......694
Showing OOBM ARP information.......694
`show oobm ipv6`.......694
show oobm ipv6 (for stacked switches).......695
show oobm ipv6 member (for stacked switches).......696
show oobm ip detail (for stacked switches).......696
Application server commands.......697
Application client commands.......698

Chapter 38 Websites.......701

Chapter 39 Support and other resources.......702

Accessing Hewlett Packard Enterprise Support.......702
Accessing updates.......702
Customer self repair.......702
Remote support.......703
Warranty information.......703
Regulatory information.......703
Documentation feedback.......704

Chassis Redundancy (HPE 5400R Switches).. 705

Overview of chassis management redundancy.. 705
Nonstop switching with redundant management modules.. 705
How the management modules interact.. 705
About using redundant management.. 705
Transition from no redundancy to nonstop switching.. 706
About setting the rapid switchover stale timer.. 706
About directing the standby module to become active.. 706
Preferred management module.. 706
`redundancy active-management`.. 706
`redundancy preferred-active-management`.. 707
`show redunancy`.. 709
Determining active module.. 710
Diagram of the decision process.. 711
Hotswapping management modules.. 711
Hotswapping out the active management module.. 712
Management module switchover.. 712
Software version mismatch between active and hotswapped module.. 714
Other software version mismatch conditions.. 714
About turning off redundant management.. 714
Disable management module redundancy with two modules present.. 715
Disable management module redundancy with only one module present.. 715
Active management module commands.. 715
Viewing modules.. 716
Syncing commands.. 716
Using the WebAgent for redundant management.. 716
Enabling or disabling redundant management.. 717
Transitioning from `no redundancy` to nonstop switching 720
Setting the Rapid Switchover Stale Timer.. 720
Directing the standby module to become active.. 721
Directing the standby module to become active.. 722
Setting the active management module for next boot.. 723
Resetting the management module.. 725
Viewing management information.. 726
Viewing information about the management and fabric modules.. 727
Viewing information about the redundancy role of each management module.. 727
Viewing which software version is in each flash image.. 727
Viewing system software image information for both management modules.. 728
Viewing the status of the switch and its management modules.. 728
Standby management module commands.. 729
Viewing redundancy status on the standby module.. 729
Viewing the flash information on the standby module.. 730
Viewing the version information on the standby module.. 730
Setting the default flash for boot.. 731
Booting the active management module from the current default flash.. 731
`boot` command 732
`Boot` and `reload` commands with OSPFv2 or OSPFv3 enabled 734
Displaying module events.. 736
Viewing log events.. 736
Copying crash file information to another file.. 737
Viewing saved crash information.. 738
Enabling and disabling fabric modules.. 738
Nonstop switching features.. 739
Nonstop switching with VRRP.. 739

Example nonstop routing configuration....740
Nonstop forwarding with RIP....741
Nonstop forwarding with OSPFv2 and OSPFv3....741
About downloading a new software version....744
File synchronization after downloading....744
Potential software version mismatches after downloading....745
Downloading a software version serially if the management module is corrupted....746
Unsupported zl modules....746
Hot swapping of management modules....747
Rapid routing switchover and stale timer....747
Task Usage Reporting....747
Help text....747
Command tab....748
Command ouput....749

Smart Rate Technology....751
Show Smart Rate port....751
Rate-Limiting — GMB features when Fast-Connect SmartRate ports are configured....752
Error messages....753
Speed-duplex....753
Limitations on 5Gbps ports....753
Error messages....753

HPE Networking 6th Generation Switch ASIC....755
Introduction....755
Commands....755
Configuration setup....755
V3 to V2 compatibility....756
Show commands....757
Version 2 — version 3 blade compatibility on the 5400R switch....759
Allow V2 command....759
Show commands....759

MAC Address Management....761
Overview....761
Determining MAC addresses....761
Viewing the MAC addresses of connected devices....761
Viewing the switch's MAC address assignments for VLANs configured on the switch....762
Viewing the port and VLAN MAC addresses....763

Chapter 1
About this guide

This guide provides information on how to configure, manage, and monitor basic switch operation.

Applicable products

This guide applies to these products:

Aruba 3800 Switch Series (J9573A, J9574A, J9575A, J9576A, J9584A)

Aruba 3810 Switch Series (JL071A, JL072A, JL073A, JL074A, JL075A, JL076A)

Aruba 5400R zl2 Switch Series (J9821A, J9822A, J9850A, J9851A, JL001A, JL002A, JL003A, JL095A)

Switch prompts used in this guide

Examples in this guide are representative and may not match your particular switch/environment. The following table explains the types of command prompts that may be used in examples, along with information on what each prompt indicates.

Prompt	Explanation
`switch#`	`#` indicates manager context (authority).
`switch>`	`>` indicates operator context (authority).
`switch(config)#`	`(config)` indicates the config context.
`switch(vlan-x)#`	`(vlan-x)` indicates the vlan context of config, where *x* represents the VLAN ID. For example: `switch(vlan-128)#`.
`switch(eth-x)#`	`(eth-x)` indicates the interface context of config, where `x` represents the interface. For example: `switch(eth-48)#`.
`switch-Stack#`	`Stack` indicates stacking is enabled.
`switch-Stack(config)#`	`Stack(config)` indicates the config context while stacking is enabled.
`switch-Stack(stacking)#`	`Stack(stacking)` indicates the stacking context of config while stacking is enabled.
`switch-Stack(vlan-x)#`	`Stack(vlan-x)` indicates the vlan context of config while stacking is enabled, where *x* represents the VLAN ID. For example: `switch-Stack(vlan-128)#`.
`switch-Stack(eth-x/y)#`	`Stack(eth-x/y)` indicates the interface context of config, in the form `(eth-<member-in-stack>/<interface>)`. For example: `switch(eth-1/48)#`

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

For successful time protocol setup and specific configuration details, contact your system administrator regarding your local configuration. The HPE Aruba OS switch utilizes the Network Time Protocol (NTP)

NTP

NTP synchronizes the time of day among a set of distributed time servers and clients in order to correlate events when receiving system logs and other time-specific events from multiple network devices. NTP uses the User Datagram Protocol (UDP) as its transport protocol.

All NTP communications use Coordinated Universal Time (UTC). An NTP server usually receives its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server, and then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of each other.

NTP uses a stratum to describe the distance between a network device and an authoritative time source:

- A stratum 1 time server is directly attached to an authoritative time source (such as a radio or atomic clock or a GPS time source).
- A stratum 2 NTP server receives its time through NTP from a stratum 1 time server.

Before synchronizing, NTP compares the time reported by several network devices and does not synchronize with one that is significantly different, even if it is a stratum 1.

You an use the security features of NTP to avoid the accidental or malicious setting of incorrect time. One such mechanism is available: an encrypted authentication mechanism.

Though similar, the NTP algorithm is more complex and accurate than the Simple Network Time Protocol (SNTP).

Enabling this feature results in synchronizing the system clock; therefore, it may affect all sub-systems that rely on system time.

NTP related commands

The following commands allow the user to configure NTP or show NTP configurations.

timesync

Syntax

```
[no]timesync [timep |sntp | timep-or-sntp | ntp]
```

Description

Use this command to configure the protocol for network time synchronization.

Parameters and options

no

Deletes all timesync configurations on the device.

timep

Updates the system clock using TIMEP.

sntp

Updates the system clock using SNTP.

timep-or-sntp

Updates the system clock using TIMEP or SNTP (default).

ntp

Updates the system clock using NTP

timesync

```
Switch(config)# timesync
sntp                    Update the system clock using SNTP.
timep                   Update the system clock using TIMEP.
timep-or-sntp           Update the system clock using TIMEP or SNTP.
ntp                     Update the system clock using NTP.
```

timesync ntp

Syntax

`timesync ntp`

Description

Use this command to update the system clock using NTP.

ntp

Syntax

`[no] ntp [broadcast|unicast]`

Description

This command selects the operating mode of the NTP client. Defaults to broadcast.

Parameters and options

no

Using `no ntp` disables NTP and removes all NTP configurations on the device.

no ntpExample

```
switch(config)# no ntp
This will delete all NTP configurations on this device. Continue [y/n]?
```

broadcast

Sets ntp client to operate in broadcast mode.

unicast

Sets ntp client to operate in unicast mode.

[no] ntp

This command disables NTP and removes all NTP configurations on the device.

Syntax

```
[no] ntp [authentication <key-id>
 | broadcast | enable | max-association
<integer> | server
<IP-ADDR> | trap
<trap-name> | unicast]
```

Description

Disable NTP and removes the entire NTP configuration.

Options

authentication

Configure NTP authentication.

broadcast

Operate in broadcast mode.

enable

Enable/disable NTP.

max-association

Maximum number of Network Time Protocol (NTP) associations.

server

Configure a NTP server to poll for time synchronization.

trap

Enable/disable NTP traps.

unicast

Operate in unicast mode.

Example

```
switch(config)# no ntp
This will delete all NTP configurations on this device. Continue [y/n]?
```

ntp enable

Syntax

```
ntp enable
```

Description

Use this command to enable or disable NTP on the switch.

Restrictions

Validation	Error/Warning/Prompt
If timeSync is in SNTP or Timep when NTP is enabled.	`Timesync is not configured to NTP.`
When timesync is NTP and ntp is enabled and we try to change timesync to SNTP.	`Disable NTP before changing timesync to SNTP or TIMEP`

Enable ntp

```
switch(config)# ntp
enable        Enable/disable NTP.
```

ntp authentication

Syntax

ntp authentication key-id *<KEY-ID>* [authentication-mode *<MODE>* key-value *<KEY-STRING>*] [trusted]

Description

This command is used for authentication of NTP server by the NTP client.

Parameters and options

key-id *<KEY-ID>*

Sets the key-id for the authentication key.

authentication-mode

Sets the NTP authentication mode

key-value <KEY-STRING>

Sets the key-value for the authentication key.

[trusted]

Sets the authentication key as trusted.

ntp authentication

```
Switch(config)# ntp
Authentication      Configure NTP authentication.


Switch(config)# ntp authentication
key-id              Set the key-id for this authentication key.


Switch(config)# ntp authentication key-id
<1-4294967295>      Set the authentication key-id.


Switch(config)# ntp authentication key-id 1
authentication-mode  Set the NTP authentication mode.
trusted              Set this authentication key as trusted.


Switch(config)# ntp authentication key-id 1
authentication-mode|trusted md5
Authenticate using MD5.


Switch(config)# ntp authentication key-id 1
authentication-mode|trusted md5key-value  Set the NTP authentication key.


Switch(config)# ntp authentication key-id 1
```

```
authentication-mode|trusted md5 key-value
KEY            Enter a string to be set as the NTP authentication key.
```

ntp max-associations

Syntax

`ntp max-associations <number>`

Description

Use this command to configure the maximum number of servers associated with this NTP client.

Parameters and options

<number>

Sets the maximum number of NTP associations, in the range of 1–8.

ntp max-associations

```
Switch(config)# ntp
max-associations      Maximum number of NTP associations.

Switch(config)# ntp max-associations
<1-8>                  Enter the number.
```

Restrictions

Validation	Error/Warning/Prompt
When the number of configured NTP servers is more than the max-associations value.	`The maximum number of NTP servers allowed is <number>.`
When the max-associations value is less than the (n) number of configured NTP servers.	`Max-associations value cannot be less than the number of NTP servers configured.`

ntp server

Syntax

`[no] ntp server <IP-ADDR|IPv6-ADDR> [key <KEY-ID>] [oobm] [max-poll <MAX-POLL-VAL>] [min-poll <MIN-POLL-VAL>][burst | iburst] [version <1-4>]`

Description

This command is used to configure the NTP servers. Configure a maximum of 8 NTP servers.

Parameters and options

no

Removes the unicast NTP configurations on the device.

IP-ADDR

Sets the IPv4 address of the NTP server.

IPv6-ADDR

Sets the IPv6 address of the NTP server.

KEY-ID

Specifies the authentication key.

`oobm`

Specifies that the NTP Unicast server is accessible over an OOBM interface.

MIN-POLL-VAL

Configures the minimum time intervals in seconds. Range is 4–17.

MAX-POLL-VAL

Configures the maximum time intervals in power of 2 seconds. Range is 4–17 (e.g., 5 would translate to 2 raised to 5 or 32).

`burst`

Enables burst mode.

`iburst`

Enables initial burst mode.

`version`

Sets version 1–4.

Restrictions

Validation	Error/Warning/Prompt
If authentication key-id not configured	`Authentication key-id has not been configured.`
If Key-id is not marked as trusted	`Key-id is not trusted.`
When min poll value is more than max poll value	`NTP max poll value should be more than min poll value.`

ntp server configuration

```
Switch(config)# ntp
server          Allow the software clock to be synchronized by an NTP
time server.
broadcast       Operate in broadcast mode.
unicast         Operate in unicast mode.


Switch(config)# ntp server
IP-ADDR         IPv4 address of the NTP server.
IPV6-ADDR       IPv6 address of the NTP server.

Switch(config)# ntp server <IP-ADDR>
Key             Specify the authentication key.


Switch(config)# ntp server <IP-ADDR> key key-id
Max-poll        Configure the maximum time intervals in seconds.

Switch(config)# ntp server <IP-ADDR> key key-id max-poll
<4-17>          Enter an integer number.
```

```
Switch(config)# ntp server <IP-ADDR> key key-id
Min-poll        Configure the minimum time intervals in seconds.


Switch(config)# ntp server <IP-ADDR> key key-id min-poll
<4-17>          Enter an integer number.


Switch(config)# ntp server <IP-ADDR> key key-id prefer max-poll
<max-poll-val> min-poll <min-poll-val>
iburst          Enable initial burst (iburst) mode.
burst           Enable burst mode.


Switch(config)# ntp server IP-ADDR key key-id prefer maxpoll <number>
minpoll <number> iburst
```

ntp server key-id

Syntax

```
ntp server <IP-ADDR |IPV6-ADDR>
key—id <key-id> [max-poll
<max-poll-val>] [min-poll
<min-poll-val>] [burst | iburst]
```

Description

Configure the NTP server. *<IP-ADDR>* indicates the IPv4 address of the NTP server. *<IPV6-ADDR>* indicates the IPv6 address of the NTP server.

Options

burst

Enables burst mode.

iburst

Enables initial burst (iburst) mode.

key-id

Set the authentication key to use for this server.

max-poll *<max-poll-val>*

Configure the maximum time intervals in seconds.

min-poll *<min-poll-val>*

Configure the minimum time intervals in seconds.

ntp ipv6-multicast

Syntax

```
ntp ipv6-multicast
```

Description

Use this command to configure NTP multicast on a VLAN interface.

Restrictions

Validation	Error/Warning/Prompt
If ipv6 is not enabled on vlan interface	IPv6 address not configured on the VLAN.

ntp ipv6–multicast

```
Switch(vlan-2)# ntp
ipv6-multicast       Configure the interface to listen to the NTP multicast
packets.
```

debug ntp

Syntax

`debug ntp [event|packet]`

Description

Use this command to display debug messages for NTP.

Parameters and options

`event`

Displays event log messages related to NTP.

`packets`

Displays NTP packet messages.

debug ntp

```
Switch(config)# debug ntp
event                  Display event log messages related to NTP.
packet                 Display NTP packet messages.
```

ntp trap

Syntax

`[no] ntp trap <TRAP-NAME>`

Description

Use this command to configure NTP traps.

Parameters and options

`no`

Disables NTP traps.

TRAP-NAME

Specifies the NTP trap name.

Specifiers

Specify trap names as follows:

```
ntp-mode-change
ntp-stratum-change
ntp-peer-change
```

```
ntp-new-association
ntp-remove-association
ntp-config-change
ntp-leapsec-announced
ntp-alive-heartbeat
```

Usage

The traps defined below are generated as the result of finding an unusual condition while parsing an NTP packet or a processing a timer event. Note that if more than one type of unusual condition is encountered while parsing the packet or processing an event, only the first one will generate a trap. Possible trap names are:

- 'ntpEntNotifModeChange' The notification to be sent when the NTP entity changes mode, including starting and stopping (if possible).

- 'ntpEntNotifStratumChange' The notification to be sent when stratum level of NTP changes.

- 'ntpEntNotifSyspeerChanged' The notification to be sent when a (new) syspeer has been selected.

- 'ntpEntNotifAddAssociation' The notification to be sent when a new association is mobilized.

- 'ntpEntNotifRemoveAssociation' The notification to be sent when an association is demobilized.

- 'ntpEntNotifConfigChanged' The notification to be sent when the NTP configuration has changed.

- 'ntpEntNotifLeapSecondAnnounced' The notification to be sent when a leap second has been announced.

- 'ntpEntNotifHeartbeat' The notification to be sent periodically (as defined by ntpEntHeartbeatInterval) to indicate that the NTP entity is still alive.

show ntp statistics

Syntax

`show ntp statistics`

Description

Use this command to show NTP statistics.

show ntp statistics

```
Switch(config)# show ntp statistics

NTP Global statistics information

NTP In Packets                : 100
NTP Out Packets               : 110
NTP Bad Version Packets       : 4
NTP Protocol Error Packets    : 0
```

show ntp status

Syntax

`show ntp status`

Description

Use this command to show the status of the NTP.

show ntp status

```
Switch(config)# show ntp status

NTP Status information
NTP Status              : Disabled               NTP Mode        : Broadcast
Synchronization Status : Synchronized           Peer Dispersion : 8.01 sec
Stratum Number          : 2                      Leap Direction  : 1
Reference Assoc Id      : 1                      Clock Offset    : 0.0000 sec
Reference               : 192.0.2.1              Root Delay      : 0.00 sec
Precision               : 2**7                   Root Dispersion : 15.91 sec
NTP Uptime              : 01d 09h 15m            Time Resolution : 1
Drift                   : 0.000000000 sec/sec

System Time             : Tue Aug 25 04:59:11 2015
Reference Time          : Mon Jan  1 00:00:00 1990
```

show ntp authentication

Syntax

`show ntp authentication`

Description

Use this command to show the authentication status of the NTP.

show ntp authentication

```
Switch(config)# show ntp authentication

NTP Authentication Information

Key-ID     Auth Mode    Trusted
--------   -----------  -------
67            md5         yes
7             md5         no
```

show ntp associations

Syntax

`show ntp associations`

Description

Use this command to show the NTP associations configured for your system.

show ntp associations

```
Switch(config)# show ntp associations

                          NTP Associations Entries

Address           St   T  When Poll  Reach   Delay    Offset   Dispersion
--------------    ---  -- ---- ----- ------ -------   -------  ----------
121.0.23.1        16   u   -   1024    0      0.000     0.000    0.000
231.45.21.4       16   u   -   1024    0      0.000     0.000    0.000
55.21.56.2        16   u   -   1024    0      0.000     0.000    0.000
```

```
23.56.13.1         3    u 209  1024    377      54.936   -6.159    12.688
91.34.255.216      4    u 132  1024    377      1.391     0.978    3.860
```

show ntp associations detail

Syntax

`show ntp associations detail <IP ADDR>`

Description

Use this command to show the detailed status of NTP associations configured for your system.

Parameters and options

IP-ADDR

Specify the IPv4 address of the NTP server.

show ntp association detail

```
Switch(config)# show ntp association detail <IP ADDR>

NTP association information

IP address                 : 172.31.32.2                  Peer Mode         : Server
Status                     : Configured, Insane, Invalid Peer Poll Intvl  : 64
Stratum                    : 5                            Root Delay        : 137.77
sec
Ref Assoc ID               : 0                            Root Dispersion   : 142.75
Association Name           : NTP Association 0            Reach             : 376
Reference ID               : 16.93.49.4                   Delay             : 4.23 sec
Our Mode                   : Client                       Offset            : -8.587
sec
Our Poll Intvl             : 1024                         Precision         : 2**19
Dispersion                 : 1.62 sec
Association In Packets     : 60
Association Out Packets    : 60
Association Error Packets  : 0
Origin Time                : Fri Jul  3 11:39:40 2015
Receive Time               : Fri Jul  3 11:39:44 2015
Transmit Time              : Fri Jul  3 11:39:44 2015

-----------------------------------------------------------------------------
Filter Delay =   4.23    4.14    2.41    5.95    2.37    2.33    4.26    4.33
Filter Offset = -8.59   -8.82   -9.91   -8.42  -10.51  -10.77  -10.13  -10.11
```

Validation Rules

Validation	Error/Warning/Prompt
If access-list name is not valid.	Please enter a valid access-list name.
If the authentication method is being set to two-factor authentication, various messages display.	If both the public key and username/password are not configured: Public key and username/password should be configured for a successful two-factor authentication. If public key is configured and username is not configured: Username and password should be configured for a successful two-factor authentication. If the username is configured and public key is not configured: Public key should be configured for a successful two-factor authentication. If “ssh-server” certificate is not installed at the time of enabling certificate-password authentication: The “ssh-server” certificate should be installed for a successful two-factor authentication.
If the authentication method is set to two-factor while installing the public key, a message displays.	The client public keys without username will not be considered for the two-factor authentication for the SSH session.
If the username and the key installation user for that privilege do not match, a message displays and installation is not allowed. This will also happen when the authentication method is set for two-factor.	The username in the key being installed does not match the username configured on the switch.
If the maximum number of <username : TA profile> associations is reached for a given TA profile, a message displays.	Maximum number of username associations with a TA profile is 10.
If secondary authentication type for two-factor authentication chosen is not "none", a message displays.	Not legal combination of authentication methods.
If the authentication method is anything other than two-factor and the two-factor authentication method options are set, a message displays.	Not legal combination of authentication methods.
If two-factor authentication is set and user tries to SSH into another system using “ssh <ip | hostname>” command, a message displays.	SSH client is not supported when the two-factor authentication is enabled.

Table Continued

Validation	Error/Warning/Prompt
If timeSync is in SNTP or TimeP when NTP is enabled.	Timesync is not configured to NTP.
If timesync is NTP and NTP is enabled and we try to change timesync to SNTP.	Disable NTP before changing timesync to SNTP or TimeP.
If we try to configure NTP servers more than the configured max-associations value.	The maximum number of NTP servers allowed is 2.
If we have ‘n’ NTP servers configured and we try to configure a max-associations value less than (n) number of NTP servers already configured.	Max-associations value cannot be less than the number of NTP servers configured.
If authentication key-id is not configured.	Authentication key-id %d has not been configured.
If key-id is not marked as trusted.	Key-id %d is not trusted.
If min poll value is more than max poll value.	NTP max poll value should be more than min poll value.
If ipv6 is not enabled on vlan interface.	IPv6 address not configured on the VLAN.

Event log messages

Event	Message
RMON_AUTH_TWO_FACTOR_AUTHEN_STATUS	W 01/01/15 18:24:03 03397: auth: %s. Examples: W 01/01/15 18:24:03 03397: auth: Public key and username/password should be configured for the successful two-factor authentication. W 01/01/15 18:24:03 03397: auth: Username and password should be configured for the successful two-factor authentication. W 01/01/15 18:24:03 03397: auth: Public key should be configured for the successful two-factor authentication. I 01/01/15 18:24:03 03397: auth: The validation of certificate of SSH user (user1) is successful.
RMON_SSH_KEY_TWO_FACTOR_EN	W 01/01/15 18:24:03 03399: ssh: %s. Examples: W 01/01/15 18:24:03 03399: ssh: The client public keys without username will not be considered for the two-factor authentication for SSH session. W 01/01/15 18:24:03 03399: ssh: The privilege level for the user with the SSH key conflicts with the user configured.
RMON_SSH_TWO_FACTOR_AUTH_FAIL	W 01/01/15 18:24:03 03398: ssh: %s. Examples: W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in public key authentication. W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in username/password authentication. W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in validating the client certificate. W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed as “ssh-server” certificate is not installed.

Elements of time synchronization

Time synchronization contains several elements. These include:

- **Protocol** — SNTP or TimeP. The switch offers `TimeP` and `SNTP` (Simple Network Time Protocol) and a `timesync` command for changing the time protocol selection (or turning off time protocol operation.)
- **Authentication modes** — Broadcast or Unicast for SNTP, and DHCP or Manual for TimeP
- **Status** — Enabled or Disabled. Simply selecting a time synchronization protocol does not enable that protocol on the switch. You must also enable the protocol itself by setting the appropriate parameter (enabled or disabled).

Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time. In addition, the switch retains the parameter settings for both time protocols, even if you change from one protocol to the other. Thus, if you select a time protocol, the switch uses the parameters you last configured for the selected protocol.

Time synchronization protocols

Use the `timesync` command to set the time synchronization protocol, either `SNTP` or `TimeP`.

- **SNTP**—To run SNTP as the switch's time synchronization protocol, you must also select SNTP as the time synchronization method using the CLI `timesync` command, or the menu interface **Time Sync Method** parameter.
- **TimeP**—You can manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one, designated TimeP server. This option enhances security by specifying which time server to use.

timesync

Syntax

`timesync [timep|sntp]`

Description

The `timesync` command configures the network time protocol for `sntp` or `timep`.

Parameters and options

`sntp`

Sets the time protocol to `SNTP`.

`TimeP`

Sets the time protocol to `TIMEP`.

timesync [timep | sntp]

```
switch# timesync timep
```

```
switch# timesync sntp
```

Setting a time protocol on the switch

Procedure

1. Select a time synchronization protocol: `SNTP` or `TimeP` (the default). See **timesync** on page 39.
2. Enable the protocol. Choose one:
   - SNTP: `Broadcast` or `Unicast`
   - TimeP: `DHCP` or `Manual`

3. Configure the remaining parameters for the time protocol you selected.
4. View the configuration.

The SNTP protocol

SNTP provides the following operating modes:

- **Broadcast mode**

  The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address; see the documentation provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable Poll Interval expires three consecutive times without an update received from the first-detected server. If the Poll Interval (configurable up to 720 seconds) expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval (configurable up to 720 seconds) expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.

  NOTE: To use Broadcast mode, the switch and the SNTP server must be in the same subnet.

- **Unicast mode**

  Directs the switch to poll a specific server periodically for SNTP time synchronization.The default value between each polling request is 720 seconds, but can be configured.At least one manually configured server IP address is required.

  NOTE: At least one `key-id` must be configured as `trusted`, and it must be associated with one of the SNTP servers. To edit or remove the associated `key-id` information or SNTP server information, SNTP authentication must be disabled.

  The switch periodically requests a time update, for the purposes of time synchronization, from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI `sntp` server command.) This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast. The default value between each polling request is 720 seconds, but can be configured. At least one manually configured server IP address is required.

  When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured, with either the server address parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured `Poll Interval` time has expired.

  If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Selecting and configuring SNTP

Procedure

1. Use the `SNTP` command to specify whether the switch operates in broadcast or unicast mode. With no mode specified, the setting defaults to broadcast.

Prerequisites

- Configure at least one `key-id` as `trusted`, and then associate it with one of the SNTP servers (see **SNTP authentication trusted keys** on page 49
- Configure the appropriate parameters, such as poll interval, server address and version
- To edit or remove the associated `key-id` information or SNTP server information, disable SNTP authentication

sntp

> To enable authentication, you must configure either unicast or broadcast mode. After authentication is enabled, changing the mode from unicast to broadcast or vice versa is not allowed; you must disable authentication and then change the mode.

To set the SNTP mode or change from one mode to the other, enter the appropriate command.

Syntax

`sntp`

Description

This command configures SNTP, including specifying whether the switch operates in broadcast or unicast mode.

Parameters and options

`Disabled`

The Default. SNTP does not operate, even if specified by the Menu interface **Time Sync Method** parameter or the CLI `timesync` command.

`Unicast`

Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address.

`Broadcast`

Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.

Poll interval (seconds)

In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update.

In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.

Value is between 30 to 720 seconds.

Server Address

Used only when the **SNTP Mode** is set to `Unicast`. Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI.

Server Version

Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 to 7.

Priority

Specifies the order in which the configured servers are polled for getting the time.

Value is between 1 and 3.

oobm

For switches that have a separate out-of-band management port, specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)

sntp broadcast|unicast output

```
switch# sntp broadcast
```

```
switch# sntp unicast
```

Enabling SNTP in Broadcast mode

Because the switch provides an SNTP polling interval (default: 720 seconds), you need only **timesync** on page 39 and **sntp** on page 41 commands for minimal SNTP broadcast configuration.

Figure 1: SNTP in Broadcast Mode on page 42 shows time synchronization in the factory default configuration, TimeP.

Procedure

1. To view the current time synchronization, enter `show sntp`.
2. Use the timesync command to set SNTP as the time synchronization mode:

   `timesync sntp`
3. Use the SNTP command to enable SNTP for Broadcast mode:

   `sntp broadcast`
4. View the SNTP configuration again to verify the configuration.

Figure 1: *SNTP in Broadcast Mode*

```
HP Switch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] :720

HP Switch(config)# timesync sntp

HP Switch(config)# sntp broadcast

HP Switch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] :720
```

show sntp displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.

show sntp again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.

Configuring SNTP in unicast mode

As with broadcast mode, configuring SNTP for unicast mode enables SNTP. For unicast operation, however, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing unicast server with another. To add a second or third server, you must use the CLI.

The following is an example of a full SNTP unicast operation.

Procedure

1. Select the SNTP protocol:

   ```
   switch(config)# timesync sntp
   ```

2. Set the mode to unicast:

   ```
   switch(config)# sntp unicast
   ```

3. Specify the SNTP server and set the server priority:

   ```
   switch(config)# sntp server priority 1 10.28.227.141
   ```

   This specifies the SNTP server and accepts the current SNTP server version (default: 3).

```
switch(config)# show sntp
SNTP Configuration
SNTP Authentication : Disabled
Time Sync Mode: Timep
SNTP Mode : disabled
Poll Interval (sec) [720] : 720
Source IP Selection: Outgoing Interface
switch(config)# timesync sntp
switch(config)# sntp broadcast
switch(config)# show sntp
SNTP Configuration
SNTP Authentication : Disabled
Time Sync Mode: Sntp
SNTP Mode : Broadcast
Poll Interval (sec) [720] : 720
Source IP Selection: Outgoing Interface
```

If the SNTP server you specify uses SNTP v4 or later, use the `sntp server` command to specify the correct version number. For example, suppose you learned that SNTP v4 was in use on the server you specified above (IP address 10.28.227.141.) You would use the following commands to delete the server IP address , re-enter it with the correct version number for that server.

```
switch(config)# sntp server priority 1 10.28.227.141 4
switch(config)# show sntp
SNTP Configuration
SNTP Authentication : Disabled
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720
Source IP Selection: Outgoing Interface
Priority SNTP Server Address  Version Key-id
```

```
-------- ------------------  ------- ----------
1        10.28.227.141         4       0
```

Figure 2: *SNTP in unicast mode*

```
HP Switch(config)# show sntp

 SNTP Configuration

  Time Sync Mode: Sntp
  SNTP Mode : Unicast
  Poll Interval (sec) [720] : 720


  Priority SNTP Server Address                         Protocol Version
  -------- ----------------------------------------- ----------------
  1        2001:db8::215:60ff:fe79:8980              7
  2        10.255.5.24                               3
  3        fe80::123%vlan10                          3
```

In this example, the **Poll Interval** and the **Protocol Version** appear at their default settings.

Both IPv4 and IPv6 addresses are displayed.

Note: Protocol Version appears only when there is an IP address configured for an SNTP server.

If the SNTP server you specify uses SNTP v4 or later, use the `sntp server` command to specify the correct version number. For example, suppose SNTP v4 is in use on the server you specified above (IP address 10.28.227.141.) Use the SNTP commands shown in the following figure to delete the server IP address, and then re-enter it with the correct version number for that server.

Figure 3: *Specifying the SNTP protocol version number*

```
HP Switch(config)# no sntp server 10.28.227.141
HP Switch(config)# sntp server 10.28.227.141 4
HP Switch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 600

  IP Address      Protocol Version
  -------------   ----------------
  10.28.227.141   4
```

Deletes unicast SNTP server entry.

Re-enters the unicast server with a non-default protocol version.

show sntp displays the result.

Viewing SNTP parameters

Viewing SNTP server addresses using the CLI

The System Information screen in the menu interface displays only one SNTP server address, even if the switch is configured for two or three servers.

show management

Syntax

`show management`

Description

Displays all configured SNTP servers on the switch.

Viewing SNTP server addresses using the GUI

```
switch# show management

Status and Counters - Management Address Information
Time Server Address : fe80::215:60ff:fe7a:adc0%vlan10

Priority   SNTP Server Address Protocol Version
-------- ---------------------------------------------- ----------------
1        2001:db8::215:60ff:fe79:8980 7
2        10.255.5.24 3
3        fe80::123%vlan10 3

Default Gateway : 10.0.9.80

VLAN Name    MAC Address         | IP Address
------------ ------------------- + -------------------
DEFAULT_VLAN 001279-88a100       | Disabled
VLAN10       001279-88a100       | 10.0.10.17
```

Enabling SNTP client authentication

The command `sntp authentication` enables SNTP client authentication on the switch. If SNTP authentication is not enabled, SNTP packets are not authenticated.

Enabling SNTP authentication allows network devices such as HPE switches to validate the SNTP messages received from an NTP or SNTP server before updating the network time. NTP or SNTP servers and clients must be configured with the same set of authentication keys so that the servers can authenticate the messages they send and clients (switches) can validate the received messages before updating the time.

This feature provides support for SNTP client authentication on switches, which addresses security considerations when deploying SNTP in a network.

Requirements to enable SNTP client authentication

You must configure all of the the following items to enable SNTP client authentication on the switch.

SNTP client Authentication Support Requirements

- Timesync mode must be SNTP. Use the `timesync sntp` command. SNTP is disabled by default.
- SNTP must be in unicast or broadcast mode.
- The MD5 authentication mode must be selected.
- An SNTP authentication key-identifier (`key-id`) must be configured on the switch and a value (`key-value`) must be provided for the authentication key. A maximum of 8 sets of `key-id` and `key-value` can be configured on the switch.
- Among the keys that have been configured, one key or a set of keys must be configured as trusted. Only trusted keys will be used for SNTP authentication.
- If the SNTP server requires authentication, one of the trusted keys has to be associated with the SNTP server.
- SNTP client authentication must be enabled on the switch. If client authentication is disabled, packets are processed without authentication. All of the above steps are necessary to enable authentication on the client.

SNTP server authentication support

The following must be performed on the SNTP server:

- The same authentication key-identifier, trusted key, authentication mode and key-value that were configured on the SNTP client must also be configured on the SNTP server.
- SNTP server authentication must be enabled on the server. If any of the parameters on the server are changed, the parameters have to be changed on all the SNTP clients in the network as well. The authentication check will fail on the clients otherwise, and the SNTP packets will be dropped.

NOTE

SNTP server is not supported on HPE products.

If any of the parameters on the server are changed, the parameters have to be changed on all the SNTP clients in the network as well. The authentication check fails on the clients otherwise, and the SNTP packets are dropped.

Viewing all SNTP authentication keys that have been configured on the switch

Enter the `show sntp authentication` command.

Show SNTP authentication command output

```
switch(config)# show sntp authentication

  SNTP Authentication Information

  SNTP Authentication : Enabled

  Key-ID   Auth Mode   Trusted
  -------- ---------- --------
  55       MD5         Yes
  10       MD5         No
```

SNTP poll interval

This parameter is different from the *poll interval* parameter used for the TimeP operation. Enabling SNTP mode also enables the SNTP poll interval.

sntp poll-interval

Syntax

`sntp poll-interval <30-720>`

Description

Configures the poll interval to specify the amount of time between updates of the system clock using SNTP. Defaults to 720 seconds, and the range is 30 to 720 seconds.

Changing an SNTP poll interval to 300 seconds

```
switch# sntp 300
```

SNTP unicast time polling with multiple SNTP servers

When you use the Menu interface to configure an SNTP server IP address, the new address writes over the current primary address, if one is configured.

When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured with either the `Server Address` parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured *Poll Interval* time has expired.

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

SNTP server priority

Set the server priority to choose the order in which to poll configured servers.

sntp server priority

Syntax

`[no] sntp server priority <ip-address>`

Description

Polls for the current time among configured SNTP servers.

Parameters and options

no

Deletes a server address. If there are multiple addresses and you delete one of them, the switch re-orders the address priority.

`server priority <1-3>`

Specifies the polling order of the configured SNTP servers. Value is between 1 and 3.

<IP-ADDRESS>

Supports bot IPv4 and IPv6 addresses.

Set the server priority

To set one server to priority 1 and another to priority 2:

```
switch# sntp server priority 1 10.28.22.141
switch# sntp server priority 2 2001:db8::215:60ff:fe79:8980
```

Delete a server address

To delete the primary address and automatically convert the secondary address to primary:

```
switch(config)# no sntp server 10.28.227.141
```

SNTP software version

sntp server <version>

Syntax

```
sntp server [<IP-ADDRESS>] [<VERSION>]
```

Description

Specifies the SNTP software version to use. Assigned on a per-server basis.

Parameters and options

<IP-ADDRESS>

SNTP server ip-address

<VERSION>

The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 to 7.

SNTP server address

Required only for unicast mode. Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI.

sntp server <ip-address>

Syntax

```
sntp server <ip-address>
```

Description

Specifies the IP address of the SNTP server for use in unicast mode.

Parameters and options

<ip-address>

An IPv4 or IPv6 address of an SNTP server.

Adding SNTP server addresses

You can configure one SNTP server address using either the Menu interface or the CLI. To configure a second and third address, you must use the CLI. To configure these remaining two addresses, you would do the following:

Creating additional SNTP server addresses with the CLI

```
HP-5406zl(config)# no sntp server priority 1 2001:db8::215:60ff:fe79:8980
HP-5406zl(config)# no sntp server priority 2 10.255.5.24
```

NOTE

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

SNTP authentication trusted keys

Trusted keys are used in SNTP authentication. In unicast mode, you must associate a key with a specific NTP/SNTP server. That key is used for authenticating the SNTP packet.

In unicast mode, a specific server is configured on the switch so that the SNTP client communicates with the specified server to get the date and time.

In broadcast mode, the SNTP client switch checks the size of the received packet to determine if it is authenticated. If the broadcast packet is authenticated, the key-id value is checked to see if the same key-id value is configured on the SNTP client switch. If the switch is configured with the same key-id value, and the key-id value is configured as "trusted," the authentication succeeds. Only trusted key-id value information is used for SNTP authentication.

If the packet contains key-id value information that is not configured on the SNTP client switch, or if the received packet contains no authentication information, it is discarded. The SNTP client switch expects packets to be authenticated if SNTP authentication is enabled.

When authentication succeeds, the time in the packet is used to update the time on the switch.

trusted

Syntax

`trusted`

Description

Parameters and options

Configuration files and the `include-credentials` command

You can use the `include-credentials` command to store security information in the running-config file. This allows you to upload the file to a TFTP server and then later download the file to the switches on which you want to use the same settings.

The authentication key values are shown in the output of the `show running-config` and `show config` commands only if the `include-credentials` command was executed.

When SNTP authentication is configured and `include-credentials` has not been executed, the SNTP authentication configuration is not saved.

The following example shows an enabled SNTP authentication with a key-id of 55.

Configuration file with SNTP authentication information

```
 switch(config) # show config
Startup configuration:
timesync sntp
sntp broadcast
sntp 50
sntp authentication
sntp server priority 1 10.10.10.2.3 key-id 55
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
```

In this example, the `include-credentials` command has not been executed and is not present in the configuration file. The configuration file is subsequently saved to a TFTP server for later use. The SNTP authentication information is not saved and is not present in the retrieved configuration files, as shown in the following example.

Retrieved configuration file when include credentials is not configured

```
switch(config) # copy tftp startup-config 10.2.3.44 config1
Switch reboots ...
Startup configuration
timesync sntp
sntp broadcast
sntp 50 sntp server priority 1 10.10.10.2.3
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4
```

The SNTP authentication line and the Key-ids are not displayed. Reconfigure SNTP authentication.

If `include-credentials` is configured, the SNTP authentication configuration is saved in the configuration file. When the `show config` command is entered, all of the information that has been configured for SNTP authentication displays, including the key-values.

Figure 4: *Saved SNTP Authentication information when include-credentials is configured*

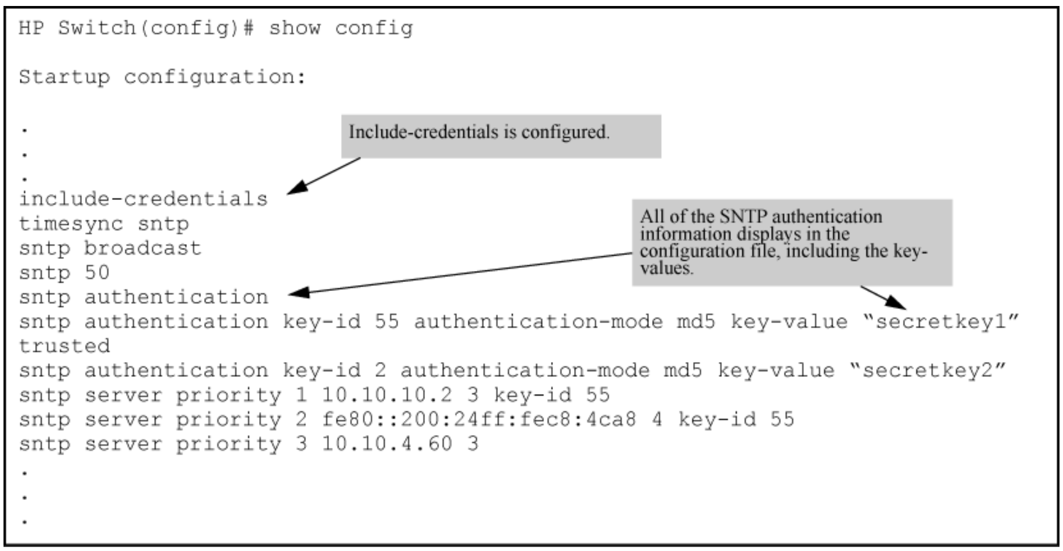

```
HP Switch(config)# show config

Startup configuration:

.
.
.
include-credentials
timesync sntp
sntp broadcast
sntp 50
sntp authentication
sntp authentication key-id 55 authentication-mode md5 key-value "secretkey1"
trusted
sntp authentication key-id 2 authentication-mode md5 key-value "secretkey2"
sntp server priority 1 10.10.10.2 3 key-id 55
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
sntp server priority 3 10.10.4.60 3
.
.
.
```


Configuring the key-identifier, authentication mode, and key-value

Configures the `key-id`, `authentication-mode`, and `key-value`, which are required for authentication. It is executed in the global configuration context.

At least one `key-id` must be configured as `trusted`, and it must be associated with one of the SNTP servers. To edit or remove the associated `key-id` information or SNTP server information, SNTP authentication must be disabled.

A numeric key identifier in the range of 1-4,294,967,295 (2^{32}) that identifies the unique key value. It is sent in the SNTP packet.

The secret key that is used to generate the message digest. Up to 32 characters are allowed for *`key-string`* .

NOTE

For the 3800 switches, when the switch is in enhanced secure mode, commands that take a secret key as a parameter have the echo of the secret typing replaced with asterisks. The input for <key-string> is prompted for interactively.

`encrypted-key <key-string>`

Set the SNTP authentication key value using a base64–encoded aes-256 encrypted string.

sntp authentication

Syntax

```
sntp authentication key-id <KEY-ID> authentication-mode md5 key-value <key-string>
trusted [encrypted-key <key-string>]
```

Description

Configures a key-id, authentication-mode (MD5 only), and key-value, which are required for authentication.

Parameters and options

KEY-ID

A numeric key identifier in the range of 1-4,294,967,295 (2^{32}) that identifies the unique key value. It is sent in the SNTP packet.

`key-value <KEY-STRING>`

The secret key that is used to generate the message digest. Up to 32 characters are allowed for `key-string` .

Disabling key-id

sntp authentication key-id

Syntax

```
no sntp authentication key-id <KEY-ID>
```

Description

The `no` version of the command deletes the authentication key.

Default: No default keys are configured on the switch.

Setting parameters for SNTP authentication

```
switch# sntp authentication key-id 55 authentication-mode md5 key-value secretkey1
```

Configuring a key-id as trusted

Trusted keys are used during the authentication process. You can configure the switch with up to eight sets of key-id/key-value pairs. Select one, specific set for authentication; this is done by configuring the set as `trusted`. The `key-id` itself must already be configured on the switch. To enable authentication, at least one `key-id` must be configured as `trusted`.

- Trusted keys are used in SNTP authentication.
- If the packet contains key-id value information that is not configured on the SNTP client switch, or if the received packet contains no authentication information, it is discarded. The SNTP client switch expects packets to be authenticated if SNTP authentication is enabled.
- When authentication succeeds, the time in the packet is used to update the time on the switch.
- In unicast mode: The trusted key is associated with a specific NTP/SNTP server, and configured on the switch so that the SNTP client communicates with the server to get the date and time. The key is used for authenticating the SNTP packet.
- In : The SNTP client switch checks the size of the received packet to determine if it is authenticated. If the broadcast packet is authenticated, the key-id value is checked to see if the same key-id value is configured on the SNTP client switch. If the switch is configured with the same key-id value, and the key-id value is configured as "trusted," the authentication succeeds. Only trusted key-id value information is used for SNTP authentication.

sntp authentication key-id trusted

Syntax

```
[no] sntp authentication key-id <KEY-ID> trusted
```

Description

Trusted keys are used during the authentication process. You can configure the switch with up to eight sets of key-id/key-value pairs. Select one, specific set for authentication; this is done by configuring the set as `trusted`. The `key-id` itself must already be configured on the switch.

Parameters and options

no

The `no` version of the command indicates the key is unreliable (not trusted).

Default: No key is trusted by default.

`key-id <KEY-ID>`

`trusted`

To enable authentication, configure at least one `key-id` as `trusted`.

Associating a key with an SNTP server

sntp server

Syntax

```
[no] sntp server priority <1-3> <IP-ADDRESS> <VERSION-NUM> <KEY-ID>
<1-4,294,967,295>
```

Description

Configures a `key-id` to be associated with a specific server. The key itself must be configured on the switch.

The `no` version of the command disassociates the key from the server. This does not remove the authentication key.

Default: No key is associated with any server by default.

Parameters and options

`priority <1-3>`

Specifies the order in which the configured servers are polled for getting the time.

`<IP-ADDRESS>`

The IP address of the server. Supports IPv4 or IPv6.

`version-num`

Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 - 7.

`<KEY-ID>`

Optional command. The key identifier sent in the SNTP packet. This `key-id` is associated with the SNTP server specified in the command.

Associating a `key-id` with a specific server

```
switch(config)# sntp server priority 1 10.10.19.5 2 key-id 55
```

sntp server priority

Syntax

[no] sntp server priority 1-3 [<*IP-ADDRESS*>]<*VERSION-NUM*>[<*KEY-ID*> <*1-4,294,967,295*>]

Description

Configures a key to be associated with a specific server. The key itself must already be configured on the switch. Default: No key is associated with any server by default.

Parameters and options

no

Disassociates the key from the server. This does not remove the authentication key.

priority

Specifies the order in which the configured servers are polled for getting the time.

version-num

Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 - 7.

key-id

Optional command. The key identifier sent in the SNTP packet. This `key-id` is associated with the SNTP server specified in the command.

Associating a key-id with a specific server

```
switch# sntp server priority 1 10.10.19.5 2 key-id 55
```

Enabling and disabling SNTP client authentication

The `sntp authentication` command enables SNTP client authentication on the switch. If SNTP authentication is not enabled, SNTP packets are not authenticated.

sntp authentication

Syntax

`[no] sntp authentication`

Description

Enables the SNTP client authentication. SNTP client authentication defaults to disabled.:

Parameters and options

no

Disables authentication.

Viewing SNTP authentication configuration information

show sntp

Syntax

`show sntp authentication`

Description

The `show sntp` command displays SNTP configuration information, including any SNTP authentication keys that have been configured on the switch.

show sntp authentication

To display all the SNTP authentication keys that have been configured on the switch, enter the `show sntp authentication` command.

```
switch(config) # show sntp authentication
SNTP Authentication Information
SNTP Authentication: Enabled
Key-ID                 Auth Mode              Trusted
-------               -----------             -------
55                        MD5                  YES
10                        MD5                   NO
```

Show SNTP authentication command output

```
switch(config)# show sntp authentication

  SNTP Authentication Information

  SNTP Authentication : Enabled

  Key-ID  Auth Mode  Trusted
  ------- ---------- --------
  55      MD5        Yes
  10      MD5        No
```

Viewing all SNTP authentication keys that have been configured on the switch

SNTP configuration information

```
switch(config)# show sntp

 SNTP Configuration

  SNTP Authentication : Enabled
  Time Sync Mode: Sntp
  SNTP Mode : Unicast
  Poll Interval (sec) [720] : 720

  Priority SNTP Server Address                       Protocol Version KeyId
  -------- --------------------------------------- ---------------- -----
  1        10.10.10.2                              3                55
  2        fe80::200:24ff:fec8:4ca8                3                55
```

SNTP Statistics command output

To display the statistical information for each SNTP server, enter the `sntp statistics` command. The number of SNTP packets that have failed authenrtication is displayed for each SNTP server address.

```
switch(config) # show sntp statistics
SNTP statistics
Received Packets:   0
Sent Packets:       3
Dropped Packets:    0

SNTP Server Address                   Auth Failed Pkts
-------------------                   ----------------
10.10.10.1                                0
fe80::200:24ff:fec8:4ca8                  0
```

The `show sntp` command displays SNTP configuration information, including any SNTP authentication keys that have been configured on the switch.

SNTP configuration information

```
switch# show sntp

 SNTP Configuration

  SNTP Authentication : Enabled
  Time Sync Mode: Sntp
  SNTP Mode : Unicast
  Poll Interval (sec) [720] : 720

  Priority SNTP Server Address                       Protocol Version KeyId
  -------- --------------------------------------- ---------------- -----
  1        10.10.10.2                                3                55
  2        fe80::200:24ff:fec8:4ca8                  3                55
```

`show sntp authentication` command output

To display all the SNTP authentication keys that have been configured on the switch, enter the `show sntp authentication` command.

```
switch(config) # show sntp authentication
SNTP Authentication Information
SNTP Authentication: Enabled
Key-ID                Auth Mode             Trusted
-------              -----------            -------
55                       MD5                  YES
10                       MD5                   NO
```

Displays all SNTP authentication keys configured on the switch.

```

switch(config) # show sntp authentication
SNTP Authentication Information
SNTP Authentication: Enabled

Key-ID  Auth Mode  Trusted
-------  ----------- -------
55 MD5 YES
10 MD5 NO
```

Viewing statistical information for each SNTP server

To display the statistical information for each SNTP server, enter the `show sntp statistics` command.

The number of SNTP packets that have failed authentication is displayed for each SNTP server address.

show sntp statistics

```
switch(config)# show sntp statistics
SNTP Statistics

  Received Packets : 0
  Sent Packets : 3
  Dropped Packets : 0

  SNTP Server Address                      Auth Failed Pkts
  ---------------------------------------- ----------------
  10.10.10.1                                   0
  fe80::200:24ff:fec8:4ca8                     0
```

To display the statistical information for each SNTP server, enter the `show sntp statistics` command.

show sntp statistics

Syntax

`show sntp statistics`

Description

Shows the number of SNTP packets that have failed authentication for each SNTP server address.

SNTP authentication statistical information

Shows the statistical information for each SNTP server. The number of SNTP packets that have failed authentication is displayed for each SNTP server address.

```
switch(config) # show sntp statistics
SNTP statistics
Received Packets: 0
Sent Packets: 3
Dropped Packets: 0
SNTP Server Address    Auth Failed Pkts
-------------------    ----------------
10.10.10.1          0
fe80::200:24ff:fec8:4ca8   0
```

```
switch# show sntp statistics
SNTP Statistics

  Received Packets : 0
  Sent Packets : 3
  Dropped Packets : 0

  SNTP Server Address                      Auth Failed Pkts
  ---------------------------------------- ----------------
  10.10.10.1                                   0
  fe80::200:24ff:fec8:4ca8                     0
```

SNTP messages in the event log

If an SNTP time change of more than three seconds occurs, the switch's Event Log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Storing security information in the running-config file

Enter the `include-credentials` command.

The TimeP Protocol

Enabling TimeP as the time protocol means configuring it for either DHCP or manual mode.

To run TimeP as the time synchronization protocol, you must also select TimeP as the time synchronization method by using the CLI `timesync` command or the menu interface **Time Sync Method** parameter.

Procedure

1. To view the current time synchronization, enter `show timep`.
2. Use the timesync command to set TimeP as the time synchronization mode:

   `timesync timep`
3. Use the ip timep command to enable timep for dhcp or manual mode:

   `ip timep dhcp|manual`
4. View the SNTP configuration again to verify the configuration.

Enabling TimeP mode

Enabling the TimeP mode configures it for either broadcast or unicast. Run TimeP as the switch's time synchronization protocol and select TimeP as the time synchronization method by using the CLI `timesync` command (or the menu interface **Time Sync Method** parameter).

Procedure

1. View the current time synchronization using `show sntp`.
2. Set TimeP as the synchronization mode using `timesync sntp`.
3. Enable TimeP for DHCP mode using `sntp broadcast`.
4. View the TimeP configuration using `show sntp`.

Figure 5: *Enabling TimeP operation in DHCP mode*

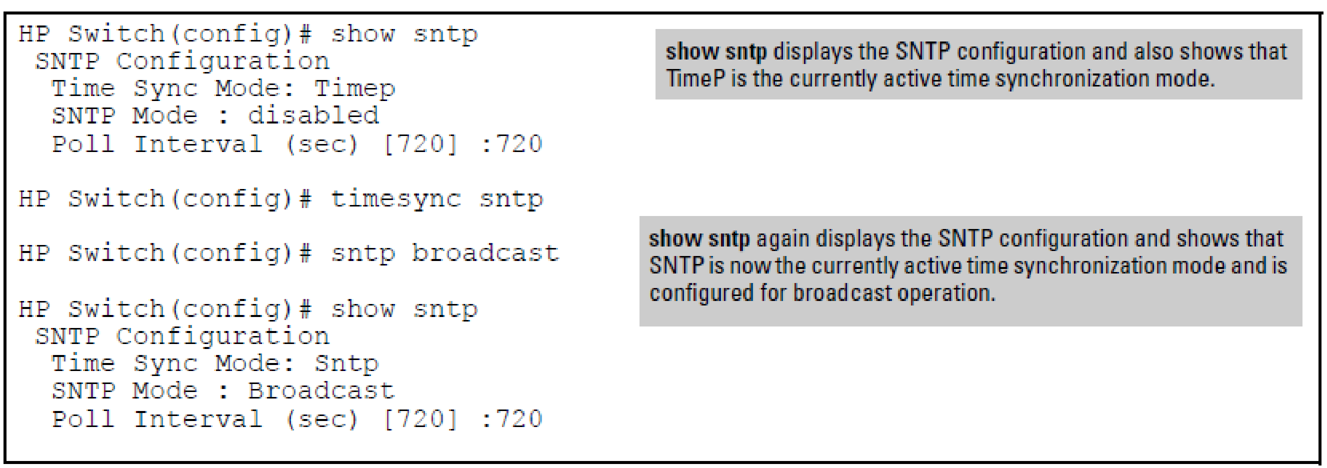

timesync timep

Syntax

`timesync timep`

Description

Selects `TimeP` as the time synchronization method.

TimeP in DHCP mode

Because the switch provides a TimeP polling interval (default: 720 minutes), you need the **timesync timep** on page 57 and `ip timep` commands only, for a minimal TimeP DHCP configuration.

ip timep dhcp

Syntax

`ip timep dhcp`

Description

Configuring TimeP for DHCP operation

```
switch# show timep

Timep Configuration

Time Sync Mode: Sntp
TimeP Mode : Disabled
Poll Interval (min) [720] : 720

switch# timesync timep

switch# ip timep dhcp

switch# show timep

Timep Configuration
Time Sync Mode: Timep
TimeP Mode : DHCP Poll Interval (min): 720
```

Enabling TimeP for DHCP

Suppose time synchronization is configured for SNTP. Following this example to enable TimeP for DHCP.

Procedure

1. View the current time synchronization.
2. `show timep` displays the TimeP configuration and also shows that SNTP is the currently active time synchronization mode.
3. Select TimeP as the time synchronization mode.
4. Enable TimeP for DHCP mode.
5. View the TimeP configuration.
6. `show timep` again displays the TimeP configuration and shows that TimeP is now the currently active time synchronization mode.

```
switch(config)# show timep

 Timep Configuration

  Time Sync Mode: Sntp
```

```
  TimeP Mode : Disabled
  Poll Interval (min) [720] : 720

switch(config)# timesync timep

switch(config)# ip timep dhcp

switch(config)# show timep

 Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : DHCP Poll Interval (min): 720
```

Viewing, enabling, and modifying the TimeP protocol(Menu)

Procedure

1. From the Main Menu, select:

   2. Switch Configuration

   1. System Information

   Figure 6: *System Information screen (default values)*

   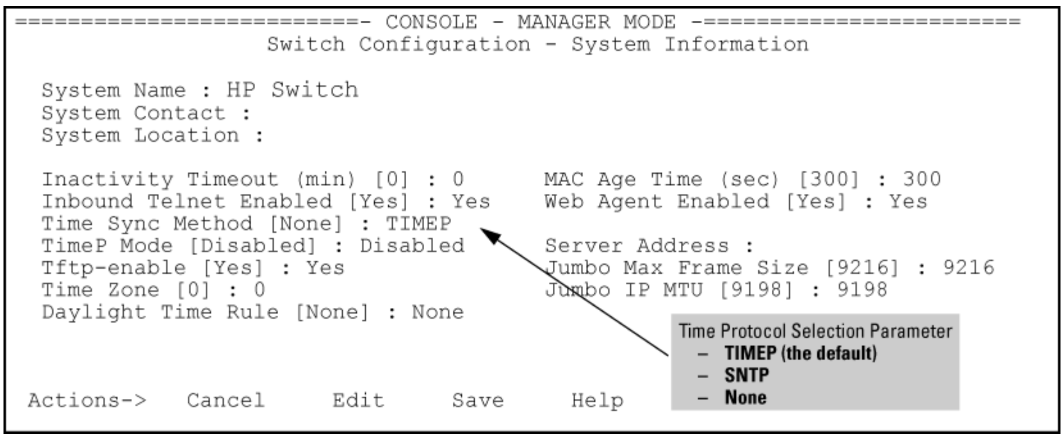

2. Press **[E]** (for **Edit**.)

   The cursor moves to the **System Name** field.
3. Use the right arrow key to move the cursor to the **Time Sync Method** field.
4. If **TIMEP** is not already selected, use the **Space** bar to select **TIMEP**, then press the right arrow key once to display and move to the **TIMEP Mode** field.
5. Do one of the following:
   a. Use the **Space** bar to select the **DHCP** mode.
      - Press the right arrow key to move the cursor to the **Poll Interval** field.
      - Go to step 6.
   b. **Enabling TIMEP or DHCP**

      ```
      Time Sync Method [None] :   TIMEP
      TimeP Mode [Disabled] :     DHCP
      Poll Interval (min) [720] : 720
      Time Zone [0] :             0
      Daylight Time Rule [None] : None
      ```

      Use the **Space** bar to select the **Manual** mode.

c. Press the right arrow key to move the cursor to the **Server Address** field.
d. Enter the IP address of the TimeP server you want the switch to use for time synchronization.

NOTE

This step replaces any previously configured TimeP server IP address.

e. Press the right arrow key to move the cursor to the **Poll Interval** field, then go to step 6.

6. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.
7. Select **[Enter]** to return to the **Actions** line, then select **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

TimeP operation in manual mode

As with DHCP mode, configuring `timep` for Manual Mode enables `timep`; but for manual operation, you must also specify the IP address of the `timep` server. (The switch allows only one `timep` server.)

timesync timep

Syntax

```
timesync timep
```

Description

Activates TimeP in manual mode with a specified TimeP server. By default, SNTP traffic goes through the data ports.

ip timep

Syntax

```
ip timep manual<IP-ADDR>
```

Description

Activate TimeP in manual mode with a specified TimeP server. (By default, SNTP traffic goes through the data ports.)

Parameters and options

manual

<IP-ADDR>

Enabling TimeP in manual mode

Select TimeP and configure it for manual operation using a TimeP server address of 10.28.227.141, and the default poll interval (720 minutes, assuming the TimeP poll interval is already set to the default).

Procedure

1. Select TimeP:

   ```
   switch(config)# timesync timep
   ```
2. Activate TimeP in manual mode:

   ```
   switch(config)# ip timep manual 10.28.227.141
   ```
3. Review the TimeP status:

   ```
   switch(config)# show timep
   ```

   show timep output

```
switch(config)# show timep
 Timep Configuration

  Time Sync Mode: Timep
  TimeP Mode :  Manual                Server Address : 10.28.227.141
  Poll Interval (min) : 720
```

Viewing, enabling, and modifying the TimeP protocol (Menu)

Procedure

1. From the Main Menu, select **2. Switch Configuration**, and then select **1. System Information**.

   Figure 7: *System Information screen (default values)*

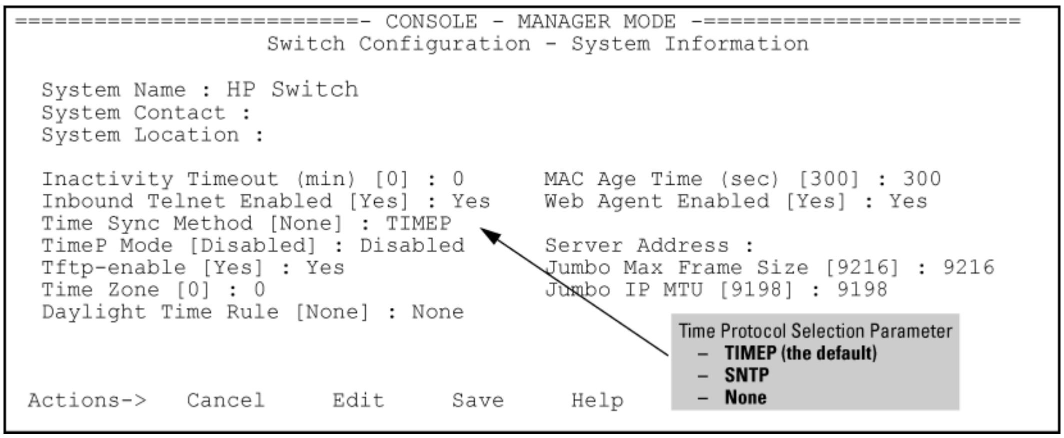

2. Press **[E]** (for **Edit**.)

   The cursor moves to the **System Name** field.
3. Use the right arrow key to move the cursor to the **Time Sync Method** field.
4. If **TIMEP** is not already selected, use the **Space** bar to select **TIMEP**, then press the right arrow key once to display and move to the **TIMEP Mode** field.
5. Do one of the following:

   a. Use the **Space** bar to select the **DHCP** mode.

   - Press the right arrow key to move the cursor to the **Poll Interval** field.
   - Go to step 6.

```
Time Sync Method [None] :   TIMEP
TimeP Mode [Disabled] :     DHCP
Poll Interval (min) [720] : 720
Time Zone [0] :             0
Daylight Time Rule [None] : None
```

   b. Use the **Space**bar to select the **Manual** mode.

   - Press the right arrow key to move the cursor to the **Server Address** field.
   - Enter the IP address of the TimeP server you want the switch to use for time synchronization.

     NOTE: This step replaces any previously configured TimeP server IP address.

   - Press the right arrow key to move the cursor to the **Poll Interval** field, then go to step 6.
6. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.
7. Select **[Enter]** to return to the **Actions** line, then select **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

Current TimeP configuration

Using different `show` commands, you can display either the full TimeP configuration or a combined listing of all TimeP, SNTP, and VLAN IP addresses configured on the switch.

show timep

Syntax

`show timep`

Description

Lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol. (If the TimeP Mode is set to `Disabled` or `DHCP`, the Server field does not appear.)

TimeP configuration when TimeP is the selected Time synchronization method

If you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, `show timep` lists the following:

```
switch(config)# show timep

 Timep Configuration

  Time Sync Mode: Timep
  TimeP Mode [Disabled] : DHCP    Server Address : 10.10.28.103
  Poll Interval (min) [720] : 720
```

TimeP configuration when TimeP is not the selected time synchronization method

If SNTP is the selected time synchronization method, `show timep` still lists the TimeP configuration even though it is not currently in use. Even though, in this example, SNTP is the current time synchronization method, the switch maintains the TimeP configuration.

```
switch(config)# show timep

 Timep Configuration

  Time Sync Mode: Sntp
  TimeP Mode [Disabled] : Manual   Server Address : 10.10.28.100
  Poll Interval (min) [720] : 720
```

show management

Syntax

`show management`

Description

Examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch plus the IP addresses and default gateway for all VLANs configured on the switch.

Show IP addressing for all configured time servers and VLAN

```
switch(config)# show management

 Status and Counters - Management Address Information
```

```
Time Server Address : 10.10.28.100

Priority SNTP Server Address                              Protocol Version
-------- ---------------------------------------------- ----------------
1        10.10..28.101                                    3
2        10.255.5.24                                      3
3        fe80::123%vlan10                                 3


Default Gateway : 10.0.9.80

VLAN Name    MAC Address         | IP Address
------------ ------------------- + -------------------
DEFAULT_VLAN 001279-88a100       | 10.30.248.184
VLAN10 001279-88a100             | 10.0.10.17
```

Change from one TimeP server to another

To change from one TimeP server to a different server, use the `no ip timep` command to disable TimeP mode then reconfigure TimeP in manual mode with the new server IP address.

TimeP poll interval

ip timep

Syntax

`ip timep [dhcp|manual] interval [1-9999]`

Description

Specifies how long the switch waits between time polling intervals. The default is 720 minutes and the range is 1 to 9999 minutes. (This parameter is separate from the `poll interval` parameter used for SNTP operation.)

Disable time synchronization protocols

Disabling TimeP in manual mode

no ip timep

Syntax

`[no] ip timep`

Description

Disables TimeP.

Parameters and options

`no`

To change from one TimeP server to another, you must use the `no ip timep` command to disable TimeP mode, the reconfigure TimeP in manual mode with the new server IP address.

Disabling time synchronization

Either of these methods can be used to disable time synchronization without changing the Timep or SNTP configuration.

no timesync

Syntax

`[no] timesync`

Description

Disables time synchronization by changing the `Time Sync Mode` configuration to `Disabled`. This halts time synchronization without changing your TimeP configuration.The recommended method for disabling time synchronization is to use the `timesync` command.

TimeP with time synchronization disabled

Suppose TimeP is running as the switch's time synchronization protocol, with `DHCP` as the TimeP mode, and the factory-default polling interval. You would halt time synchronization with this command:

```
switch(config)# no timesync
```

If you then viewed the TimeP configuration, you would see the following:

```
switch(config)# show timep

 Timep Configuration
 Time Sync Mode: Disabled
 TimeP Mode : DHCP Poll Interval (min): 720
```

Disabling timsync using the GUI

Procedure

1. Set the `Time Synch Method` parameter to `None`.
2. Press **[Enter]**, then **[S]** (for **Save**.)

Disabling the TimeP mode

no ip timep

Syntax

`no ip timep`

Description

Disables TimeP by changing the TimeP mode configuration to `Disabled` and prevents the switch from using it as the time synchronization protocol, even if it is the selected `Time Sync Method` option.

Disabling time synchronization by disabling the TimeP mode parameter

If the switch is running TimeP in DHCP mode, `no ip timep` changes the TimeP configuration as shown below and disables time synchronization. Even though the TimeSync mode is set to TimeP, time synchronization is disabled because `no ip timep` has disabled the TimeP mode parameter.

```
switch(config)# no ip timep

switch(config)# show timep

 Timep Configuration
```

```
Time Sync Mode: Timep
TimeP Mode : Disabled
```

Disabling time synchronization without changing the SNTP configuration

timesync

Syntax

`[no] timesync`

Description

Recommended method for disabling time synchronization. Halts time synchronization without changing your SNTP configuration.

Halt time synchronization

Suppose SNTP is running as the switch's time synchronization protocol, with `broadcast` as the SNTP mode and the factory-default polling interval. You would halt time synchronization with this command:

```
switch(config)# no timesync
```

If you then viewed the SNTP configuration, you would see the following:

SNTP with time synchronization disabled

```
switch(config)# show sntp
SNTP Configuration
SNTP Authentication : Disabled
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720
```

Disabling SNTP mode

Procedure

1. To view the current time synchronization, enter `show sntp`.
2. Use the sntp command to disable sntp mode:

   `no sntp`
3. View the SNTP configuration again to verify the configuration.

Disabling SNTP Mode

If you want to prevent the SNTP from being used even if it is selected by `timesync` (or the Menu interface's **Time Sync Method** parameter), configure the SNTP mode as disabled.

no sntp

Syntax

`[no] sntp`

Description

Disables SNTP by changing the SNTP mode configuration to `Disabled`.

Disabling time synchronization by disabling the SNTP mode

If the switch is running SNTP in unicast mode with an SNTP server at 10.28.227.141 and a server version of 3 (the default), `no sntp` changes the SNTP configuration as shown below and disables time synchronization on the switch.

```
switch(config)# no sntp
switch(config)# show sntp

SNTP Configuration
SNTP Authentication : Disabled
Time Sync Mode: SNTP
SNTP Mode : disabled
Poll Interval (sec) [720] : 719
Source IP Selection: Outgoing Interface

Priority SNTP Server Address                     Version Key-id
-------- --------------------------------------- ------- ----------
1        2001:db8::215:60ff:fe79:8980            7       0
2        10.255.5.24                             3       0
```

Deleting an SNTP server

Syntax

`[no] sntp server priority <PRIORITY> <IP-ADDRESS>`

Description

Deletes the specified SNTP server.

Deleting an SNTP server when only one server is configured disables SNTP unicast operation.

Disabling SNTP by deleting a server

sntp server priority

Syntax

`[no] sntp server priority <PRIORITY> <IP-ADDRESS> version key-id <KEY_ID>`

Description

Disabling SNTP by deleting the specified SNTP server. Uses the `no` version of the command to disable SNTP.

Disabling time synchronization in DHCP mode by disabling the TimeP mode parameter

The `[no] ip timep` command changes the TimeP configuration for both DHCP and manual modes, as shown below, and disables time synchronization. Even though the TimeSync mode is set to TimeP, time synchronization is disabled because the `no ip timep` command has disabled the TimeP mode parameter.

`ip timep`

Syntax

`[no] ip timep`

Description

To change from one TimeP server to another, you must use the `no ip timep` command to disable TimeP mode, then reconfigure TimeP in manual mode with the new server IP address.

Disabling TimeP in manual mode

```
Timep Configuration

Time Sync Mode: Sntp
TimeP Mode : Disabled
Poll Interval (min) [720] : 720

switch# timesync timep

switch# ip timep manual

switch# show timep

Timep Configuration
Time Sync Mode: Timep
TimeP Mode : DHCP Poll Interval (min): 720
```

Disabling TimeP in DHCP mode

```
switch# no ip timep

switch# show timep

Timep Configuration
Time Sync Mode: Timep
TimeP Mode : Disabled
```

Other time protocol commands

Features that apply to both SNTP and TimeP protocols.

Show management command

show management

Syntax

`show management`

Description

This command shows the switch addresses available for management, and the time server if the switch uses one. It can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.

Display showing IP addressing for all configured time servers and VLANs

```
switch(config)# show management
Status and Counters - Management Address Information

Time Server Address : 10.10.28.100

  Priority SNTP Server Address Protocol Version
  -------- ------------------- ----------------
  1        10.10.28.101             3
  2        10.255.5.24              3

  Default Gateway     : 10.0.9.80

  VLAN Name    MAC Address     | IP Address
  ------------ --------------- + ---------------
  DEFAULT_VLAN 001871-c42f00   | 10.30.248.184
  VLAN10       001871-c42f00   | 10.0.10.17

 Internet (IPv6) Service

  Interface Name   : DEFAULT_VLAN
  IPv6 Status      : Disabled

  Interface Name   : VLAN10
  IPv6 Status      : Disabled
```

Show SNTP command

In the factory-default configuration (where TimeP is the selected time synchronization method), `show sntp` still lists the SNTP configuration, even though it is not currently in use.

show sntp

Syntax

`show sntp [authentication|statistics]`

Description

Shows configured time protocol and servers. Lists both the time synchronization method (`TimeP`, `SNTP`, or `None`) and the SNTP configuration, even if SNTP is not the selected time protocol. Configure the switch with SNTP as the time synchronization method, and then enable SNTP in broadcast mode with the default poll interval, `show sntp`.

Parameters and options

`Authentication`

Displays all the configured SNTP authentication information.

`Statistics`

Displays SNTP protocol statistics.

Figure 8: *SNTP configuration when SNTP is not the selected time synchronization method*

```
HP Switch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] :720

HP Switch(config)# timesync sntp

HP Switch(config)# sntp broadcast

HP Switch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] :720
```

show sntp displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.

show sntp again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.

`show sntp authentication` command with authentication disabled

To display all the SNTP authentication keys that have been configured on the switch, enter the `show sntp authentication` command.

```
switch(config) # show sntp authentication
SNTP Authentication Information
SNTP Authentication: Enabled

Key-ID  Auth Mode  Trusted
------- ----------- -------
55      MD5      YES
10      MD5      NO
```

To display the statistical information for each SNTP server, enter the `sntp statistics` command. The number of SNTP packets that have failed authentication is displayed for each SNTP server address.

```
switch(config) # show sntp statistics
SNTP statistics
Received Packets: 0
Sent Packets: 3
Dropped Packets: 0
SNTP Server Address    Auth Failed Pkts
-------------------    ----------------
10.10.10.1             0
fe80::200:24ff:fec8:4ca8   0
```

Show TimeP command

Using different `show` commands, you can display either the full TimeP configuration or a combined listing of all TimeP, SNTP, and VLAN IP addresses configured on the switch.

show

Syntax

`show timep | management`

Description

Displays the timep and management information for the switch.

Parameters and options

timep

Lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol. (If the TimeP Mode is set to `Disabled` or `DHCP`, the Server field does not appear.)

management

Helps you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch plus the IP addresses and default gateway for all VLANs configured on the switch.

TimeP configuration when TimeP is the selected Time synchronization method

If you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, `show timep` lists the following:

```
switch# show timep

 Timep Configuration

 Time Sync Mode: Timep
 TimeP Mode [Disabled] : DHCP    Server Address : 10.10.28.103
 Poll Interval (min) [720] : 720
```

TimeP configuration when TimeP is not the selected time synchronization method

If SNTP is the selected time synchronization method, `show timep` still lists the TimeP configuration even though it is not currently in use. Even though, in this example, SNTP is the current time synchronization method, the switch maintains the TimeP configuration (see data in bold below):

```
switch# show timep

 Timep Configuration

  Time Sync Mode: Sntp
  TimeP Mode [Disabled] : Manual   Server Address : 10.10.28.100
  Poll Interval (min) [720] : 720
```

Display showing IP addressing for all configured time servers and VLANs

```
switch# show management

 Status and Counters - Management Address Information

  Time Server Address : 10.10.28.100

  Priority SNTP Server Address                            Protocol Version
```

```
-------- ---------------------------------------------- ----------------
1        10.10..28.101                                  3
2        10.255.5.24                                    3
3        fe80::123%vlan10                               3


Default Gateway : 10.0.9.80

VLAN Name    MAC Address         | IP Address
------------ ------------------- + -------------------
DEFAULT_VLAN 001279-88a100       | 10.30.248.184
VLAN10 001279-88a100             | 10.0.10.17
```

Viewing current resource usage

showqous

Syntax

`showqos|access-list|policyresources`

Description

Displays the resource usage of the policy enforcement engine on the switch by software feature. For each type of resource, the amount still available and the amount used by each software feature is shown.

Parameters and options

`show resources`

This output allows you to view current resource usage and, if necessary, prioritize and reconfigure software features to free resources reserved for less important features.

`qos|access-list|openflow|policy`

Display the same command output and provide different ways to access task-specific information. See the *OpenFlow administrators guide.*

Unavailable resources

The resource usage on a switch configured for ACLs, QoS, RADIUS-based authentication, and other features:

- The "Rules Used" columns show that ACLs, VT, mirroring, and other features (for example, Management VLAN) have been configured globally or per-VLAN, because identical resource consumption is displayed for each port range in the switch. If ACLs were configured per-port, the number of rules used in each port range would be different.
- The switch is also configured for VT and is either blocking or throttling routed traffic with a high rate-of-connection requests.

- Varying ICMP rate-limiting configurations on ports 1 to 24, on ports 25 to 48, and on slot A, have resulted in different meter usage and different rule usage listed under QoS. Global QoS settings would otherwise result in identical resource consumption on each port range in the switch.
- There is authenticated client usage of IDM resources on ports 25 to 48.

Figure 9: *Viewing current QoS resource usage on a series 3500yl switch*

```
HP Switch# show qos resources

 Resource usage in Policy Enforcement Engine

       |    Rules     |  Rules Used
 Ports |  Available   |  ACL  |  QoS  |  IDM  |  VT   | Mirror | Other |
 ------+--------------+-------+-------+-------+-------+--------+-------|
 1-24  |         3014 |    15 |    11 |     0 |     1 |      0 |     3 |
 25-48 |         3005 |    15 |    10 |    10 |     1 |      0 |     3 |
 A     |         3017 |    15 |     8 |     0 |     1 |      0 |     3 |

       |    Meters    |  Meters Used
 Ports |  Available   |  ACL  |  QoS  |  IDM  |  VT   | Mirror | Other |
 ------+--------------+-------+-------+-------+-------+--------+-------|
 1-24  |          250 |       |     5 |     0 |       |        |     0 |
 25-48 |          251 |       |     4 |     0 |       |        |     0 |
 A     |          253 |       |     2 |     0 |       |        |     0 |

       | Application  |
       | Port Ranges  |  Application Port Ranges Used
 Ports |  Available   |  ACL  |  QoS  |  IDM  |  VT   | Mirror | Other |
 ------+--------------+-------+-------+-------+-------+--------+-------|
 1-24  |         3014 |     2 |     0 |     0 |       |      0 |     0 |
 25-48 |         3005 |     2 |     0 |     0 |       |      0 |     0 |
 A     |         3017 |     2 |     0 |     0 |       |      0 |     0 |

 0 of 8 Policy Engine management resources used.
 Key:
 ACL = Access Control Lists
 QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
 IDM = Identity Driven Management
 VT  = Virus Throttling blocks
 Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
 Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

 Resource usage includes resources actually in use, or reserved for future
 use by the listed feature.  Internal dedicated-purpose resources, such as
 port bandwidth limits or VLAN QoS priority, are not included.
```

Viewing information on resource usage

Cause

The switch allows you to view information about the current usage and availability of resources in the Policy Enforcement engine, including the following software features:

- Access control lists (ACL)
- Quality-of-service (QoS), including device and application port priority, ICMP rate-limiting, and QoS policies
- Dynamic assignment of per-port or per-user ACLs and QoS through RADIUS authentication designated as “IDM”, with or without the optional identity-driven management (IDM) application
- Virus throttling (VT) using connection-rate filtering
- Mirroring policies, including switch configuration as an endpoint for remote intelligent mirroring
- Other features, including:

    - Management VLAN
    - DHCP snooping
    - Dynamic ARP protection
    - Jumbo IP-MTU

When insufficient resources are available

Cause

The switch has ample resources for configuring features and supporting:

- RADIUS-authenticated clients (with or without the optional IDMapplication)
- VT and blocking on individual clients.

NOTE

Virus throttling does not operate on IPv6 traffic.

If the resources supporting these features become fully subscribed:

- The current feature configuration, RADIUS-authenticated client sessions, and VT instances continue to operate normally.
- The switch generates an event log notice to say that current resources are fully subscribed.
- Currently engaged resources must be released before any of the following actions are supported:
    - Modifying currently configured ACLs, IDM, VT, and other software features, such as Management VLAN, DHCP snooping, and dynamic ARP protection.

      You can modify currently configured classifier-base QoS and mirroring policies if a policy has not been applied to an interface. However, sufficient resources must be available when you apply a configured policy to an interface.
    - Acceptance of new RADIUS-based client authentication requests (displayed as a new resource entry for IDM.)

      Failure to authenticate a client that presents valid credentials may indicate that insufficient resources are available for the features configured for the client in the RADIUS server. To troubleshoot, check the event log.
    - Throttling or blocking of newly detected clients with high rate-of-connection requests (as defined by the current VT configuration.)

      The switch continues to generate Event Log notifications (and SNMP trap notification, if configured) for new instances of high-connection-rate behavior detected by the VT feature.

Policy enforcement engine

Cause

The policy enforcement engine is the hardware element in the switch that manages QoS, mirroring, and ACL policies, as well as other software features, using the rules that you configure. Resource usage in the policy enforcement engine is based on how these features are configured on the switch:

- Resource usage by dynamic port ACLs and VT is determined as follows:

    - Dynamic port ACLs configured by a RADIUS server (with or without the optional IDM application) for an authenticated client determine the current resource consumption for this feature on a specified slot. When a client session ends, the resources in use for that client become available for other uses.
    - A VT configuration (connection-rate filtering) on the switch does not affect switch resources unless traffic behavior has triggered either a throttling or blocking action on the traffic from one or more clients. When the throttling action ceases or a blocked client is unblocked, the resources used for that action are released.
- When the following features are configured globally or per-VLAN, resource usage is applied across all port groups or all slots with installed modules:
    - ACLs
    - QoS configurations that use the following commands:
        - QoS device priority (IP address) through the CLI using the `qos device-priority` command
        - QoS application port through the CLI using `qos tcp-port` or `qos udp-port`
        - VLAN QoS policies through the CLI using `service-policy`
    - Management VLAN configuration
    - DHCP snooping
    - Dynamic ARP protection
    - Remote mirroring endpoint configuration
    - Mirror policies per VLAN through the CLI using `monitor service`
    - Jumbo IP-MTU
- When the following features are configured per-port, resource usage is applied only to the slot or port group on which the feature is configured:
    - ACLs or QoS applied per-port or per-user through RADIUS authentication
    - ACLs applied per-port through the CLI using the `ip access-group` or `ipv6 traffic-filter` commands
    - QoS policies applied per port through the CLI using the `service-policy`command
    - Mirror policies applied per-port through the CLI using the `monitor all service` and `service-policy`commands
    - ICMP rate-limiting through the CLI using the `rate-limit icmp`command
    - VT applied to any port (when a high-connection-rate client is being throttled or blocked)

Usage notes for show resources output

Cause

- A 1:1 mapping of internal rules to configured policies in the switch does not necessarily exist. As a result, displaying current resource usage is the most reliable method for keeping track of available resources. Also, because some internal resources are used by multiple features, deleting a feature configuration may not increase the amount of available resources.
- Resource usage includes resources actually in use or reserved for future use by the listed features.
- "Internal dedicated-purpose resources" include the following features:
    - Per-port ingress and egress rate limiting through the CLI using `rate-limit in/out`
    - Per-port ingress and egress broadcast rate limiting through the CLI using `rate-limit bcast/mcast`
    - Per-port or per-VLAN priority or DSCP through the CLI using `qos priority` or `qos dscp`
    - Per protocol priority through the CLI using `qos protocol`
- For chassis products (for example, the 5400zl or 8212zl switches), 'slots' are listed instead of 'ports,' with resources shown for all installed modules on the chassis.
- The "Available" columns display the resources available for additional feature use.
- The "IDM" column shows the resources used for RADIUS-based authentication with or without the IDM option.
- "Meters" are used when applying either ICMP rate-limiting or a QoS policy with a rate-limit class action.

Services

The `services` command requires a slot-name parameter followed by an option. Options permitted in this command depend on the context (operator, manager, or configure).

Show services

Syntax

`show services <SLOT-ID>[details | device]`

Description

Show services modules information.

Parameters

Slot-id

Show services modules information

Options

***<SLOT-ID>* details**

Display application information for the specified slot.

***<SLOT-ID>* device**

Display the current configuration of the devices.

Show services

```
switch# show services

        Installed Services
Slot    Index Description                          Name
------  --------------------------------      ------------------
H,L       1. Services zl Module                  services-module
L         2. HP ProCurve MSM765 zl Int-Ctlr        msm765-applicati
H         3. Threat Management Services zl Module  tms-module
```

No parameters

This `no parameters` command lists only installed modules which have applications running that provide a pass-through CLI feature.

show services

Syntax

`show services`

Description

Show services of only installed modules.

Show services

```
switch# show services

        Installed Services

Slot      Index Description                           Name
H,L         1. Services zl Module                       services-module
L          2. HP ProCurve MSM765 zl Int-Ctlr           msm765-applicati
H          3.Threat Management Services zl Module      tms-module
```

Show services locator

Syntax

`show services <SLOT-ID>[details | device]`

Description

Show services information.

Parameters

details

Display application information for the specified slot.

device

Display the current configuration of the devices.

Options

Slot-id

Display summary table for the specified slot.

Show services f

```
switch# show services f
Status and Counters - Services Module F Status
 HP Services zl Module J9840A
Versions             :
 Current status      : running
For more information, use the show commands in services context
```

Show servers f details

```
switch# show services f details
 Status and Counters - Services Module F Status
 HP Services zl Module J9840A
 Versions             :
 Current status       : running

Description                                  Version             Status
-------------------------------------------- ------------------- ---------
```

```
Services zl Module                                                    hardware
HP MSM775 zl Premium Controller     J9840A                            installed

 For more information, use the show commands in services context
```

Show services f status

```
Status and Counters - Services Module F Status
HP Services zl Module J9840A
Versions            :
Current status      : running
Description                               Version              Status
----------------------------------------- -------------------- ----------
Services zl Module                                             hardware
HP Adv Services v2 zl Module w/ HDD         J9857A             installed

For more information, use the show commands in services context
```

Show services device

Adding the keyword “device” displays information about whether certain external devices are enabled or disabled. This command is equivalent to the “services <slot> device” command with no additional parameters.

show services device

Syntax

`show services` *`slot-id`* `device`

Description

- USB port (x86–side) May be one of:
  - “disabled” (normal state)
  - “enabled” – enabled once the x86 boots into the OS, but disabled before OS boot to prevent inadvertently booting to an inserted USB key.
  - “boot” – enabled all the time, both for and after x86 OS boot.
- ShutdownFront-panel shutdown/reset button:
  - “enabled” – default state
  - “disabled” – for increased physical security
- PXE (PXE-boot)Not displayed for all modules.

Show services device

```
switch# show services d device
Services Module Device Configuration
Device          |                 State
----------------|--------------------
USB             |             disabled
Shutdown        |              enabled
PXE             |              enabled
```

Requesting a reboot

Syntax

`services <SLOT>boot[product|PXE|service|USB]`

Description

This command requests a reboot (graceful shutdown and restart) of the x86.

Parameters

product

Boot to the Product OS.

PXE

Boot to the PXE or Product OS (if supported).

service

Boot to the Service OS.

USB

Boot to the USB or Product OS (if supported).

If no parameters are given, the switch attempts to boot to the same OS (product, service, or USB) that was enabled before the command was given. If the `services <slot> boot product|usb` command is given on a non-permitted module, one of the following error messages is returned:

Services b boot

```
switch# services b boot product
Command not supported for the Services module in slot B.

switch# services b boot pxe
Command not supported for the Services module in slot B.

switch# services b boot usb
Command not supported for the Services module in slot B.
```

Services in Operator/Manager/Configure context

This top-level command requires a slot-name parameter followed by a subcommand. Permitted subcommands depend on one of the three contexts: operator, manager, or configure.

Services (operator)

Syntax

`services <SLOT-ID>[<INDEX>| locator | name <NAME>]`

Description

Displays applications installed and running for the services module in the Operator context.

Parameters

Integer

Index of the services CLI to access.

Locator

Control services module locator LED.

Name

Name of the services CLI to access.

Options

<SLOT-ID>

Device slot identifier for the services module.

<SLOT-ID> <INDEX>

Configure parameters for the installed application.

<SLOT-ID> locator

Controls services module locator LED.

<SLOT-ID> name <NAME>

Configure parameters for the installed application.

Services (manager)

Syntax

```
services <SLOT-ID>[<INDEX> | boot | locator | name <NAME> | reload | serial |
shutdown]
```

Description

Display applications installed and running for the services module or change the module's state (reload or shutdown).

Parameters

Boot

Reboot the services module.

Integer

Index of the services CLI to access.

Locator

Control services module locator LED.

Name

Name of the services CLI to access.

Reload

Reset the services module.

Serial

Connect to application via serial port.

Shutdown

Shutdown (halt) the services module.

Options

slot-id

Device slot identifier for the services module.

<slot-id> <index>

Configure parameters for the installed application.

<slot-id> boot

Reboot the services module.

<slot-id> locator

Controls services module locator LED.

<slot-id> name <name>

Configure parameters for the installed application.

<slot-id> reload

Reset the services module.

<slot-id> serial

Connect to services module via serial port.

<slot-id> shutdown

Shutdown (halt) the services module.

Services (configure)

Syntax

```
[no] services [<SLOT-ID> <INDEX> boot | locator | name <NAME> | reload | serial |
shutdown] services <slot-id> device [shutdown | usb]
```

Description

Configure parameters for the services module or change the module's state (reload or shutdown).

Parameters and options

`slot-id`

Device slot identifier for the services module.

`<SLOT-ID> <INDEX>`

Configure parameters for the installed application.

`<SLOT-ID> boot`

Reboot the services module.

`<SLOT-ID> locator`

Controls services module locator LED.

`<SLOT-ID> name<NAME>`

Configure parameters for the installed application.

`<SLOT-ID> reload`

Reset the services module.

***<SLOT-ID>* serial**

Connect to services module via serial port.

***<SLOT-ID>* shutdown**

Shutdown (halt) the services module.

Enable or disable devices.

Enable or disable devices. This command must be run from the configure context.

no services

Syntax

`no services <SLOT> device [PXE|shutdown|USB|CF]`

Parameters

PXE

Enable or disable booting from PXE (if supported).

shutdown

Enable or disable the shutdown or reset button.

USB

Enable or Disable the USB after boot.

CF

Enable or disable the Compact Flash or SD1 card.

Accessing CLI-passthrough

Accessing the CLI-passthrough feature on modules that support the feature. Feature can be reported by the `show services` command given with no additional parameters.

services

Syntax

`services <SLOT>[<INDEX>|<NAME>]`

Description

Parameters

ASCII-STR

Enter an ASCII string.

Show services

```
switch# show services

Installed Services

Slot    Index   Description                               Name
H,L      1.    Services zl Module                       services-module
```

```
L       2.    HP ProCurve MSM765 zl Int-Ctlr            msm765-applicati
H       3.    Threat Management Services zl Module      tms-module
```

Show services set locator module

This command sets the Module Locator LED to either solid-on, off or slow-blink for a specified duration of time or to turn it off before the previously-specified time has passed. Options are permitted in this command for the Operator.

command name

Syntax

`show services <SLOT>[blink <1-1440>|off|on]`

Parameters

blink

Blink the locator LED. Default 30 mins. Range <1-1440>.

off

Turn the locate led off.

on

Turn the locate led on.

show services d

```
switch# show services d locator blink
```

Reloading services module

command name

Syntax

`services <SLOT> reload`

Description

Reloads the services module and is similar to the command `services<slot> boot` with no additional parameters given.

Connection to the application via a serial port

You are entering a mode on this product that is Hewlett Packard Enterprise Confidential and Proprietary. This mode, the commands and functionality specific to this mode, and all output from this mode are Hewlett Packard Enterprise Confidential and Proprietary. You may use this mode only by specific permission of, and under the direction of, an Hewlett Packard Enterprise support engineer or Hewlett Packard Enterprise technical engineer. Unauthorized or improper use of this mode will be considered by Hewlett Packard Enterprise to be unauthorized modification of the product, and any resulting defects or issues are not eligible for coverage under the Hewlett Packard Enterprise product warranty or any Hewlett Packard Enterprise support or service. UNAUTHORIZED OR IMPROPER USE OF THIS MODE CAN MAKE THE PRODUCT COMPLETELY INOPERABLE.

SvcOS login: <CTRL-Z>

command name

Syntax

```
services <SLOT>serial
```

Description

Starts a serial-passthrough session to the x86.

Shutdown the services module.

command name

Syntax

```
services <SLOT>shutdown
```

Description

Similar to services <*slot*>boot with no additional parameters given. This command is similar in that it attempts a graceful shutdown of the x86 except that this command does not restart the x86. If the graceful-shutdown attempt fails, no follow-up attempt is made to do a hard shutdown.

Transceiver status

The following information is displayed for each installed transceiver:

- Port number on which transceiver is installed.
- Type of transceiver.
- Product number — Includes revision letter, such as A, B, or C. If no revision letter follows a product number, this means that no revision is available for the transceiver.
- Part number — Allows you to determine the manufacturer for a specified transceiver and revision number.

Operating notes

- For a non- switches installed transceiver (see line 23 **Figure 10: Example of show tech transceivers command** on page 85), no transceiver type, product number, or part information is displayed. In the Serial Number field, non-operational is displayed instead of a serial number.
- The following error messages may be displayed for a non-operational transceiver:

- ```
  Unsupported Transceiver. (SelfTest Err#060)
  Check: http://www.hpe.com/rnd/device_help/2_inform for more info.
  ```
- ```
  This switch only supports revision B and above transceivers.
  Check: http://www.hpe.com/rnd/device_help/2_inform for more info.
  ```
- `Self test failure.`
- `Transceiver type not supported in this port.`
- `Transceiver type not supported in this software version.`
- ```
  Not an HP Switch Transceiver.
  Go to: http://www.hpe.com/rnd/device_help/2_inform for more info.
  ```

show interfaces transceivers

Syntax

`show interfaces transceivers`

Description

Figure 10: Example of show tech transceivers command on page 85 shows sample output from the `show tech transceivers` command. The Part # column enables you to determine the manufacturer for a specified transceiver and revision number.

- Remotely identify transceiver type and revision number without having to physically remove an installed transceiver from its slot.
- Display real-timestatus information about all installed transceivers, including non-operational transceivers.

Figure 10: *Example of `show tech transceivers` command*

```
HP Switch# show tech transceivers

Transceiver Technical Information:
 Port # |    Type     | Prod # | Serial #          | Part #
 -------+-------------+--------+-------------------+----------
 21     | 1000SX      | J4858B | CN605MP23K        |
 22     | 1000LX      | J4859C | H117E7X           | 2157-2345
 23     | ??          | ??     | non operational   |
 25     | 10GbE-CX4   | J8440A | US509RU079        |
 26     | 10GbE-CX4   | J8440A | US540RU002        |
 27     | 10GbE-LR    | J8437B | PPA02-2904:0017   | 2157-2345
 28     | 10GbE-SR    | J8436B | 01591602          | 2158-1000
 29     | 10GbE-ER    | J8438A | PPA03-2905:0001   |

The following transceivers may not function correctly:
 Port #          Message
 --------        -------------------------------------
 Port 23         Self test failure.
```

Configuring the type of a module

module type

Syntax

`module <module-num> type <module-type>`
```

Description

Allows you to configure the type of the module.

Clearing the module configuration

Syntax

`no module <SLOT>`

Description

Allows removal of the module configuration in the configuration file after the module has been removed. Enter an integer between 1 and 12 for *slot*.

- This command can be used to swap a module for a different type.
- This command will save the changes to both the running and startup configuration without a user issuing a 'write memory'

Example

```
switch# no module 3
```

Configuring transceivers and modules that have not been inserted

Transceivers

Previously, a port had to be valid and verified for the switch to allow it to be configured. Transceivers are removable ports and considered invalid when not present in the switch, so they cannot be configured unless they are already in the switch. For switches, the verification for allowable port configurations performed by the CLI is removed and configuration of transceivers is allowed even if they are not yet inserted in the switch.

Modules

You can create or edit configuration files (as text files) that can be uploaded to the switch without the modules having been installed yet. Additionally, you can pre-configure the modules with the CLI `module` command.

The same `module` command used in an uploaded configuration file is used to define a module that is being pre-configured. The validation performed when issued through the CLI is still performed just as if the command was executed on the switch, in other words, as if the module were actually present in the switch.

NOTE

You cannot use this method to change the configuration of a module that has already been configured. The slot must be empty and the configuration file must not have a configuration associated with it.

Clearing the module configuration

Because of the hot-swap capabilities of the modules, when a module is removed from the chassis, the module configuration remains in the configuration file. `[no] module slot` allows you to remove the module configuration information from the configuration file.

This does not change how hot-swap works.

Power consumption

NOTE

The `show system power-supply detailed` command is only supported on the 5400R and 3810M switches.

show system power-supply

Syntax

`show system power-supply [detailed | fahrenheit]`

Description

Shows power supply information in either full detail or full detail in Fahrenheit only. Default temperature is displayed in degrees Celsius.

Command context

`manager and operator`

Parameters

`detailed`

Shows detailed switch power supply sensor information.

`fahrenheit`

Shows detailed switch power supply sensor information with temperatures in degrees Fahrenheit.

Usage

- The `show system power-supply detailed` command shows detailed information for the local power supplies only.
- The `show system power-supply detailed` command shows detailed information for power supplies in the `powered` state only.

Examples

Use of the command `show system power-supply` shows the power supply status for all active switches.

```
Switch# show system power-supply

Power Supply Status:

 PS#    Model      Serial          State          AC/DC  + V    Wattage
 ---- --------- ------------ ----------------- -------------- ----------
   1   J9828A    IN30G4D009   Powered           AC 120V/240V     700
   2   J9828A    IN30G4D00C   Powered           AC 120V/240V     700
   3                          Not Present       -- ---------       0
   4   J9830A    IN43G4G05H   Powered           AC 120V/240V    2750

   3 /  4 supply bays delivering power.
   Total power: 4150 W
```

Use of the command `show system power-supply detailed` shows the power supply status in detail for all active switches.

```
Switch# show system power-supply detailed
```

```
Status and Counters - Power Supply Detailed Information

 PS# Model   Serial      State        Status
 --- ------- ----------- ------------ -------------------------------------
 1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                      AC MAIN Voltage      : 209 Volts
                                      Power Supplied       : 31 Watts
                                      Power Capacity       : 700 Watts
                                      Inlet Temp (C/F)     : 27.0C/80.6F
                                      Internal Temp (C/F)  : 30.5C/86.0F
                                      Fan 1 Speed          : 1600 RPM (47%)
                                      Fan 2 Speed          : 1600 RPM (47%)

 2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                      AC MAIN Voltage      : 209 Volts
                                      Power Supplied       : 21 Watts
                                      Power Capacity       : 700 Watts
                                      Inlet Temp (C/F)     : 27.7C/80.6F
                                      Internal Temp (C/F)  : 32.5C/89.6F
                                      Fan 1 Speed          : 1600 RPM (47%)
                                      Fan 2 Speed          : 1600 RPM (47%)

 3                       Not Present

 4   J9830A  IN43G4G05H  Powered      AC Power Consumption : 90 Watts
                                      AC MAIN/AUX Voltage  : 210/118 Volts
                                      Power Supplied       : 16 Watts
                                      Power Capacity       : 2750 Watts
                                      Inlet Temp (C/F)     : 30.9C/86.0F
                                      Internal Temp (C/F)  : 65.6C/149.0F
                                      Fan 1 Speed          : 2000 RPM (37%)
                                      Fan 2 Speed          : 1950 RPM (36%)



  3 /  4 supply bays delivering power.
  Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply fahrenheit` shows the power supply status in Fahrenheit for all active switches.

```
Switch# show system power-supply detailed fahrenheit
Power Supply Status:
 Mem PS# Model   Serial      State        Status
 --- --- ------- ----------- ------------ -----------------------------------
 1   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                          AC MAIN/AUX Voltage : 118/208 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

 1   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
                                          AC Input Voltage    : 208 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

 1   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                          AC Input Voltage    : 119 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

 1   4                           Not Present
```

```
2   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                          AC MAIN/AUX Voltage : 118/208 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

2   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
                                          AC Input Voltage    : 208 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

2   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                          AC Input Voltage    : 119 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)
2   4                            Not Present
------------------------------------------------------------------------------
   6 / 8 supply bays delivering power.
   Total Input Power: 378 Watts
```

Use of the command `show system power-supply detailed` shows the power supply status all active switches including a nonpowered J9830A PSU.

```
switch# show system power-supply detailed

 Status and Counters - Power Supply Detailed Information

  PS# Model   Serial      State        Status
  --- ------- ----------- ------------ -------------------------------------
  1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 31 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.0C/80.6F
                                       Internal Temp (C/F)  : 30.5C/86.0F
                                       Fan 1 Speed          : 1600 RPM
                                       Fan 2 Speed          : 1600 RPM

  2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 21 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.7C/80.6F
                                       Internal Temp (C/F)  : 32.5C/89.6F
                                       Fan 1 Speed          : 1600 RPM
                                       Fan 2 Speed          : 1600 RPM

  3                       Not Present

  4   J9830A  IN43G4G05H  Aux Not Powered


   2 /  4 supply bays delivering power.
   Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply` shows the power supply status all active switches with power supply #2 showing permanent failure.

```
switch# show system power-supply

Power Supply Status:
```

```
PS#    Model       Serial          State         AC/DC   + V    Wattage
---- --------- ------------ ---------------- -------------- ---------
  1                        Not Present       -- ---------       0
  2    J9829A    IN30G4D00C  Permanent Failure  AC 120V/240V     1100
  3    J9829A    IN30G4D00D  Powered           -- ---------       1100
  4    J9829A    IN43G4G05H  Powered           AC 120V/240V      1100

  3 /  4 supply bays delivering power.
  Total power: 3300 W
```

Table 1: *Field key for output of* `show system power-supply detailed`

Field	Description
AC Power Consumption	Actual power consumed from AC input
AC MAIN/AUX Voltage	Actual voltage measured on AC Input: • Two voltages are displayed for PS#4, as the J9830A includes two AC input IEC connectors. • Most power-supplies contain a single AC Input IEC connector and are labeled MAIN.
Power Supplied	Actual voltage being supplied from the power-supply to the switch for general power and PoE.
Power Capacity	The maximum power that the power-supply can provide to the switch.
Inlet Temp (C/F)	The thermal sensor at the inlet of the power-supply - shown in both Celsius and Fahrenheit
Internal Temp (C/F)	The thermal sensor internal to the power-supply (will vary depending upon the model) - shown in both Celsius and Fahrenheit. NOTE: There is no "Output Temperature Sensor" on either the 5400R or 3810M switches.
Fan Speed	Shows the current fan speed in RPM and the percent of total fan speed utilization. For PSUs that contain more than one fan, a separate line will be included for each.
Currently Supplying	A summary of the total power being supplied and the total capacity (same summary as seen on the command `show system power-supply`).

Fans

There are three fan types:

- Power supply fans
- Fan-tray fans
- Stacking switch fans

show system

Syntax

```
show system [chassislocate | information | temperature]
```

Description

Shows global system information and operational parameters for the switch.

Command context

`manager and operator`

Parameters

`chassislocate`

Shows the chassis locator LED status. Possible values are ON, Off, and Blink. When the status is On or Blink, the number of minutes that the Locator LED will continue to be on or to blink is displayed.

`information`

Displays global system information and operational parameters for the switch.

`temperature`

Shows system temperature and settings.

Usage

- To show system fans, see **show system fan**.
- To show chassis power supply and settings, see **show system power-supply** .
- To show system fans for VSF members, see **show system fans**.

Examples

Locating the system chassis by LED blink using the `show system chassislocate` command.

```
HP Switch(config)# show system chassislocate

Chassis Locator LED: ON 5 minutes 5 seconds

HP Switch(config)# show system chassislocate

Chassis Locator LED: BLINK 10 minutes 6 seconds

HP Switch(config)# show system chassislocate

Chassis Locator LED: OFF
```

Showing the general switch system information by using the `show system` command.

```
HP Switch(config)# show system

 Status and Counters - General System Information

  System Name        : HP Switch Switch
  System Contact     :
  System Location    :

  MAC Age Time (sec) : 300

  Time Zone          : 0
  Daylight Time Rule : None


  Software revision  : T.13.XX          Base MAC Addr      : 001635-b57cc0
  ROM Version        : K.12.12          Serial Number      : LP621KI005

  Up Time            : 51 secs          Memory    - Total   : 152,455,616
  CPU Util (%)       : 3                            Free    : 110,527,264

  IP Mgmt  - Pkts Rx : 0                Packet    - Total   : 6750
             Pkts Tx : 0                Buffers     Free    : 5086
                                                    Lowest  : 5086
                                                    Missed  : 0
```

show system fans

Syntax

`show system fans`

Description

Shows the state, status, and location of system fans.

Command context

`manager and operator`

Usage

- Command can be executed using various command contexts. See examples for use of command context PoEP and VSF.

Examples

The state of all system fans is shown by using the command `show system fans`.

```
switch# show system fans

Fan Information
  Num   | State        | Failures | Location
-------+-------------+----------+----------
Fan-1   | Fan OK       |    0     | Fan Tray
Fan-2   | Fan OK       |    0     | Fan Tray
Fan-3   | Fan OK       |    0     | Fan Tray
Fan-4   | Fan OK       |    0     | Fan Tray
Fan-5   | Fan OK       |    0     | Fan Tray
Fan-6   | Fan OK       |    0     | Fan Tray
Fan-7   | Fan Removed  |    0     | PS 1
Fan-8   | Fan Failed   |    2     | PS 2
Fan-9   | Fan OK       |    0     | PS 3
Fan-10  | Fan OK       |    0     | PS 4

1 / 10 Fans in Failure State
1 / 10 Fans have been in Failure State
```

The state of all system fans within the PoEP context is shown by using the command `show system fans`.

```
switch(PoEP)# show system fans

Fan Information
  Num  | State       | Failures | Location
-------+-------------+---------------------
Fan-1  | Fan OK      |   0      | Chassis
Fan-2  | Fan OK      |   0      | Chassis
Fan-3  | Fan OK      |   0      | Chassis
Fan-4  | Fan Removed |   0      | PS 1
Fan-5  | Fan Failed  |   2      | PS 2

1 / 5 Fans in Failure State
1 / 5 Fans have been in Failure State
```

The state of all stacked switch system fans is shown by using the command `show system fans` within the `stacked` context.

```
Switch(stacked)# show system fans

Fan Information
Member  1

  Num  | State       | Failures | Location
-------+-------------+---------------------
Sys-1  | Fan OK      |   0      | Chassis
Sys-2  | Fan OK      |   0      | Chassis
Sys-3  | Fan OK      |   0      | Chassis
Sys-4  | Fan Removed |   0      | PS 1
Sys-5  | Fan Failed  |   2      | PS 2

0 / 5 Fans in Failure State
0 / 5 Fans have been in Failure State

Member  2

  Num  | State       | Failures | Location
-------+-------------+---------------------
Sys-1  | Fan OK      |   0      | Chassis
Sys-2  | Fan OK      |   0      | Chassis
Sys-3  | Fan OK      |   0      | Chassis
Sys-4  | Fan OK      |   0      | PS 1
Sys-5  | Fan OK      |   0      | PS 2

0 / 5 Fans in Failure State
0 / 5 Fans have been in Failure State
```

The state of all VSF switch members system fans is shown by using the command `show system fans` from within the VSF context.

```
VSF-Switch# show system fans

Fan Information
VSF-Member  1

  Num  | State       | Failures | Location
-------+-------------+----------+---------
Sys-1  | Fan OK      |   0      | Fan Tray
Sys-2  | Fan OK      |   0      | Fan Tray
Sys-3  | Fan OK      |   0      | Fan Tray
Sys-4  | Fan OK      |   0      | Fan Tray
Sys-5  | Fan Removed |   0      | PS 1
```

```
Sys-6   | Fan Failed  |    2      | PS 2

1 / 6 Fans in Failure State
1 / 6 Fans have been in Failure State

VSF-Member  2

  Num   | State         | Failures | Location
-------+-------------+----------+---------
Sys-1   | Fan OK        |    0      | Fan Tray
Sys-2   | Fan OK        |    0      | Fan Tray
Sys-3   | Fan OK        |    0      | Fan Tray
Sys-4   | Fan OK        |    0      | Fan Tray
Sys-5   | Fan OK        |    0      | PS 1
Sys-6   | Fan OK        |    0      | PS 2

0 / 6 Fans in Failure State
0 / 6 Fans have been in Failure State
HP-VSF-Switch#
```

show system power-supply

Syntax

`show system power-supply [detailed | fahrenheit]`

Description

Shows power supply information in either full detail or full detail in Fahrenheit only. Default temperature is displayed in degrees Celsius.

Command context

`manager and operator`

Parameters

`detailed`

Shows detailed switch power supply sensor information.

`fahrenheit`

Shows detailed switch power supply sensor information with temperatures in degrees Fahrenheit.

Usage

- The `show system power-supply detailed` command shows detailed information for the local power supplies only.
- The `show system power-supply detailed` command shows detailed information for power supplies in the `powered` state only.

Examples

Use of the command `show system power-supply` shows the power supply status for all active switches.

```
Switch# show system power-supply

Power Supply Status:

 PS#    Model      Serial         State          AC/DC   + V    Wattage
 ---- ---------- ------------ ----------------- --------------- ----------
   1    J9828A    IN30G4D009   Powered           AC 120V/240V      700
   2    J9828A    IN30G4D00C   Powered           AC 120V/240V      700
```

```
  3                          Not Present       -- ---------        0
  4   J9830A    IN43G4G05H   Powered           AC 120V/240V     2750

  3 /  4 supply bays delivering power.
  Total power: 4150 W
```

Use of the command `show system power-supply detailed` shows the power supply status in detail for all active switches.

```
Switch# show system power-supply detailed

 Status and Counters - Power Supply Detailed Information

  PS# Model   Serial      State        Status
  --- ------- ----------- ------------ -------------------------------------
  1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 31 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.0C/80.6F
                                       Internal Temp (C/F)  : 30.5C/86.0F
                                       Fan 1 Speed          : 1600 RPM (47%)
                                       Fan 2 Speed          : 1600 RPM (47%)

  2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 21 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.7C/80.6F
                                       Internal Temp (C/F)  : 32.5C/89.6F
                                       Fan 1 Speed          : 1600 RPM (47%)
                                       Fan 2 Speed          : 1600 RPM (47%)

  3                           Not Present

  4   J9830A  IN43G4G05H  Powered      AC Power Consumption : 90 Watts
                                       AC MAIN/AUX Voltage  : 210/118 Volts
                                       Power Supplied       : 16 Watts
                                       Power Capacity       : 2750 Watts
                                       Inlet Temp (C/F)     : 30.9C/86.0F
                                       Internal Temp (C/F)  : 65.6C/149.0F
                                       Fan 1 Speed          : 2000 RPM (37%)
                                       Fan 2 Speed          : 1950 RPM (36%)



   3 /  4 supply bays delivering power.
   Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply fahrenheit` shows the power supply status in Fahrenheit for all active switches.

```
Switch# show system power-supply detailed fahrenheit
Power Supply Status:
 Mem PS# Model   Serial      State        Status
 --- --- ------- ----------- ------------ -----------------------------------
 1   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                          AC MAIN/AUX Voltage : 118/208 Volts
                                          Inlet/Internal Temp : 85.6F/87.7F
                                          Fan 1 Speed (util)  : 1650RPM (20%)
                                          Fan 2 Speed (util)  : 1600RPM (19%)

 1   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
```

```
                                         AC Input Voltage    : 208 Volts
                                         Inlet/Internal Temp : 85.6F/87.7F
                                         Fan 1 Speed (util)  : 1650RPM (20%)
                                         Fan 2 Speed (util)  : 1600RPM (19%)

 1   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                         AC Input Voltage    : 119 Volts
                                         Inlet/Internal Temp : 85.6F/87.7F
                                         Fan 1 Speed (util)  : 1650RPM (20%)
                                         Fan 2 Speed (util)  : 1600RPM (19%)

 1   4                       Not Present

 2   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                         AC MAIN/AUX Voltage : 118/208 Volts
                                         Inlet/Internal Temp : 85.6F/87.7F
                                         Fan 1 Speed (util)  : 1650RPM (20%)
                                         Fan 2 Speed (util)  : 1600RPM (19%)

 2   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
                                         AC Input Voltage    : 208 Volts
                                         Inlet/Internal Temp : 85.6F/87.7F
                                         Fan 1 Speed (util)  : 1650RPM (20%)
                                         Fan 2 Speed (util)  : 1600RPM (19%)

 2   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                         AC Input Voltage    : 119 Volts
                                         Inlet/Internal Temp : 85.6F/87.7F
                                         Fan 1 Speed (util)  : 1650RPM (20%)
                                         Fan 2 Speed (util)  : 1600RPM (19%)
 2   4                       Not Present
-------------------------------------------------------------------------------
   6 / 8 supply bays delivering power.
   Total Input Power: 378 Watts
```

Use of the command `show system power-supply detailed` shows the power supply status all active switches including a nonpowered J9830A PSU.

```
switch# show system power-supply detailed

 Status and Counters - Power Supply Detailed Information

  PS# Model   Serial      State        Status
  --- ------- ----------- ------------ -------------------------------------
  1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 31 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.0C/80.6F
                                       Internal Temp (C/F)  : 30.5C/86.0F
                                       Fan 1 Speed          : 1600 RPM
                                       Fan 2 Speed          : 1600 RPM

  2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 21 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.7C/80.6F
                                       Internal Temp (C/F)  : 32.5C/89.6F
                                       Fan 1 Speed          : 1600 RPM
                                       Fan 2 Speed          : 1600 RPM

  3                       Not Present
```

```
  4   J9830A  IN43G4G05H  Aux Not Powered


   2 /  4 supply bays delivering power.
   Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply` shows the power supply status all active switches with power supply #2 showing permanent failure.

```
switch# show system power-supply

Power Supply Status:

 PS#    Model      Serial          State         AC/DC  + V    Wattage
 ---- --------- ------------ ----------------- -------------- ---------
   1                          Not Present       -- ---------        0
   2    J9829A    IN30G4D00C   Permanent Failure  AC 120V/240V     1100
   3    J9829A    IN30G4D00D   Powered           -- ---------      1100
   4    J9829A    IN43G4G05H   Powered           AC 120V/240V     1100

   3 /  4 supply bays delivering power.
   Total power: 3300 W
```

Table 2: *Field key for output of* `show system power-supply detailed`

Field	Description
AC Power Consumption	Actual power consumed from AC input
AC MAIN/AUX Voltage	Actual voltage measured on AC Input: • Two voltages are displayed for PS#4, as the J9830A includes two AC input IEC connectors. • Most power-supplies contain a single AC Input IEC connector and are labeled MAIN.
Power Supplied	Actual voltage being supplied from the power-supply to the switch for general power and PoE.
Power Capacity	The maximum power that the power-supply can provide to the switch.
Inlet Temp (C/F)	The thermal sensor at the inlet of the power-supply - shown in both Celsius and Fahrenheit
Internal Temp (C/F)	The thermal sensor internal to the power-supply (will vary depending upon the model) - shown in both Celsius and Fahrenheit. NOTE: There is no "Output Temperature Sensor" on either the 5400R or 3810M switches.
Fan Speed	Shows the current fan speed in RPM and the percent of total fan speed utilization. For PSUs that contain more than one fan, a separate line will be included for each.
Currently Supplying	A summary of the total power being supplied and the total capacity (same summary as seen on the command `show system power-supply`).

Fan failures and SNMP traps

Power supply fan-fault

Power supply events indicating an internal fan-fault are reported by SNMP traps issued up to 10 seconds after the corresponding power supply fan fault occures.

For single event power supply fan faults, a corresponding SNMP trap is issued.

Single event and corresponding SNMP trap issued for a power supply fan-fault and recovery

Shown is a fan-fault (fan 1 of 2) and spontaneous recovery a few seconds after within power supply in bay number 2. The event is issued as informative (I).

```
I 11/30/16 14:01:59 02778 chassis: AM1: Internal power supply 2: Fan 1 OK.
```

Fan-tray fan-fault events

For single event fan-tray fan-faults, the corresponding SNMP traps are issued only for fans within the fan-tray.

Single event and corresponding SNMP trap issued for a fan-tray fan-fault and recovery

Shown is a fan-tray fan-fault (fan number 3) and spontaneous recovery a few seconds after. The event is issued as "Informative" (I).

```
I 11/30/16 14:03:08 00070 chassis: AM1: Fan OK: Fan: 3  Failures: 1
```

Shown is a fan-tray fan-fault (fan number 3) failure. The event is issued as a "Warning" (w).

```
W 11/30/16 14:02:38 00070 chassis: AM1: Fan failure: Fan: 3 Failures: 1
```

System boot diagnostics

Power-on self-test (POST) is the diagnostic testing sequence that runs on start-up to verify hardware functionality. The command `show system post` provides information by slot or interface module, which aids in troubleshooting issues at startup.

`show system post`

Syntax

`show system post <SLOT>`

Description

Shows detailed POST information by slot or interface module, which aids in troubleshooting issues at start-up.

Command context

`manager`

Parameters

<SLOT>

Specifies the slot number for detailed POST information.

Example

Show the detailed POST information on slot 1.

```
switch#  show system post 1

Slot 1 POST results

Failed Results:

Timestamp          Test Type  Port  Error Code
------------------ ---------  ---- -----------
01/25/15 11:02:32  Loopback   2     0x1010060
01/25/15 11:02:50  Loopback   7     0x1032000
01/25/15 11:02:59  POE        10    0xFFFFFFF
01/25/15 11:02:50  MACSEC     21    0x1082000

All Ports except the ones listed above have passed the following self-tests
1. Loopback
2. POE
3. MACSEC

Note: This is just a reference. POE and MACSEC tests may not show up in cases
where there is no support.
```

Field descriptions

Table 3: *Field descriptions*

Field	Description
Timestamp	Time and date of the POST failure result
Test Type	There are 3 types of tests: **Loopback** Packet loopback test **PoE** PoE controller test **MACSEC** Loopback test with MACSEC enabled
Port	Slot port number
Error Code	Reason behind failure

show system post member

Syntax

`show system post member <MEMBER-ID>`

Description

Shows the detailed POST results for the specified stack member.

Command context

`manager`

Parameters

<MEMBER-ID>

Specifies the member ID for the stack.

Examples

Show the POST results for the stack member 1.

```
switch-stack# show system post member 1

POST results for member 1

Failed Results:

Timestamp         Test Type   Port   Error Code
----------------- ---------   ----   -----------
01/25/15 11:02:32  Loopback   1/2    0x1010060
01/25/15 11:02:50  Loopback   1/7    0x1032000
01/25/15 11:02:59  POE        1/10   0xFFFFFFF
01/25/15 11:02:50  MACSEC     1/21   0x1082000


All Ports except the ones listed above have passed the following self-tests
1. Loopback
2. POE
3. MACSEC


Note: This is just a reference. POE and MACSEC tests may not show up in cases
where there is no support.
```

show system post vsf member

Syntax

`show system post vsf member <MEMBER-ID>`

Description

Shows the detailed POST results for the specified VSF stack member.

Command context

`manager`

Parameters

<MEMBER-ID>

Specifies the member ID of the VSF stack.

Examples

Show the VSF member 1 detailed POST results.

```
switch-VSF# show system post vsf member 1

POST results for vsf member 1

Failed Results:

Timestamp         Test Type   Port   Error Code
----------------- ---------   ----   -----------
```

```
01/25/15 11:02:32  Loopback   1/2    0x1010060
01/25/15 11:02:50  Loopback   1/7    0x1032000
01/25/15 11:02:59  POE        1/10   0xFFFFFFF
01/25/15 11:02:50  MACSEC     1/21   0x1082000


All Ports except the ones listed above have passed the following self-tests
1. Loopback
2. POE
3. MACSEC


Note: This is just a reference. POE and MACSEC tests may not show up in cases
where there is no support.
```

Viewing port status and configuration

show interfaces

Syntax

`show interfaces [brief | config | <PORT-LIST>]`

Description

Shows interface information for ports or trunk groups in brief or configuration detail.

Command context

`operator or manager`

Parameters

brief

Shows the current operating status for all ports or trunk groups on the switch in brief detail.

config

Shows the configuration data for all ports or trunk groups on the switch.

<PORT-LIST>

Specifies the list of ports for which status information will be shown.

<TRUNK-GROUP>

Specifies the trunk group for which status information will be shown. The status information shown consists of total transmit and receive counters and weighted average rate for the trunk group specified. The weighted average rate is calculated in 5 minute intervals.

Usage

Both external and internal ports are supported on the same module. Internal ports have an “I” suffix in the output of the show command to indicate that they are internal ports.

- “10GbE-INT” – Internal 10G data-plane ports (1i-2i, 4i-5i)
- “1GbE-INT” – Internal 1G control-plane port (3i)
- Port 3i always shows as link-down.

Examples

Show brief information and status of all interfaces.

```
switch# show interfaces brief

Status and Counters - Port Status
                | Intrusion                           MDI   Flow  Bcast
Port  Type      | Alert     Enabled Status Mode       Mode  Ctrl  Limit
----- --------- + --------- ------- ------ ---------- ----- ----- ------
B1    100/1000T | No        Yes     Down   Auto-10-100 Auto  off   0
```

```
B2    100/1000T | No        Yes     Down   1000FDx     Auto  off   0
B3    100/1000T | No        Yes     Down   1000FDx     Auto  off   0
B4    100/1000T | No        Yes     Down   1000FDx     Auto  off   0
B5    100/1000T | No        Yes     Down   1000FDx     Auto  off   0
B6    100/1000T | No        Yes     Down   1000FDx     Auto  off   0
```

Show the configuration of the interfaces currently available.

```
switch# show interfaces config

Port Settings
  Port  Type      | Enabled Mode         Flow Ctrl MDI
  ----- --------- + ------- ------------ --------- ----
  B1    100/1000T | Yes     Auto-10-100  Disable   Auto
  B2    100/1000T | Yes     Auto         Disable   Auto
  B3    100/1000T | Yes     Auto         Disable   Auto
  B4    100/1000T | Yes     Auto         Disable   Auto
  B5    100/1000T | Yes     Auto         Disable   Auto
  B6    100/1000T | Yes     Auto         Disable   Auto
```

Show brief information and status of interfaces for ports D1i, D2i, and D3i.

```
switch# show interfaces brief d1i-d3i

Status and Counters - Port Status
                  | Intrusion                          MDI  Flow Bcast
Port   Type       | Alert     Enabled Status Mode       Mode Ctrl Limit
------ ---------- + --------- ------- ------ ---------- ---- ---- -----
D1i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
D2i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
D3i    1GbE-INT   | No        Yes     Down   1000FDx    NA   off  0
```

Show the brief information and status of interfaces for ports B1 through internal port B3i.

```
switch# show interfaces brief b1-b3i

Status and Counters - Port Status
                  | Intrusion                          MDI  Flow Bcast
 Port   Type      | Alert     Enabled Status Mode       Mode Ctrl Limit
 ------ ---------- + --------- ------- ------ ---------- ---- ---- -----
 B1     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B2     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B3     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B4     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B5     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B6     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B7     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B8     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B9     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B10    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B11    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B12    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B1i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
 B2i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
 B3i    1GbE-INT   | No        Yes     Up     1000FDx    NA   off  0
```

Show detailed interface information for port trunk 1.

```
switch# show interface trk1

Status and Counters - Port Counters for port Trk1
  Name  :  Trk1
  MAC Address      : 3464a9-b1b8bf
  Link Status      : Up
```

```
  Totals (Since boot or last clear) :
   Bytes Rx         : 777,713,956          Bytes Tx         : 596,853,141
   Unicast Rx       : 1,154,693            Unicast Tx       : 0
   Bcast/Mcast Rx   : 48,563               Bcast/Mcast Tx   : 607,474,910
  Errors (Since boot or last clear) :
   FCS Rx           : 0                    Drops Tx         : 0
   Alignment Rx     : 0                    Collisions Tx    : 0
   Runts Rx         : 0                    Late Colln       : 0
   Giants Rx        : 0                    Excessive Colln  : 0
   Total Rx Errors  : 0                    Deferred Tx      : 0
  Others (Since boot or last clear) :
   Discard Rx       : 0                    Out Queue Len    : 0
   Unknown Protos   : 0
  Rates (5 minute weighted average) :
   Total Rx(Kbps)   : 76,800                 Total Tx(Kbps): 76,800
   Unicast Rx (Pkts/sec) : 21              Unicast Tx (Pkts/sec) : 0
   B/Mcast Rx (Pkts/sec) : 114,878         B/Mcast Tx (Pkts/sec) : 114,900
   Utilization Rx   : 00.76 %
Utilization Tx   : 00.76 %
```

Viewing transceiver information

NOTE

40G QSFP+ BIDI transceiver support is available only on switches running KB software.

This features provides the ability to view diagnostic monitoring information for transceivers with Diagnostic Optical Monitoring (DOM) support. The following table indicates the support level for specific transceivers:

Product #	Description	Support[1]
J8436A	10GbE X2–SC SR Optic	V
J8437A	10GbE X2–SC LR Optic	V
J8440B	10GbE X2-CX4 Xcver	NA
J8440C	10GbE X2-CX4 Xcver	NA
J4858A	Gigabit-SX-LC Mini-GBIC	V
J4858B	Gigabit-SX-LC Mini-GBIC	V
J4858C	Gigabit-SX-LC Mini-GBIC	V (some)
J9054B	100-FX SFP-LC Transceiver	N
J8177C	Gigabit 1000Base-T Mini-GBIC	NA

Table Continued

Product #	Description	Support[1]
J9150A	10GbE SFP+ SR Transceiver	D
J9151A	10GbE SFP+ LR Transceiver	D
J9152A	10GbE SFP+ LRM Transceiver	D
J9153A	10GbE SFP+ ER Transceiver	D
J9144A	10GbE X2-SC LRM Transceiver	D
J8438A	10Gbe X2-SC ER Transceiver	D
JH233A	40G QSFP+ MPO eSR4 Transceiver	V
JH232A	40G QSFP+ LC LR4 SM Transceiver	V
JL308A	40G QSFP+BIDI	V
JH231A	40G QSFP+ MPO SR4 Transceiver	V

[1] Support indicators:

- V - Validated to respond to DOM requests
- N - No support of DOM
- D - Documented by the component suppliers as supporting DOM
- NA - Not applicable to the transceiver (copper transceiver)

NOTE

Not all transceivers support Digital Optical Monitoring. If DOM appears in the Diagnostic Support field of the `show interfaces transceiver detail` command, or the hpicfTransceiverMIB hpicfXcvrDiagnostics MIB object, DOM is supported for that transceiver.

The port VLAN tagged status

The `show interfaces status` command displays port status, configuration mode, speed, type and tagged or untagged information.

Tagged values can be:

- VLAN ID: When the VLAN number is displayed, the port is a member of a single tagged VLAN.
- multi: When “multi” is displayed, the port is a member of multiple tagged VLANs.
- no: When “no” is displayed, the port is not a member of any tagged VLAN.

Untagged values can be:

- VLAN-ID: When the VLAN number is displayed, the port is a member of a single untagged VLAN.
- multi: When "multi" is displayed, the port is added to multiple untagged VLANs.
- no: When "no" is displayed, the port is not a member of any tagged VLAN.

If the port is part of a trunk, then the trunk_VLAN membership is displayed in the Tagged and Untagged columns.

show interfaces

```
switch(config)# show interfaces status
Port Name Status Config-mode Speed Type Tagged Untagged
-------- -------- ------ ----------- ------- --------- ------ ---------
A1 Up Auto 1000FDx 100/1000T 2 1
A2 Down 10HDx 10HDx 100/1000T multi 2
A3 Down 100HDx 100HDx 100/1000T multi 3
A4 Down 10FDx 10FDx 100/1000T 5 4
A5-Trk1 Down 100FDx 100FDx 100/1000T No No
A6 Down Auto 1000FDx 100/1000T No 6
A7 Down Auto-10 10HDx 100/1000T No 7
```

Dynamically updating the show interfaces command

command name

Syntax

`show interfaces display`

Description

Uses the `display` option to initiate the dynamic update of the `show interfaces` command, with the output being the same as the `show interfaces` command.

Usage

Select **Back** to exit the display.

show interfaces display

```
switch# show interfaces display
```

When using the **display** option in the CLI, the information stays on the screen and is updated every 3 seconds, as occurs with the display using the menu feature. The update is terminated with **CTRL-C**.

You can use the arrow keys to scroll through the screen when the output does not fit in one screen.

Figure 11: *`show interfaces display` command with dynamically updating output*

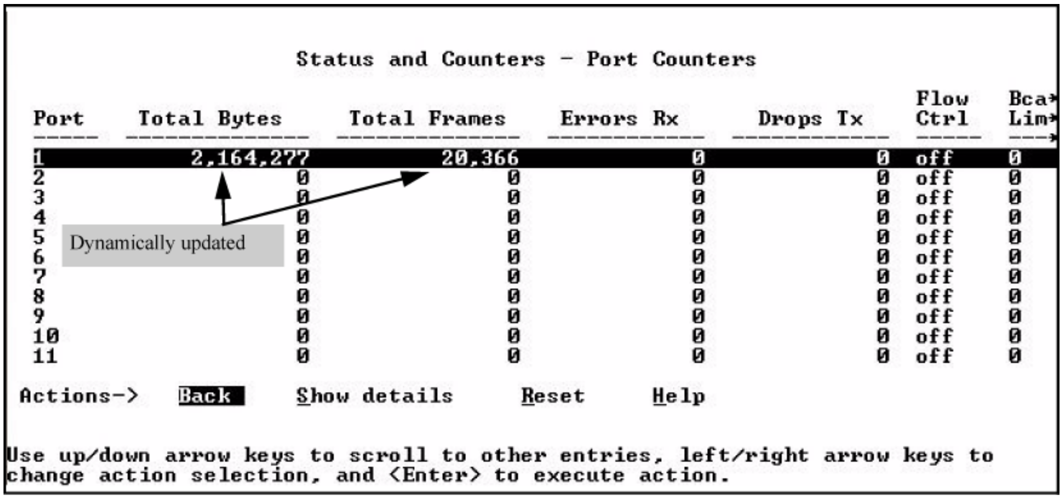

Customizing the show interfaces command

You can create `show` commands displaying the information that you want to see in any order you want by using the option.

show interfaces custom

Syntax

`show interfaces custom <PORT-LIST> <COLUMN-LIST>`

Description

Select the information that you want to display. Supported columns are shown in the following section.

Parameters and options

port

port identifier, such as A2.

type

Port type, such as 100/1000T

status

Port status, up or down.

speed

Connection speed and duplex, such as 1000FDX

mode

Configured mode, auto, auto-100, 100FDX

mdi

MDI mode, auto, MDIX

flow

Flow control, on or off

name

Friendly port name

vlanid

The vlan id this port belongs to, or "tagged" if it belongs to more than one vlan.

enabled

Port is or is not enabled, yes or no.

intrusion

Intrusion alert status, no.

bcast

Broadcast limit, 0

You can specify the column width by entering a colon after the column name, then indicating the number of characters to display. In **Example of the custom show interfaces command** on page 108, the Name column displays only the first four characters of the name. All remaining characters are truncated.

Each field has a fixed minimum width to be displayed. If you specify a field width smaller than the minimum width, the information is displayed at the minimum width. For example, if the minimum width for the Name field is 4 characters and you specify Name:2, the Name field displays 4 characters.

You can enter parameters in any order. There is a limit of 80 characters per line; if you exceed this limit an error displays.

Example of the custom `show interfaces` command

```
switch# show int custom 1-4 port name:4 type vlan intrusion speed enabled mdi

 Status and Counters - Custom Port Status

                                       Intrusion
  Port Name       Type       VLAN  Alert     Speed   Enabled MDI-mode
  ---- ---------- ---------- ----- --------- ------- ------- --------
  1    Acco       100/1000T  1     No        1000FDx Yes     Auto
  2    Huma       100/1000T  1     No        1000FDx Yes     Auto
  3    Deve       100/1000T  1     No        1000FDx Yes     Auto
  4    Lab1       100/1000T  1     No        1000FDx Yes     Auto
```

show interface smartrate

Syntax

`show interface <PORT-LIST> smartrate`

Description

The option `smartrate` has been added to the `show interface <PORT-LIST>` command. This option is used to display port diagnostics on a Smart Rate port only. If the command is run on a non - Smart Rate port, a message similar to `Port A1: This command is only applicable to Smart Rate ports` will display.

show interface port utilization

Syntax

`show interface port-utilization`

Description

Use the `show interface port-utilization` command to view a real-time rate display for all ports on the switch. **show interface port-utilization command** on page 109 shows a sample output from this command.

- For each port on the switch, the command provides a real-time display of the rate at which data is received (Rx) and transmitted (Tx) in terms of kilobits per second (KBits/s), number of packets per second (Pkts/s), and utilization (Util) expressed as a percentage of the total bandwidth available.
- The `show interfaces` *<PORT-LIST>* command can be used to display the current link status and the port rate average over a 5 minute period. Port rates are shown in bits per second (bps) for ports up to 1 Gigabit; for 10 Gigabit ports, port rates are shown in kilobits per second (Kbps.)

show interface port-utilization command

```
HP Switch(config)# show interfaces port-utilization
 Status and Counters - Port Utilization

                                Rx                           Tx
 Port       Mode      | --------------------------- | ---------------------------
                      | Kbits/sec   Pkts/sec   Util  | Kbits/sec  Pkts/sec   Util
 ---------  --------- + ---------- ---------- ----- + ---------- ---------- -----
 B1         1000FDx   | 0          0          0     | 0          0          0
 B2         1000FDx   | 0          0          0     | 0          0          0
 B3         1000FDx   | 0          0          0     | 0          0          0
 B4         1000FDx   | 0          0          0     | 0          0          0
 B5         1000FDx   | 0          0          0     | 0          0          0
 B6         1000FDx   | 0          0          0     | 0          0          0
 B7         100FDx    | 624        86         00.62 | 496        0          00.49
```

Enabling or disabling ports and configuring port mode

You can configure one or more of the following port parameters.

interface

Syntax

`interface` *<PORT-LIST>* `[disable|enable]`

Description

Disables or enables the port for network traffic. Does not use the `no` form of the command. Defaults to `enable`. You can substitute `int` for `interface` (for example, `int` *<PORT-LIST>* .)

Parameters and options

`speed-duplex [auto-10|10-full|10-half|100-full|100-half|auto|auto-100|1000-full]`

Specifies the port's data transfer speed and mode. Does not use the `no` form of the command. Default: `auto`.

The 10/100 auto-negotiation feature allows a port to establish a link with a port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.

Configure port C5 for auto-10-100

```
switch# int c5 speed-duplex auto-10-100
```

Configure ports C1 through C3 and port C6 for 100Mbps full-duplex

```
switch# int c1-c3,c6 speed-duplex 100-full
```

Similarly, to configure a single port with the above command settings, you could either enter the same command with only the one port identified or go to the context level for that port and then enter the command. For example, to enter the context level for port C6 and then configure that port for 100FDx:

```
switch# int e c6
switch(eth-C6#) speed-duplex 100-full
```

If port C8 was disabled, and you wanted to enable it and configure it for 100FDx with flow-control active, you could do so with either of the following command sets:

Figure 12: *Two methods for changing a port configuration*

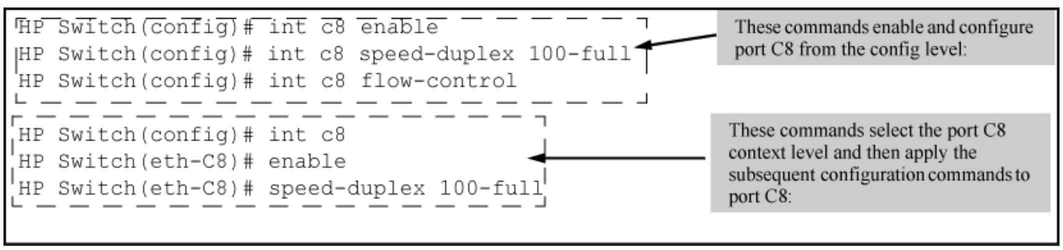

Basic USB port commands

This feature allows configuration of the USB port with either the CLI or SNMP.

Syntax

`show usb-port`

Description

To display the status of the USB port: Displays the status of the USB port. It can be enabled, disabled, or not present.

`show usb-port` command output

```
switch# show usb-port

  USB port status: enabled
  USB port power status: power on    (USB device detected in port)
```

One of the following messages indicates the presence or absence of the USB device:

- Not able to sense device in USB port
- USB device detected in port
- no USB device detected in port

The autorun feature works only when a USB device is inserted and the USB port is enabled.

usb-port

Syntax

`usb-port`

`no usb-port`

Description

Enables the USB port.

The `no` form of the command disables the USB port and any access to the device.

Command context

`Config`

show usb-port

Syntax

`show usb-port`

Description

Displays the status of the USB port. It can be enabled, disabled, or not present.

Command context

`operator`

Usage

One of the following messages indicates the presence or absence of a USB device:

- `Not able to sense device in USB port`
- `USB device detected in port`
- `no USB device detected in port`

Example

Display USB port status.

```
switch# show usb-port

  USB port status: enabled
  USB port power status: power on    (USB device detected in port)
```

Enabling or disabling flow control

You must enable flow control on both ports in a given link. Otherwise, flow control does not operate on the link and appears as `Off` in the `show interfaces brief` port listing, even if flow control is configured as enabled on the port in the switch. Also, the port (speed-duplex) mode must be set to `Auto` (the default).

To disable flow control on some ports, while leaving it enabled on other ports, just disable it on the individual ports you want to exclude.

interface flow-control

Syntax

```
interface <PORT-LIST> flow-control
```

Description

Enables or disables flow control packets on the port. Default: Disabled.

Parameters

no

The `no` form of the command disables flow control on the individual ports.

Examples

```
no interface <PORT-LIST> flow-control
```

Usage

1. You want to enable flow control on ports A1-A6.
2. Later, you decide to disable flow control on ports A5 and A6.
3. As a final step, you want to disable flow control on all ports.

Assuming that flow control is currently disabled on the switch, you would use these commands:

Configuring flow control for a series of ports

```
switch# int a1-a6 flow-control
switch# show interfaces brief

 Status and Counters - Port Status

                     | Intrusion                           MDI  Flow Bcast
  Port   Type        | Alert     Enabled Status Mode       Mode Ctrl Limit
  ------ ---------- + --------- ------- ------ ---------- ---- ---- -----
  A1     10GbE-T    | No        Yes     Up     1000FDx    NA   on   0
  A2     10GbE-T    | No        Yes     Up     10GigFD    NA   on   0
  A3     10GbE-T    | No        Yes     Up     10GigFD    NA   on   0
  A4     10GbE-T    | No        Yes     Up     10GigFD    NA   on   0
  A5     10GbE-T    | No        Yes     Up     10GigFD    NA   on   0
  A6     10GbE-T    | No        Yes     Up     10GigFD    NA   on   0
  A7     10GbE-T    | No        Yes     Down   10GigFD    NA   off  0
  A8     10GbE-T    | No        Yes     Up     10GigFD    NA   off  0
```

Example continued from Configuring flow control for a series of ports

```
switch# no int a5-a6 flow-control
switch# show interfaces brief

 Status and Counters - Port Status

                     | Intrusion                           MDI  Flow Bcast
  Port   Type        | Alert     Enabled Status Mode       Mode Ctrl Limit
  ------ ---------- + --------- ------- ------ ---------- ---- ---- -----
  A1     10GbE-T    | No        Yes     Up     1000FDx    NA   on   0
  A2     10GbE-T    | No        Yes     Down   10GigFD    NA   on   0
  A3     10GbE-T    | No        Yes     Down   10GigFD    NA   on   0
  A4     10GbE-T    | No        Yes     Down   10GigFD    NA   on   0
  A5     10GbE-T    | No        Yes     Down   10GigFD    NA   off  0
```

```
  A6      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A7      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A8      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
```

Example continued from Example continued from Configuring flow control for a series of ports

```
switch# no int a1-a4 flow-control
switch# show interfaces brief

 Status and Counters - Port Status

                     | Intrusion                              MDI  Flow Bcast
  Port    Type       | Alert      Enabled Status Mode         Mode Ctrl Limit
  ------ ---------- + ---------- ------- ------ ---------- ---- ---- -----
  A1      10GbE-T    | No         Yes      Down    1000FDx     NA   off  0
  A2      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A3      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A4      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A5      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A6      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A7      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
  A8      10GbE-T    | No         Yes      Down    10GigFD     NA   off  0
```

Configuring auto-MDIX

interface mdix-mode

Syntax

`interface <PORT-LIST> mdix-mode [auto-mdix|mdi|mdix]`

Description

The auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables.

Parameters

auto-mdix

The automatic, default setting. This configures the port for automatic detection of the cable (either straight-through or crossover.)

mdi

The manual mode setting that configures the port for connecting to either a PC or other MDI device with a crossover cable, or to a switch, hub, or other MDI-X device with a straight-through cable.

mdix

The manual mode setting that configures the port for connecting to either a switch, hub, or other MDI-X device with a crossover cable, or to a PC or other MDI device with a straight-through cable.

show interfaces config

Syntax

`show interfaces config`

Description

Lists the current per-port Auto/MDI/MDI-X configuration.

show interfaces brief

Syntax

`show interfaces brief`

Description

- Where a port is linked to another device, this command lists the MDI mode the port is currently using.
- In the case of ports configured for `Auto` (`auto-mdix`), the MDI mode appears as either `MDI` or `MDIX`, depending upon which option the port has negotiated with the device on the other end of the link.
- In the case of ports configured for `MDI` or `MDIX`, the mode listed in this display matches the configured setting.
- If the link to another device was up, but has gone down, this command shows the last operating MDI mode the port was using.
- If a port on a given switch has not detected a link to another device since the last reboot, this command lists the MDI mode to which the port is currently configured.

`show interfaces config` displays the following data when port A1 is configured for `auto-mdix`, port A2 is configured for `mdi`, and port A3 is configured for `mdix`:

Example of displaying the current MDI configuration

```
switch# show interfaces config

 Port Settings

  Port   Type       | Enabled Mode         Flow Ctrl MDI
  ------ ---------- + ------- ------------ --------- ----
  A1     10GbE-T    | Yes     Auto         Disable   Auto
  A2     10GbE-T    | Yes     Auto         Disable   MDI
  A3     10GbE-T    | Yes     Auto         Disable   MDIX
  A4     10GbE-T    | Yes     Auto         Disable   Auto
  A5     10GbE-T    | Yes     Auto         Disable   Auto
  A6     10GbE-T    | Yes     Auto         Disable   Auto
  A7     10GbE-T    | Yes     Auto         Disable   Auto
  A8     10GbE-T    | Yes     Auto         Disable   Auto
```

Example of displaying the current MDI operating mode

```
switch# show interfaces brief

 Status and Counters - Port Status

                    | Intrusion                           MDI  Flow Bcast
  Port   Type       | Alert     Enabled Status Mode       Mode Ctrl Limit
  ------ ---------- + --------- ------- ------ ---------- ---- ---- -----
  A1     10GbE-T    | No        Yes     Up     1000FDx    MDIX off  0
  A2     10GbE-T    | No        Yes     Down   10GigFD    MDI  off  0
  A3     10GbE-T    | No        Yes     Down   10GigFD    MDIX off  0
  A4     10GbE-T    | No        Yes     Down   10GigFD    Auto off  0
  A5     10GbE-T    | No        Yes     Down   10GigFD    Auto off  0
  A6     10GbE-T    | No        Yes     Down   10GigFD    Auto off  0
  A7     10GbE-T    | No        Yes     Down   10GigFD    Auto off  0
  A8     10GbE-T    | No        Yes     Down   10GigFD    Auto off  0
```

Viewing port configuration (Menu)

The menu interface displays the configuration for ports and (if configured) any trunk groups.

From the Main Menu, select:

Procedure

1. **Status and Counters**
2. Select **Port Status**

Figure 13: *Switch port status screen*

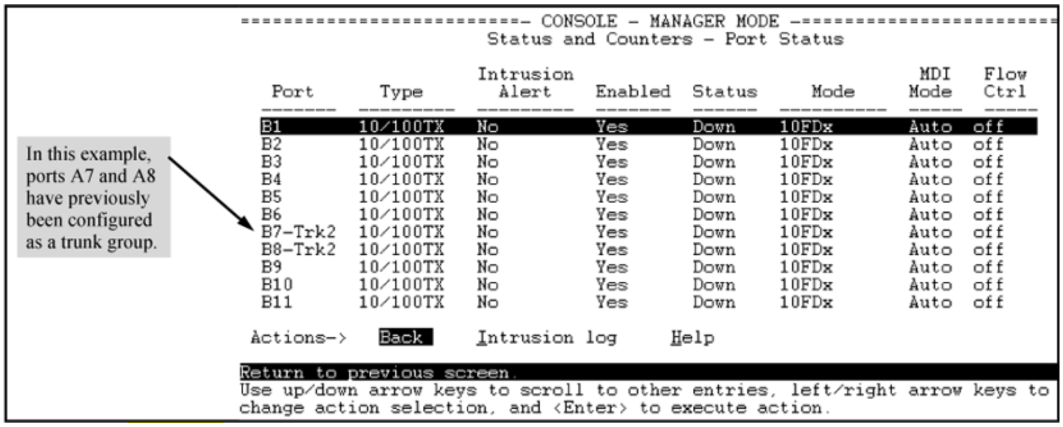

```
==========================- CONSOLE - MANAGER MODE -========================
                    Status and Counters - Port Status

                          Intrusion                                MDI    Flow
  Port      Type          Alert      Enabled  Status   Mode        Mode   Ctrl
  -------   ---------     ---------  -------  ------   ----------  -----  -----
 B1         10/100TX      No         Yes      Down     10FDx       Auto   off
 B2         10/100TX      No         Yes      Down     10FDx       Auto   off
 B3         10/100TX      No         Yes      Down     10FDx       Auto   off
 B4         10/100TX      No         Yes      Down     10FDx       Auto   off
 B5         10/100TX      No         Yes      Down     10FDx       Auto   off
 B6         10/100TX      No         Yes      Down     10FDx       Auto   off
 B7-Trk2    10/100TX      No         Yes      Down     10FDx       Auto   off
 B8-Trk2    10/100TX      No         Yes      Down     10FDx       Auto   off
 B9         10/100TX      No         Yes      Down     10FDx       Auto   off
 B10        10/100TX      No         Yes      Down     10FDx       Auto   off
 B11        10/100TX      No         Yes      Down     10FDx       Auto   off

 Actions->   Back     Intrusion log     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```


Configuring ports (Menu)

The menu interface uses the same screen for configuring both individual ports and port trunk groups.

From the Main Menu:

Procedure

1. Select **Switch Configuration...**
2. Select **Port/Trunk Settings**

   Figure 14: *Port/trunk settings with a trunk group configured*

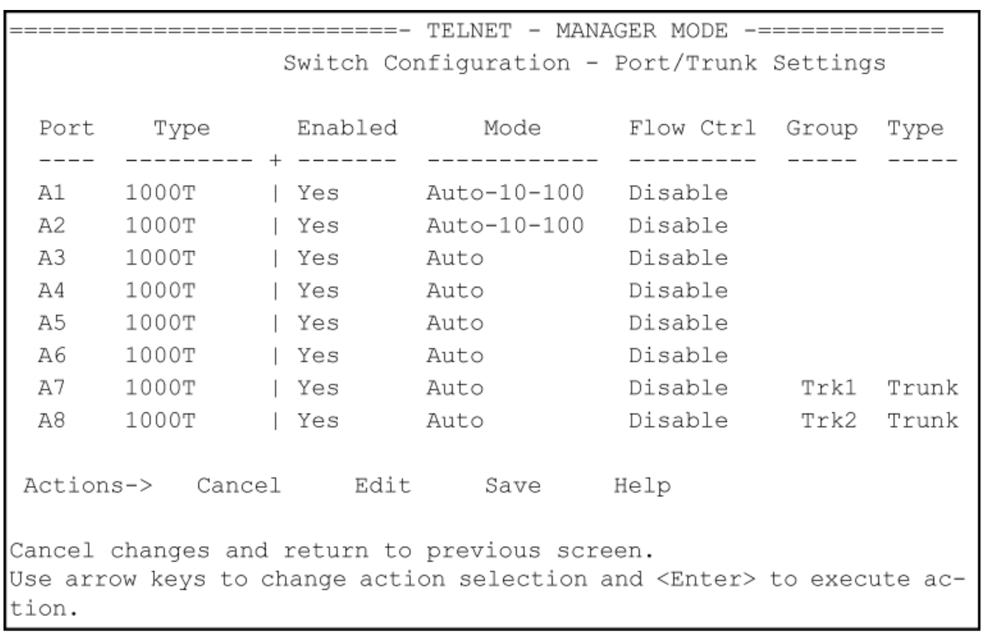

```
===========================- TELNET - MANAGER MODE -=============
                   Switch Configuration - Port/Trunk Settings

  Port    Type        Enabled     Mode          Flow Ctrl  Group  Type
  ----  --------- + -------  ------------  ---------  -----  -----
  A1    1000T     | Yes      Auto-10-100   Disable
  A2    1000T     | Yes      Auto-10-100   Disable
  A3    1000T     | Yes      Auto          Disable
  A4    1000T     | Yes      Auto          Disable
  A5    1000T     | Yes      Auto          Disable
  A6    1000T     | Yes      Auto          Disable
  A7    1000T     | Yes      Auto          Disable     Trk1   Trunk
  A8    1000T     | Yes      Auto          Disable     Trk2   Trunk

 Actions->    Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute ac-
tion.
```

3. On the keyboard, press `[E]` (for **E**dit.)

The cursor moves to the `Enabled` field for the first port.

4. When you have finished making changes to the above parameters, press `[Enter]`, and then press `[S]` (for `Save`.)

Configuring friendly port names

interface name

Syntax

`interface <PORT-LIST> name <port-name-string>`

Description

Assigns a port name to *<PORT-LIST>*.

Parameter

no

Deletes the port name from *<PORT-LIST>*.

Configuring a single port name

Suppose that you have connected port A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

Example of configuring a friendly port name

```
switch# int A3 name Bill_Smith@10.25.101.73
switch# write mem
switch# show name A3

 Port Names
  Port : A3
   Type : 10/100TX
   Name : Bill_Smith@10.25.101.73
```

Configuring the same name for multiple ports

Suppose that you want to use ports A5 through A8 as a trunked link to a server used by a drafting group. In this case you might configure ports A5 through A8 with the name "Draft-Server:Trunk."

Example of configuring one friendly port name on multiple ports

```
switch# int a5-a8 name Draft-Server:Trunk
switch# write mem
switch# show name a5-a8

 Port Names

  Port : A5
   Type : 10GbE-T
   Name : Draft-Server:Trunk
 Port : A6
   Type : 10GbE-T
   Name : Draft-Server:Trunk
 Port : A7
```

```
   Type : 10GbE-T
   Name : Draft-Server:Trunk
 Port : A8
   Type : 10GbE-T
   Name : Draft-Server:Trunk
```

Viewing friendly port names with other port data

show name

Syntax

`show name`

Description

Displays a listing of port numbers with their corresponding friendly port names and also quickly shows you which ports do not have friendly name assignments. (`show name` data comes from the running-config file.)

show interface

Syntax

`show interface <PORT-NUMBER>`

Displays the friendly port name, if any, along with the traffic statistics for that port. (The friendly port name data comes from the running-config file.)

show config

Syntax

`show config`

Description

Includes friendly port names in the per-port data of the resulting configuration listing. (`show config` data comes from the startup-config file.)

Listing all ports or selected ports with their friendly port names

show name

Syntax

`show name <PORT-LIST>`

Description

Lists the friendly port name with its corresponding port number and port type. The `show name` command without a port list shows this data for all ports on the switch.

Example of friendly port name data for all ports on the switch

```
switch# show name
Port Names
```

```
 Port   Type      Name
 ------ --------- ----------------------------------------------------------
 A1     10GbE-T
 A2     10GbE-T
 A3     10GbE-T   Bill_Smith@10.25.101.73
 A4     10GbE-T
 A5     10GbE-T   Draft-Server:Trunk
 A6     10GbE-T   Draft-Server:Trunk
 A7     10GbE-T   Draft-Server:Trunk
 A8     10GbE-T   Draft-Server:Trunk
```

Example of friendly port name data for specific ports on the switch

```
switch# show name A3-A5

 Port Names

  Port : A3
  Type : 10GbE-T
  Name : Bill_Smith@10.25.101.73
 Port : A4
  Type : 10GbE-T
  Name :
 Port : A5
  Type : 10GbE-T
  Name : Draft-Server:Trunk
```

Including friendly port names in per-port statistics listings

show interface

Syntax

`show interface <PORT-NUMBER>`

Description

Includes the friendly port name with the port's traffic statistics listing. A friendly port name configured to a port is automatically included when you display the port's statistics output.

If you configure port A1 with the name "O'Connor_10.25.101.43," the `show interface` output for this port appears similar to the following:

Usage

```
switch(config)# show interfaces 1
switch(config)# show interface 1 hc
switch(config)# show int queues 1
```

Example `show interfaces 1`

```
switch(config)# show interfaces 1

Status and Counters - Port Counters for port 1

  Name  :
  MAC Address      : 1cc1de-4d8d7f
  Link Status      : Up
  Port Enabled     : <Yes/No>
  Totals (Since boot or last clear) :
```

```
   Bytes Rx         : 218,939                Bytes Tx         : 99,019
   Unicast Rx       : 538                    Unicast Tx       : 290
   Bcast/Mcast Rx   : 1,132                  Bcast/Mcast Tx   : 166
  Errors (Since boot or last clear) :
   FCS Rx           : 0                      Drops Tx         : 0
   Alignment Rx     : 0                      Collisions Tx    : 0
   Runts Rx         : 0                      Late Colln Tx    : 0
   Giants Rx        : 0                      Excessive Colln : 0
   Total Rx Errors : 0                       Deferred Tx      : 0
  Others (Since boot or last clear) :
   Discard Rx       : 0                      Out Queue Len    : 0
   Unknown Protos   : 0
  Rates (5 minute weighted average) :
   Total Rx (bps) : 0                        Total Tx (bps) : 0
   Unicast Rx (Pkts/sec) : 0                 Unicast Tx (Pkts/sec) : 0
   B/Mcast Rx (Pkts/sec) : 0                 B/Mcast Tx (Pkts/sec) : 0
   Utilization Rx  :     0 %                 Utilization Tx  :     0 %
```

Example `show interface 1 hc`

```
switch(config)# show interface 1 hc
Status and Counters - Port Counters for port 1

  Name  :
  MAC Address      : 1cc1de-4d8d7f
  Link Status      : Up
  Port Enabled     : <Yes/No>
  Totals (Since boot or last clear) :
   Bytes Rx         : 221,044                Bytes Tx         : 100,043
   Unicast Rx       : 545                    Unicast Tx       : 294
   Bcast/Mcast Rx   : 1,143                  Bcast/Mcast Tx   : 168
  Errors (Since boot or last clear) :
   FCS Rx           : 0                      Drops Tx         : 0
   Alignment Rx     : 0                      Collisions Tx    : 0
   Runts Rx         : 0                      Late Colln Tx    : 0
   Giants Rx        : 0                      Excessive Colln : 0
   Total Rx Errors : 0                       Deferred Tx      : 0
  Others (Since boot or last clear) :
   Discard Rx       : 0                      Out Queue Len    : 0
   Unknown Protos   : 0
  Rates (5 minute weighted average) :
   Total Rx (bps) : 0                        Total Tx (bps) : 0
   Unicast Rx (Pkts/sec) : 0                 Unicast Tx (Pkts/sec) : 0
   B/Mcast Rx (Pkts/sec) : 0                 B/Mcast Tx (Pkts/sec) : 0
switch(config)#
```

Example `show int queues 1`

```
switch(config)# show int queues 1
 Status and Counters - Port Counters for port 1

  Name  :
  MAC Address       : 1cc1de-4d8d7f
  Link Status       : Down
  Port Enabled      : <Yes/No>
  Port Totals (Since boot or last clear) :
   Rx Packets       : 1,754               Tx Packets       : 482
   Rx Bytes         : 229,698             Tx Bytes         : 103,963
   Rx Drop Packets : 0                    Tx Drop Packets : 0
   Rx Drop Bytes    : 0                   Tx Drop Bytes    : 0
```

```
 Egress Queue Totals (Since boot or last clear) :
     Tx Packets          Dropped Packets    Tx Bytes          Dropped Bytes
Q1  0                   0                  0                 0
Q2  0                   0                  0                 0
Q3  313                 0                  59,681            0
Q4  0                   0                  0                 0
Q5  0                   0                  0                 0
Q6  0                   0                  0                 0
Q7  0                   0                  0                 0
Q8  169                 0                  44,282            0
```

NOTE

For a given port, if a friendly port name does not exist in the running-config file, the Name line in the above command output appears as `Name  :  not assigned`.

Searching the configuration for ports with friendly port names

This option tells you which friendly port names have been saved to the startup-config file. (`show config` does not include ports that have only default settings in the startup-config file.)

show config

Syntax

`show config`

Description

Includes friendly port names in a listing of all interfaces (ports) configured with non-default settings. Excludes ports that have neither a friendly port name nor any other non-default configuration settings.

If you configure port A1 with a friendly port name:

Figure 15: *Listing of the startup-config file with a friendly port name configured*

```
HP Switch(config)# int A1 name Print_Server@10.25.101.43
HP Switch(config)# write mem
HP Switch(config)# int A2 name Herbert's_PC

HP Switch(config)# show config

 Startup configuration:
; J9091A Configuration Editor; Created on release K.15.05.xxxx
hostname "HPSwitch"
 interface AQ
    name "Print_Server@10.25.101.43"
 exit

snmp-server community "public" Unrestricted
.
.
.
```

This command sequence saves the friendly port name for port A1 in the startup-config file. The name entered for port A2 is not saved because it was executed after **write memory**.

Configuring uni-directional link detection

interface link-keepalive

Syntax

`interface <PORT-LIST> link-keepalive`

Description

Enables UDLD on a port or range of ports. Default: UDLD disabled

Parameters and options

no

To disable this feature, enter the `no` form of the command.

link-keepalive interval *<INTERVAL>*

Determines the time interval to send UDLD control packets. The *interval* parameter specifies how often the ports send a UDLD packet. You can specify from 10 to 100, in 100-ms increments, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Default: 50 (5 seconds)

link-keepalive retries *<NUM>*

Determines the maximum number of retries to send UDLD control packets. The `num` parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 to 10.

Default: 5

link-keepalive vlan *<VID>*

Assigns a VLAN ID to a UDLD-enabled port for sending tagged UDLD control packets.Under default settings, untagged UDLD packets can still be transmitted and received on tagged only ports; however, a warning message is logged.

The `no` form of the command disables UDLD on the specified ports.

Default: UDLD packets are untagged; tagged-only ports transmit and receive untagged UDLD control packets

Enabling UDLD

UDLD is enabled on a per-port basis.

When at least one port is UDLD-enabled, the switch will forward out UDLD packets that arrive on non-UDLD-configured ports out of all other non-UDLD configured ports in the same vlan. That is, UDLD control packets will "pass through" a port that is not configured for UDLD. However, UDLD packets will be dropped on any blocked ports that are not configured for UDLD.

Enable UDLD on port a1

```
switch# interface al link-keepalive
```

Enter the appropriate port range to enable the feature on a trunk group

```
switch#interface al-a4 link-keepalive
```

Changing the keepalive interval

By default, ports enabled for UDLD send a link health-check packet once every 5 seconds. You can change the interval to a value from 10 to 100 deciseconds, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Change packet interval to seven seconds

```
switch# link-keepalive interval 70
```

Changing the keepalive retries

By default, a port waits 5 seconds to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 to 10.

Change the maximum number of attempts to four

```
switch# link-keepalive retries 4
```

Configuring UDLD for tagged ports

The default implementation of UDLD sends the UDLD control packets untagged, even across tagged ports. If an untagged UDLD packet is received by a non-Hewlett Packard Enterprise switch, that switch may reject the packet. To avoid such an occurrence, you can configure ports to send out UDLD control packets that are tagged with a specified VLAN.

enable ports to receive and send UDLD control packets tagged with a specific VLAN ID

```
switch# interface llink-keepalive vlan 22
```

- You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.
- If a VLAN ID is not specified, UDLD control packets are sent out of the port as untagged packets.
- To re-assign a VLAN ID, re-enter the command with the new VLAN ID number. The new command overwrites the previous command setting.
- When configuring UDLD for tagged ports, you may receive a warning message if there are any inconsistencies with the VLAN configuration of the port. See **Viewing port status and configuration** on page 102 for potential problems.

Viewing UDLD information

show link-keepalive

Syntax

`show link-keepalive`

Description

Displays all the ports that are enabled for `link-keepalive`.

Parameters

statistics

Displays detailed statistics for the UDLD-enabled ports on the switch.

clear link-keepalive

Syntax

`clear link-keepalive statistics`

Description

Clears UDLD statistics. This command clears the packets sent, packets received, and transitions counters in the `show link-keepalive statistics` display.

Parameters

statistics

Displays detailed statistics for the UDLD-enabled ports on the switch.

Viewing summary information on all UDLD-enabled ports

Enter the `show link-keepalive` command.

show link-keepalive command

Figure 16: *`show link-keepalive`*

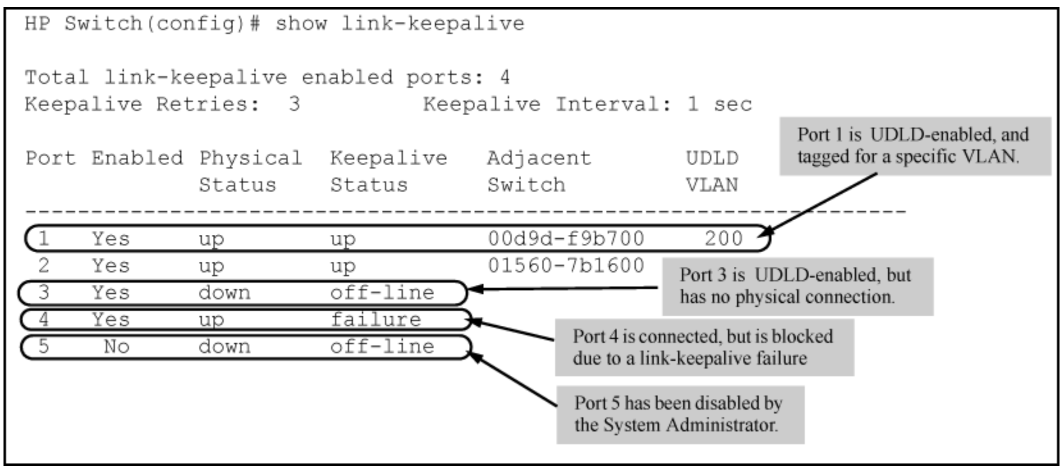

```
HP Switch(config)# show link-keepalive

Total link-keepalive enabled ports: 4
Keepalive Retries:  3         Keepalive Interval: 1 sec

Port Enabled Physical  Keepalive   Adjacent      UDLD
             Status    Status      Switch        VLAN
-------------------------------------------------------------------
 1   Yes     up        up          00d9d-f9b700   200
 2   Yes     up        up          01560-7b1600
 3   Yes     down      off-line
 4   Yes     up        failure
 5    No     down      off-line
```


Viewing detailed UDLD information for specific ports

Enter the `show link-keepalive statistics` command.

show link-keepalive command

Figure 17: *`show link-keepalive statistics`*

Ports 1 and 2 are UDLD-enabled and show the number of health check packets sent and received on each port.

```
HP Switch(config)# show link-keepalive statistics

Port:                    1
Current State:           up          Neighbor MAC Addr:    0000a1-b1c1d1
Udld Packets Sent:       1000        Neighbor Port:        5
Udld Packets Received:   1000        State Transitions:    2
Port Blocking:           no          Link-vlan:            1

Port:                    2
Current State:           up          Neighbor MAC Addr:    000102-030405
Udld Packets Sent:       500         Neighbor Port:        6
Udld Packets Received:   450         State Transitions:    3
Port Blocking:           no          Link-vlan:            200

Port:                    3
Current State:           off line    Neighbor MAC Addr:    n/a
Udld Packets Sent:       0           Neighbor Port:        n/a
Udld Packets Received:   0           State Transitions:    0
Port Blocking:           no          Link-vlan:            1

Port:                    4
Current State:           failure     Neighbor MAC Addr:    n/a
Udld Packets Sent:       128         Neighbor Port:        n/a
Udld Packets Received:   50          State Transitions:    8
Port Blocking:           yes         Link-vlan:            1
```

Port 4 is shown as blocked due to a link-keepalive failure

Port status and Port parameters

Connecting transceivers to fixed-configuration devices

Cause

If the switch either fails to show a link between an installed transceiver and another device or demonstrates errors or other unexpected behavior on the link, check the port configuration on both devices for a speed and/or duplex (mode) mismatch.

- To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface or `show interfaces brief` in the CLI.
- To display information about the transceivers installed on a switch, enter the `show tech receivers` command in the CLI.

Enabled

`Yes` (default): The port is ready for a network connection.

`No`: The port will not operate, even if properly connected in a network. Use this setting, for example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes.

Status (read-only

`Up`: The port senses a link beat.

Mode

The port's speed and duplex (data transfer operation) setting.

`10/100/1000Base-T Ports:`

- `Auto-MDIX` (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI.)

  To see what the switch negotiates for the auto setting, use the CLI `show interfaces brief` command or the `3. Port Status` option under `1. Status and Counters` in the menu interface.
- `MDI`: Sets the port to connect with a PC using a crossover cable (manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- `MDIX`: Sets the port to connect with a PC using a straight-through cable (manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- `Auto-10`: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled.) Hewlett Packard Enterprise recommends auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.)
- `10HDx`:10 Mbps, half-duplex
- `10FDx`: 10 Mbps, full-duplex
- `Auto-100`: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.
- `Auto-10-100`: Allows the port to establish a link with the port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.
- `Auto-1000`: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features.
- `100Hdx`: Uses 100 Mbps, half-duplex.
- `100Fdx`: Uses 100 Mbps, full-duplex

Gigabit Fiber-Optic Ports (Gigabit-SX, Gigabit-LX, and Gigabit-LH):

- `1000FDx`: 1000 Mbps (1 Gbps), full-duplex only
- `Auto` (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port.

Gigabit Copper Ports:

- `1000FDx`: 1000 Mbps (1 Gbps), full-duplex only
- `Auto` (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port.

10-Gigabit CX4 Copper Ports:

- Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed.

10-Gigabit SC Fiber-Optic Ports (10-GbE SR, 10-GbE LR, 10-GbE ER):

- Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed.

NOTE

Conditioning patch cord cables are not supported on 10-GbE.

Auto-MDIX

The switch supports Auto-MDIX on 10Mb, 100Mb, and 1 Gb T/TX (copper) ports. (Fiber ports and 10-gigabit ports do not use this feature.)

- `Automdix`: Configures the port for automatic detection of the cable type (straight-through or crossover.)
- `MDI`: Configures the port to connect to a switch, hub, or other MDI-X device with a straight-through cable.
- `MDIX`: Configures the port to connect to a PC or other MDI device with a straight-through cable.

flow-control

- `Disabled` (default): The port does not generate flow control packets, and drops any flow control packets it receives.
- `Enabled`: The port uses 802.3x link layer flow control, generates flow-control packets, and processes received flow-control packets.

With the port mode set to `Auto` (the default) and flow control enabled, the switch negotiates flow control on the indicated port. If the port mode is not set to `Auto`, or if flow control is disabled on the port, flow control is not used. Note that flow control must be enabled on both ends of a link.

Broadcast limit

Specifies the percentage of the theoretical maximum network bandwidth that can be used for broadcast traffic. Any broadcast traffic exceeding that limit will be dropped. Zero (0) means the feature is disabled.

The broadcast-limit command operates at the port context level to set the broadcast limit for a port on the switch.

NOTE

This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.

Error messages associated with the show interfaces command

`Requesting too many fields (total characters exceeds 80)`

Total length of selected data exceeds one line

`Field name is misspelled`

Invalid input: *input*

`Mistake in specifying the port list`

Module not present for port or invalid port: *input*

`The port list is not specified`

Incomplete input: custom

Using pattern matching with the `show interfaces custom` command

Cause

If you have included a pattern matching command to search for a field in the output of the `show int custom` command, and the `show int custom` command produces an error, the error message may not be visible and the output is empty. For example, if you enter a command that produces an error (such as vlan is misspelled) with the pattern matching `include` option, the output may be empty:

```
switch# show int custom 1-3 name [vlun|include vlan1]
```

Try the `show int custom` command first to ensure there is output, and then enter the command again with the pattern matching option.

You can substitute `int` for `interface`; that is: show `int custom`.

Auto-MDIX configurations

Copper ports on the switch can automatically detect the type of cable configuration (MDI or MDI-X) on a connected device and adjust to operate appropriately.

This means you can use a "straight-through" twisted-pair cable or a "crossover" twisted-pair cable for any of the connections—the port makes the necessary adjustments to accommodate either one for correct operation. The following port types on your switch support the IEEE 802.3ab standard, which includes the "Auto MDI/MDI-X" feature:

- 10/100-TX xl module ports
- 100/1000-T xl module ports
- 10/100/1000-T xl module ports

Using the above ports:

- If you connect a copper port using a straight-through cable on a switch to a port on another switch or hub that uses MDI-X ports, the switch port automatically operates as an MDI port.
- If you connect a copper port using a straight-through cable on a switch to a port on an end node—such as a server or PC—that uses MDI ports, the switch port automatically operates as an MDI-X port.

Switch auto-MDIX supports operation in forced speed and duplex modes.

For more information on this subject, see the IEEE 802.3ab standard reference. For more information on MDI-X, see the installation and getting started guide.

Manual override

If you require control over the MDI/MDI-X feature, you can set the switch to either of these non-default modes:

- Manual MDI
- Manual MDI-X

Table 4: *Cable types for auto and manual MDI/MDI-X settings*

Setting	MDI/MDI-X device type	
	PC or other MDI device type	Switch, hub, or other MDI-X device
Manual MDI	Crossover cable	Straight-through cable
Manual MDI-X	Straight-through cable	Crossover cable
Auto-MDI-X (the default)	Either crossover or straight-through cable	

The AutoMDIX features apply only to copper port switches using twisted-pair copper Ethernet cables.

About using friendly port names

Optional: This feature enables you to assign alphanumeric port names of your choosing to augment automatically assigned numeric port names. This means you can configure meaningful port names to make it easier to identify the source of information listed by some `show` commands. (Note that this feature **augments** port numbering, but **does not replace** it.)

Configuring and operating rules for friendly port names

- At either the global or context configuration level, you can assign a unique name to a port. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the `show name`*<PORT-LIST>*, `show config`, and `show interface` *`port-number`* commands. They do not appear in the output of other `show` commands or in Menu interface screens. (See **Viewing friendly port names with other port data** on page 117.)
- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.
- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.
- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the `write memory` command.)

Uni-directional link detection (UDLD)

Uni-directional link detection (UDLD) monitors a link between two switches and blocks the ports on both ends of the link if the link fails at any point between the two devices. This feature is particularly useful for detecting failures in fiber links and trunks. **Figure 18: UDLD** on page 128 shows an example.

Figure 18: *UDLD*

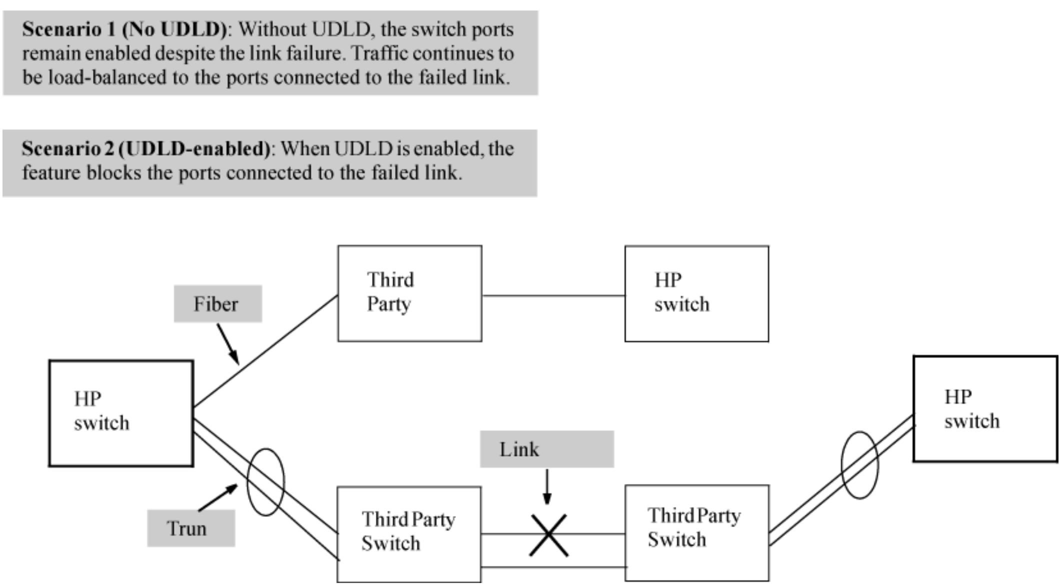

In this example, each switch load balances traffic across two ports in a trunk group. Without the UDLD feature, a link failure on a link that is not directly attached to one of the switches remains undetected. As a result, each switch continue to send traffic on the ports connected to the failed link. When UDLD is enabled on the trunk ports on each switch, the switches detect the failed link, block the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Similarly, UDLD is effective for monitoring fiber optic links that use two uni-direction fibers to transmit and receive packets. Without UDLD, if a fiber breaks in one direction, a fiber port may assume the link is still good (because the other direction is operating normally) and continue to send traffic on the connected ports. UDLD-enabled ports; however, will prevent traffic from being sent across a bad link by blocking the ports in the event that either the individual transmitter or receiver for that connection fails.

Ports enabled for UDLD exchange health-check packets once every five seconds (the link-keepalive interval.) If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for four more intervals. If the port still does not receive a health-check packet after waiting for five intervals, the port concludes that the link has failed and blocks the UDLD-enabled port.

When a port is blocked by UDLD, the event is recorded in the switch log or via an SNMP trap (if configured); and other port blocking protocols, like spanning tree or meshing, will not use the bad link to load balance packets. The port will remain blocked until the link is unplugged, disabled, or fixed. The port can also be unblocked by disabling UDLD on the port.

Configuring UDLD

Consult the release notes and current manuals for required software versions and to determine if your switch model interoperates with UDLD.

When UDLD enabled on at least one port , UDLD packet received on UDLD disabled port will be re-forwarded out on all other UDLD disabled ports on the same VLAN as per the below conditions.

Procedure

1. If the incoming port itself is already blocked on the VLAN it will be dropped right away, and no re-forwarding will be done.
2. UDLD packet will be re-forwarded to other UDLD disabled ports of the same VLAN that are in forwarding state(non blocked ports).

Prerequisites

Procedure

1. When configuring UDLD, keep the following considerations in mind:
   a. UDLD is configured on a per-port basis and must be enabled at both ends of the link. See the note below for a list of switches that support UDLD.
   b. To configure UDLD on a trunk group, you must configure the feature on each port of the group individually. Configuring UDLD on a trunk group's primary port enables the feature on that port only.
   c. Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.

Uplink failure detection

Uplink Failure Detection (UFD) is a network path redundancy feature that works in conjunction with NIC teaming functionality. UFD continuously monitors the link state of the ports configured as links-to-monitor (LtM), and when these ports lose link with their partners, UFD will disable the set of ports configured as links-to-disable (LtD.) When an uplink port goes down, UFD enables the switch to auto-disable the specific downlinks connected to the NICs. This allows the NIC teaming software to detect link failure on the primary NIC port and fail over to the secondary NIC in the team.

NIC teams must be configured for switch redundancy when used with UFD, that is, the team spans ports on both Switch A and Switch B. The switch automatically enables the downlink ports when the uplink returns to service. For an example of teamed NICs in conjunction with UFD, see **Figure 19: Teamed NICs in conjunction with UFD** on page 130.) For an example of teamed NICs with a failed uplink, see **Figure 20: Teamed NICs with a failed uplink** on page 130.

For UFD functionality to work as expected, the NIC teaming must be in Network Fault Tolerance (NFT) mode.

Figure 19: *Teamed NICs in conjunction with UFD*

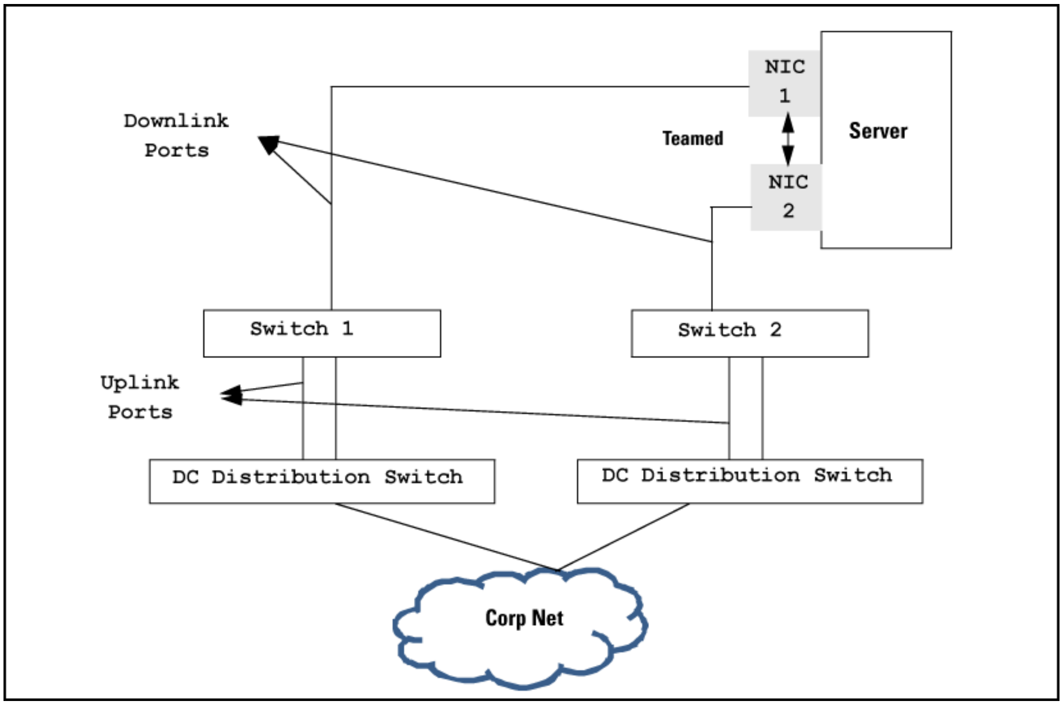

Figure 20: *Teamed NICs with a failed uplink*

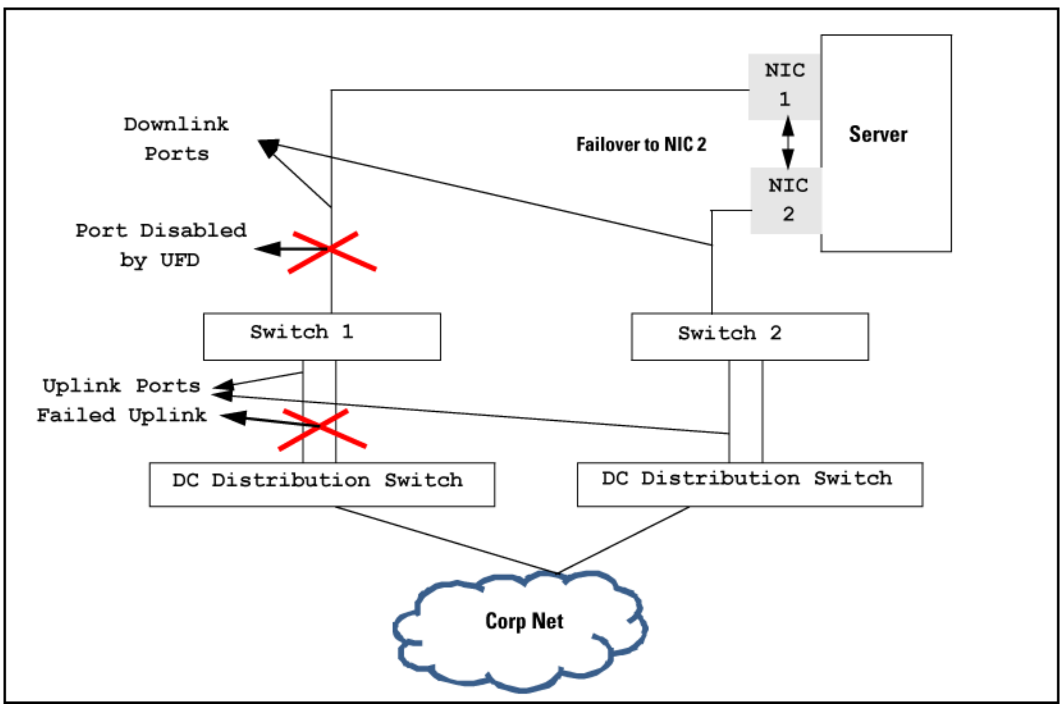

Configuration Guidelines for UFD

Below is a list of configuration guidelines to be followed for UFD. These are applicable only to blade switches where there is a clear distinction between downlink and uplink ports.

1. UFD is required only when uplink-path redundancy is not available on the blade switches.
2. An LtM can be either one or more uplink ports or one or more multi-link trunk group of uplink ports.
3. Ports that are already members of a trunk group are not allowed to be assigned to an LtM or LtD.
4. A trunk group configured as an LtM can contain multiple uplink ports, but no downlink ports or ISL (Inter-Switch-Link) ports.
5. A port cannot be added to a trunk group if it already belongs to an LtM or LtD.
6. An LtD can contain one or more ports, and/or one or more trunks
7. A trunk group configured as an LtD can contain multiple downlink ports, but no uplink ports or ISL (Inter-Switch-Link) ports.

A common API will be provided for higher layers, like CLI and SNMP, which will determine if a port-list can be an LtM or LtD. The API will handle the platform specific details and ensure a uniform code flow for blade and other switch families.

HPE switches do not have a clear distinction between uplink and downlink ports so some of the points listed above may not be applicable.

UFD enable/disable

uplink-failure-detection

Syntax

`uplink-failure-detection`

Description

Used to globally enable UFD. The [no] option globally disables UFD.

UFD track data configuration

uplink-failure-detection-track

syntax

`uplink-failure-detection-track` *`track-id`* `links-to-disable` *`port-list`* `links-to-monitor` *`port-list`*

Description

The above command is used to configure ports given as LtM and ports given as LtD for track-id. This command will accept trunk interfaces as well.

Parameters and options

`no`

Use at the beginning of any parameter to remove tracking or monitoring, respectively.

`ufd track-id`

Use with the <TRACK-ID> option.

From within track-id context:

`links-to-monitor`

Use with the <PORT-LIST> option.

`links-to-disable`

Use with the <PORT-LIST> option.

<TRACK-ID>

Use with ufd track-id parameter

<PORT-LIST>

Use with links-to monitor and links-to-disable parameters within track-id context.

uplink-failure-detection-track

Configure ports 18,19,20 as LtM and ports 1,2,3 as LtD for track-id 10:

```
switch(config)# uplink-failure-detection-track 10 links-to-monitor 18,19,20 links-
to-disable 1,2,3
```

Remove any track data associated with track-id 10.

```
switch(config)# no uplink-failure-detection-track 10
```

Remove port 18 as LtM and port 1 as LtD from track-id 10. This command can be issued from track-id context as well.

```
switch(config)# no uplink-failure-detection-track 10 links-to-monitor 18 links-to-
disable 1
```

UFD minimum uplink threshold configuration

uplink-failure-detection-track

Syntax

`uplink-failure-detection-track` *`track-id`* `minimum-uplink-threshold` *`threshold value`*

Description

Configures the minimum uplink threshold value to a number which is the same as the number of LtM ports that must fail to trigger the disabling of LtD ports. This number of LtM ports must be up to enable the LtD ports if in disable state.

Parameters

failure-count

Specify the number of monitored links that must fail before disabling links-to-disable ports.

all

Set the failure-count equal to the number of links-to-monitor ports configured. Default is `all`.

Options

<NUMBER>

The number of ports to be set as links-to-monitor ports failure count.

Usage

Inside a track-id context:

`monitor-threshold` *`<threshold value>`* `|` *`<all>`*

show uplink-failure-detection

show uplink-failure-detection

Syntax

`show uplink-failure-detection`

uplink failure detection information

```
switch(config)# show uplink-failure-detection

Uplink Failure Detection Information
UFD Enabled    : Yes
Track | Monitored    Links to     LtM    LtD            LtM      LtD
ID    | Links        Disable      State  State          Lacp Key Lacp Key
----- + ------------ ------------ ------ -------------- -------- --------
1     | Dyn1         Dyn2         Up     Up             100      200
2     |                           Down   Auto-Disabled  300      400
3     | 1            D3           Up     Up
10    | 2,3          D4,D5        Down   Auto-Disabled
11    | Trk1         D6           Up     Up
```

UFD operating notes

- A port cannot be added to a trunk group if it already belongs to an LtM or LtD.
- Ports that are already members of a trunk group cannot be assigned to an LtM or LtD.
- Trunks that are configured as LtM or LtD cannot be deleted.

Configuring ports as LtM and LtD for track 3

```
switch# uplink-failure-detection track 3 links-to-monitor 5,6,7
links-to-disable 8,9,10
```

Removing a LtM port and an LtD port for track 3

```
switch# no uplink-failure-detection track 3 links-to-monitor 5
links-to-disable 8
```

Error log

Cause

UFD will log messages in the following scenarios

- Admin status change.
- When an LtM looses link to its partner and as a result number of LtM ports down becomes equal or greater then the LtM failure count,UFD will disable the LtD.
- When an LtM returns to service and as a result the number of LtM ports down becomes lesser than the LtM failure count, UFD auto-enables the LtD.

Invalid port error messages

Cause

- When a user specifies an invalid LtM port, a message similar to the following is displayed.

```
Invalid port(s) specified as links-to-monitor.
```

- When a user specifies an invalid LtD port, a message similar to the following is displayed.

```
Invalid port(s) specified as links-to-disable.
```

- When user specifies an invalid threshold value an error message similar to the following is displayed.

```
Invalid threshold value.
```

- When user tries to configure threshold value greater then number of LtM ports configured an error message similar to the following is displayed.

```
Invalid port(s) specified as links-to-disable.
```

- When a user specifies an invalid LtD port an error message similar to the following is displayed.

```
Invalid port(s) specified as links-to-disable.
```

Port Shutdown with Broadcast Storm

A LAN broadcast storm arises when an excessively high rate of broadcast packets flood the LAN. Occurrence of LAN broadcast storm disrupts traffic and degrades network performance. To prevent LAN traffic from being disrupted, the use of fault-finder commands trigger a port disablement when a broadcast storm is detected. Commands can be used only to support broadcast traffic and not multicast and unicast types of traffic.

The waiting period range for re-enabling ports is 0 to 604800 seconds. The default waiting period to re-enable a port is zero which prevents the port from automatic re-enabling.

Avoid port flapping when choosing the waiting period by considering the time to re-enable carefully.

Configuration Commands

fault-finder broadcast-storm

Syntax

fault-finder broadcast-storm [ethernet] <*PORT-LIST*> action [warn|warn-and-disable <*SECONDS*>][percent <*PERCENT*>|pps <*RATE*>

Descripton

Use the following commands to configure the broadcast-storm on a port.

Parameters and options

ethernet

—

action

—

warn

—

warn-and-disable

—

percent

—

pps

—

<PORT-LIST>

—

<SECONDS>

—

<PERCENT>

—

<RATE>

—

Usage

To remove the current configuration of broadcast-storm on a port, use:

```
no fault-finder broadcast-storm [ethernet] <PORT-LIST>
```

Configuration example 1

```
switch(config)# fault-finder broadcast-storm [ethernet] <AI> action [warn-and-
disable <65535>]<percent 10>
```

Configuration example 2

```
switch(config)# fault-finder broadcast-storm [ethernet] <A2> action [warn-and-
disable] pps <100>
```

Configuration example 3

```
switch(config)# fault-finder broadcast-storm [ethernet] <A22> action [warn] pps
<100>
```

Viewing broadcast-storm configuration

show fault-finder broadcast-storm

Syntax

```
show fault-finder broadcast-storm [ethernet <PORT-LIST>]
```

Description

Display the broadcast-storm-control configuration.

Parameters

broadcast-storm

Configure broadcast storm control.

pps

Rising threshold level in number of broadcast packets per second.

Percent

Rising threshold level as a percentage of bandwidth of the port. The percentage is calculated on 64 byte packet size.

warn

Log the event only.

warn-and-disable

Log the event and disable the port.

seconds

Re-enable the port after waiting for the specified number of seconds. Default is not to re-enable.

show example 1

Port

A1

Bcast storm

Yes

Port status

Down

Rising threshold

10%

Action

warnanddisable

Disable timer

65535

Disable timer left

—

Show example 2

```
switch(config)# show fault-finder broadcast-storm
```

Port

A1

Bcast storm

Yes

Port status

Down

Rising threshold

200 pps

Action

warnanddisable

Disable timer

10

Disable timer left

9

Show example 3

`switch(config)# show fault-finder broadcast-storm A1`

Port

A1

Bcast storm

No

Port status

Up

Rising threshold

—

Action

none

Disable timer

—

Disable timer left

—

Show example 4

`switch(config)# show fault-finder broadcast-storm`

Port

A1

Bcast storm

Yes

Port status

Up

Rising threshold

75%

Action

warn

Disable timer

—

Disable timer left

—

Broadcast-storm event logs

Depending on the configuration of broadcast storm control, several of the following messages can be logged:

- FFI: port <ID>-Administrator action required to re-enable.
- FFI: port <ID>-Excessive Broadcasts. Broadcast-storm control threshold <configured value> percent exceeded.
- FFI: port <ID>-Excessive Broadcasts. Broadcast-storm control threshold <configured value>pps exceeded.
- FFI: port <ID>-Port disabled by Fault-finder.
- ports:Fault-Finder(<FF ID>) has disabled port A1 for 100 Seconds.

The following messages can be logged after the port is enabled:

- ports: port <ID> timer (<FF ID>) has expired.
- ports: port <ID> is now on-line.

Event log

```
l 01/01/90 00:35:20 00025 ip: DEFAULT_VLAN: ip address 10.100.38.231/24 configured
on vlan 1
l 01/01/90 00:35:20 00083 dhcp: updating IP address and subnet mask
l 01/01/90 00:35:05 00076 ports: port A1 is now on-line
l 01/01/90 00:35:02 00900 ports: port A1 timer (71) has expired
W 01/01/90 00:34:13 00026 ip: DEFAULT_VLAN: ip address 10.100.38.231/24 removed
from vlan 1
l 01/01/90 00:34:12 00077 ports: port A1 is now off-line
l 01/01/90 00:34:12 00898 ports:Fault-Finder(71) has disabled port A1 for 5seconds
M 01/01/90 00:34:12 02673 FFI: port A1-Port disabled by Fault-finder.
W 01/01/90 00:34:12 02676 FFI: port A1-Excessive Broadcasts. Broadcast-storm
control threshold 4 percent exceeded.
---- Reverse event Log listing: Events Since Boot ----
I 01/01/90 00:08:44 00025 ip: DEFAULT_VLAN: ip address 10.100.38.231/24 configured
on vlan 1
I 01/01/90 00:08:44 00083 dhcp: updating IP address and subnet mask
I 01/01/90 00:08:11 00076 ports: port A1 is now on-line
I 01/01/90 00:08:08 00900 ports: port A1 timer (71) has expired
W 01/01/90 00:06:29 00026 ip: DEFAULT_VLAN: ip address 10.100.38.231/24 removed
from vlan 1
I 01/01/90 00:06:28 00077 ports: port A1 is now off-line
I 01/01/90 00:06:28 00898 ports:Fault-Finder(71) has disabled port A1 for 100
seconds
M 01/01/90 00:06:28 02673 FFI: port A1-Port disabled by Fault-finder.
W 01/01/90 00:06:28 02675 FFI: port A1-Excessive Broadcasts. Broadcast-storm
control threshold 10 pps exceeded.
```

PoE

PoE technology allows IP telephones, wireless LAN access points, and other appliances to receive power and transfer data over existing ethernet LAN cabling. For more information about PoE technology, see the PoE planning and implementation guide, which is available on the HPE Networking website at

http://www.hpe.com/networking/support.

PoE terminology

Power-over-ethernet (PoE) and Power-over-ethernet plus (PoE+ or POEP) operate similarly in most cases. The CLI commands are the same for a PoE module or a PoE+ zl module. Any differences between PoE and PoE+ operation are noted; otherwise, the term "PoE" is used to designate both PoE and PoE+ functionality.

Planning and implementing a PoE configuration

This section provides an overview of some considerations for planning a PoE application. For additional information on this topic, refer to the PoE planning and implementation guide which is available on the Networking web site at **http://www.hpe.com/networking/support**.

Some of the elements you may want to consider for a PoE installation include:

- Port assignments to VLANs
- Use of security features
- Power requirements

This section can help you to plan your PoE installation. If you use multiple VLANs in your network, or if you have concerns about network security, you should read the first two topics. If your PoE installation comes close to (or is likely to exceed) the system's ability to supply power to all devices that may request it, then you should also read the third topic. (If it is unlikely that your installation will even approach a full utilization of the PoE power available, then you may find it unnecessary to spend much time on calculating PoE power scenarios.)

Power requirements

To get the best PoE performance, you should provide enough PoE power to exceed the maximum amount of power that is needed by all the PDs that are being used.

By connecting an external power supply you can optionally provision more PoE wattage per port and or supply the switch with redundant 12V power to operate should an internal power supply fail.

By installing a second power supply in the 5406zl or a third power supply in a 5412zl chassis, depending on how many PoE ports are being supplied with power, the switch can have redundant power if one power supply fails. A Power Supply Shelf (external power supply) can also be connected to the 5400zl switches to provide extra or redundant PoE power.

For example, if the 5406zl has two 24-port PoE modules (J8702A) installed, and all ports are using 15.4 watts, then the total wattage used is 739.2 watts (48 x 15.4.) To supply the necessary PoE wattage a J8713A power supply is installed in one of the power supply slots.

To gain redundant power, a second J8713A must be installed in the second power supply slot. If the first power supply fails, then the second power supply can supply all necessary power.

See the PoE planning and implementation guide for detailed information about the PoE/PoE+ power requirements.

Assigning PoE ports to VLANs

If your network includes VLANs, you may want to assign various PoE-configured ports to specific VLANs. For example, if you are using PoE telephones in your network, you may want to assign ports used for telephone access to a VLAN reserved for telephone traffic.

Applying security features to PoE configurations

You can utilize security features built into the switch to control device or user access to the network through PoE ports in the same way as non-PoE ports.

Using Port Security, you can configure each switch port with a unique list of MAC addresses for devices that are authorized to access the network through that port. For more information, see the Access security guide for your switch.

Assigning priority policies to PoE traffic

You can use the configurable QoS (Quality of Service) features in the switch to create prioritization policies for traffic moving through PoE ports.

Table 5: *Classifiers for prioritizing outbound packets*

Priority	QoS classifier
1	UDP/TCP application type (port)
2	Device priority (destination or source IP address)
3	IP type of service (ToS) field (IP packets only)
4	VLAN priority
5	Incoming source-port on the switch
6	Incoming 802.1 priority (present in tagged VLAN environments)

For more on this topic, see the advanced traffic management guide.

PoE operation

Using the commands described in this chapter, you can:

- Enable or disable PoE operation on individual ports.
- Monitor PoE status and performance per module.
- Configure a non-default power threshold for SNMP and Event Log reporting of PoE consumption on either all PoE ports on the switch or on all PoE ports in one or more PoE modules.
- Specify the port priority you want to use for provisioning PoE power in the event that the PoE resources become oversubscribed.

Power-sourcing equipment (PSE) detects the power needed by a powered device (PD) before supplying that power, a detection phase referred to as "searching." If the PSE cannot supply the required amount of power, it does not supply any power. For PoE using a Type 1 device, a PSE will not supply any power to a PD unless the PSE has at least 17 watts available. For example, if a PSE has a maximum available power of 382 watts and is already supplying 378 watts, and is then connected to a PD requiring 10 watts, the PSE will not supply power to the PD.

For PoE+ using Type 2 devices, the PSE must have at least 33 watts available. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

PoE configuration options

In the default configuration, PoE support is enabled on the ports in a PoE module installed on the switch. The default priority for all ports is **low** and the default power notification threshold is **80%**. Using the CLI, you can:

- Disable or re-enable PoE operation on individual PoE ports
- Enable support for pre-standard devices
- Change the PoE priority level on individual PoE ports
- Change the threshold for generating a power level notice
- Manually allocate the amount of PoE power for a port by usage, value, or class
- Allocate PoE power based on the link-partner's capabilities via LLDP

The ports support standard networking links and PoE links. You can connect either a non-PoE device or a PD to a port enabled for PoE without reconfiguring the port.

PD support

To best utilize the allocated PoE power, spread your connected PoE devices as evenly as possible across modules. Depending on the amount of power delivered to a PoE module, there may or may not always be enough power available to connect and support PoE operation on all ports in the module. When a new PD connects to a PoE module and the module does not have enough power left for that port, if the new PD connects to a port "X" that has a:

Higher

PoE priority than another port "Y" that is already supporting another PD, the power is removed from port "Y" and delivered to port "X." In this case the PD on port "Y" loses power and the PD on port "X" receives power.

Lower

priority than all other PoE ports currently providing power to PDs, power is not supplied to port "X" until one or more PDs using higher priority ports are removed.

In the default configuration (usage), when a PD connects to a PoE port and begins operating, the port retains only enough PoE power to support the PD's operation. Unused power becomes available for supporting other PD connections. However, if you configure the `poe-allocate-by` option to either value or class, all of the power configured is allocated to the port.

For PoE (not PoE+), while 17 watts must be available for a PoE module on the switch to begin supplying power to a port with a PD connected, 17 watts per port is not continually required if the connected PD requires less power. For example, with 20 watts of PoE power remaining available on a module, you can connect one new PD without losing power to any connected PDs on that module. If that PD draws only 3 watts, 17 watts remain available, and you can connect at least one more PD to that module without interrupting power to any other PoE devices connected to the same module. If the next PD you connect draws 5 watts, only 12 watts remain unused. With only 12 unused watts available, if you then connect yet another PD to a higher-priority PoE port, the lowest-priority port on the module loses PoE power and remains unpowered until the module once again has 17 or more watts available.

For PoE+, there must be 33 watts available for the module to begin supplying power to a port with a PD connected. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

Disconnecting a PD from a PoE port makes that power available to any other PoE ports with PDs waiting for power. If the PD demand for power becomes greater than the PoE power available, power is transferred from the lower-priority ports to the higher-priority ports. (Ports not currently providing power to PDs are not affected.)

PoE power priority

If a PSE can provide power for all connected PD demand, it does not use its power priority settings to allocate power. However, if the PD power demand oversubscribes the available power, the power allocation is prioritized to the ports that present a PD power demand. This causes the loss of power from one or more lower-priority ports to meet the power demand on other, higher-priority ports. This operation occurs regardless of the order in which PDs connect to the module's PoE-enabled ports.

Power allocation is prioritized according to the following methods:

Priority class

Assigns a power priority of **low** (the default), **high**, or **critical**to each enabled PoE port.

Port-number priority

A lower-numbered port has priority over a higher-numbered port within the same configured priority class, for example, port A1 has priority over port A5 if both are configured with **high** priority.

Assigning PoE priority with two or more modules

Ports across two or more modules can be assigned a class priority of low (the default), high, or critical. For example, A5, B7, and C10 could all be assigned a priority class of **Critical**. When power is allocated to the ports on a priority basis, the **Critical** priority power requests are allocated to module A first, then Module B, then C, and so on. Next, the **High** priority power requests are allocated, starting with module A, then B, then C, and the remaining modules in order. Any remaining power is allocated in the same manner for the **Low** priority ports, beginning with module A though the remaining modules. If there is not enough PoE power for all the PDs connected to PoE modules in the switch, power is allocated according to priority class across modules.

All ports on module C are prioritized as Critical.

```
switch# interface c1-c24 power-over-ethernet
    critical
```

All ports on module A are prioritized as Low.

```
switch# interface a1-a24 power-over-ethernet
    low
```

There are 48 PDs attached to all ports of modules A and C (24 ports each module); however, there is enough PoE power for only 32 ports (8.5 watts × 32 ports=273 watts.) The result is that all the **Critical** priority ports on module C receive power, but only 8 ports on module A receive power.

On module A, the port A1 has the highest priority of the ports in that module if all ports are in the same priority class, which is the case for this example. Since a minimum 17 + 5 watts of power is allocated per PoE module for PoE, port A1 will always receive PoE power. If another port on module A had a higher priority class than port A1, that port would be allocated the power before port A1.

For PoE+ modules there must be a minimum of 33 + 5 watts of power allocated per PoE+ module.

About configuring PoE

In the default configuration,PoE support is enabled on the ports in a PoE module installed on the switch. The default priority for all ports is **low** and the default power notification threshold is **80%**.

Using the CLI, you can:

- Disable or re-enable PoE operation on individual PoE ports.
- Enable support for pre-standard devices.
- Change PoE priority level on individual PoE ports.
- Change the threshold for generating a power level notice.
- Manually allocate the amount of PoE power for a port by usage, value, or class.
- Allocate PoE power based on the link-partner's capabilities via LLDP.

For a given level, ports are prioritized by port number in ascending order. For example, if ports A1 to A24 have a priority level of critical, port A1 has priority over ports A2 to A24.

If there is not enough power available to provision all active PoE ports at a given priority level, the lowest-numbered port at that level is provisioned first. For chassis switches, the lowest-numbered port at that level starting with module A, then B, C, and so on is provisioned. PoE priorities are invoked only when all active PoE ports cannot be provisioned (supplied with PoE power.)

In chassis switches, you can use one command to set the same priority level on PoE ports in multiple modules. For example, to configure the priority to **High** for ports c5 to c10, C23 to C24, D1 to D10, and D12, you could use this command:

```
switch# interface c5-c10,c23-c24,d1-d10,d12 power-over-ethernet high
```

PoE priority example

Table 6: *PoE priority operation on a PoE module*

Port	Priority setting	Configuration command and resulting operation with PDs connected to ports C3 through C24
C3 - C17	**Critical**	In this example, the following CLI command sets ports C3 to C17 to **Critical**: `switch# interface c3-c17 power-over-ethernet critical` The **critical** priority class always receives power. If there is not enough power to provision PDs on all ports configured for this class, no power goes to ports configured for **high** and **low** priority. If there is enough power to provision PDs on only some of the critical-priority ports, power is allocated to these ports in ascending order, beginning with the lowest-numbered port in the class, which, in this case, is port 3.
C18 - C21	**high**	In this example, the following CLI command sets ports C19 to C22 to **high**: `switch# interface c19-c22 power-over-ethernet high` The **high** priority class receives power only if all PDs on ports with a **critical** priority setting are receiving power. If there is not enough power to provision PDs on all ports with a **high** priority, no power goes to ports with a low priority. If there is enough power to provision PDs on only some of the high-priority ports, power is allocated to these ports in ascending order, beginning, in this example, with port 18, until all available power is in use.

Table Continued

Port	Priority setting	Configuration command and resulting operation with PDs connected to ports C3 through C24
C22 - C24	**low**	In this example, the CLI command sets ports C23 to C24 to **low**[1]: `switch# interface c23-c24 power-over-ethernet low` This priority class receives power only if all PDs on ports with **high** and **critical** priority settings are receiving power. If there is enough power to provision PDs on only some low- priority ports, power is allocated to the ports in ascending order, beginning with the lowest-numbered port in the class (port 22, in this case), until all available power is in use.
C1 - C2	**N/A**	In this example, the CLI command disables PoE power on ports C1 to C2: `switch# no interface c1-c2 power-over-ethernet` There is no priority setting for the ports in this example.

[1] In the default PoE configuration, the ports are already set to **low** priority. In this case, the command is not necessary.

Disabling or re-enabling PoE port operation

interface

Syntax

```
interface <PORT-LIST> power-over-ethernet
```

Description

Re-enables PoE operation on *<PORT-LIST>* and restores the priority setting in effect when PoE was disabled on *<PORT-LIST>*.

Default: All PoE ports are initially enabled for PoE operation at Low priority. If you configure a higher priority, this priority is retained until you change it.

For PoE, disabling all ports allows the 22 watts of minimum PoE power or the 38 watts for PoE+ power allocated for the module to be recovered and used elsewhere. You must disable ALL ports for this to occur.

Options

no

The `no` form of the command disables PoE operation on *<PORT-LIST>*

<PORT-LIST>

—

Enabling support for pre-standard devices

power-over-ethernet

Syntax

```
power-over-ethernet pre-std-detect
```

Description

Detects and powers pre-802.3af standard devices. The switches covered in this guide also support some pre-802.3af devices. The default setting for the `pre-std-detect` PoE parameter has changed. In earlier software, the default setting is "on." In K.15.02 and later software, the default setting is "off."

Options

no

—

Configuring the PoE port priority

interface

Syntax

```
interface <PORT-LIST> power-over-ethernet [critical|high|low]
```

Description

Reconfigures the PoE priority level on *<PORT-LIST>*. For a given level, ports

Parameters

critical

Specifies the highest-priority PoE support for *<PORT-LIST>*. The active PoE ports at this level are provisioned before the PoE ports at any other level are provisioned.

high

Specifies the second priority PoE support for *<PORT-LIST>*. The active PoE ports at this level are provisioned before the `Low` priority PoE ports are provisioned.

low

(Default) Specifies the third priority PoE support for *<PORT-LIST>*. The active PoE ports at this level are provisioned only if there is power available after provisioning any active PoE ports at the higher priority levels.

Controlling PoE allocation

int

Syntax

```
int <PORT-LIST> poe-allocate-by [usage|class|value]
```

Description

Allows you to manually allocate the amount of PoE power for a port by either its class or a defined value.

The default option for PoE allocation is usage, which is what a PD attached to the port is allocated. You can override this value by specifying the amount of power allocated to a port by using the class or value options.

Parameters

no

usage

(Default) The automatic allocation by a PD. The allowable PD requirements are lower than those specified for PSEs to allow for power losses along the Cat-5 cable.

class

Uses the power ramp-up signature of the PD to identify the device power class.

value

A user-defined level of PoE power allocated for that port.

Table 7: *Power classes and their values*

Power class	Value
0	Depends on cable type and PoE architecture. Maximum power level output of 15.4 watts at the PSE.This is the default class; if there is not enough information about the load for a specific classification, the PSE classifies the load as class 0 (zero.)
1	Requires at least 4 watts at the PSE.
2	Requires at least 7 watts at the PSE.
3	15.4 watts
4	For PoE+Maximum power level output of 30 watts at the PSE.

PoE port allocation by class

To allocate by class for ports 6 to 8:

```
switch# int 6-8 PoE-allocate-by class
```

Manually configuring PoE power levels

You can specify a power level (in watts) allocated for a port by using the `value` option. This is the maximum amount of power that will be delivered.

Procedure

1. To configure a port by value, first set the PoE allocation by entering the `poe-allocate-by value` command:

   ```
   switch(config) # int A6 poe-allocate-by value
   ```

   or in interface context:

   ```
   switch(eth-A6) # poe-allocate-by value
   ```
2. Then select a value:

   ```
   switch(config) # int A6 poe-value 15
   ```

   or in interface context:

```
switch(eth-A6) # poe-value 15
```

3. To view the settings, enter the `show power-over-ethernet` command, shown below.

Figure 21: *Displaying PoE allocation by value and the maximum power delivered*

```
HP Switch(config)# show power-over-ethernet A6

 Status and Counters - Port Power Status for port A6

  Power Enable       : Yes
                                          LLDP Detect       : enabled
  Priority           : low                Configured Type   :
  AllocateBy         : value              Value             : 15 W
  Detection  Status  : Delivering         Power Class       : 2

  Over Current Cnt   : 0                  MPS Absent Cnt    : 0
  Power Denied Cnt   : 0                  Short Cnt         : 0

  Voltage            : 55.1 V             Current           : 154 mA
  Power              : 8.4 W
```

Maximum power delivered

Detection status: fault

Symptom

A **fault** occurs, as shown in Figure **Figure 22: Showing PoE power value set too low for the PD** on page 147.

Figure 22: *Showing PoE power value set too low for the PD*

```
HP Switch(config)# int A7 poe-value 4

HP Switch(config)# show power-over-ethernet A7

 Status and Counters - Port Power Status for port A7

  Power Enable       : Yes
                                          LLDP Detect       : enabled
  Priority           : low                Configured Type   :
  AllocateBy         : value              Value             : 4 W
  Detection  Status  : fault              Power Class       : 2

  Over Current Cnt   : 1                  MPS Absent Cnt    : 0
  Power Denied Cnt   : 2                  Short Cnt         : 0

  Voltage            : 55.1 V             Current           : 154 mA
  Power              : 8.4 W
```

Cause

Setting the PoE maximum value to less than what the PD requires.

Action

1. Increase the PoE maximum value.

Configuring PoE redundancy (chassis switches only)

PoE redundancy occurs automatically when enabled. The switch keeps track of power use and does not supply PoE power to additional PoE devices trying to connect if that results in the switch not having enough power in reserve for redundancy.

power-over-ethernet redundancy

Syntax

`power-over-ethernet rdundancy [n+1|full]`

Description

Allows you to set the amount of power held in reserve for redundancy.

Parameters

`no`	Means that all available power can be allocated to PDs. Default: No PoE redundancy enforced.
`n+1`	One of the power supplies is held in reserve for redundancy. If a single power supply fails, no powered devices are shut down. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.
`full`	Half of the available power supply is held in reserve for redundancy. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.

Changing the threshold for generating a power notice

power-over-ethernet slot

Syntax

`power-over-ethernet [slot` *`<SLOT-ID-RANGE> <threshold 1 - 99>`*`]`

Description

This command configures the notification threshold for PoE power usage on either a global or per-module (slot) basis.

Parameters and options

slot *<SLOT-ID-RANGE>*

Specifies the PoE usage level (as a percentage of the PoE power available on a module) at which the switch generates a power usage notice. This notice appears as an SNMP trap and a corresponding Event Log message and occurs when a PoE module's power consumption crosses the configured threshold value. That is, the switch generates a notice whenever the power consumption on a module either exceeds or drops below the specified percentage of the total PoE power available on the module.

Without the `slot PoE` *`<SLOT-ID-RANGE>`* option, the switch applies one power threshold setting on all PoE modules installed in the switch.

<THRESHOLD 1–99>

—

Enabling or disabling ports for allocating power using LLDP

int poe-lldp-detect

Syntax

```
int <PORT-LIST> poe-lldp-detect [enabled|disabled]
```

Description

Enables or disables ports for allocating PoE power based on the link-partner's capabilities via LLDP.

Default: Enabled

Enable LLDP detection

```
switch(config) # int A7 poe-lldp-detect enabled
```

Interface context

```
switch(eth-A7) # poe-lldp-detect enabled
```

Enabling PoE detection via LLDP TLV advertisement

lldp config

Syntax

```
lldp config <port-number>
```

Description

Configure the lldp for the desired port by number.

Parameter

`basicTlvEnable`

Enables the basic advertised TLV for each port by number.

Options

<management_addr>

Use the option *<management_addr>* to specify specific devices to enable TLV advertisement.

Usage

```
lldp config <port_num> basicTlvEnable <management_addr>
```

Negotiating power using the DLL

When a PD requests power on a PoE port, LLDP interacts with PoE to see if there is enough power to fulfill the request. Power is set at the level requested. If the PD goes into power-saving mode, the power supplied is reduced; if the need for power increases, the amount supplied increases. PoE and LLDP interact to meet the current power demands.

int poe-lldp-detect

Syntax

```
int <PORT-LIST> poe-lldp-detect [enabled|disabled]
```

Description

Allows the data link layer to be used for power negotiation between a PD on a PoE port and LLDP.

Default: Disabled

Enable LLDP

```
switch(config) # int 7 PoE-lldp-detect enabled
```

Interface context

```
switch(eth-7) # PoE-lldp-detect enabled
```

NOTE

Detecting PoE information via LLDP affects only power delivery; it does not affect normal Ethernet connectivity.

Port with LLDP configuration information obtained from the device

```
switch(config)# show power-over-ethernet brief
Status and Counters - Port Power Status
System Power Status     : No redundancy
PoE Power Status        : No redundancy

Available: 273 W  Used: 0 W Remaining: 273 W

Module A Power
Available: 273 W  Used: 0 W Remaining: 273 W

PoE    | Power  Power     Alloc Alloc  Actual Configured Detection   Power
Port   | Enable Priority  By    Power  Power  Type       Status      Class
------  ------ --------- ----- ----- ------ ----------- ----------- ---
Pre-std
Detect
------
A1     | Yes    low       usage  17 W  0.0 W            Searching   0
off
A2     | Yes    low       usage  17 W  0.0 W            Searching   0
off
A3     | Yes    critical usage  17 W  0.0 W            Searching   0
off
A4     | Yes    critical usage  17 W  0.0 W            Searching   0
off
A5     | Yes    critical usage  17 W  0.0 W            Searching   0
off
A6     | Yes    high      usage  17 W  0.0 W            Searching   0
off
A7     | Yes    high      usage  17 W  0.0 W            Searching   0
off
```

```
A8      | Yes    high     usage  17 W  0.0 W             Searching   0
off
A9      | Yes    low      usage  17 W  0.0 W             Searching   0
off
A10     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A11     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A12     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A13     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A14     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A15     | Yes    low      usage  17 W  0.0 W             Searching   0
off
A16     | Yes    low      usage  17 W  0.0 W             Searching   0
```

Figure 23: *Port with LLDP configuration*

```
HPswitch(config)# show power-over-ethernet brief

Status and Counters - Port Power Status

  System Power Status    : No redundancy
  PoE Power Status       : No redundancy

  Available: 300 W  Used: 0 W  Remaining: 300 W

  Module A Power
  Available: 300 W  Used: 5 W  Remaining: 295 W

  PoE    | Power  Power    Alloc Alloc Actual Configured  Detection   Power
  Port   | Enable Priority By    Power Power  Type        Status      Class
  ------ + ------ -------- ----- ----- ------ ----------- ----------- -----
  A1     | Yes    low      usage 17 W  5.0 W  Phone1      Delivering  1
  A2     | Yes    low      usage 17 W  0.0 W              Searching   0
  A3     | Yes    low      usage 17 W  0.0 W              Searching   0
  A4     | Yes    low      usage 17 W  0.0 W              Searching   0
  A5     | Yes    low      usage 17 W  0.0 W              Searching   0
  A6     | Yes    low      usage 17 W  0.0 W              Searching   0
```

Initiating advertisement of PoE+ TLVs

lldp config

Syntax

`lldp config <PORT-LIST>lldp config dot3TlvEnable poe_config`

Description

Enables advertisement of data link layer power using PoE+ TLVs.The TLV is processed only after the physical layer and the data link layer are enabled. The TLV informs the PSE about the actual power required by the device.

Default: Enabled

Temporary PoE+ power drop

Symptom

A temporary PoE+ power drop occurs.

Cause

If LLDP is disabled at runtime, and a PD is using PoE+ power that has been negotiated through LLDP, there is a temporary power drop; the port begins using PoE+ power through the PLC. This event is recorded in the Event Log.When LLDP is enabled again, it causes a temporary power drop. This event is also recorded in the Event Log.

Sample event log messages:

```
W 08/04/10 13:35:50 02768 ports: Port A1 PoE power dropped.
Exceeded physical classification for a PoE Type1 device (LLDP process disabled)
W 08/04/10 13:36:31 02771 ports: Port A1 PoE power dropped.
Exceeded physical classification due to change in classification type (LLDP
process enabled)
```

Viewing PoE when using LLDP information

show lldp config

Syntax

`show lldp config <PORT-LIST>`

Description

Displays the LLDP port configuration information, including the TLVs advertised.

LLDP port configuration information with PoE

Figure **Figure 25: Local power information** on page 153 shows an example of the local device power information using the `show lldp info local-device <PORT-LIST>` command.

Figure 24: *LLDP port configuration information with PoE*

```
HPSwitch(config)# show lldp config 4

 LLDP Port Configuration Detail

  Port : 4
  AdminStatus [Tx_Rx] : Tx_Rx
  NotificationEnabled [False] : False
  Med Topology Trap Enabled [False] : False


  TLVS Advertised:
   * port_descr
   * system_name
   * system_descr
   * system_cap

   * capabilities
   * network_policy
   * location_id
   * poe

   * macphy_config
   * poeplus_config

  IpAddress Advertised:
```

Figure 25: *Local power information*

```
HPswitch(config) show lldp info local-device A1

LLDP Local Port Information Detail

  Port     : A1
  PortType : local
  PortId   : 1
  PortDesc : A1
  Pvid     : 1

Poe Plus Information Detail

    Poe Device Type              : Type2 PSE
    Power Source                 : Primary
    Power Priority               : low
    PD Requested Power Value     : 20 Watts
    PSE Actual Power Value       : 20 Watts
```

Figure **Figure 26: Remote power information** on page 154 shows an example of the remote device power information using the `show lldp info remote-device <PORT-LIST>` command.

Figure 26: *Remote power information*

```
HPswitch(config) show lldp info remote-device A3

LLDP Remote Device Information Detail

  Local Port    : A3
  ChassisType   : mac-address
  ChassisId     : 00 16 35 ff 2d 40
  PortType      : local
  PortId        : 23
  SysName       : HPSwitch
  System Descr  : HP Switch 3500-24, revision K.14.xx
  PortDescr     : 23
  Pvid          : 55

  System Capabilities Supported  : bridge, router
  System Capabilities Enabled    : bridge

  Remote Management Address
     Type     : ipv4
     Address  : 10.0.102.198


  Poe Plus Information Detail

    Poe Device Type              : Type2 PD
    Power Source                 : Only PSE
    Power Priority               : low
    PD Requested Power Value     : 20 Watts
    PSE Actual Power Value       : 20 Watts
```

Viewing the global PoE power status of the switch

show power-over-ethernet

Syntax

`show power-over-ethernet [brief] [ethernet <PORT-LIST>] [slot <SLOT-ID-RANGE>]`

Description

Displays the switch's global PoE power status, including:

- Total Available Power

  Lists the maximum PoE wattage available to provision active PoE ports on the switch. This is the amount of usable power for PDs.
- Total Failover Power

  Lists the amount of PoE power available in the event of a single power supply failure. This is the amount of power the switch can maintain without dropping any PDs.
- Total Redundancy Power

  Indicates the amount of PoE power held in reserve for redundancy in case of a power supply failure.
- Total Remaining Power

  The amount of PoE power still available.

Parameters

brief

Displays PoE information for each port.

Options

<PORT-LIST>

Displays PoE information for the ports in PORT-LIST.

<SLOT-ID-RANGE>

Displays PoE information for the selected slots. Enter the `all` option to display the PoE information for all slots.

Show power-over-ethernet

The command `show power-over-ethernet` displays data similar to that shown in **Figure 27: show power-over-ethernet command output** on page 155.

Figure 27: *`show power-over-ethernet` command output*

```
HP Switch(config)# show power-over-ethernet

 Status and Counters - System Power Status

  Pre-standard Detect     : On
  System Power Status     : No redundancy
  PoE Power Status        : No redundancy

 Chassis power-over-ethernet:

  Total Available Power   :  600 W
  Total Failover Power    :  300 W
  Total Redundancy Power  :    0 W
  Total used Power        :    9 W +/- 6W
  Total Remaining Power   :  591 W

 Internal Power
        1   300W/POE  /Connected.
        2   300W/POE  /Connected.
        3   Not Connected.
        4   Not Connected.
 External Power
        EPS1   /Not Connected.
        EPS2   /Not Connected.
```

Viewing PoE status on all ports

show power-over-ethernet

Syntax

`show power-over-ethernet brief`

Description

Displays the port power status.

- PoE Port

  Lists all PoE-capable ports on the switch.
- Power Enable

  Shows `Yes` for ports enabled to support PoE (the default) and `No` for ports on which PoE is disabled.
- Power Priority

Lists the power priority (Low, High, and Critical) configured on ports enabled for PoE.

- Alloc by

  Displays how PoE is allocated (usage, class, value)
- Alloc Power

  The maximum amount of PoE power allocated for that port (expressed in watts.)Default: 17 watts for PoE; 33 watts for PoE+.
- Actual Power

  The power actually being used on that port.
- Configured Type

  If configured, shows the user-specified identifier for the port. If not configured, this field is empty.
- Detection Status:
  - Searching: The port is trying to detect a PD connection.
  - Delivering: The port is delivering power to a PD.
  - Disabled: On the indicated port, either PoE support is disabled or PoE power is enabled but the PoE module does not have enough power available to supply the port's power needs.
  - Fault: The switch detects a problem with the connected PD.
  - Other Fault: The switch has detected an internal fault that prevents it from supplying power on that port.
- Power Class Shows the 802.3af power class of the PD detected on the indicated port.

Table 8: *Power Classes*

Power class	Description
0	0.44 to 12.95 watts can be drawn by the PD. Default class.
1	0.44 to 3.84 watts
2	3.84 to 6.49 watts
3	6.49 to 12.95 watts
4	For PoE+; up to 25.5 watts can be drawn by the PD

Show power-over-ethernet brief

`show power-over-ethernet brief` displays this output:

Figure 28: *`show power-over-ethernet brief` command output*

```
HP Switch(config)# show power-over-ethernet brief

 Status and Counters - Port Power Status

  System Power Status     : No redundancy
  PoE Power Status        : No redundancy

  Available: 600 W  Used: 9 W  Remaining: 591 W

  Module A Power
  Available: 408 W  Used: 9 W  Remaining: 399 W

  PoE    | Power   Power     Alloc Alloc  Actual Configured  Detection   Power
  Port   | Enable  Priority  By    Power  Power  Type        Status      Class
  ------ + ------- --------- ----- ------ ------ ----------- ----------- ------
  A1     | Yes     low       usage 17 W   0.0 W              Searching   0
  A2     | Yes     low       usage 17 W   0.0 W              Searching   0
  A3     | Yes     low       usage 17 W   0.0 W              Searching   0
  A4     | Yes     low       usage 17 W   0.0 W              Searching   0
  A5     | Yes     low       usage 17 W   0.0 W              Searching   0
  A6     | Yes     low       usage 17 W   8.4 W              Delivering  2
  A7     | Yes     low       usage 17 W   0.0 W              Searching   0
  A8     | Yes     low       usage 17 W   0.0 W              Searching   0
  A9     | Yes     low       usage 17 W   0.0 W              Searching   0
```

You can also show the PoE information by **slot**:

Figure 29: *Showing the PoE information by slot*

```
HP Switch(config)# show power-over-ethernet slot A

 Status and Counters - System Power Status for slot A

  Maximum Power        : 408 W          Operational Status  : On
  Power In Use         :   9 W +/- 6 W  Usage Threshold (%) : 80
```

Viewing the PoE status on specific ports

show power-over-ethernet

Syntax

`show power-over-ethernet <PORT-LIST>`

Description

Displays the following PoE status and statistics (since the last reboot) for each port in *<PORT-LIST>*:

Power Enable	Shows `Yes` for ports enabled to support PoE (the default) and `No` for ports on which PoE is disabled. For ports on which power is disabled, this is the only field displayed by `show power-over-ethernet <PORT-LIST>` .
Priority	Lists the power priority (Low, High, and Critical) configured on ports enabled for PoE.

Table Continued

Allocate by	How PoE is allocated (usage, class, value.)
Detection Status	**Searching** The port is available to support a PD. **Delivering** The port is delivering power to a PD. **Disabled** PoE power is enabled on the port but the PoE module does not have enough power available to supply the port's power needs. **Fault** The switch detects a problem with the connected PD. **Other Fault** The switch has detected an internal fault that prevents it from supplying power on that port.
Over Current Cnt	Shows the number of times a connected PD has attempted to draw more than 15.4 watts for PoE or 24.5 watts for PoE+. Each occurrence generates an Event Log message.
Power Denied Cnt	Shows the number of times PDs requesting power on the port have been denied because of insufficient power available. Each occurrence generates an Event Log message.
Voltage	The total voltage, in volts, being delivered to PDs.
Power	The total power, in watts, being delivered to PDs.
LLDP Detect	Port is enabled or disabled for allocating PoE power, based on the link-partner's capabilities via LLDP.
Configured Type	If configured, shows the user-specified identifier for the port. If not configured, the field is empty.
Value	The maximum amount of PoE power allocated for that port (expressed in watts.) Default: 17 watts for PoE; 33 watts for PoE+

Table Continued

Power Class	Shows the power class of the PD detected on the indicated port. Classes include: **0** 0.44 to 12.95 watts **1** 0.44 to 3.84 watts **2** 3.84 to 6.49 watts **3** 6.49 to 12.95 watts **4** For PoE+; up to 25.5 watts can be drawn by the PD
MPS Absent Cnt	Shows the number of times a detected PD has no longer requested power from the port. Each occurrence generates an Event Log message. ("MPS" refers to the "maintenance power signature.")
Short Cnt	Shows the number of times the switch provided insufficient current to a connected PD.
Current	The total current, in mA, being delivered to PDs.

PoE status of ports

If you want to view the PoE status of ports A6 and A7, you would use **`show power-over-ethernet A6-A7`** to display the data:

Figure 30: *`show power-over-ethernet PORT-LIST` output*

```
HP Switch(config)# show power-over-ethernet A6-A7

 Status and Counters - Port Power Status for port A6

  Power Enable       : Yes
                                       LLDP Detect        : enabled
  Priority           : low             Configured Type    :
  AllocateBy         : value           Value              : 17 W
  Detection  Status  : Delivering      Power Class        : 2

  Over Current Cnt   : 0               MPS Absent Cnt     : 0
  Power Denied Cnt   : 0               Short Cnt          : 0

  Voltage            : 55.1 V          Current            : 154 mA
  Power              : 8.4 W

 Status and Counters - Port Power Status for port A7

  Power Enable       : yes
                                       LLDP Detect        : disabled
  Priority           : low             Configured Type    :
  AllocateBy         : value           Value              : 17 W
  Detection  Status  : Searching       Power Class        : 0

  Over Current Cnt   : 0               MPS Absent Cnt     : 0
  Power Denied Cnt   : 0               Short Cnt          : 0

  Voltage            : 0 V             Current            : 0 mA
  Power              : 0 W
```

Configuring thresholds for generating a power notice

You can configure one of the following thresholds:

A global power threshold that applies to all modules on the switch.

This setting acts as a trigger for sending a notice when the PoE power consumption on any PoE module installed in the switch crosses the configured global threshold level. (Crossing the threshold level in either direction—PoE power usage either increasing or decreasing—triggers the notice.) The default setting is 80%.

A per-slot power threshold that applies to an individual PoE module installed in the designated slot.

This setting acts as a trigger for sending a notice when the module in the specified slot exceeds or goes below a specific level of PoE power consumption.

setting global notification

Suppose slots A, B, and C each have a PoE module installed. In this case, executing the following command sets the global notification threshold to 70% of available PoE power.

```
switch# power-over-ethernet threshold
      70
```

With this setting, if module B is allocated 100 watts of PoE power and is using 68 watts, and then another PD is connected to the module in slot B that uses 8 watts, the 70% threshold of 70 watts is exceeded. The switch sends an SNMP trap and generates this Event Log message:

```
Slot B POE usage has exceeded threshold of 70%.
```

If the switch is configured for debug logging, it also sends the Event Log message to the configured debug destinations.

On any PoE module, if an increasing PoE power load (1) exceeds the configured power threshold—which triggers the log message and SNMP trap—and then (2) later decreases and drops below the threshold again, the switch generates another SNMP trap, plus a message to the Event Log and any configured Debug destinations.

PoE/PoE+ allocation using LLDP

LLDP with PoE

When using PoE, enabling `poe-lldp-detect` allows automatic power configuration if the link partner supports PoE. When LLDP is enabled, the information about the power usage of the PD is available, and the switch can then comply with or ignore this information. You can configure PoE on each port according to the PD (IP phone, wireless device, and so on) specified in the LLDP field. The default configuration is for PoE information to be ignored if detected through LLDP.

Detecting PoE information via LLDP affects only power delivery; it does not affect normal Ethernet connectivity.

LLDP with PoE+

PoE+ with LLDP Overview

The DLC for PoE provides more exact control over the power requirement between a PSE and PD. The DLC works in conjunction with the PLC and is mandatory for any Type-2 PD that requires more than 12.95 watts of input power.

NOTE

DLC is defined as part of the IEEE 802.3at standard.

You can implement the power negotiation between a PSE and a PD at the physical layer or at the data link layer. After the link is powered at the physical layer, the PSE can use LLDP to query the PD repeatedly to discover the power needs of the PD. Communication over the data link layer allows finer control of power allotment, which makes it possible for the PSE to supply dynamically the power levels needed by the PD. Using LLDP is optional for the PSE but mandatory for a Type 2 PD that requires more than 12.95 watts of power.

If the power needed by the PD is not available, that port is shut off.

PoE allocation

There are two ways LLDP negotiates power with a PD:

- Using LLDP MED TLVs

  Disabled by default. Enable using the `int <PORT-LIST> PoE-lldp-detect [enabled|disabled]` command, as shown below.

  LLDP MED TLVs sent by the PD are used to negotiate power only if the LLDP PoE+ TLV is disabled or inactive; if the LLDP PoE+ TLV is sent as well (not likely), the LLDP MED TLV is ignored.
- Using LLDP PoE+ TLVs

  Enabled by default. The LLDP PoE+ TLV is always advertised unless it has been disabled (enable it by using the `lldp config <PORT-LIST> dot3TlvEnable poeplus_config` command.)

  It always takes precedence over the LLDP MED TLV.

Enabling `PoE-lldp-detect` allows the data link layer to be used for power negotiation. When a PD requests power on a PoE port, LLDP interacts with PoE to see if there is enough power to fulfill the request. Power is set at the level requested. If the PD goes into power-saving mode, the power supplied is reduced; if the need for power increases, the amount supplied is increased. PoE and LLDP interact to meet the current power demands.

Operation note

The advertisement of power with TLVs for LLDP PoE+ is enabled by default. If LLDP is disabled at runtime and a PD is using PoE+ power that has been negotiated through LLDP, there will be a temporary power drop. The port will begin using PoE+ power through the PLC. This event is recorded in the event log. An example message would look like the following:

```
W 08/04/10 13:35:50 02768 ports: Port A1 PoE power dropped. Exceeded physical
classification for a PoE Type1 device (LLDP process disabled)
```

When LLDP is enabled again, it causes a temporary power drop. This event is also recorded in the event log. An example message looks like the following:

```
W 08/04/10 13:36:31 02771 ports: Port A1 PoE power dropped. Exceeded physical
classification due to change in classification type (LLDP process enabled)
```

Port trunking overview

Port trunking allows you to assign up to eight physical links to one logical link (trunk) that functions as a single, higher-speed link providing dramatically increased bandwidth. This capability applies to connections between backbone devices as well as to connections in other network areas where traffic bottlenecks exist. A trunk group is a set of up to eight ports configured as members of the same port trunk. The ports in a trunk group do not have to be consecutive. For example:

Figure 31: *Conceptual example of port trunking*

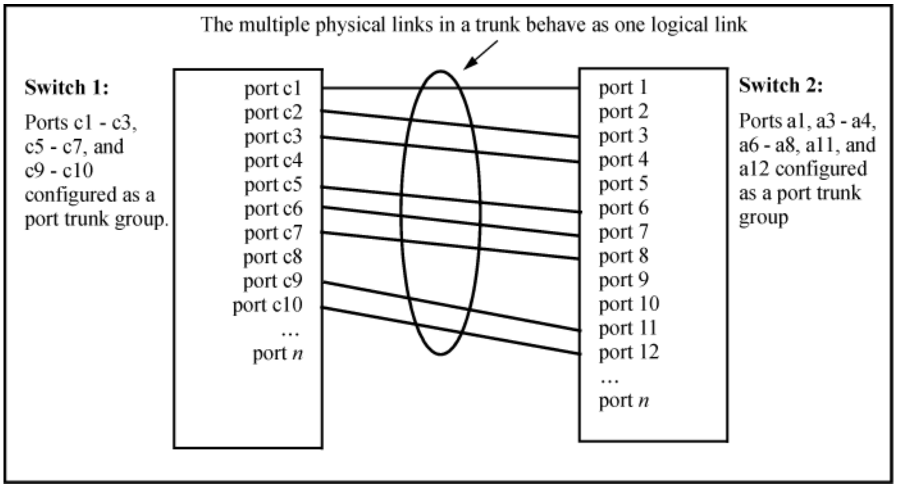

With full-duplex operation in a eight-port trunk group, trunking enables the following bandwidth capabilities:

Port trunk connections and configuration

All port trunk links must be point-to-point connections between a switch and another switch, router, server, or workstation configured for port trunking. No intervening, non-trunking devices are allowed. It is important to note that ports on both ends of a port trunk group must have the same mode (speed and duplex) and flow control settings.

Link connections

The switch does not support port trunking through an intermediate, non-trunking device such as a hub, or using more than onemedia type in a port trunk group. Similarly, for proper trunk operation, all links in the same trunk group must have the same speed, duplex, and flow control.

Port security restriction

Port security does not operate on a trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch resets the port security parameters for those ports to the factory-default configuration.

To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Viewing and configuring port trunk groups

You can list the trunk type and group for all ports on the switch or for selected ports. You can also list LACP-only status information for LACP-configured ports.

To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Viewing static trunk type and group for all ports or for selected ports

show trunks

Syntax

```
show trunks <PORT-LIST>
```

Description

Omitting the *<PORT-LIST>* parameter results in a static trunk data listing for all LAN ports in the switch.

Static trunk group

In a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in **Figure 32: Listing specific ports belonging to static trunks** on page 163 and **Example of a show trunk listing without specifying ports** on page 164 for displaying port data for ports belonging to static trunks.

Using a port list specifies, for switch ports in a static trunk group, only the ports you want to view. In this case, the command specifies ports A5 through A7. However, because port A6 is not in a static trunk group, it does not appear in the resulting listing:

Figure 32: *Listing specific ports belonging to static trunks*

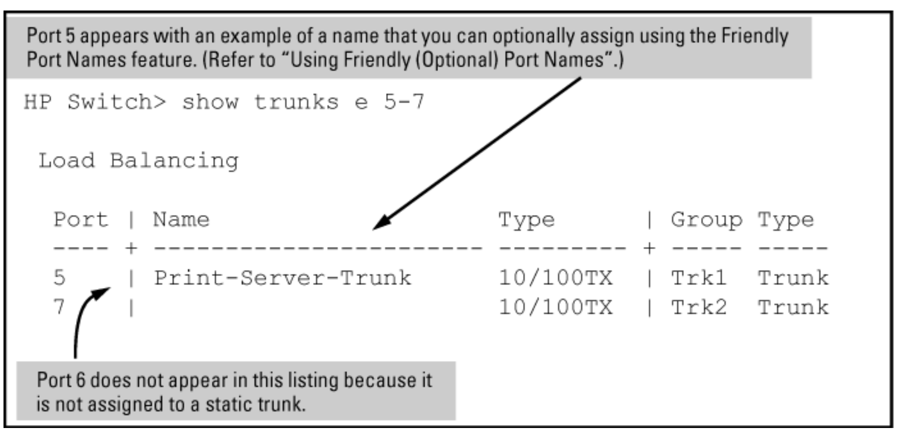

The `show trunks <PORT-LIST>` command in the above example includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In **Example of a show trunk listing without specifying ports** on page 164, the command does not include a port list, so the switch lists all ports having static trunk membership.

Example of a show trunk listing without specifying ports

```
switch# show trunks

 Load Balancing

 Port | Name                    Type      | Group Type
 ---- + ----------------------- --------- + ----- -----
 4    | Print-Server-Trunk      10/100TX  | Trk1  Trunk
 5    | Print-Server-Trunk      10/100TX  | Trk1  Trunk
 7    |                         10/100TX  | Trk2  Trunk
 8    |                         10/100TX  | Trk2  Trunk
```

Viewing static LACP and dynamic LACP trunk data

show lacp

Syntax

`show lacp`

Description

Lists data for only the LACP-configured ports.

Example of a show LACP listing

Ports A1 and A2 have been previously configured for a static LACP trunk.

```
switch# show lacp

       LACP     Trunk    Port     LACP     LACP     Admin  Oper
Port   Enabled  Group    Status   Partner  Status   Key    Key
----   -------  -------  -------  -------  -------  ------ ------
A1     Active   Trkl     Up       Yes      Success  0      250
A2     Active   Trkl     Up       Yes      Success  0      250
A3     Active   A3       Down     No       Success  0      300
A4     Passive  A4       Down     No       Success  0      0
A5     Passive  A5       Down     No       Success  0      0
A6     Passive  A6       Down     No       Success  0      0
```

Configuring a static trunk or static LACP trunk group

Configure port trunking before you connect the trunked links between switches. Otherwise, a broadcast storm could occur. If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured.

trunk

Syntax

trunk *<PORT-LIST>* *<trk1 | trk2 | ..... trkN>* [trunk | lacp | dt-lacp | dt-trunk]

Description

Configures the specified static trunk type.

Static trunk group

This example uses ports C4 to C6 to create a non-protocol static trunk group with the group name `Trk2`.

```
switch# trunk c4-c6 trk2 trunk
```

Removing ports from a static trunk group

Removing a port from a trunk can create a loop and cause a broadcast storm. When you remove a port from a trunk where spanning tree is not in use, Switch recommends that you first disable the port or disconnect the link on that port.

no trunk

Syntax

`no trunk <PORT-LIST>`

Description

Removes the specified ports from an existing trunk group.

remove ports from an existing trunk group

To remove ports C4 and C5 from an existing trunk group:

```
switch# no trunk c4-c5
```

Enabling dynamic LACP trunk groups

An individual trunk can have up to eight links, with additional standby links if you are using LACP.

interface lacp active

Syntax

`interface <PORT-LIST> lacp active`

Description

Configure trunk group types. Configures *<PORT-LIST>* as LACP active. If the ports at the other end of the links on *<PORT-LIST>* are configured as LACP passive, this command enables a dynamic LACP trunk group on *<PORT-LIST>*.

Enable a dynamic LACP trunk group

This example uses ports C4 and C5 to enable a dynamic LACP trunk group.

```
switch# interface c4-c5 lacp active
```

Remove ports from a dynamic LACP trunk group

To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP `Active` and LACP `passive` without first removing LACP operation from the port.)

> Unless spanning tree is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where spanning tree is not in use, Hewlett Packard Enterprise recommends that you first disable the port or disconnect the link on that port.

no interface lacp

Syntax

`no interface <PORT-LIST> lacp`

Description

Removes *<PORT-LIST>* from any dynamic LACP trunk and returns the ports in *<PORT-LIST>* to passive LACP.

Using no interface lacp

Port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, do the following:

```
switch# no interface c6 lacp
switch# interface c6 lacp passive
```

If the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Set the LACP key

During dynamic link aggregation using LACP, ports with the same key are aggregated as a single trunk.

lacp

Syntax

`lacp [active|passive|key 0-65535]`

Enable LACP and configure an LACP key

```
switch# int A2-A3 lacp active
switch# int A2-A3 lacp key 500

switch# show lacp

                                LACP

         LACP      Trunk    Port               LACP     Admin  Oper
  Port   Enabled   Group    Status   Partner   Status   Key    Key
  ----   -------   -------  -------  -------   -------  ------ ------
  A2     Active    A2       Down     No        Success  500    500
  A3     Active    A3       Down     No        Success  500    500
```

Interface configured with a different LACP key

```
switch# int A5 lacp active
switch# int A5 lacp key 250

switch# show lacp

LACP

       LACP      Trunk     Port                LACP      Admin   Oper
Port   Enabled   Group     Status    Partner   Status    Key     Key
----   -------   -------   -------   -------   -------   ------  ------
A1     Active    Dyn1      Up        Yes       Success   100     100
A2     Active    Dyn1      Up        Yes       Success   100     100
A3     Active    Dyn1      Up        Yes       Success   100     100
A4     Active    Dyn1      Up        Yes       Success   100     100
A5     Active    A5        Up        No        Success   250     250
```

Viewing and configuring a static trunk group (Menu)

Prerequisites

To avoid a broadcast storm, configure port trunking **before** you connect the trunked links to another switch, routing switch, or server.(If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Enabling or Disabling Ports and Configuring Port_Mode".)

This procedure uses the Port/Trunk Settings screen to configure a static port trunk group on the switch.

Procedure

1. Review the Prerequisites.
2. From the Main Menu, select **Switch Configuration ...** , and then select **Port/Trunk Settings**.
3. On the keyboard, press **[E]** (for `Edit`), and then use the arrow keys to access the port trunk parameters.

   Figure 33: *Menu screen for configuring a port trunk group*

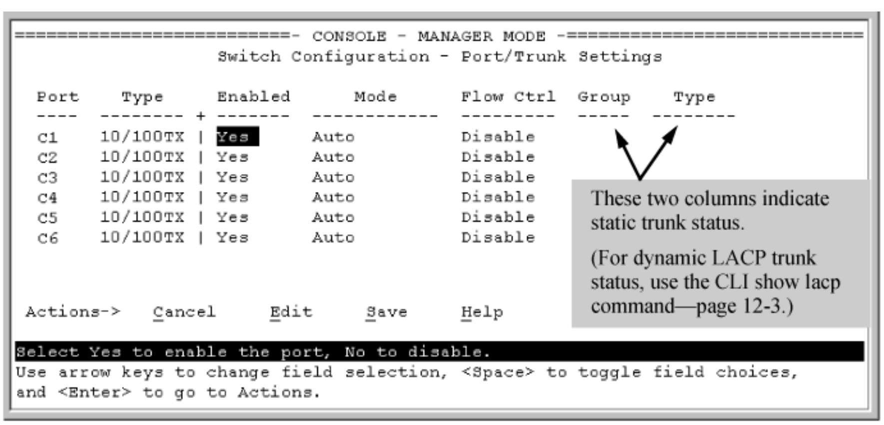

4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group assignment (`Trk1`, `Trk2`, and so on) for the selected port.
   a. For proper trunk operation, all ports in a trunk must have the same media type and mode (such as 10/100TX set to 100FDx, or 100FX set to 100FDx.) The flow control settings must also be the same for all ports in a given trunk.

b. You can configure the trunk group with up to eight ports per trunk. If multiple VLANs are configured, all ports within a trunk will be assigned to the same VLAN or set of VLANs. (With the 802.1Q VLAN capability built into the switch, more than one VLAN can be assigned to a trunk. See the *Advanced Traffic Management Guide*.

(To return a port to a non-trunk status, keep pressing the Space bar until a blank appears in the highlighted Group value for that port.)

Figure 34: *Configuration for a Two-Port Trunk Group*

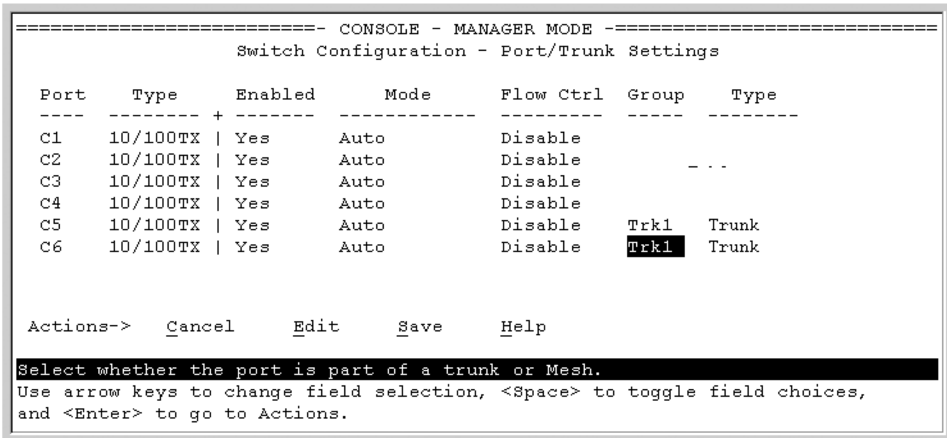

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled      Mode        Flow Ctrl  Group    Type
  ----  --------- + -------  ------------  ---------  -----  --------
  C1    10/100TX | Yes      Auto          Disable
  C2    10/100TX | Yes      Auto          Disable
  C3    10/100TX | Yes      Auto          Disable
  C4    10/100TX | Yes      Auto          Disable
  C5    10/100TX | Yes      Auto          Disable    Trk1   Trunk
  C6    10/100TX | Yes      Auto          Disable    Trk1   Trunk


 Actions->   Cancel     Edit     Save     Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

6. Move the cursor to the Type column for the selected port and use the Space bar to select the trunk type:

   a. LACP

   b. Trunk (the default type if you do not specify a type)

   All ports in the same trunk group on the same switch must have the same Type (`LACP` or `Trunk`.)

7. When you are finished assigning ports to the trunk group, press **[Enter]**, then **[S]** (for `Save`) and return to the Main Menu. (It is not necessary to reboot the switch.)

   During the Save process, traffic on the ports configured for trunking is delayed for several seconds. If the Spanning Tree Protocol is enabled, the delay may be up to 30 seconds.

8. Connect the trunked ports on the switch to the corresponding ports on the opposite device. If you previously disabled any of the trunked ports on the switch, enable them now. (See "Viewing Port Status and Configuring Port Parameters")

Check the Event Log to verify that the trunked ports are operating properly.

Enable L4-based trunk load balancing

Use this command with the `L4-based` option to enable load balancing on Layer 4 information when it is present.

trunk-load-balance

Syntax

`trunk-load-balance [L3-based|L4-based]`

Description

When the `L4-based` option is configured, enables load balancing based on Layer 4 information if it is present. If it is not present, Layer 3 information is used if present; if Layer 3 information is not present, Layer 2 information is used. The configuration is executed in global configuration context and applies to the entire switch.

Defaults to L3-based load balancing.

Parameters

L3–based

Load balance on Layer 3 information if present, or Layer 2 information.

L4–based

Load balance on Layer 4 port information if present, or Layer 3 if present, or Layer 2.

Enabling L4-based trunk load balancing

Figure 35: *Enabling L4-based trunk load balancing*

```
HPswitch(config)# trunk-load-balance L4-based
```

Figure 36: *Output when L4-based trunk load balancing is enabled*

```
HPswitch(config)# show trunk

Load Balancing Method: L4-based

  Port | Name                             Type      | Group  Type
  ---- + -------------------------------- --------- + ------ --------
  41                                      100/1000T   Trk1   Trunk
  42                                      100/1000T   Trk1   Trunk
```

Figure 37: *Running config file when L4-based trunk load balancing is enabled*

```
HP Switch(config) # show running-config

Running configuration:

; J9091A Configuration Editor; Created on release #K.15.02.0001x

hostname "Switch"
module 1 type J8702A
module 5 type J9051A
module 7 type J8705A
module 10 type J8708A
module 12 type J8702A
trunk-load-balance L4-based
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,G1-G24,J1-J4,L1-L24
   ip address dhcp-bootp
   tagged EUP
   no untagged EDP
   exit
snmp-server community "public" unrestricted
```

If L4 trunk load balancing is enabled, a line appears in the running-config file. If it is not enabled, nothing appears as this is the default and the default values are not displayed.

Viewing trunk load balancing

The `show trunks load-balance interface` command displays the port on which the information will be forwarded out for the specified traffic flow with the specified source and destination address.

show trunks

Syntax

```
show trunks load-balance interface <trunk-id> mac <src-addr> <dest-addr> [ip <src-addr> <dest-addr> <src-tcp-port> <src-upd-port> <dest-tcp-port> <dest-udp-port> inbound-port <port-num>
```

Description

Displays the port on which the information will be forwarded out for the specified traffic flow with the specified source and destination address.

Options

trunk-id

The trunk id (trk1, trk2, etc.)

mac *src-addr dest-addr*

The source MAC address and the destination MAC address.

ip *src-addr dest-addr*

The source IPv4 /IPv6 address and the destination IPv4/IPv6 address.

[*src-tcp-port*|*src-udp-port*]

The source TCP port or the source UDP port.

[*dest-tcp-port*|*dest-udp-port*]

The destination TCP port or the destination UDP port.

inbound-port *port-num*

the port number of which the traffic is received.

Information about the forwarding port

```
switch# show trunks load-balance interface trk1 mac 424521-498421 534516-
795463 inbound-port a5
Traffic in this flow will be forwarded out port 23 based on the confiugred L2
load balancing.
```

Operating notes

The port cannot be determined if:

- All the ports in the trunk are down.
- The MAC address is all zeros.
- The source MAC address is broadcast or multicast.

Distributed trunking

Configure ISC ports

Configure the ISC ports before you configure the trunks for distributed trunking.

switch-interconnect

Syntax

```
[no] switch-interconnect <PORT-LIST>|<TRK1|TRK2|...TRKn>
```

Configures an InterSwitch-Connection (ISC) port. Override an ISC configuration by configuring the command with a different value.

A port that is already part of a trunk cannot be configured as an ISC interface.

Parameters

no

Removes the ISC interface configuration.

<PORT-LIST>|<trk1|trk2|...trkN>

The interconnect interface that connects two distributed trunking switches. It can be a physical port, manual LACP trunk, or manual non-protocol trunk.

Configuring distributed trunking ports

Distributed trunking ports must be configured manually.

trunk

Syntax

`trunk <PORT-LIST> <trk1|trk2|...trkN> trunk <PORT-LIST>|lacp | dt-lacp | dt-trunk`

Description

Configures distributed trunking on a switch. Use either the `dt-lacp` or `dt-trunk` option.

The trunk groups and trunk types must be identical in both switches. For example, if Switch Local is configured with `trk1` and uses the `dt-lacp` option, Switch Remote also must be configured with `trk1` and use the `dt-lacp` option to form a distributed trunk. Similarly, if Switch Local is configured with `trk2` and uses the `dt-trunk` option, Switch Remote must be configured with `trk2` and use the `dt-trunk` option to form the distributed trunk.

DT requires that the platforms at both ends of the DT-link be the same and running the same software version.

Parameters

no

The `no` form of the command removes the distributed trunking configuration on the switch.

ISC port configuration

Figure 38: Configuring distributed trunking on page 171 shows an ISC port being configured for the local switch and the remote switch.

Figure 38: *Configuring distributed trunking*

```
HP Switch Local(config)# switch-interconnect a7
HP Switch Remote(config)# switch-interconnect a8

HP Switch Local(config)# trunk a9-a10 trk10 dt-lacp
HP Switch Remote(config)# trunk a5-a6 trk10 dt-lacp

HP Switch Local(config)# trunk a1-a2 trk20 dt-trunk
HP Switch Remote(config)# trunk a3-a4 trk20 dt-trunk
```

Configuring peer-keepalive links

distributed-trunking

Syntax

```
distributed-trunking [hold-timer3-10|peer-keepalive <DESTINATION> ip-address|vlan
<VID> [interval <400-10000>][timeout <3-20>] [udp-port <1024-49151>]
```

Description

Distributed trunking uses a VLAN interface between DT peers to transmit periodic peer-keepalive messages. This command configures the peer-keepalive parameters for distributed trunking.

Parameters and options

`no`

The `no` form of the command removes the distributed trunking configuration on the switch.

`hold-timer`

Configures the hold time in seconds. The range is 3–10 seconds, and defaults to 3.

`peer-keepalive`

—

`<DESTINATION>`

The destination IPv4 address to be used by DT switches to send peer-keepalive messages to the peer DT switch when the ISC is down.

`vlan` ***`<VID>`***

The VID of the VLAN used exclusively for sending and receiving peer-keepalive messages.

`interval`

The interval between peer-keepalive messages (in milliseconds), in the range of 400–10000 milliseconds. Defaults to 1000 milliseconds.

`timeout`

The peer-keepalive timeout in seconds, in the range of 3–10 seconds. Defaults to 5 seconds.

`udp-port`

The source UDP port to be used for transmitting peer-keepalive HELLO messages, in the range of 1024–49151.

Viewing distributed trunking information

show lacp distributed

Syntax

```
show lacp disributed
```

Description

Displays information about distributed trunks and LACP status.

Example

```
HP Switch Local (config#): show lacp distributed
                               Distributed LACP
Local Port Status:

         LACP        Trunk      Port       LACP        LACP       Admin     Oper
```

```
Port     Enabled     Group      Status     Partner     Status    Key       Key
-----    ---------   -------   ---------   ---------   ------   -------   ---
A9       Active      Trk10       Up          Yes        Sucess    350      350
A10      Active      Trk10       Up          Yes        Sucess    350      350

Remote Port Status
         LACP        Trunk      Port       LACP        LACP      Oper
Port     Enabled     Group      Status     Partner     Status    Key
-----    ---------   -------   ---------   ---------   ------   -----
A5       Active      Trk10      Up          Yes        Sucess     200
A6       Active      Trk10      Up          Yes        Sucess     200
```

show distributed-trunk

Syntax

```
show distributed-trunk consistency-parameters global
```

Description

This command displays configured features on VLANs that have dt - lacp or dt - trunk ports as member port. This command also displays VLAN memberships and loop - protect status of a given DT trunk. You can use this command to determine if there is any mismatch in the configuration parameters on VLANs configured for DT ports or on DT interfaces.

show distributed-trunk

```
show distributed-trunk consistency-parameters global

                              Local              Peer
                              -----              -----
Image Version                 K.15.XX            K.15.XX
IP Routing                    Enabled            Enabled
Peer-keepalive interval (ms)  1000               1000

IGMP enabled VLANs on Local : 1-10, 100-110, 501 ,600
610 ,800
IGMP enabled VLANs on Peer : 1-10, 100-110, 501 ,600

DHCP-snooping enabled VLANs on Local : 1,2
DHCP-snooping enabled VLANs on Peer : 1

Loop-protect enabled VLANs on Local : 1,4
Loop-protect enabled VLANs on Peer : 1,5

MLD enabled VLANs on Local : 1-10
MLD enabled VLANs on Peer : 1-10
```

Example

```
Show distributed-trunk
consistency-parameters trunk <trk1...trkN>
Allowed VLANs on Local : 1-10, 100-110, 501 ,600
610 ,800
Allowed VLANs on Peer : 1-10, 100-110, 501 ,600
610 ,800

Name            Local Value             Peer Value
---------------------- --------------------------------- ------------------------
Loop-protect     Enabled                  Enabled
```

Viewing peer-keepalive configuration

Viewing switch interconnect

Syntax

`show switch-interconnect`

Description

Displays information about switch interconnect settings.

Figure 39: *Switch-interconnect settings*

```
HPSwitch(config)# show switch-interconnect

Port           :Trk2
Status         :Up
Active VLANs   :2,3,4,30
```

Port trunk operations

The switches covered in this guide offer these options for port trunking:

- LACP: IEEE 802.3ad—
- Trunk: Non-Protocol—

Up to 144 trunk groups are supported on the switches. The actual maximum depends on the number of ports available on the switch and the number of links in each trunk. (Using the link aggregation control protocol—LACP—option, you can include standby trunked ports in addition to the maximum of eight actively trunking ports.) The trunks do not have to be the same size; for example, 100 two-port trunks and 11 eight-port trunks are supported.

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, and so on) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, Switch recommends that you leave the port Mode settings at `Auto` (the default.) LACP also operates with `Auto-10`, `Auto-100`, and `Auto-1000 (if negotiation selects FDx)`, and `10FDx`, `100FDx`, and `1000FDx` settings. (The 10-gigabit ports available for some switch models allow only the `Auto` setting.)

Fault tolerance

If a link in a port trunk fails, the switch redistributes traffic originally destined for that link to the remaining links in the trunk. The trunk remains operable as long as there is at least one link in operation. If a link is restored, that link is automatically included in the traffic distribution again. The LACP option also offers a standby link capability, which enables you to keep links in reserve for service if one or more of the original active links fails. (See **Trunk group operation using LACP** on page 182.)

Trunk configuration methods

Dynamic LACP trunk

The switch automatically negotiates trunked links between LACP-configured ports on separate devices, and offers one dynamic trunk option: LACP. To configure the switch to initiate a dynamic LACP trunk with another device, use the `interface` command in the CLI to set the default LACP option to `active` on the ports you want to use for the trunk. For example, the following command sets ports C1 to C4 to `LACP active`:

```
switch(config) int c1-c4 lacp active
```

The preceding example works if the ports are not already operating in a trunk. To change the LACP option on ports already operating as a trunk, you must first remove them from the trunk. For example, if ports C1 to C4 are LACP-active and operating in a trunk with another device, you would do the following to configure them to LACP-passive:

```
switch# no int c1-c4 lacp
```

Removes the ports from the trunk:

```
switch# int c1-c4 lacp passive
```

Dynamic LACP Standby Links

Dynamic LACP trunking enables you to configure standby links for a trunk by including more than eight ports in a dynamic LACP trunk configuration. When eight ports (trunk links) are up, the remaining link(s) will be held in standby status. If a trunked link that is “Up” fails, it will be replaced by a standby link, which maintains your intended bandwidth for the trunk. In this example, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows Standby status, while the remaining eight links are “Up”.

```
switch# show lacp
                              LACP
       LACP      Trunk    Port                LACP    Admin    Oper
Port   Enabled   Group    Status    Partner   Status  Key      Key
----   -------   -----    ------    -------   ------  ----     -----
A1     Active     Dyn1     Up         Yes     Success  100      100
A2     Active     Dyn1     Up         Yes     Success  100      100
A3     Active     Dyn1     Up         Yes     Success  100      100
A4     Active     Dyn1     Up         Yes     Success  100      100
A5     Active     Dyn1     Up         Yes     Success  100      100
A6     Active     Dyn1     Up         Yes     Success  100      100
A7     Active     Dyn1     Up         Yes     Success  100      100
A8     Active     Dyn1     Up         Yes     Success  100      100
A9     Active     Dyn1    Standby     Yes     Success  100      100
```

Viewing LACP Local Information

```
switch# show lacp local
LACP Local Information.
System ID: 001871-b98500
               LACP                Tx      Rx Timer
Port   Trunk   Mode   Aggregated   Timer   Expired
----- ------- ------- ----------- ------ ---------
A2       A2    Active    Yes        Fast    No
A3       A3    Active    Yes        Fast    No
```

Viewing LACP Peer Information

Use the `show lacp peer` command to display information about LACP peers. The System ID represents the MAC address of a partner switch. It will be zero if a partner is not found.

```
switch# show lacp peer
LACP Peer Information.
System ID: 001871-b98500
Local  Local                         Port       Oper     LACP     Tx
Port   Trunk   System ID       Port  Priority   Key      Mode     Timer
------ ------ -------------- ----- --------- ------- -------- -----
```

```
A2     A2     123456-654321   2        0        100     Passive   Fast
A3     A3     234567-456789   3        0        100     Passive   Fast
```

Viewing LACP Counters

Use the show lacp counters command to display statistical information about LACP ports.

Note on the Marker Protocol. Data traffic can be dynamically redistributed in port channels. This may occur when a link is added or removed, or there is a change in load-balancing. Traffic that is redistributed in the middle of a traffic flow could potentially cause mis-ordered data packets.

LACP uses the marker protocol to prevent data packets from being duplicated or reordered due to redistribution. Marker PDUs are sent on each port-channel link. The remote system responds to the marker PDU by sending a marker responder when it has received all the frames received on this link prior to the marker PDU. When the marker responders are received by the local system on all member links of the port channel, the local system can redistribute the packets in the traffic flow correctly.

For the switches covered in this guide, the marker BPDUs are not initiated, only forwarded when received, resulting in the Marker fields in the output usually displaying zeros.

```
switch# show lacp counters
LACP Port Counters.
               LACP      LACP      Marker    Marker    Marker     Marker
Port   Trunk   PDUs Tx   PDUs Rx   Req. Tx   Req. Rx   Resp. Tx   Resp. Rx   Error
----   ------  -------   -------   -------   -------   --------   --------   -----
A2       A2    1234      1234       0          0         0         0          0
A3       A3    1234      1234      0          0         0         0          0
```

Using keys to control dynamic LACP trunk configuration

The `lacp key` option provides the ability to control dynamic trunk configuration. Ports with the same key will be aggregated as a single trunk.

There are two types of keys associated with each port, the Admin key and the Operational key. The Operational key is the key currently in use. The Admin key is used internally to modify the value of the Operational key. The Admin and Operational key are usually the same, but using static LACP can alter the Operational key during runtime, in which case the keys would differ.

The `lacp key` command configures both the Admin and Operational keys when using dynamic LACP trunks. It only configures the Admin key if the trunk is a static LACP trunk. It is executed in the interface context.

Static trunk

The switch uses the links you configure with the Port/Trunk Settings screen in the menu interface or the `trunk` command in the CLI to create a static port trunk. The switch offers two types of static trunks: LACP and Trunk.

Table 9: *Trunk types used in static and dynamic trunk groups*

Trunking method	LACP	Trunk
Dynamic	Yes	No
Static	Yes	Yes

Table 10: *Trunking options for LACP and Trunk protocols*

<table>
<tr><th>Protocol</th><th>Trunking Options</th></tr>
<tr><td>LACP (802.3ad)</td><td>Provides dynamic and static LACP trunking options.
Dynamic LACP — Use the switch-negotiated dynamic LACP trunk when:
The port on the other end of the trunk link is configured for Active or Passive LACP.
You want fault-tolerance for high-availability applications. If you use an eight-link trunk, you can also configure one or more additional links to operate as standby links that will activate only if another active link goes down.
Static LACP — Use the manually configured static LACP trunk when:
The port on the other end of the trunk link is configured for a static LACP trunk.
You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group.
You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled.
You want to use a monitor port on the switch to monitor an LACP trunk.</td></tr>
<tr><td>Trunk(non-protocol)</td><td>Provides manually configured, static-only trunking to:
Most Switch and routing switches not running the 802.3ad LACP protocol.
Windows NT and HP-UX workstations and servers
Use the Trunk option when:
The device to which you want to create a trunk link is using a non-802.3ad trunking protocol.
You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol.
You want to use a monitor port on the switch to monitor traffic on a trunk.</td></tr>
</table>

Operating port trunks

Media:

For proper trunk operation, all ports on both ends of a trunk group must have the same media type and mode (speed and duplex.) (For the switches, Switch recommends leaving the port Mode setting at `Auto` or, in networks using Cat 3 cabling, `Auto-10`.)

Port Configuration

The default port configuration is `Auto`, which enables a port to sense speed and negotiate duplex with an auto-enabled port on another device. Switch recommends that you use the `Auto` setting for all ports you plan to use for trunking. Otherwise, you must manually ensure that the mode setting for each port in a trunk is compatible with the other ports in the trunk.

Recommended port mode setting for LACP

```
switch# show interfaces config

 Port Settings

  Port  Type      | Enabled Mode         Flow Ctrl MDI
  ----- --------- + ------- ------------ --------- ----
  1     10/100TX  | Yes     Auto         Enable    Auto
  2     10/100TX  | Yes     Auto         Enable    MDI
```

All of the following operate on a per-port basis, regardless of trunk membership:

- Enable/Disable
- Flow control (Flow Ctrl)

LACP is a full-duplex protocol.

Trunk configuration:

All ports in the same trunk group must be the same trunk type (LACP or trunk.) All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.

A trunk appears as a single port labeled`Dyn1`(for an LACP dynamic trunk) or `Trk1` (for a static trunk of type LACP, Trunk) on various menu and CLI screens.

For spanning-tree or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for spanning-tree or VLAN operation.)

Traffic distribution:

All of the switch trunk protocols use the SA/DA (source address/destination address) method of distributing traffic across the trunked links.

Spanning Tree:

802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch.) 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch.) For each Spanning Tree instance, you can adjust Spanning Tree parameters on a per-port basis.

A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named Trk1, they are listed in the Spanning Tree display as Trk1 and do not appear as individual ports in the Spanning Tree displays.

When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.

NOTE

A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI `show spanning-tree` display, but not in the Spanning Tree Operation display of the Menu interface.

If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.

Figure 40: *Example of a port trunk in a Spanning Tree listing*

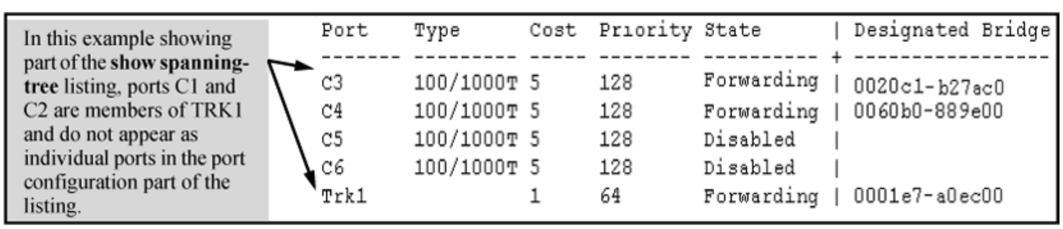

IP multicast protocol (IGMP):

A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN.

A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or `show ip igmp` listing.

VLANs:

Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.

NOTE

For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled.

Port security

Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the `show port-security` listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you see the following message and the command is not executed:

<PORT-LIST> Command cannot operate over a logical port.

Monitor port

NOTE

A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.

Show port-security log

Syntax

`show port-security intrusion-log`

Description

show port-security intrusion-log

```
switch(config)# sh port-security intrusion-log

Status and Counters - Intrusion Log

Port   MAC Address   Date / Time
------ ------------- --------------------------
23     000087-c78b49 11/19/14 11:09:30
23     000087-c78041 11/19/14 11:12:29
23     000087-c781c1 11/19/14 11:14:08
```

Static or dynamic trunk group overview

Configure port trunking before you connect the trunked links between switches.

Failure to configure port trunking before connecting the trunked links between switches can result in a broadcast storm. If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until after you have configured the trunk.

An individual trunk can have up to eight links, with additional standby links if you are using LACP. You can configure trunk group types as follows:

Trunk Type	Trunk Group Membership	
	TrkX (static)	DynX (dynamic)
LACP	Yes	Yes
Trunk	Yes	No

Enabling a dynamic LACP trunk group

In the default port configuration, all ports on the switch are set to disabled. To enable the switch to automatically form a trunk group that is dynamic on both ends of the link, the ports on one end of a set of links must be LACP Active. The ports on the other end can be either LACP Active or LACP Passive. The `active` command enables the switch to automatically establish a (dynamic) LACP trunk group when the device on the other end of the link is configured for LACP Passive. **Figure 41: Criteria for automatically forming a dynamic LACP trunk** on page 180 provides an example.

Figure 41: *Criteria for automatically forming a dynamic LACP trunk*

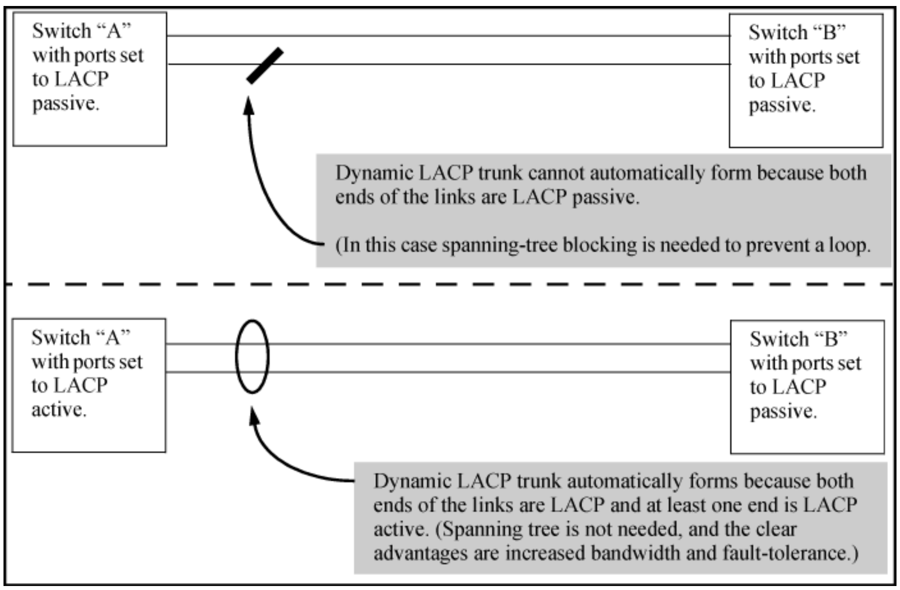

Dynamic LACP standby links

Dynamic LACP trunking enables you to configure standby links for a trunk by including more than eight ports in a dynamic LACP trunk configuration. When eight ports (trunk links) are up, the remaining links are held in standby status. If a trunked link that is "Up" fails, it is replaced by a standby link, which maintains your intended bandwidth for the trunk. In this example, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows `Standby` port status, while the remaining eight links show `Up` port status.

A dynamic LACP trunk with one standby link

```
switch# show lacp

                              LACP

      LACP      Trunk     Port                LACP      Admin  Oper
Port  Enabled   Group     Status    Partner   Status    Key    Key
----  -------   -------   -------   -------   -------   ------ ------
Al    Active    Dyn1      Up        Yes       Success   100    100
```

```
A2     Active     Dyn1       Up        Yes       Success   100    100
A3     Active     Dyn1       Up        Yes       Success   100    100
A4     Active     Dyn1       Up        Yes       Success   100    100
A5     Active     Dyn1       Up        Yes       Success   100    100
A6     Active     Dyn1       Up        Yes       Success   100    100
A7     Active     Dyn1       Up        Yes       Success   100    100
A8     Active     Dyn1       Up        Yes       Success   100    100
A9     Active     Dyn1       Standby   Yes       Success   100    100
```

Viewing LACP local information

Example of LACP local information

```
switch# show lacp local

LACP Local Information.

System ID: 001871-b98500

              LACP                     Tx     Rx Timer
  Port Trunk  Mode      Aggregated  Timer  Expired
  ---- ------ --------- ----------- ------ --------
  A2   A2     Active    Yes         Fast   No
  A3   A3     Active    Yes         Fast   No
```

Viewing LACP peer information

Use the `show lacp peer` command to display information about LACP peers. The System ID represents the MAC address of a partner switch. It will be zero if a partner is not found.

Example of LACP peer information

```
switch# show lacp peer

LACP Peer Information.

System ID: 001871-b98500

  Local  Local                         Port      Oper    LACP     Tx
  Port   Trunk  System ID       Port   Priority  Key     Mode     Timer
  ------ ------ --------------- ----- ---------- -------- --------- -----
  A2     A2     123456-654321   2     0          100     Passive  Fast
  A3     A3     234567-456789   3     0          100     Passive  Fast
```

Viewing LACP counters

Use the `show lacp counters` command to display statistical information about LACP ports.

NOTE

Data traffic can be dynamically redistributed in port channels. This may occur when a link is added or removed, or there is a change in load-balancing. Traffic that is redistributed in the middle of a traffic flow could potentially cause mis-ordered data packets.

LACP uses the marker protocol to prevent data packets from being duplicated or reordered due to redistribution. Marker PDUs are sent on each port-channel link. The remote system responds to the marker PDU by sending a marker responder when it has received all the frames received on this link prior to the marker PDU. When the marker responders are received by the local system on all member links of the port channel, the local system can redistribute the packets in the traffic flow correctly.

For the switches covered in this guide, the marker BPDUs are not initiated, only forwarded when received, resulting in the Marker fields in the output usually displaying zeros.

Example of LACP counters output

```
switch# show lacp counters

LACP Port Counters.


  LACP         LACP       Marker     Marker    Marker    Marker
  Port Trunk   PDUs Tx    PDUs Rx    Req. Tx   Req. Rx   Resp. Tx Resp. Rx Error
  ---- ------ ---------- ---------- --------- --------- --------- --------- ---------
  A2   A2     1234       1234       0         0         0         0         0
  A3   A3     1234       1234       0         0         0         0         0
```

Trunk group operation using LACP

The switch can automatically configure a dynamic LACP trunk group, or you can manually configure a static LACP trunk group.

NOTE

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, and so on) and the same speed and enforces speed and duplex conformance across a trunk group. For most installations, Switch recommends that you leave the port mode settings at `Auto` (the default.) LACP also operates with `Auto-10`, `Auto-100`, and `Auto-1000` (if negotiation selects FDx), and `10FDx`, `100FDx`, and `1000FDx` settings.

LACP trunk status commands include:

Trunk display method	Static LACP trunk	Dynamic LACP trunk
CLI `show lacp` command	Included in listing.	Included in listing.
CLI `show trunk` command	Included in listing.	Not included.
Port/Trunk Settings screen in menu interface	Included in listing.	Not included

Thus, to display a listing of dynamic LACP trunk ports, you must use the `show lacp` command.

In most cases, trunks configured for LACP on the switches operate as follows:

Table 11: *LACP trunk types*

LACP port trunk configuration	Operation
Dynamic LACP	This option automatically establishes an 802.3ad-compliant trunk group, with LACP for the port Type parameter and DynX for the port Group name, where X is an automatically assigned value from 1 to 144, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 144 trunk groups in any combination of static and dynamic trunks.) NOTE Dynamic LACP trunks operate only in the default VLAN (unless GVRP is enabled and `Forbid` is used to prevent the trunked ports from joining the default VLAN.) Thus, if an LACP dynamic port forms using ports that are not in the default VLAN, the trunk automatically moves to the default VLAN unless GVRP operation is configured to prevent this from occurring. In some cases, this can create a traffic loop in your network. Under the following conditions, the switch automatically establishes a dynamic LACP port trunk group and assigns a port Group name: • The ports on both ends of each link have compatible mode settings (speed and duplex.) • The port on one end of each link must be configured for LACP Active and the port on the other end of the same link must be configured for either LACP Passive or LACP Active. For example: 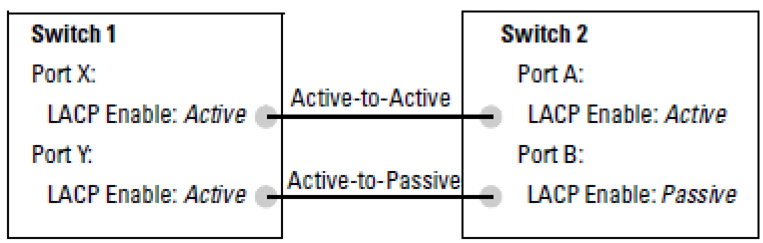 Either of the above link configurations allows a dynamic LACP trunk link. Backup Links: A maximum of eight operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more additional (backup) links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing eight-port dynamic LACP trunk, ensure that the ports in the standby link are configured as either active-to-active or active-to-passive between switches. Displaying dynamic LACP trunk data: To list the configuration and status for a dynamic LACP trunk, use the CLI `show lacp` command. NOTE The dynamic trunk is automatically created by the switch and is not listed in the static trunk listings available in the menu interface or in the CLI `show trunk` listing.
Static LACP	Provides a manually configured, static LACP trunk to accommodate these conditions:

LACP port trunk configuration	Operation
	• The port on the other end of the trunk link is configured for a static LACP trunk. • You want to configure non-default Spanning Tree or IGMP parameters on an LACP trunk group. • You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. • You want to use a monitor port on the switch to monitor an LACP trunk. The trunk operates if the trunk group on the opposite device is running one of the following trunking protocols: • Active LACP • Passive LACP • Trunk This option uses LACP for the port Type parameter and TrkX for the port Group parameter, where X is an automatically assigned value in a range corresponding to the maximum number of trunks the switch allows. Displaying static LACP trunk data : To list the configuration and status for a static LACP trunk, use the CLI `show lacp` command. To list a static LACP trunk with its assigned ports, use the CLI `show trunk` command or display the menu interface Port/Trunk Settings screen.Static LACP does not allow standby ports.

Default port operation

In the default configuration, LACP is disabled for all ports. If LACP is not configured as Active on at least one end of a link, the port does not try to detect a trunk configuration and operates as a standard, untrunked port. To display this data for a switch, execute the following command in the CLI:

```
switch# show lacp
```

Table 12: *LACP port status data*

Status name	Meaning
Port Numb	Shows the physical port number for each port configured for LACP operation (C1, C2, C3 ….) Unlisted port numbers indicate that the missing ports that are assigned to a static trunk group are not configured for any trunking.
LACP Enabled	Active: The port automatically sends LACP protocol packets. Passive: The port does not automatically send LACP protocol packets and responds only if it receives LACP protocol packets from the opposite device. A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports does not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device. NOTE: In the default switch configuration, LACP is disabled for all ports.

Table Continued

Status name	Meaning
Trunk Group	TrkX: This port has been manually configured into a static LACP trunk. Trunk group same as port number: The port is configured for LACP, but is not a member of a port trunk.
Port Status	Up: The port has an active LACP link and is not blocked or in standby mode. Down: The port is enabled, but an LACP link is not established. This can indicate, for example, a port that is not connected to the network or a speed mismatch between a pair of linked ports. Disabled: The port cannot carry traffic. Blocked: LACP, Spanning Tree has blocked the port. (The port is not in LACP standby mode.) This may be caused by a (brief) trunk negotiation or a configuration error, such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support. NOTE: Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. Standby: The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or "standby" unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a standby port, if available, to replace the failed port.
LACP Partner	Yes: LACP is enabled on both ends of the link. No: LACP is enabled on the switch, but either LACP is not enabled or the link has not been detected on the opposite device.
LACP Status	Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link. Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore is not able to send LACP packets across the link. This can be caused, for example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.

LACP operating notes and restrictions

802.1X (Port-based access control) configured on a port

To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables 802.1X on that port.

```
switch# aaa port-access authenticator b1
LACP has been disabled on 802.1x port(s.)
switch#
```

The switch does not allow you to configure LACP on a port on which port access (802.1X) is enabled. For example:

```
switch# int b1 lacp passive
Error configuring port  port-number  : LACP and 802.1x cannot
be run together.
switch#
```

To restore LACP to the port, you must first remove the 802.1X configuration of the port and then re-enable LACP active or passive on the port.

Port security

To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables port security on that port. For example:

```
switch# port-security a17 learn-mode static address-
limit 2 LACP has been disabled on secured port(s.)
switch#
```

The switch does not allow you to configure LACP on a port on which port security is enabled. For example:

```
switch# int a17 lacp passive
Error configuring port A17: LACP and port security cannot be
run together.
switch#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Changing trunking methods

To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP trunks

When a port is configured for LACP (active or passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

NOTE

Static LACP allows ports with different speed to be part of the same trunk.

Dynamic LACP trunks

You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the `trunk` command. (See "Using the CLI To Configure a Static or Dynamic Trunk Group")

VLANs and dynamic LACP

A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use `Forbid` to prevent the ports from joining the default VLAN.)

If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.

Blocked ports with older devices.

Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. The LACP status of the blocked ports is shown as "Failure."

If one of the other ports becomes disabled, a blocked port replaces it (Port Status becomes "Up".) When the other port becomes active again, the replacement port goes back to blocked (Port Status is "Blocked".) It can take a few seconds for the switch to discover the current status of the ports.

Figure 42: *Blocked ports with LACP*

```
HP Switch(eth-B1-B8)# show lacp

                              LACP

   PORT     LACP       TRUNK      PORT       LACP       LACP
   NUMB     ENABLED    GROUP      STATUS     PARTNER    STATUS
   ----     -------    -------    -------    -------    -------
   B1       Active     Dyn1       Up         Yes        Success
   B2       Active     Dyn1       Up         Yes        Success
   B3       Active     Dyn1       Up         Yes        Success
   B4       Active     Dyn1       Up         Yes        Success
   B5       Active     Dyn1       Blocked    Yes        Failure
   B6       Active     Dyn1       Blocked    Yes        Failure
   B7       Active     B7         Down       No         Success
   B8       Active     B8         Down       No         Success
```

If there are ports that you do not want on the default VLAN, ensure that they cannot become dynamic LACP trunk members. Otherwise a traffic loop can unexpectedly occur. For example:

Figure 43: *A dynamic LACP trunk forming in a VLAN can cause a traffic loop*

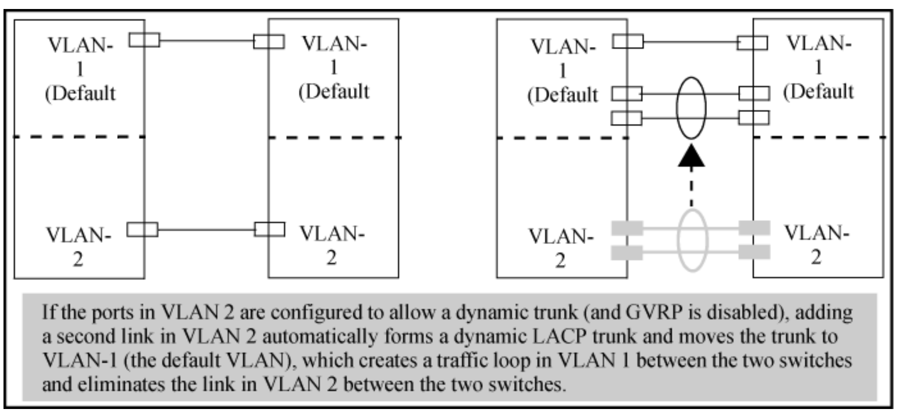

Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP

If Spanning Tree, IGMP, or both are enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-duplex, different port speeds, or both not allowed in LACP trunks

Theports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx.) The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking. (10-gigabit ports operate only at FDx.)

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx.) However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If the port is a 10-gigabit port.
- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/static LACP interoperation

A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links are ignored.

Trunk group operation using the "trunk" option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

When a trunk group is configured with the `trunk` option, the switch automatically sets the trunk to a priority of "4" for Spanning Tree operation (even if Spanning Tree is currently disabled.) This appears in the running-config file as `spanning-tree Trkn priority 4`. Executing `write memory` after configuring the trunk places the same entry in the startup-config file.

Use the `trunk` option to establish a trunk group between a switch and another device, where the other device's trunking operation fails to operate properly with LACP trunking configured on the switches.

Viewing trunk data on the switch

Static trunk group

Appears in the menu interface and the output from the CLI `show trunk` and `show interfaces` commands.

Dynamic LACP trunk group

Appears in the output from the CLI `show lacp` command.

Interface option	Dynamic LACP trunk group	Static LACP trunk group	Static non-protocol
Menu interface	No	Yes	Yes
CLI `show trunk`	No	Yes	Yes
CLI `show interfaces`	No	Yes	Yes
CLI `show lacp`	Yes	Yes	No
CLI `show spanning-tree`	No	Yes	Yes
CLI `show igmp`	No	Yes	Yes
CLI `show config`	No	Yes	Yes

Outbound traffic distribution across trunked links

The two trunk group options (LACP and trunk) use SA/DA pairs for distributing outbound traffic over trunked links. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and may also send traffic from the same source address to a different destination address through the same link or a different link, depending on the mapping of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through links depending on the path assignment.

The load-balancing is done on a per-communication basis. Otherwise, traffic is transmitted across the same path. That is, if Client A attached to Switch 1 sends five packets of data to Server A attached to Switch 2, the same link is used to send all five packets. The SA/DA address pair for the traffic is the same. The packets are not evenly distributed across any other existing links between the two switches; they all take the same path.

Figure 44: *Example of single path traffic through a trunk*

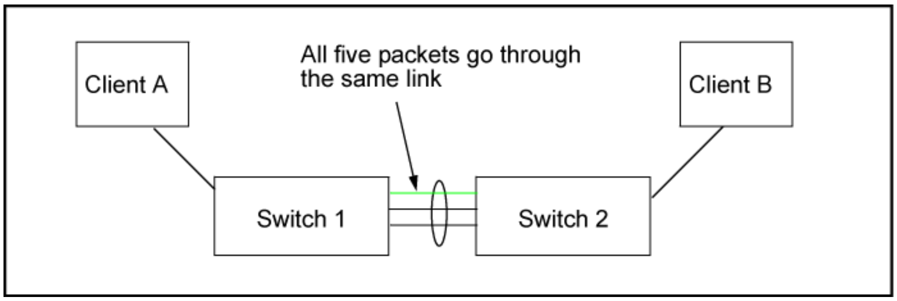

The actual distribution of the traffic through a trunk depends on a calculation using bits from the SA/DA. When an IP address is available, the calculation includes the last five bits of the IP source address and IP destination address; otherwise, the MAC addresses are used. The result of that process undergoes a mapping that determines which link the traffic goes through. If you have only two ports in a trunk, it is possible that all the traffic will be sent through one port even if the SA/DA pairs are different. The more ports you have in the trunk, the more likely it is that the traffic will be distributed among the links.

When a new port is added to the trunk, the switch begins sending traffic, either new traffic or existing traffic, through the new link. As links are added or deleted, the switch redistributes traffic across the trunk group.

Figure 45: *Example of port-trunked network*

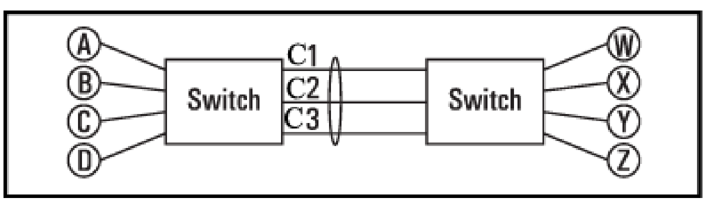

Table 13: *Example of link assignments in a trunk group (SA/DA distribution)*

Source	Destination	Link
Node A	Node W	1
Node B	Node X	2
Node C	Node Y	3
Node D	Node Z	1

Table Continued

Source	Destination	Link
Node A	Node Y	2
Node B	Node W	3

Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while other links in the same trunk have unused bandwidth capacity, even if the assignments were evenly distributed across the links in a trunk.

Trunk load balancing using Layer 4 ports

Trunk load balancing using Layer 4 ports allows the use of TCP/UDP source and destination port number for trunk load balancing. This is in addition to the current use of source and destination IP address and MAC addresses. Configuration of Layer 4 load balancing would apply to all trunks on the switch. Only non-fragmented packets will have their TCP/UDP port number used by load balancing. This ensures that all frames associated with a fragmented IP packet are sent through the same trunk on the same physical link.

The priority for using Layer 4 packets when this feature is enabled is as follows:

1. If the packet protocol is an IP packet and has Layer 4 port information, use Layer 4.
2. If the packet protocol is an IP packet and does **not** have Layer 4 information, use Layer 3 information.
3. If the packet is **not** an IP packet, use Layer 2 information.

Distributed trunking overview

The IEEE standard 802.3ad requires that all links in a trunk group originate from the same switch. Distributed trunking uses a proprietary protocol that allows two or more port trunk links distributed across two switches to create a trunk group. The grouped links appear to the downstream device as if they are from a single device. This allows third party devices such as switches, servers, or any other networking device that supports trunking to interoperate with the distributed trunking switches (DTSs) seamlessly. Distributed trunking provides device-level redundancy in addition to link failure protection.

DTSs are connected by a special interface called the InterSwitch-Connect (ISC) port. This interface exchanges information so that the DTSs appear as a single switch to a downstream device, as mentioned above. Each distributed trunk (DT) switch in a DT pair must be configured with a separate ISC link and peer-keepalive link. The peer-keepalive link is used to transmit keepalive messages when the ISC link is down to determine if the failure is a link-level failure or the complete failure of the remote peer.

The downstream device is a distributed trunking device (DTD.) The DTD forms a trunk with the DTSs. The connecting links are DT links and the ports are DT ports. A distributed trunk can span a maximum of two switches.

Before you configure the switch, Hewlett Packard Enterprise recommends that you review the **Distributed trunking restrictions** on page 199 for a complete list of operating notes and restrictions.

NOTE

DT is not supported between different platforms. The generic application of the DT protocol across series is not supported. For example:

- DT is not supported between an Aruba 3810M switch and an Aruba 5400R switch.
- DT is not supported on different platforms that make it generic for the Aruba 3800 switch.

You can group together distributed trunks by configuring two individual dt-lacp/dt-trunk trunks with the same trunk group name in each switch. The DT ports are grouped dynamically after the configuration of distributed trunking.

Before you configure the switch, Hewlett Packard Enterprise recommends that you review the **Distributed trunking restrictions** on page 199 for a complete list of operating notes and restrictions.

In **Figure 46: Example of distributed trunking with three different distributed trunks with three servers** on page 191, three different distributed trunks with three different servers have one common ISC link. Each trunk spans only two DTSs, which are connected at the ISC ports so they can exchange information that allows them to appear as one device to the server.

Figure 46: *Example of distributed trunking with three different distributed trunks with three servers*

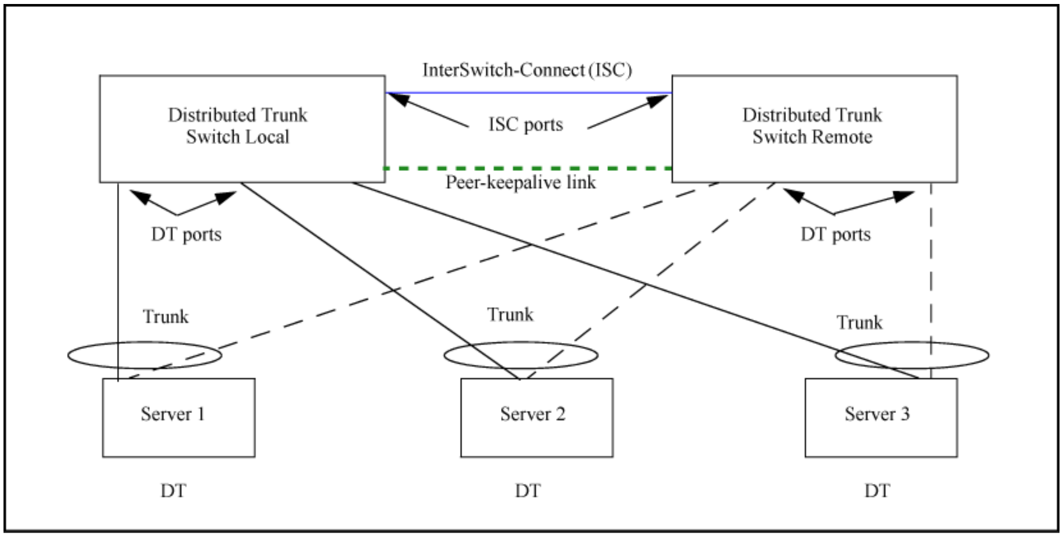

An example of distributed trunking switch-to-switch in a square topology is shown in **Figure 47: Distributed trunking switch-to-switch square topology** on page 192.

Figure 47: *Distributed trunking switch-to-switch square topology*

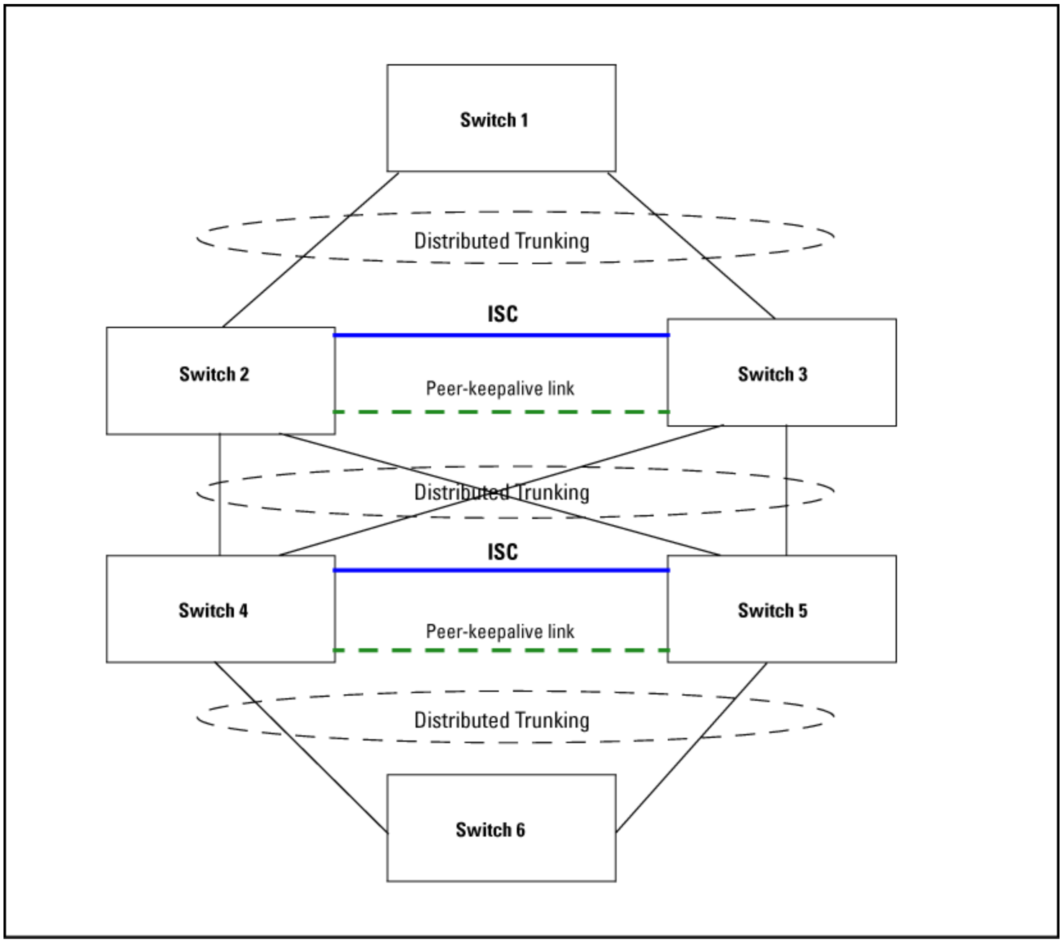

Distributed trunking interconnect protocol

Distributed trunking uses the distributed trunking interconnect protocol (DTIP) to transfer DT-specific configuration information for the comparison process and to synchronize MAC and DHCP snooping binding data between the two DT peer switches.

NOTE

For DHCP snooping to function correctly in a DT topology, the system time must be the same on both switches, and the ISC must be trusted for DHCP snooping.

Configuring distributed trunking

The following parameters must be configured identically on the peer devices or undesirable behavior in traffic flow may occur:

- The ISC link must have a VLAN interface configured for the same VLAN on both DT switches.
- VLAN membership for all DT trunk ports should be the same on both DT switches in a DT pair.
- IGMP-snooping or DHCP-snooping configuration on a DT VLAN should be the same on both DT switches. For example, for a DT, if IGMP-snooping or DHCP-snooping is enabled on a VLAN that has a DT port as a member port of the VLAN, the same must be configured on the peer DT on the same VLAN.
- Loop-protection configuration on a DT VLAN should be the same for both DT switches.

Configuring peer-keepalive links

Distributed trunking uses UDP-based peer-keepalive messages to determine if an ISC link failure is at the link level or the peer has completely failed. The following operating rules must be followed to use peer-keepalive links:

- An IP address must be configured for a peer-keepalive VLAN interface and the same IP address must be configured as a peer-keepalive destination on the peer DT switch.
- There must be logical Layer 3 connectivity between the two IP addresses configured for the peer-keepalive VLAN interface.
- Only peer-keepalive messages are sent over the peer-keepalive VLAN (Layer 3 link.) These messages indicate that the DT switch from which the message originates is up and running. No data or synchronization traffic is sent over the peer-keepalive VLAN.
- STP cannot run on peer-keepalive links.
- The peer-keepalive VLAN can have only one member port. If you attempt to assign a second member port to this VLAN, or if you attempt to configure a VLAN that has more than one member port as a peer-keepalive VLAN, this message displays:

  ```
  A keepalive VLAN can only have one member port.
  ```
- A port cannot be a member of a regular VLAN and a peer-keepalive VLAN. An error message displays:

  ```
  A port cannot simultaneously be a member of a keepalive and a non-keepalive VLAN.
  ```
- The DEFAULT VLAN cannot be a peer-keepalive VLAN. An error message displays:

  ```
  The default VLAN cannot be configured as a keepalive VLAN.
  ```

NOTE

If you are upgrading your software from a version prior to K.15.05.**xxxxx** with a configuration that violates any of the above operating rules, the following message displays:

```
DT: Keepalive mis-configuration detected. Reconfigure the keepalive VLAN.
```

You must then manually correct the configuration.

DT switches have an operational role that depends on the system MAC address. The bridge with the lowest system MAC address acts as the DT primary device; the other device is the DT secondary device. These roles are used to determine which device forwards traffic when the ISC link is down.

Figure 48: *ISC link failure with peer-keepalive*

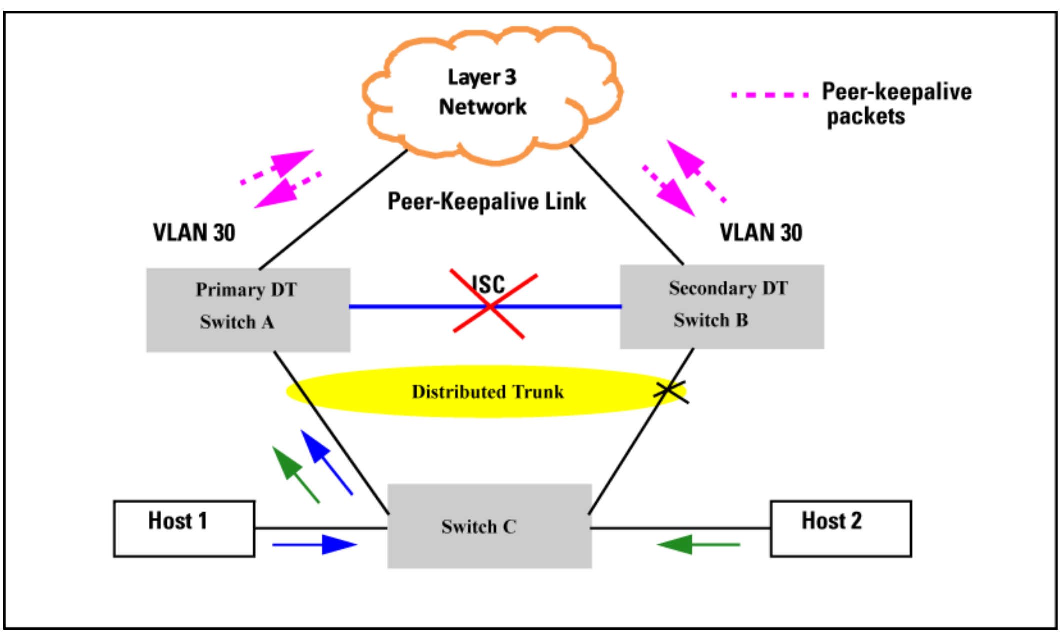

Peer-keepalive messages are sent by both the DT switches as soon as the switches detect that the ISC link is down. Peer-keepalive message transmission (sending and receiving) is suspended until the peer-keepalive hold timer expires. When the hold timer expires, the DT switches begin sending peer-keepalive messages periodically while receiving peer-keepalive messages from the peer switch. If the DT switch fails to receive any peer-keepalive messages for the timeout period, it continues to forward traffic, assuming that the DT peer switch has completely failed.

Conversely, if the failure is because the ISC link went down and the secondary DT switch receives even one peer-keepalive message from the primary peer, the secondary switch disables all its DT ports. The primary switch always forwards the traffic on its DT ports even if it receives peer-keepalive messages from the secondary DT switch.

In both situations, if the ISC link or the DT switch becomes operational, both the DT peers sync the MAC addresses learned during the failover and continue to forward traffic normally. The peer-keepalive timers and operation is halted.

Maximum DT trunks and links supported

Table 14: *Maximum supported DT trunks and links*

Description	Max number
Maximum number of groups (DT trunks) in a DT switch (that is, maximum number of servers supported)	144
Maximum number of switches that can be aggregated	2
Maximum number of physical links that can be aggregated in a single switch from a server (that is, maximum number of ports that can be in a trunk connected to a single switch)	4

From the server perspective, this means that there could be a maximum total of 60 servers connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that a single server could have a maximum of eight links (that is, four on each DT switch) in a DT trunk.

Forwarding traffic with distributed trunking and spanning tree

Refer to **Figure 49: Distributed trunking with STP forwarding unicast, broadcast, and multicast traffic** on page 195 for the following discussion about forwarding traffic when spanning tree is enabled. In this example, it is assumed that traffic is sent from a host off switch B to a server, and from the server back to the host. STP can block any one of the upstream links; in this example, STP has blocked all the links except the I1 link connected to DT1.

NOTE

STP is automatically disabled on the DT ports.

Figure 49: *Distributed trunking with STP forwarding unicast, broadcast, and multicast traffic*

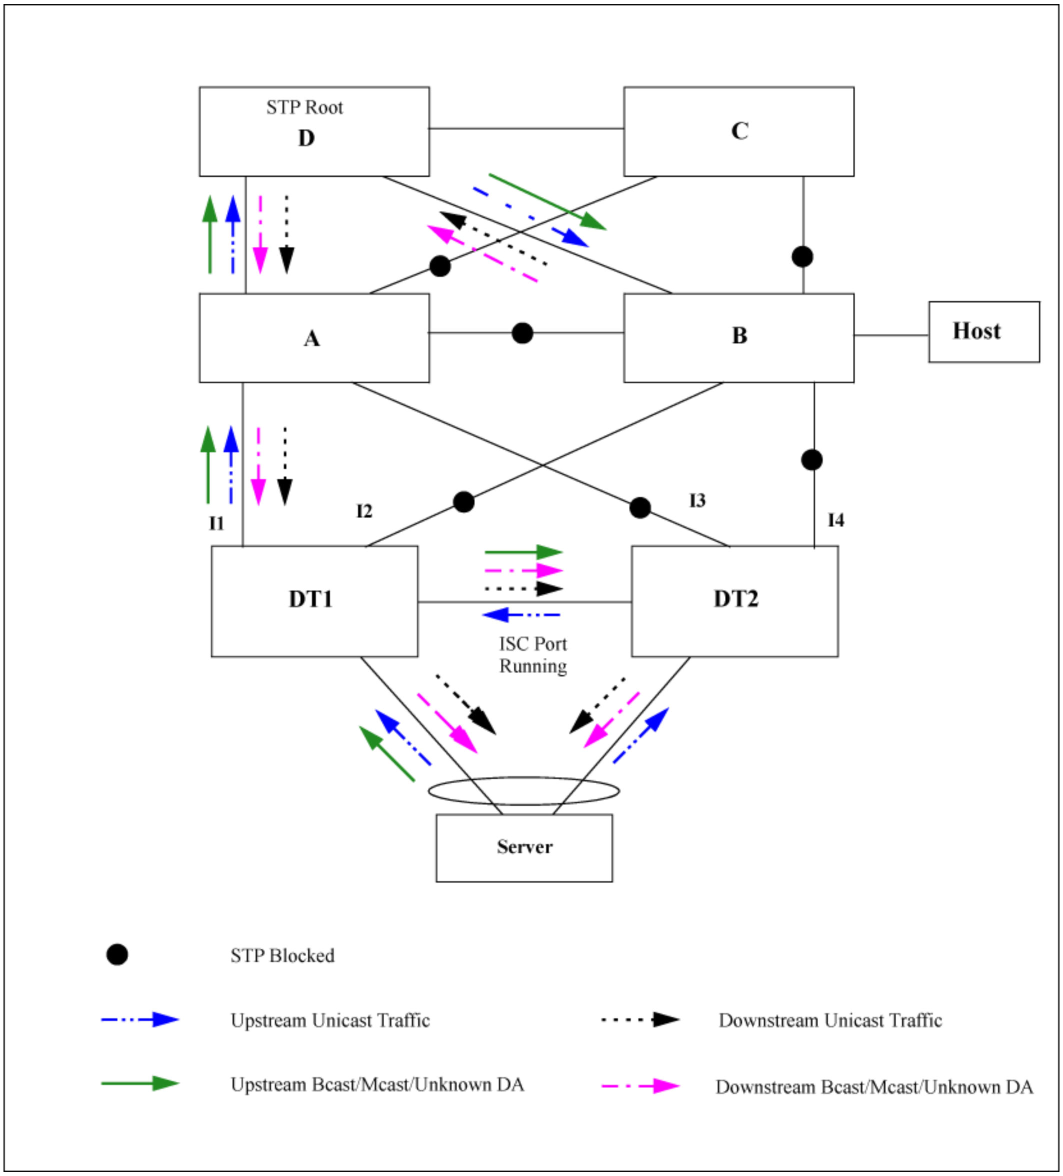

Forwarding unicast traffic

Refer to **Figure 50: Unicast traffic flow across DT switches** on page 196 for the following discussion about forwarding traffic with switch-to-switch distributed trunking. Traffic from Host X or Y that is destined for Host F is always forwarded by Switch A over one of its standard 802.1AX trunk links to either Switch B or Switch C. When either Switch B or Switch C receives incoming traffic from Switch A, the traffic is directly forwarded to Switch F without traversing the ISC link.

Traffic from Host Y to Host D may go over the ISC if Switch A sends it to Switch C instead of sending it to Switch B.

Figure 50: *Unicast traffic flow across DT switches*

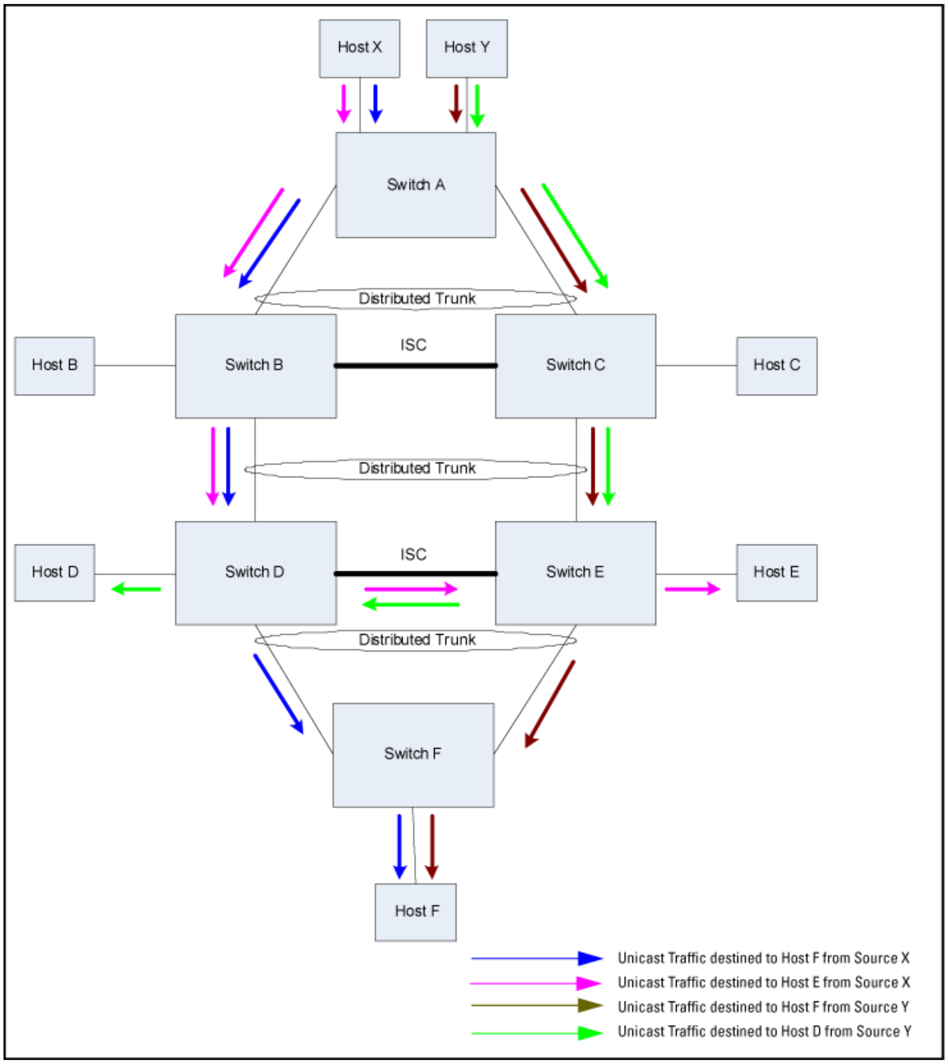

Forwarding broadcast, multicast, and unknown traffic

In the example shown in **Figure 51: Broadcast/multicast/unknown traffic flow access DT switches** on page 197, multicast/broadcast/unknown traffic from Host X or Y is always forwarded by Switch A over one of its standard 802.3ad trunk links to either Switch B or C. Switch B or C forwards the traffic on all the links including the ISC port, but not on the port that the traffic was received on. The peer DT switch (B or C) that receives broadcast/ multicast/unknown traffic over the ISC port does not forward the packets to any of the DT trunks; the packet is sent only over the non-DT ports. The one exception is if the DT trunk on the peer aggregation device is down, then traffic received over the ISC is forwarded to the corresponding DT trunk.

Figure 51: *Broadcast/multicast/unknown traffic flow access DT switches*

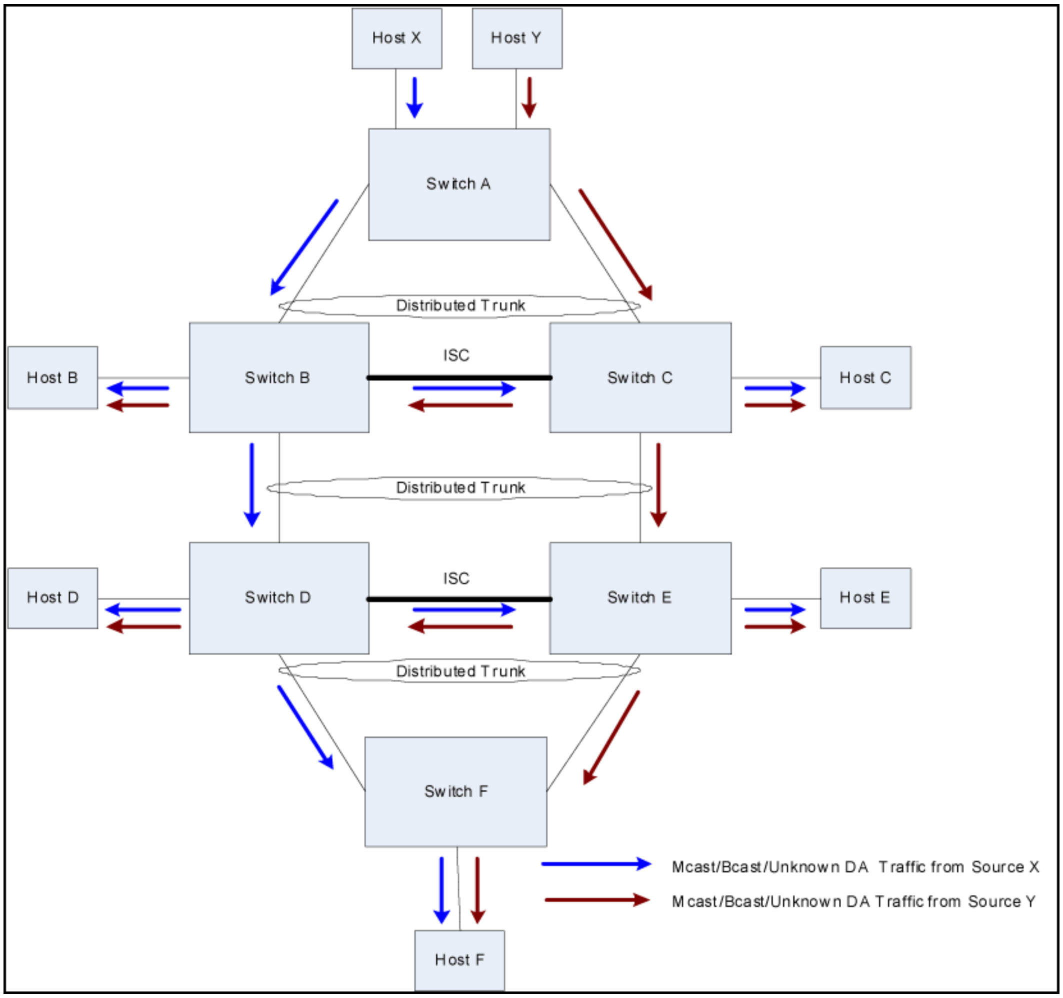

IP routing and distributed trunking

In switch-to-switch distributed trunking, the peer DT switches behave like independent Layer 3 devices with their own IP addresses in each active VLAN. If a DT switch receives a packet destined for the peer DT switch, it switches the packet through the ISC link. Interfaces on a VLAN using DT typically use a single default gateway pointing to only one of the DT switches in a DT pair.

The example in the following figure shows Layer 3 (IP unicast) forwarding in a DT topology. The packet is sent as follows:

1. Switch A selects the link (using the trunk hash) to the DT pair. The packet is sent to the selected link DT_SW_1.
2. When DT_SW_1 receives the packet, it determines, based on the MAC address, that the packet must be sent over the ISC link to DT_SW_2.
3. When the packet arrives, DT_SW_2 performs a lookup and determines that the packet needs to be sent to Switch B.

Figure 52: *Layer 3 forwarding (IP unicast) in DT topology*

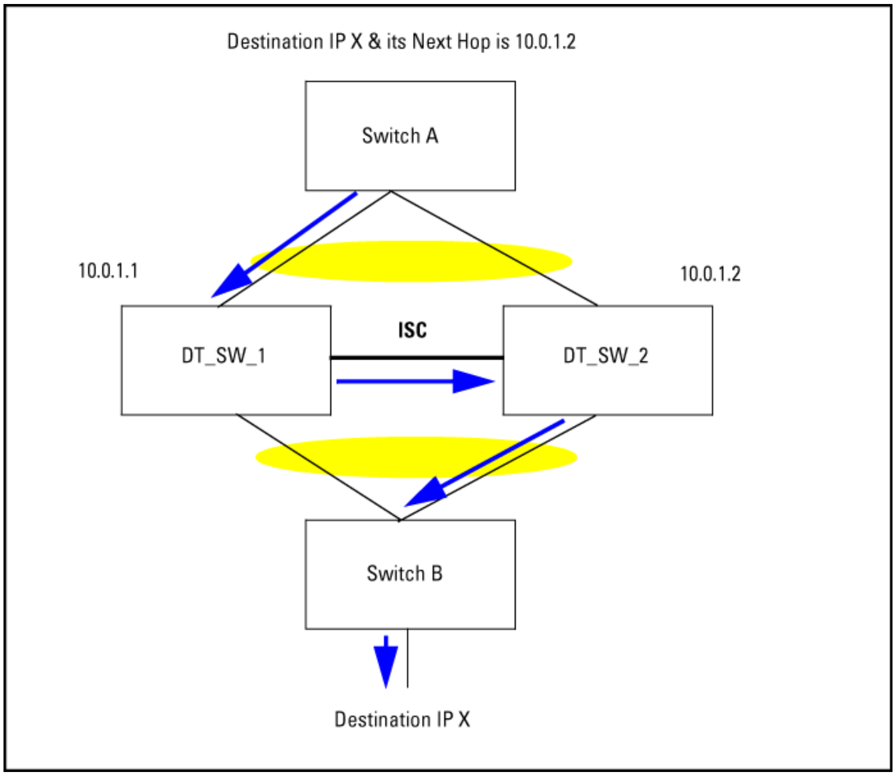

Another example in the next figure shows Layer 3 (IP unicast) forwarding in a DT topology. The packet is sent as follows:

1. Host 2 sends a packet to Switch C.
2. Switch C performs a lookup in the routing table and determines that the default gateway IP address is 10.0.0.1.
3. Layer 2 lookup determines that the outgoing interface is the DT port.
4. Hashing determines that the trunk member chosen is DT_SW_2 and the packet is sent there.
5. DT_SW_2 determines that the packet needs to be sent over the ISC link to DT_SW_1 based on the MAC address.
6. DT_SW_1 performs a lookup and determines that the packet goes to Switch A.

The packet is only forwarded if the outgoing interface is not a DT port, or if the outgoing DT port does not have an active interface on the peer switch.

Figure 53: *Layer 3 forwarding (IP unicast) in DT topology*

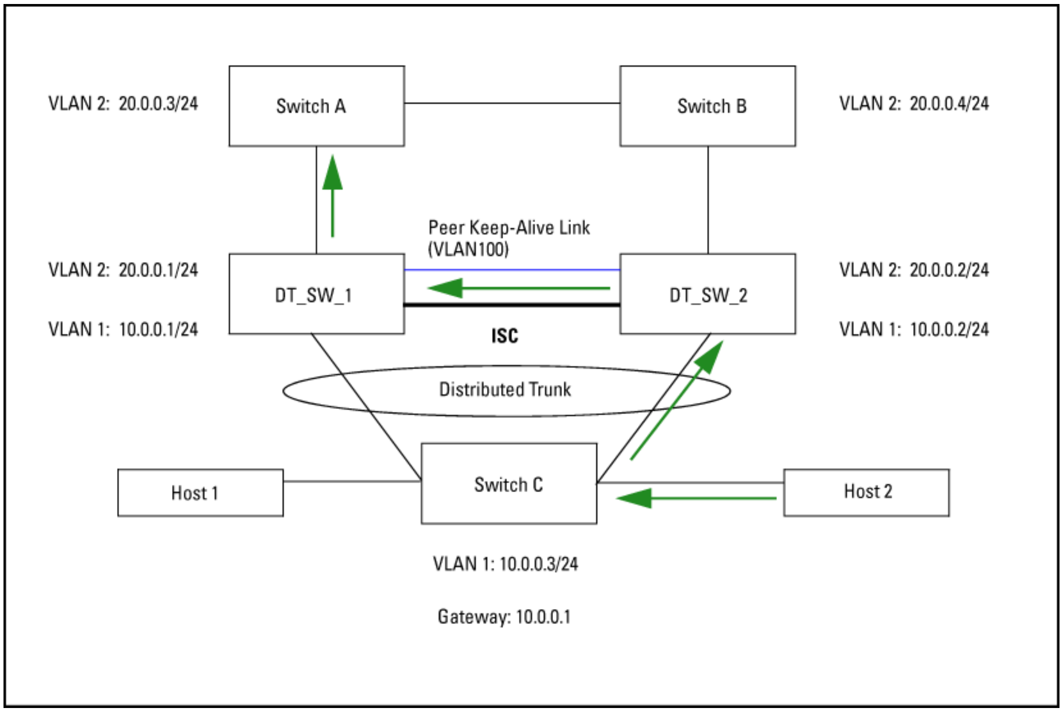

Distributed trunking restrictions

There are several restrictions with distributed trunking:

- Beginning with software version K.15.07, the switch will not allow both Distributed Trunking and MAC-based mirroring to function simultaneously. The switch will respond as follows:
  - If the user attempts to configure both, an error message will appear.
  - When a switch is updated from older software to K.15.07, if the older config file has both Distributed Trunking and MAC-based mirroring, the switch will automatically remove the MAC-based mirroring lines from the config file, and will give an explanatory error message.
- If a switch is running K.15.07 and an existing config file that has both Distributed Trunking and MAC-based mirroring is loaded onto the switch, the switch will automatically remove the MAC-based mirroring lines from the config file, and will give an explanatory error message.
- All DT linked switches must be running the same software version.
- The port trunk links should be configured manually (using manual LACP or manual trunks.) Dynamic linking across switches is not supported.
- A distributed trunk can span a maximum of two switches.
- DT is not supported between different platforms such as the HP 3800 switch and the HP 3500 switch. The generic application of the DT protocol across series is not supported.
- A maximum total of 144 servers can be connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that there can be a maximum of eight ports (four aggregated links for each DT switch) included in a DT trunk.
- Only one ISC link is supported per switch, with a maximum of 60 DT trunks supported on the switch. The ISC link can be configured as a manual LACP trunk, non-protocol trunk, or as an individual link. Dynamic LACP trunks are not supported as ISCs.
- An ISC port becomes a member of all VLANs that are configured on the switch. When a new VLAN is configured, the ISC ports become members of that VLAN.

- Port trunk links can be done only on a maximum of two switches that are connected to a specific server.
- Any VLAN that is in a distributed trunk must be configured on both switches. By default, the distributed trunk belongs to the default VLAN.
- There can be eight links in a distributed trunk grouped across two switches, with a limit of four links per distributed trunking switch.
- The limit of 144 manual trunks per switch includes distributed trunks as well.
- ARP protection is not supported on the distributed trunks.
- Dynamic IP Lockdown protection is not supported on the distributed trunks.
- QinQ in mixed VLAN mode and distributed trunking are mutually exclusive.
- Source Port Filter cannot be configured on an InterSwitch Connect (ISC) port.
- Features not supported include:
    - SVLANs in mixed mode on DT or ISC links
    - Meshing
    - IPv6 routing

Updating software versions with DT

For 15.14.x and later

Beginning with software release 15.14.x, when updating to a new software release on switches configured for DT (Distributed Trunking) on LACP type trunks, you must update the DT partner with the lowest Base MAC address first. When this partner returns to operation, then update the other partner. Use the `show system` command to determine the base MAC address on a given switch.

From no DT Keepalive support to shared DT Keepalive support

When updating software from a version that does not support DT Keepalive (prior to version K.15.03) to a version that supports shared DT keepalive (K.15.03 and greater), use the following procedure:

Procedure

1. Disable the ISC interface on both switches, and then upgrade the software. Assume a2 is configured as switch-interconnect.

```
switch# int a2 disable
switch# write mem
```

2. Configure one of the existing uplink VLANs as a keepalive VLAN, and then configure the destination keepalive IP address (peer's keepalive IP address) on both switches at bootup.

```
switch# distributed-trunking
peer-keepalive vlan 2
switch# distributed-trunking
peer-keepalive destination 20.0.0.2
```

3. Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
4. Enable the ISC link on both switches and then execute `write memory`. Assume a2 is configured as switch-interconnect.

```
switch# int a2 enable
switch# write mem
```

From no DT Keepalive support to dedicated point-to-point DT Keepalive support

When updating software from a software version that does not support DT keepalive (prior to version K.15.03) to a version with dedicated point-to-point keepalive (K.15.03 and greater), use the following procedure:

- Disable the ISC interface on both switches, and then upgrade the software. Assume a2 is configured as switch-interconnect.

```
switch# int a2 disable
switch# write mem
```

- At switch bootup, create a dedicated VLAN for keepalive, and assign only the keepalive link port as a member port of the VLAN. Configure the keepalive destination IP address.

```
switch# distributed-trunking
peer-keepalive vlan 2
switch# distributed-trunking
peer-keepalive destination 20.0.0.2
```

- Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
- Enable the ISC link on both switches, and then execute `write memory`. Assume a2 is configured as switch-interconnect.

```
switch# int a2 enable
switch# write mem
```

From shared DT Keepalive support to point-to-point Keepalive support

When updating software from a software version that does support shared DT keepalive (K.15.03, K.15.04) to a version that supports dedicated point-to-point keepalive (K.15.05), use the following procedure:

- Disable the ISC interface and undo the keepalive configuration on both switches. Ignore the warning message that is displayed by the keepalive command while undoing the configuration. Upgrade the software. Assume a2 is configured as switch-interconnect.

```
switch# int a2 disable
switch# no distributed-trunking
peer-keepalive vlan
switch# write mem
```

- At switch bootup, create a dedicated VLAN for keepalive and assign only the keepalive link port as a member port of the VLAN. Configure the keepalive destination IP address.

```
switch# vlan 10 (dedicated point-to-point VLAN interface)
switch(vlan-10)#
switch(vlan-10)# untagged b2 (keepalive link port)
switch(vlan-10)# ip address 10.0.0.1/24
switch(vlan-10)# exit
switch# distributed-trunking
peer-keepalive vlan 10
switch# distributed-trunking
peer-keepalive destination 10.0.0.2
```

- Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
- Enable the ISC link on both switches, and then execute `write memory`. Assume a2 is configured as switch-interconnect.

```
switch# int a2 enable
switch# write mem
```

Rate-limiting

The switches covered by this guide support configuring inbound and outbound rate-limiting for all traffic on a port and specifying bandwidth usage in terms of either percent or kilobits per second (kbps.)

You can enable rate limiting for various types of traffic. When a limit is enabled on a port, excess traffic above the configured rate is discarded. The default is no limit.

- All-traffic rate limiting is primarily used for end-node connections (i.e. at the network edge). It is not recommended for use on links to servers, routers, switches, or the network core. Rate limiting traffic on such links may interfere with important network functions.
- Broadcast rate limiting is used to protect the network from disruption by excessive broadcast traffic.
- ICMP rate limiting is primarily used for throttling denial of service attacks.
- Multicast rate limiting is used to protect the network from disruption by excessive multicast traffic. This is an Interface context command. It can be called directly from the interface context or following the `interface` *`<PORT-LIST>`* command.
- Queues rate limiting sets an outbound rate limit for each traffic queue on a selected interface.

Hewlett Packard Enterprise does not recommend applying rate-limiting to desirable traffic.

Rate-limiting is intended for use on edge ports in a network. It is not recommended for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

ICMP traffic is necessary for network routing functions. For this reason, blocking all ICMP traffic is not recommended.

For more information on all-traffic rate-limiting, see **All traffic rate-limiting** on page 220.

Configuring rate-limiting on all traffic

rate-limit

Syntax

```
int <PORT-LIST> rate-limit all [in|out] <0-100000000> percent <0-100>|kbps
```

Description

Configures a traffic rate limit (on non-trunked ports) on the link. The `no` form of the command disables rate-limiting on the specified ports.

The `rate-limit all` command controls the rate of traffic sent or received on a port by setting a limit on the bandwidth available. It includes options for:

- Rate-limiting on either inbound or outbound traffic.
- Specifying the traffic rate as either a percentage of bandwidth, or in terms of kilobits per second.

(Default: Disabled.)

Parameters

`in` or `out`

Specifies a traffic rate limit on inbound traffic passing through that port, or on outbound traffic.

`percent` or `kbps`

Specifies the rate limit as a percentage of total available bandwidth, or in kilobits per second.

> The granularity of actual limits varies across different switch models.

Viewing the current rate-limit configuration

show rate-limit

Syntax

`show rate-limit all <PORT-LIST>`

Description

The `show rate-limit all` command displays the per-port rate-limit configuration in the running-config file.

Options

`no`

—

`<PORT-LIST>`

Without *`<PORT-LIST>`* , this command lists the rate-limit configuration for all ports on the switch. With *`<PORT-LIST>`* , this command lists the rate-limit configuration for the specified ports. This command operates the same way in any CLI context.

rate-limited configuration

The following figure shows a rate-limiting configuration for the first six ports in the module in slot "A". In this instance:

- Ports A1–A4 are configured with an outbound rate limit of 200 Kbps.
- Port A5 is configured with an inbound rate limit of 20%.
- Port A6 is not configured for rate-limiting.

Figure 54: *Listing the rate-limit configuration*

```
HP-Switch# show rate-limit all a1-a6

 All-Traffic Rate Limit Maximum %

         | Inbound              Radius      | Outbound             Radius
  Port   | Limit     Mode       Override    | Limit     Mode       Override
  ------ + --------- --------- ----------- + --------- --------- -----------
  A1     | Disabled  Disabled  No-override | 200       kbps      No-override
  A2     | Disabled  Disabled  No-override | 200       kbps      No-override
  A3     | Disabled  Disabled  No-override | 200       kbps      No-override
  A4     | Disabled  Disabled  No-override | 200       kbps      No-override
  A5     | 20        %         No-override | Disabled  Disabled  No-override
  A6     | Disabled  Disabled  No-override | Disabled  Disabled  No-override
```

RADIUS-assigned rate-limit

To view **RADIUS**-assigned rate-limit information, use one of the following command options:

```
show port-access
    web-based clients <PORT-LIST> detailed
    mac-based clients <PORT-LIST> detailed
    authenticator clients <PORT-LIST> detailed
```

show running

The `show running` command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate-limiting.

The `show config` command displays this information for the configuration currently stored in the `startup-config` file. (Note that configuration changes performed with the CLI, but not followed by a `write mem` command, do not appear in the `startup-config` file.)

Figure 55: *Rate-limit settings listed in the* `show config` *output*

```
HP Switch(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.14.01

hostname "HP Switch 8212zl"
module 1 type J8705A
snmp-server community "public" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24
   ip address dhcp-bootp
   exit
interface A1
   rate-limit all out kbps 200
   exit
interface A2
   rate-limit all out kbps 200
   exit
interface A3
   rate-limit all out kbps 200
   exit
interface A4
   rate-limit all out kbps 200
   exit
interface A5
   rate-limit all in percent 200
   exit
interface A6
   rate-limit icmp percent 60
   rate-limit mcast in percent 60
   exit
```

Ports A1-A4 are configured with an outbound rate limit of 200 kbps.

Port A5 is configured with an inbound rate limit of 200 kbps.

Port A6 is configured with an inbound ICMP and multicast rate-limits of 60 percent each.

Configuring ICMP rate-limiting

ICMP rate-limiting provides a method for limiting the amount of bandwidth that may be used for inbound ICMP traffic on a switch port. This feature allows users to restrict ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be caused by worms or viruses (reducing their spread and effect.) In addition, ICMP rate-limiting preserves inbound port bandwidth for non-ICMP traffic.

This feature should not be used to remove all ICMP traffic from a network. ICMP is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior and should normally be configured to allow one to five percent of available inbound bandwidth (at 10 Mbps or 100 Mbps speeds) or 100 to 10,000 kbps (1Gbps or 10 Gbps speeds) to be used for ICMP traffic.

int rate-limit icmp

Syntax

`int <PORT-LIST> rate-limit icmp <ip-type> [kbps <0-10000000>|trap-clear]`

Description

The `rate-limit icmp` command controls inbound usage of a port by setting a limit on the bandwidth available for inbound ICMP traffic. Configures inbound ICMP traffic rate-limiting. You can configure a rate limit from either the global configuration level, or from the interface context level. (Default: Disabled.)

Parameters

percent

In the range of 1–100. Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.

kbps

In the range of 0–100000000. Specifies the rate at which to forward traffic in kilobits per second. Using 0 causes an interface to drop all incoming ICMP traffic and is not recommended.

Options

no

The `no` form of the command disables ICMP rate-limiting on the specified interfaces.

<IP-TYPE>

- `ip-all`: Set a rate limit for all ICMP traffic.
- `ipv4`: Set a rate limit for IPv4 ICMP traffic.
- `ipv6`: Set a rate limit for IPv6 ICMP traffic.
- `kbps`: Set the rate limit in kilobits per second.
- `percent`: Set the rate limit as a percentage of the port link speed.
- `trap-clear`: Clear an existing ICMP rate limiting trap condition.

configure an inbound rate limit

Either of the following commands configures an inbound rate limit of 1% on ports A3 to A5, which are used as network edge ports:

```
switch(config) # int a3-a5 rate-limit icmp percent 1
switch(eth-A3-A5) # rate-limit icmp percent 1
```

More information

For more information on ICMP rate-limiting operation, see **ICMP rate-limiting** on page 222.

Viewing the current ICMP rate-limit configuration

show rate-limit icmp

Syntax

`show rate-limit icmp <PORT-LIST>`

Description

The `show rate-limit icmp` command displays the per-interface ICMP rate-limit configuration in the running-config file.

Parameters

`show running`

Displays the currently applied setting for any interfaces in the switch configured for anyl traffic rate-limiting and ICMP rate-limiting.

`show config`

Displays this information for the configuration currently stored in the `startup-config` file. Configuration changes performed with the CLI, but not followed by a `write mem` command, do not appear in the `startup-config` file.

Options

`<PORT-LIST>`

Without *`<PORT-LIST>`* , this command lists the ICMP rate-limit configuration for all ports on the switch.

With *`<PORT-LIST>`* , this command lists the rate-limit configuration for the specified interfaces. This command operates the same way in any CLI context.

view a rate-limiting configuration

If you want to view the rate-limiting configuration on the first six ports in the module in slot "B":

Figure 56: *Listing the rate-limit configuration*

```
HP Switch(config)# show rate-limit icmp b1-b6

 Inbound ICMP Rate Limit Maximum Percentage

        |          Rate
  Port  | Mode     Limit
  ----- + -------- --------
  B1    | Disabled Disabled
  B2    | kbps     100
  B3    | %        5
  B4    | %        1
  B5    | %        1
  B6    | Disabled Disabled
```

Resetting the ICMP trap function of the port

Trap notification is enabled by default. When a trap notification is sent, it does not repeat unless the ICMP trap function is cleared.

int rate-limit

Syntax

```
int <PORT-LIST> rate-limit icmp trap-clear
```

Description

Resets the port ICMP trap function.

More information

You can also perform the reset through SNMP from a network management station or through the CLI with the `setmib` command.

setmib

Syntax

```
setmib hpIcmpRatelimitPortAlarmflag.internal-port-# [-i] 1
```

Description

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

reset the port and send a new message

An operator noticing an ICMP rate-limiting trap or Event Log message originating with port A1 on a switch could use either of the following commands to reset the port to send a new message if the condition occurs again:

```
switch(config)# int a1 rate-limit icmp trap-clear
switch# setmib hpicmpratelimitportalarmflag.1 -i 1
```

Determining the switch port number used in ICMP port reset commands

To enable excess ICMP traffic notification traps and Event Log messages, use the `setmib` command described on **ICMP rate-limiting trap** on page 225. The port number included in the command corresponds to the internal number the switch maintains for the designated port and not the port's external (slot/number) identity.

To match the port's external slot/number to the internal port number, use the `walkmib ifDescr` command, as shown in the following figure:

Figure 57: *Matching internal port numbers to external slot/port numbers*

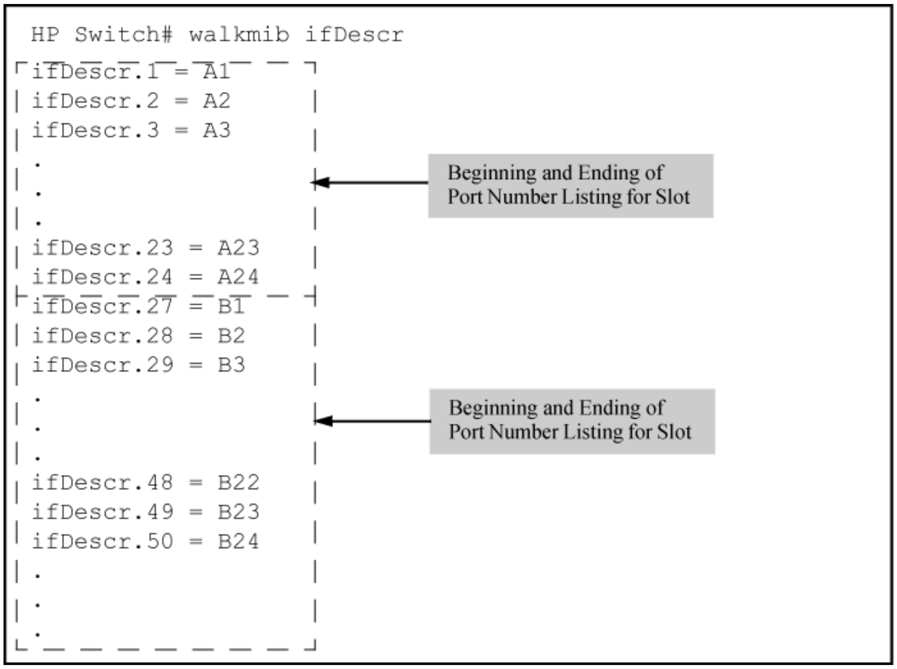

Configuring an egress/outbound broadcast limit on the switch

This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

broadcast-limit

Syntax

```
broadcast-limit 0-99
```

Description

Enables or disables broadcast limiting for outbound broadcasts on a selected port on the switch.

> You must switch to port context level before issuing the `broadcast-limit` command.

Options

<0–99>

The value selected is the percentage of traffic allowed, for example, `broadcast-limit 5` allows 5% of the maximum amount of traffic for that port. A value of zero disables broadcast limiting for that port.

port context level

Egress broadcast limiting on switches is configured on a per-port basis. You must be at the port context level for this command to work, for example:

```
switch(config) # int B1
switch(int B1) # broadcast-limit 1
```

show config

Syntax

`show config`

Description

Displays the `startup-config` file. The broadcast limit setting appears here if enabled and saved to the `startup-config` file.

Parameters and options

`running-config`

Displays the `running-config` file. The broadcast limit setting appears here if enabled. If the setting is not also saved to the `startup-config` file, rebooting the switch returns broadcast limit to the setting currently in the `startup-config` file.

---

enabling broadcast limits

The following command enables broadcast limiting of 1% of the outbound traffic rate on the selected port on the switch:

```
switch(int B1) # broadcast-limit 1
```

For a 1-Gbps port, this results in an outbound broadcast traffic rate of 10 Mbps.

Configuring inbound rate-limiting for broadcast and multicast traffic

rate-limit

Syntax

`rate-limit [bcast|mcast] in [percent <0-100>|kbps <0-100000000>]`

Description

Enables rate-limiting and sets limits for the specified inbound broadcast or multicast traffic. Only the amount of traffic specified by the percent is forwarded. You can configure rate-limiting (throttling) of inbound broadcast and multicast traffic on the switch, which helps prevent the switch from being disrupted by traffic storms if they occur on the rate-limited port. The rate-limiting is implemented as either a percentage of the total available bandwidth on the port or as kilobits per-second.

Default: Disabled

Parameters

`no`

—

`bcast`

—

`mcast`

—

`percent`

—

kbps

—

rate-limit command context

The `rate-limit` command can be executed from the global or interface context, for example:

```
switch# interface 3 rate-limit bcast in percent 10
```

or

```
switch# interface 3
switch(eth-3#) rate-limit bcast in percent 10
```

inbound broadcast rate-limit

To set a limit of 50% on inbound broadcast traffic for port 3, enter interface context for port 3 and then execute the `rate-limit` command, as shown in **Figure 58: Inbound broadcast rate-limiting of 50% on port 3** on page 210. Only 50% of the inbound broadcast traffic will be forwarded.

Figure 58: *Inbound broadcast rate-limiting of 50% on port 3*

```
HP Switch(config)# int 3
HP Switch(eth-3)# rate-limit bcast in percent 50

HP Switch(eth-3)# show rate-limit bcast
 Broadcast-Traffic Rate Limit Maximum %

  Port  | Inbound Limit Mode      Radius Override
  ----- + ------------- --------- ---------------
  1     | Disabled      Disabled  No-override
  2     | Disabled      Disabled  No-override
  3     | 50            %         No-override
  4     | Disabled      Disabled  No-override
  5     | Disabled      Disabled  No-override
```

multicast traffic rate-limit

If you rate-limit multicast traffic on the same port, the multicast limit is also in effect for that port, as shown in **Figure 59: Inbound multicast rate-limiting of 20% on port 3** on page 210. Only 20% of the multicast traffic will be forwarded.

Figure 59: *Inbound multicast rate-limiting of 20% on port 3*

```
HP Switch(eth-3)# rate-limit mcast in percent 20
HP Switch(eth-3)# show rate-limit mcast

 Multicast-Traffic Rate Limit Maximum %

  Port  | Inbound Limit Mode      Radius Override
  ----- + ------------- --------- ---------------
  1     | Disabled      Disabled  No-override
  2     | Disabled      Disabled  No-override
  3     | 20            %         No-override
  4     | Disabled      Disabled  No-override
```

disabling rate-limiting

To disable rate-limiting for a port enter the `no` form of the command, as shown in **Figure 60: Disabling inbound multicast rate-limiting for port 3** on page 211.

Figure 60: *Disabling inbound multicast rate-limiting for port 3*

```
HP Switch(eth-3)# no rate-limit mcast in

HP Switch(eth-3)# show rate-limit mcast

 Multicast-Traffic Rate Limit Maximum %

  Port  | Inbound Limit Mode      Radius Override
  ----- + ------------- --------- ---------------
  1     | Disabled      Disabled  No-override
  2     | Disabled      Disabled  No-override
  3     | Disabled      Disabled  No-override
  4     | Disabled      Disabled  No-override
```

Operating notes

- This rate-limiting option does not limit unicast traffic.
- This option does not include any form of outbound rate-limiting.

Egress per-queue rate-limiting

Egress Per-Queue Rate-Limiting is supported, including configuration on static trunks, on the Aruba 5400R and Aruba 3800 switches. (Egress Per-Queue Rate-Limiting is not supported on dynamic LACP trunks, distributed trunks, or Mesh ports.)

Overview

Egress rate-limiting permits administrators to configure the maximum percentage of traffic allowed to egress an interface for each priority queue.

- Egress per-queue rate-limiting allows configurations on both physical ports and static trunks.
- The number of queue percentages will vary based on the number of queues configured on the device (i.e. 2-queues, 4-queues, 8-queues).
- Configuration is allowed on a static trunk (manual HPE trunks and static LACP trunks), but the actual traffic enforcement occurs per-port on the individual ports belonging to the trunk.

Restrictions

- While limits on all egress traffic (`egress rate-limit all`) and limits on specific egress queues (`egress rate-limit queues`) can be configured at the same time on a given port (i.e., can be concurrent features), this may result in lower actual limits than expected. This is particularly true of queue-limits, where a packet may be dropped for the port as a whole even when the queue is below its limit.
- The egress per-queue rate-limiting is not configurable on dynamic LACP and Distributed trunks.
- Other rate-limiting features (ingress and egress) are not supported on trunked ports.

Configuration commands

show rate-limit queues

Syntax

`show rate-limit queues <PORT-LIST|TRK-LIST>`

Description

Using the `show rate-limit` command with the `queues` option added in software release 15.18 enables you to specify both individual ports and port trunk names to display the output. If nothing is specified, all physical ports and any static, non-DT trunks are displayed with their current settings previously configured with the `rate-limit queues` command. The optional *PORT-LIST* parameter limits the display output to the listed ports (and static, non-DT trunks, if any).

Command output when no port list specified

```
switch# show rate-limit queues

Outbound Queue-Based Rate-Limit %

Port   Q1  Q2  Q3  Q4  Q5  Q6  Q7  Q8
------ --- --- --- --- --- --- --- ---
A1     5   10  10  5   10  10  20  20
A2     5   10  10  5   10  10  20  20
A3     5   10  10  5   10  10  20  20
A4     5   10  10  5   10  10  20  20
A7     5   10  10  5   10  10  20  20

A22    5   10  10  5   10  10  20  20
F1     5   10  10  5   10  10  20  20

F24    5   10  10  5   10  10  20  20
Trk1   5   10  10  5   10  10  10  20
Trk6   5   10  10  5   10  10  10  20
```

Output with trunk queue set to 100 percent

```
switch# show rate-limit queues

 Outbound Queue-Based Rate-Limit %

  Port   Q1  Q2     Q3  Q4     Q5  Q6  Q7  Q8
  ------ --- ------ --- ------ --- --- --- ------
  A5     5   10     10  5      10  10  20  20
  A8     5   10     10  5      10  10  20  20
  A18    5   10     10  5      10  10  20  20
  Trk1   5   10     10  5      10  10  20  100
```

Output when port list specified

```
switch# show rate-limit queues A1-A4

Outbound Queue-Based Rate-Limit %

Port   Q1  Q2  Q3  Q4  Q5  Q6  Q7  Q8
------ --- --- --- --- --- --- --- ---
A1     5   10  10  5   10  10  20  20
A2     5   10  10  5   10  10  20  20
A3     5   10  10  5   10  10  20  20
A4     5   10  10  5   10  10  20  20
```

Output when trunk name specified

```
switch# show rate-limit queues Trk6

Outbound Queue-Based Rate-Limit %

Port   Q1  Q2  Q3  Q4  Q5  Q6  Q7  Q8
------ --- --- --- --- --- --- --- ---
Trk6   5   10  10  5   10  10  20  20
```

Guaranteed Minimum Bandwidth (GMB) for outbound traffic

Earlier software releases supported GMB configuration on a per-port basis. Now the 5400R, 3810, and 3800 switches also support GMB configuration on static trunks. (GMB configuration is not supported on dynamic LACP or distributed (DT) trunks.)

For any port, group of ports, or static trunk you can use the default minimum bandwidth settings for each outbound priority queue or a customized bandwidth profile. It is also possible to disable the feature entirely.

For application to static trunk interfaces, GMB enforcement is applied individually to each port belonging to the trunk, and not to the trunk as a whole.

By default, GMB is configured with a recommended profile for outgoing traffic that prevents higher-priority queues from starving lower-priority traffic. In the eight-queue configuration, the default values per priority queue are:

- Queue 1 (low priority): 2%
- Queue 2 (low priority): 3%
- Queue 3 (normal priority): 30%
- Queue 4 (normal priority): 10%
- Queue 5 (medium priority): 10%
- Queue 6 (medium priority): 10%
- Queue 7 (high priority): 15%
- Queue 8 (high priority): 20%

The value for each of the queues indicates the minimum percentage of port throughput that is guaranteed for that queue. If a given queue does not require its guaranteed minimum in a given service window, any extra bandwidth is allocated to the other queues, beginning with the highest-priority queue.

The actual number of queues can be two, four, or eight, depending on either the system default or the value set by the latest instance of the `qos queue-config <n-queues>` command. Per-queue values must be specified starting with queue 1 being the lowest priority and queue 8 being the highest priority. If desired, the highest-priority queue may be put into “strict” mode by specifying `strict` rather than a percentage value. In strict mode, the highest-priority queue gets all the bandwidth it needs, and any remaining bandwidth is shared among the non-strict queues based on their need and their configured bandwidth profiles. If no guaranteed minimum bandwidth is configured (i.e., the settings for all queues are 0), the traffic is serviced strictly by priority. In practice, this may cause complete starvation of some or all lower-priority queues during any periods where the output port traffic is over-subscribed.

This is an Interface context command. It can be called directly from the interface context, or following the `interface <PORT-LIST>` command. For most applications, Hewlett Packard Enterprise recommends having the same GMB profile on all the ports on a switch so that the outbound traffic profile is consistent for all outbound traffic. However, there may be instances where it may be advantageous to configure special profiles on connections to servers or to the network infrastructure (such as links to routers, other switches, or to the network core).

For more details on GMB operation, see **Guaranteed minimum bandwidth (GMB)** on page 226.

int bandwidth-min output

Syntax

`int <PORT-LIST|TRK-LIST> bandwidth-min output`

Description

Configures the default minimum bandwidth allocation for the outbound priority queue for each port in the *<PORT-LIST>*.

Parameters

no

—

Non-default GMB settings

For application to static trunk interface such as trk1, GMB enforcement is applied individually to each port belonging to the trunk, and not to the trunk as a whole.

You must specify a bandwidth percent value for all except the highest priority queue, which may instead be set to "strict" mode. The sum of the bandwidth percentages below the top queue cannot exceed 100%. (**0** can be used as a value for a queue percentage setting.)

Configuring a total of less than 100% across the outbound queue set results in unallocated bandwidth that remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority.

For example, if you configure a minimum of 10% for queues 1 to 7 and 0% for queue 8, the unallocated bandwidth is available to all eight queues in the following prioritized order:

Queue 7 (high priority)
Queue 6 (medium priority)
Queue 5 (medium priority)
Queue 4 (normal priority)
Queue 3 (normal priority)
Queue 2 (low priority)
Queue 1 (low priority)
Queue 8 (high priority)

In practice, these priorities are the result of the configured minimum of 10% for queues 1 through 7 and 0% for queue 8. However, the switch does check queue 8 periodically and services it any time the bandwidth needed in a lower-priority queue goes below its minimum.

A setting of **0** (zero percent) on a queue means that no bandwidth minimum is specifically reserved for that queue for each of the ports in the *<PORT-LIST>*.

Also, there is no benefit to setting the high-priority queue (queue 8) to **0** (zero) unless you want the medium queue (queue 4) to be able to support traffic bursts above its guaranteed minimum.

Using Strict mode

Strict mode provides the ability to configure the highest priority queue as strict. Per-queue values must be specified in priority order, with queue 1 having the lowest priority and queue 8 (or 4, or 2) having the highest priority. (The highest queue is determined by how many outbound queues are configured on the switch. Two, four, and eight queues are permitted. (See the `qos queue-config` command.) The strict queue is provided all the bandwidth it needs. Any remaining bandwidth is shared among the non-strict queues based on need and configured bandwidth profiles. (The profiles are applied to the leftover bandwidth in this case.) The total sum of percentages for non-strict queues must not exceed 100.

Configuring 0% for a queue can result in that queue being starved if any higher queue becomes over-subscribed and is then given all unused bandwidth.

The switch applies the bandwidth calculation to the link speed the port is currently using. For example, if a 10/100 Mbs port negotiates to 10 Mbps on the link, it bases its GMB calculations on 10 Mbps, not 100 Mbps.

Use `show bandwidth output<PORT-LIST|TRK-LIST>` to display the current GMB configuration. (The `show config` and `show running` commands do not include GMB configuration data.)

Outbound minimum bandwidth

To configure the following outbound minimum bandwidth availability for ports A1 through A5:

Priority of outbound port queue	Minimum bandwidth %	Effect on outbound bandwidth allocation
8	20%	Queue 8 has the first priority use of all outbound bandwidth not specifically allocated to queues 1 to 7. If, for example, bandwidth allocated to queue 5 is not being used and queues 7 and 8 become oversubscribed, queue 8 has first-priority use of the unused bandwidth allocated to queue 5.
7	15%	Queue 7 has a GMB of 15% available for outbound traffic. If queue 7 becomes oversubscribed and queue 8 is not already using all of the unallocated bandwidth, queue 7 can use the unallocated bandwidth. Also, any unused bandwidth allocated to queues 6 to queue 1 is available to queue 7 if queue 8 has not already claimed it.
6	10%	Queue 6 has a GMB of 10% and, if oversubscribed, is subordinate to queues 8 and 7 in priority for any unused outbound bandwidth available on the port.
5	10%	Queue 5 has a GMB of 10% and, if oversubscribed, is subordinate to queues 8, 7, and 6 for any unused outbound bandwidth available on the port.
4	10%	Queue 4 has a GMB of 10% and, if oversubscribed, is subordinate to queues, 8, 7, 6, and 5 for any unused outbound bandwidth available on the port.
3	30%	Queue 3 has a GMB of 30% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, and 4 for any unused outbound bandwidth available on the port.
2	3%	Queue 2 has a GMB of 3% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, 4, and 3 for any unused outbound bandwidth available on the port.
1	2%	Queue 1 has a GMB of 2% and, if oversubscribed, is subordinate to all the other queues for any unused outbound bandwidth available on the port.

Either of the following commands configures ports A1 through A5 with bandwidth settings:

```
switch(config) # int a1-a5 bandwidth-min output 2 3 30 10 10 10 15 strict
switch(eth-A1-A5) # bandwidth-min output 2 3 30 10 10 10 15 strict
```

int bandwidth-min output

Syntax

```
int <PORT-LIST|TRK-list> output <queue1_%> <queue2_%> <queue3_%> <queue_4%>
<queue5_%> <queue6_%> <queue7_%> <queue8_%>
```

Description

You can configure bandwidth minimums from either the global configuration level (as shown above) or from the port context level, however you must configure one minimum bandwidth percent setting for each outbound queue.

Parameters

no

Disables GMB for all ports in the *PORT-LIST*. In this state, which is the equivalent of setting all outbound queue minimum guarantees on a port to **0** (zero), a high level of higher-priority traffic can starve lower-priority queues, which can slow or halt lower-priority traffic in the network.

strict

Applies only to the highest-priority (last) outbund queue for each port affected by the command.

Options

<queueN_%>

A value from 0 to 100.

Minimum outbound guaranteed bandwidth

For ports in *PORT-LIST* (including static trunks) this command specifies the minimum outbound guaranteed bandwidth as a percent of the total bandwidth for each outbound queue. The queues receive service in descending order of priority. For example, to configure GMB on port A10 and trunk trk1, you would use a command with bandwidth values similar to the following:

```
switch# int a10,trk1 bandwidth-min output 2 3 30 10 10 10 15 20
```

Viewing the current GMB configuration

show bandwidth output

Syntax

```
show bandwidth output <PORT-LIST|TRK-LIST>
```

Description

This command displays the per-port GMB configuration in the `running-config` file. This command operates the same way in any CLI context. If the command lists `Disabled` for a port, there are no bandwidth minimums configured for any queue on the port.

Options

<PORT-LIST>

Without *PORT-LIST*, this command lists the GMB configuration for all ports on the switch.

With *PORT-LIST*, this command lists the GMB configuration for the specified ports.

Display the GMB configuration

To display the GMB configuration resulting from either of the above commands:

Figure 61: *Listing the GMB configuration*

```
switch# show bandwidth output a1-a5

 Outbound Guaranteed Minimum Bandwidth %

  Port   Q1  Q2     Q3  Q4     Q5  Q6  Q7  Q8
  ------ --- ------ --- ------ --- --- --- ------
  A1     2   3      30  10     10  10  15  strict
  A2     2   3      30  10     10  10  15  strict
  A3     2   3      30  10     10  10  15  strict
  A4     2   3      30  10     10  10  15  strict
  A5     2   3      30  10     10  10  15  strict
```

GMB configuration in the startup-config file

The following figure shows how the preceding listing of the GMB configuration would appear in the `startup-config` file.

Figure 62: *GMB settings listed in the `show config` output*

```
switch# show config status
Running configuration is same as the startup configuration
switch# show config

Startup configuration:
; J9821A configuration Editor; Created on release #KB.15.18.0001

hostnme "HP Switch"
module 1 type J9986A
snmp-server community "public" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24
   ip address dhcp-bootp
   exit
interface A1
   bandwidth-min output 2 3 30 10 10 10 15 strict
   exit
interface A2
   bandwidth-min output 2 3 30 10 10 10 15 strict
   exit
interface A3
   bandwidth-min output 2 3 30 10 10 10 15 strict
   exit
interface A4
   bandwidth-min output 2 3 30 10 10 10 15 strict
   exit
interface A5
   bandwidth-min output 2 3 30 10 10 10 15 strict
   exit
```

Output when trunk name specified

```
switch# show bandwidth output Trk1
```

```
Outbound Guaranteed Minimum Bandwidth %

Port   Q1  Q2     Q3  Q4     Q5  Q6  Q7  Q8
------ --- ------ --- ------ --- --- --- ------
Trk1   10  15     10  15     10  15  10  15
```

Output when no port list specified

```
switch# show bandwidth output
Outbound Guaranteed Minimum Bandwidth %

Port   Q1  Q2     Q3  Q4     Q5  Q6  Q7  Q8
 ------ --- ------ --- ------ --- --- --- ------
A1     2   3      30  10     10  10  15  20
A2     10  15     10  15     10  15  10  15
A3     2   3      30  10     10  10  15  20
A4     2   3      30  10     10  10  15  20
A5     2   3      30  10     10  10  15  20
A6     2   3      30  10     10  10  15  20
A7     2   3      30  10     10  10  15  20
A8     2   3      30  10     10  10  15  20
A9     2   3      30  10     10  10  15  20
A10    2   3      30  10     10  10  15  20
A11    2   3      30  10     10  10  15  20
A12    2   3      30  10     10  10  15  20
A13    2   3      30  10     10  10  15  20
A14    2   3      30  10     10  10  15  20
A15    2   3      30  10     10  10  15  20
A16    2   3      30  10     10  10  15  20
A17    2   3      30  10     10  10  15  20
A18    2   3      30  10     10  10  15  20
A19    2   3      30  10     10  10  15  20
A20    2   3      30  10     10  10  15  20
A21    2   3      30  10     10  10  15  20
A22    2   3      30  10     10  10  15  20
A23    2   3      30  10     10  10  15  20
A24    2   3      30  10     10  10  15  20
F1     2   3      30  10     10  10  15  20
F2     2   3      30  10     10  10  15  20
F3     2   3      30  10     10  10  15  20
F4     2   3      30  10     10  10  15  20
F5     2   3      30  10     10  10  15  20
F6     2   3      30  10     10  10  15  20
F7     2   3      30  10     10  10  15  20
F8     2   3      30  10     10  10  15  20
F9     2   3      30  10     10  10  15  20
F10    2   3      30  10     10  10  15  20
F11    2   3      30  10     10  10  15  20
F12    2   3      30  10     10  10  15  20
F13    2   3      30  10     10  10  15  20
F14    2   3      30  10     10  10  15  20
F15    2   3      30  10     10  10  15  20
F16    2   3      30  10     10  10  15  20
F17    2   3      30  10     10  10  15  20
F18    2   3      30  10     10  10  15  20
F19    2   3      30  10     10  10  15  20
F20    2   3      30  10     10  10  15  20
F21    2   3      30  10     10  10  15  20
F22    2   3      30  10     10  10  15  20
F23    2   3      30  10     10  10  15  20
F24    2   3      30  10     10  10  15  20
```

```
Trk1    10  15      10  15      10  15  10  15
Trk2    15  10      15  10      15  10  15  10
```

Validation rules

Cause

Validation	Error/Warning/Prompt
Rate-limit queues out percent command	
Valid port number?	Invalid port number
Valid trunk interface?	Invalid trunk interface
Trunk type supported?	Unsupported trunk type
Maximum bandwidth value is greater than the minimum bandwidth configured for a queue?	Invalid maximum value.
Bandwidth-min output command	
Valid trunk interface?	Invalid trunk interface
Trunk type supported?	Unsupported trunk type
Minimum bandwidth value is lesser than the maximum bandwidth configured for a queue?	Invalid minimum value.
Show rate-limit queues command	
Valid port number?	Invalid port number
Valid trunk interface?	Invalid trunk interface
Trunk type supported?	Unsupported trunk type
Show bandwidth output command	
Valid trunk interface?	Invalid trunk interface
Trunk type supported?	Unsupported trunk type

Event log

Event	Message
Invalid port number	The port number **<port num>** entered is invalid.
Invalid trunk interface	The trunk **<trunk name>** entered is invalid.

Table Continued

Event	Message
Unsupported trunk type	This command is not supported on distributed or dynamic trunks.
Invalid maximum value	The maximum bandwidth value <max value > entered should be greater than the minimum bandwidth value <min value> configured.

All traffic rate-limiting

Rate-limiting for all traffic operates on a per-port basis to allow only the specified bandwidth to be used for inbound or outbound traffic. When traffic exceeds the configured limit, it is dropped. This effectively sets a usage level on a given port and is a tool for enforcing maximum service level commitments granted to network users. This feature operates on a per-port level and is not configurable on port trunks. Rate-limiting is designed to be applied at the network edge to limit traffic from non-critical users or to enforce service agreements such as those offered by Internet Service Providers (ISPs) to provide only the bandwidth for which a customer has paid.

Rate-limiting also can be applied by a RADIUS server during an authentication client session. (See the access security guide.)

> Rate-limiting is intended for use on edge ports in a network. Hewlett Packard Enterprise does not recommend it for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

The switches also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks.

The mode using bits per second (bps) in releases before K.12.XX has been replaced by the kilobits per second (kbps) mode. Switches that have configurations with bps values are automatically converted when you update your software to the new version. However, you must manually update to kbps values an older config file that uses bps values or it will not load successfully onto a switch running later versions of the software (K.12.XX or greater.)

- The `rate-limit icmp` command specifies a rate limit on inbound ICMP traffic only (See **ICMP rate-limiting** on page 222)
- Rate-limiting does not apply to trunked ports (including meshed ports.)
- Kbps rate-limiting is done in segments of 1% of the lowest corresponding media speed.

  For example, if the media speed is 100 Kbps, the value would be 1 Mbps.

  - A 1 to 100 Kbps rate-limit is implemented as a limit of 100 Kbps
  - A limit of 101 to 199 Kbps is also implemented as a limit of 200 Kbps.
  - A limit of 201 to 299 Kbps is implemented as a limit of 300 Kbps, and so on.
- Percentage limits are based on link speed.

  For example, if a 100 Mbps port negotiates a link at 100 Mbps and the inbound rate-limit is configured at 50%, the traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, it allows no more than 5 Mbps of inbound traffic.
- Configuring a rate limit of 0 (zero) on a port blocks all traffic on that port. However, if this is the desired behavior on the port, Switch recommends that you use the `<PORT-LIST> disable` command instead of configuring a rate limit of 0.
- You can configure a rate limit from either the global configuration level or from the port context level.

configure an inbound rate limit

Either of the following commands configures an inbound rate limit of 60% on ports A3 to A5:

```
switch(config #) int a3-a5 rate-limit all in percent 60
switch(eth-A3-A5)# rate-limit all in percent 60
```

Operating notes for rate-limiting

- In general, desirable traffic should not be rate-limited.
- When going from a switch with faster links to a switch with slower links, it is better to force the speed of the port connection to be slower rather than to rate-limit the traffic.
- Rate-limiting operates on a per-port basis, regardless of traffic priority. Rate-limiting is available on all types of ports and at all port speeds configurable for these switches.
- Except for the `egress per-queue` option with static trunks on 5400R and 3800 ProVision switches, rate-limiting is not supported on trunked ports (including mesh ports.) Where trunked ports are not supported, configuring a port for rate-limiting and then adding it to a trunk suspends rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message: `<PORT-LIST> : Operation is not allowed for a trunked port.`
- Rate-limiting for inbound and outbound traffic are separate features. The rate limits for each direction of traffic flow on the same port are configured separately—even the specified limits can be different.
- Rate-limiting and hardware: The granularity of actual limits may vary across different switch models.
- Rate-limiting is visible as an outbound forwarding rate. Because inbound rate-limiting is performed on packets during packet-processing, it is not shown via the inbound drop counters. Instead, this limit is verifiable as the ratio of outbound traffic from an inbound rate-limited port versus the inbound rate. For outbound rate-limiting, the rate is visible as the percentage of available outbound bandwidth (assuming that the amount of requested traffic to be forwarded is larger than the rate-limit.)
- Operation with other features: Configuring rate-limiting on a port where other features affect port queue behavior (such as flow control) can result in the port not achieving its configured rate-limiting maximum. For example, in a situation where flow control is configured on a rate-limited port, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that is lower than the configured rate limit. In this case, the inbound traffic flow does not reach the configured rate and lower priority traffic is not forwarded into the switch fabric from the rate-limited port. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.)

  In another type of situation, an outbound port can become oversubscribed by traffic received from multiple rate-limited ports. In this case, the actual rate for traffic on the rate-limited ports may be lower than configured because the total traffic load requested to the outbound port exceeds the port's bandwidth, and thus some requested traffic may be held off on inbound.
- Traffic filters on rate-limited ports. Configuring a traffic filter on a port does not prevent the switch from including filtered traffic in the bandwidth-use measurement for rate-limiting when it is configured on the same port. For example, ACLs, source-port filters, protocol filters, and multicast filters are all included in bandwidth usage calculations.
- Monitoring (mirroring) rate-limited interfaces. If monitoring is configured, packets dropped by rate-limiting on a monitored interface are still forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)
- Optimum rate-limiting operation. Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.

Rate-limiting is applied to the available bandwidth on a port and not to any specific applications running through the port. If the total bandwidth requested by all applications is less than the configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing applications, as well as most regular network applications. Consider the following example that uses the minimum packet size:

The total available bandwidth on a 100 Mbps port "X" (allowing for Inter-packet Gap-IPG), with no rate-limiting restrictions, is:

```
(((100,000,000 bits) / 8 ) / 84) × 64 = 9,523,809 bytes per second
```

where:

- The divisor (84) includes the 12-byte IPG, 8-byte preamble, and 64-bytes of data required to transfer a 64-byte packet on a 100 Mbps link.
- Calculated "bytes-per-second" includes packet headers and data. This value is the maximum "bytes-per-second" that 100 Mbps can support for minimum-sized packets.

Suppose port "X" is configured with a rate limit of 50% (4,761,904 bytes.) If a throughput-testing application is the only application using the port and transmits 1 Mbyte of data through the port, it uses only 10.5% of the port's available bandwidth, and the rate-limit of 50% has no effect. This is because the maximum rate permitted (50%) exceeds the test application's bandwidth usage (126,642-164,062 bytes, depending upon packet size, which is only 1.3% to 1.7% of the available total.) Before rate-limiting can occur, the test application's bandwidth usage must exceed 50% of the port's total available bandwidth. That is, to test the rate-limit setting, the following must be true:

```
bandwidth usage (0.50 × 9,523,809)
```

ICMP rate-limiting

As of software version K.15.02.0004, ICMP rate-limiting and classifier-based-rate-limiting operate on the entire packet length instead of just the IP payload part of the packet. As a result, the effective metering rate is now the same as the configured rate. The rate-limiting applies to these modules:

HPE device	Product number	Minimum supported software version
HPE Switch 24-port 10/100/1000 PoE+ v2 zl Module	J9534A	K.15.02.0004
HPE Switch 20-port 10/100/1000 PoE+ / 4-port SFP v2 zl Module	J9535A	K.15.02.0004
HPE Switch 20-port 10/100/1000 PoE+ / 2-port 10-GbE SFP+ v2 zl Module	J9536A	K.15.02.0004
HPE Switch 24-port SFP v2 zl Module	J9537A	K.15.02.0004
HPE Switch 8-port 10-GbE SFP+ v2 zl Module	J9538A	K.15.02.0004
HPE 24-port 10/100 PoE+ v2 zl Module	J9547A	K.15.02.0004
HPE 20-port Gig-T / 2-port 10-GbE SFP+ v2 zl Module	J9548A	K.15.02.0004
HPE 20-port Gig-T / 4-port SFP v2 zl Module	J9549A	K.15.02.0004
HPE 24-port Gig-T v2 zl Module	J9550A	K.15.02.0004
HPE 12-port Gig-T / 12-port SFP v2 zl Module	J9637A	K.15.02.0004

ICMP rate-limiting provides a method for limiting the amount of bandwidth that may be used for inbound ICMP traffic on a switch port. This feature allows users to restrict ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be caused by worms or viruses (reducing their spread and effect.) In addition, ICMP rate-limiting preserves inbound port bandwidth for non-ICMP traffic.

CAUTION

This feature should not be used to remove all ICMP traffic from a network. ICMP is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior and should normally be configured to allow one to five percent of available inbound bandwidth (at 10 Mbps or 100 Mbps speeds) or 100 to 10,000 kbps (1Gbps or 10 Gbps speeds) to be used for ICMP traffic.

In IP networks, ICMP messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers.) Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network.

ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can separately configure both ICMP rate-limiting and all-traffic rate-limiting.

Beginning with software release K.12.xx or later, the all-traffic rate-limiting command (`rate-limit all`) and the ICMP rate-limiting command (`rate-limit icmp`) operate differently:

- All-traffic rate-limiting applies to both inbound and outbound traffic and can be specified either in terms of a percentage of total bandwidth or in terms of bits per second;
- ICMP rate-limiting applies only to inbound traffic and can be specified as only a percentage of total bandwidth.

ICMP rate-limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain.)

ICMP rate-limiting

Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. **Figure 63: ICMP rate-limiting** on page 223 shows an example of how to configure this for a small to mid-sized campus though similar rate-limit thresholds are applicable to other network environments. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. ("Normal" ICMP traffic levels should be the maximums that occur when the network is rebooting.)

Figure 63: *ICMP rate-limiting*

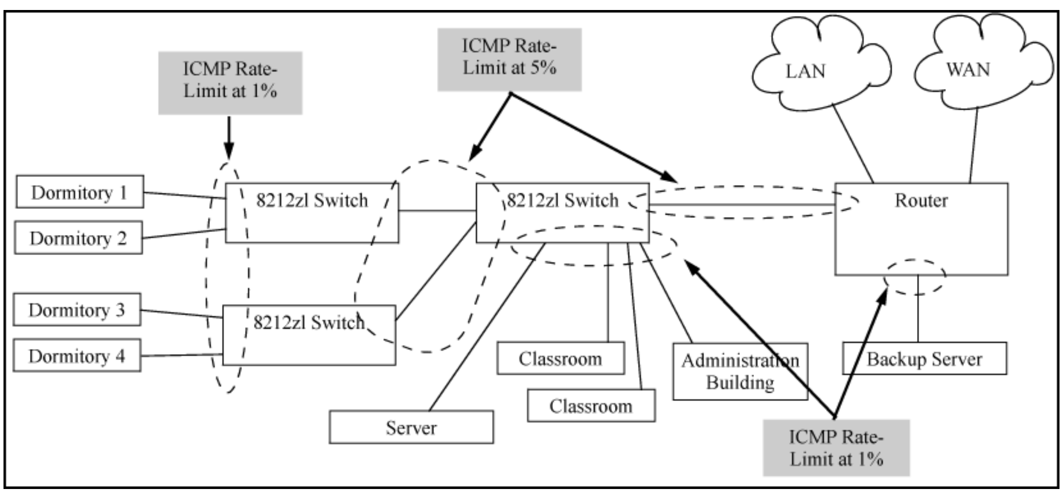

When using kbps-mode ICMP rate-limiting, the rate-limiting operates on only the IP payload part of the ICMP packet (as required by metering RFC 2698.) This means that effective metering is at a rate greater than the configured rate, with the disparity increasing as the packet size decreases (the packet to payload ratio is higher.)

Also, in kbps mode, metering accuracy is limited at low values, for example, less than 45 Kbps. This is to allow metering to function well at higher media speeds such as 10 Gbps.

Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface

ICMP and all-traffic rate-limiting can be configured on the same interface. All-traffic rate-limiting applies to all inbound or outbound traffic (including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.

If the all-traffic load on an interface meets or exceeds the currently configured all-traffic inbound rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, all excess traffic is dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has been reached.)

Example

Suppose:

- The all-traffic inbound rate-limit on port "X" is configured at 55% of the port's bandwidth.
- The ICMP traffic rate-limit on port "X" is configured at 2% of the port's bandwidth.

If at a given moment:

- Inbound ICMP traffic on port "X" is using 1% of the port's bandwidth, and
- Inbound traffic of all types on port "X" demands 61% of the ports's bandwidth,

all inbound traffic above 55% of the port's bandwidth, including any additional ICMP traffic, is dropped as long as all inbound traffic combined on the port demands 55% or more of the port's bandwidth.

Operating notes for ICMP rate-limiting

ICMP rate-limiting operates on an interface (per-port) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic.

On a given port, ICMP rate-limiting and classifier-based QoS are mutually exclusive. However, you can include ICMP rate-limiting as part of a larger classifier-QoS policy on a given port.

Interface support

ICMP rate-limiting is available on all types of ports (other than trunk ports or mesh ports), and at all port speeds configurable for the switch.

Rate-limiting is not permitted on mesh ports

Either type of rate-limiting (all-traffic or ICMP) can reduce the efficiency of paths through a mesh domain.

Except for the egress per-queue feature on 5400R and 3800 switches, rate-limiting is not supported on port trunks

All-traffic, bcast, ICMP, and mcast rate-limiting are not supported on ports configured in a trunk group.

ICMP percentage-based rate-limits are calculated as a percentage of the negotiated link speed

For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured at 5%, the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, it allows 0.5 Mbps of inbound traffic. If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured.)

ICMP rate-limiting is port-based

ICMP rate-limiting reflects the available percentage of an interface's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.

Below-maximum rates

ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.) In cases where both types of rate-limiting (`rate-limit all` and `rate-limit icmp`) are configured on the same interface, this situation is more likely to occur.

In another type of situation, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited interfaces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.

Monitoring (mirroring) ICMP rate-limited interfaces

If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface are still forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)

Optimum rate-limiting operation

Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.

Outbound traffic flow

Configuring ICMP rate-limiting on an interface does **not** control the rate of outbound traffic flow on the interface.

Testing ICMP rate-limiting

ICMP rate-limiting is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, no ICMP rate-limit can be applied. That is, an interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both `rate-limit all` and `rate-limit icmp`, the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit you need to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, you need to ensure that the ICMP traffic volume exceeds the configured maximum.

When testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP traffic.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

ICMP rate-limiting trap

If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of when the event occurred on the port.) For example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded configured limit on port A1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further; the switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port’s ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the `trap-clear` command option or the `setmib` command.

Guaranteed minimum bandwidth (GMB)

GMB provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum consideration for sending traffic out on the link to another device. This can prevent a condition where applications generating lower-priority traffic in the network are frequently or continually "starved" by high volumes of higher-priority traffic. You can configure GMB per-port.

GMB operations

The switch services per-port outbound traffic in a descending order of priority; that is, from the highest priority to the lowest priority. By default, each port offers eight prioritized, outbound traffic queues. Tagged VLAN traffic is prioritized according to the 802.1p priority the traffic carries. Untagged VLAN traffic is assigned a priority of **0** (normal.)

Table 15: *Per-port outbound priority queues*

802.1p Priority settings in tagged VLAN packets[1]	Outbound priority queue for a given port
1 (low)	1
2 (low)	2
0 (normal)	3
3 (normal)	4
4 (medium)	5
5 (medium)	6
6 (high)	7
7 (high)	8

[1]The switch processes outbound traffic from an untagged port at the "0" (normal) priority level.

You can use GMB to reserve a specific percentage of each port's available outbound bandwidth for each of the eight priority queues. This means that regardless of the amount of high-priority outbound traffic on a port, you can ensure that there will always be bandwidth reserved for lower-priority traffic.

Since the switch services outbound traffic according to priority (highest to lowest), the highest-priority outbound traffic on a given port automatically receives the first priority in servicing. Thus, in most applications, it is necessary only to specify the minimum bandwidth you want to allocate to the lower priority queues. In this case, the high-priority traffic automatically receives all unassigned bandwidth without starving the lower-priority queues.

Conversely, configuring a bandwidth minimum on only the high-priority outbound queue of a port (and not providing a bandwidth minimum for the lower-priority queues) is not recommended, because it may "starve" the lower-priority queues.

NOTE

For a given port, when the demand on one or more outbound queues exceeds the minimum bandwidth configured for those queues, the switch apportions unallocated bandwidth to these queues on a priority basis. As a result, specifying a minimum bandwidth for a high-priority queue but not specifying a minimum for lower-priority queues can starve the lower-priority queues during periods of high demand on the high priority queue. For example, if a port configured to allocate a minimum bandwidth of 80% for outbound high-priority traffic experiences a demand above this minimum, this burst starves lower-priority queues that **do not have a minimum configured**. Normally, this will not altogether halt lower priority traffic on the network, but will likely cause delays in the delivery of the lower-priority traffic.

The sum of the GMB settings for all outbound queues on a given port cannot exceed 100%.

Impacts of QoS queue configuration on GMB operation

The section on **Guaranteed Minimum Bandwidth (GMB) for outbound traffic** on page 213 assumes the ports on the switch offer eight prioritized, outbound traffic queues. This may not always be the case, however, because the switch supports a QoS queue configuration feature that allows you to reduce the number of outbound queues from eight (the default) to four queues, or two.

Changing the number of queues affects the GMB commands (`interface bandwidth-min` and `show bandwidth output`) such that they operate only on the number of queues currently configured. If the queues are reconfigured, the guaranteed minimum bandwidth per queue is automatically re-allocated according to the following percentages:

Table 16: *Default GMB percentage allocations per QoS queue configuration*

<table>
<tr><th>802.1p priority</th><th>8 queues (default)</th><th>4 queues</th><th>2 queues</th></tr>
<tr><td>1 (lowest)</td><td>2%</td><td rowspan="2">10%</td><td rowspan="4">90%</td></tr>
<tr><td>2</td><td>3%</td></tr>
<tr><td>0 (normal)</td><td>30%</td><td rowspan="2">70%</td></tr>
<tr><td>3</td><td>10%</td></tr>
<tr><td>4</td><td>10%</td><td rowspan="2">10%</td><td rowspan="4">10%</td></tr>
<tr><td>5</td><td>10%</td></tr>
<tr><td>6</td><td>15%</td><td rowspan="2">10%</td></tr>
<tr><td>7 (highest)</td><td>20%</td></tr>
</table>

For more information on queue configuration and the associated default minimum bandwidth settings, (see the advanced traffic management guide.)

Impact of QoS queue configuration on GMB commands.

Changing the number of queues causes the GMB commands (`interface bandwidth-min` and `show bandwidth output`) to operate only on the number of queues currently configured. In addition, when the `qos` `queue-config` command is executed, any previously configured `bandwidth-min output` settings are removed from the startup configuration.

Jumbo frames

The maximum transmission unitize (MTU) IP frame the switch can receive for Layer 2 frames inbound on a port. The switch drops any inbound frames larger than the MTU allowed on the port. Ports operating at a minimum of 10 Mbps on the 3500 switches and 1 Gbps on the other switches covered in this guide can accept forward frames of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. You can enable inbound jumbo frames on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and **operating** at a minimum of 10 Mbps on the 3500 switches and 1 Gbps on the other switches covered in this guide allow inbound jumbo frames of up to 9220 bytes.

Switch model	Minimum speed for jumbo traffic
3500	10 Mbps
All others in this guide	1 Gbps

Operating rules for jumbo frames

Required port speed

This feature allows inbound and outbound jumbo frames on ports operating at a minimum of 10 Mbps on the 3500 switches and 1 Gbps on the other switches.

Switch meshing

If you enable jumbo traffic on a VLAN, all meshed ports on the switch are enabled to support jumbo traffic. (On a given meshed switch, every meshed port operating at 1 Gbps or higher becomes a member of every VLAN configured on the switch.)

GVRP operation

A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.

Port adds and moves

If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.

Jumbo traffic sources

A port belonging to a jumbo-enabled VLAN can receive inbound jumbo frames through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, port 1 can receive jumbo traffic from devices on either VLAN.

Jumbo traffic-handling

- HPE Networking does not recommend configuring a voice VLAN to accept jumbo frames. Voice VLAN frames are typically small, and allowing a voice VLAN to accept jumbo frame traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo frames on all ports belonging to the VLAN.
- When the switch applies the default MTU (1522-bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in the VLAN can receive incoming frames of up to 1522 bytes. When the switch applies the jumbo MTU (9220

bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in that VLAN can receive incoming frames of up to 9220 bytes. A port receiving frames exceeding the applicable MTU drops such frames, causing the switch to generate an Event Log message and increment the "Giant Rx" counter (displayed by `show interfaces <PORT-LIST>`.)

- The switch allows flow control and jumbo frame capability to co-exist on a port.
- The default MTU is 1522 bytes (including 4 bytes for the VLAN tag.) The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag.)
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving "excessive"inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition also increments the switch's "Giant Rx" counter.
- If you do not want all ports in a given VLAN to accept jumbo frames, you can consider creating one or more jumbo VLANs with a membership comprising only the ports you want to receive jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo frames through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN.

  For example, suppose you want to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200 and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-enabled	No	No	Yes

  If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.
- Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo-enabled VLANs. This can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability, as shown in **Figure 64: Forwarding jumbo frames through non-jumbo ports** on page 229.

Figure 64: *Forwarding jumbo frames through non-jumbo ports*

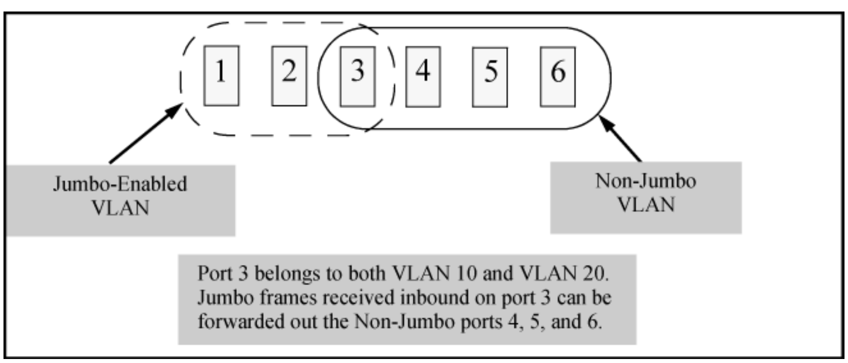

Jumbo frames can also be forwarded out non-jumbo ports when the jumbo frames received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

- If a switch belongs to a meshed domain, but does not have any VLANs configured to support jumbo traffic, the meshed ports on that switch drop any jumbo frames they receive from other devices. In this regard, if a mesh domain includes any HPE 1600M/2400M/2424M/4000M/8000M switches, along with the switches covered in this guide configured to support jumbo traffic, only the switches covered in this guide receive jumbo frames. The other switch models in the mesh will drop such frames. For more information on switch meshing, see the advanced traffic management guide.

Jumbo frame maximum size

The maximum frame size for jumbos is supported with the following proprietary MIB object:

`hpSwitchMaxFrameSize OBJECT-TYPE`

This is the value of the global `max-frame-size` supported by the switch. The default value is set to 9216 bytes.

Jumbo IP MTU

The IP MTU for jumbos is supported with the following proprietary MIB object:

`hpSwitchIpMTU OBJECT-TYPE`

This is the value of the global jumbos IP MTU (or L3 MTU) supported by the switch. The default value is set to 9198 bytes (a value that is 18 bytes less than the largest possible maximum frame size of 9216 bytes.) This object can be used only in switches that support `max-frame-size` and `ip-mtu` configuration.

Configuring jumbo frame operation

Prerequisites

Procedure

1. Determine the VLAN membership of the ports or trunks through which you want the switch to accept inbound jumbo traffic. For operation with GVRP enabled, refer to the GVRP topic under “Operating Rules”, above.
2. Ensure that the ports through which you want the switch to receive jumbo frames are operating at least at gigabit speed. (Check the Mode field in the output for the `show interfaces brief <PORT-LIST>` command.)
3. Use the `jumbo` command to enable jumbo frames on one or more VLANs statically configured in the switch. (All ports belonging to a jumbo-enabled VLAN can receive jumbo frames.
4. Execute `write memory` to save your configuration changes to the `startupconfig` file.

View the current jumbo configuration

show vlans

Syntax

`show vlans [port <PORT-LIST>] <VID>`

Description

Lists the static VLANs configured on the switch and includes a Jumbo column to indicate which VLANs are configured to support inbound jumbo traffic. All ports belonging to a jumbo-enabled VLAN can receive jumbo traffic.

Options

ports

Lists the static VLANs to which the specified ports belong, including the `Jumbo` column to indicate which VLANs are configured to support jumbo traffic.

<PORT-LIST>

Entering only one port in *<PORT-LIST>* results in a list of all VLANs to which that port belongs.

Entering multiple ports in *<PORT-LIST>* results in a superset list that includes the VLAN memberships of all ports in the list, even though the individual ports in the list may belong to different subsets of the complete VLAN listing.

<VID>

Shows port membership and jumbo configuration for the specified *vid* . (See **Figure 67: Listing the port membership and jumbo status for a VLAN** on page 232.)

show vlans

Figure 65: *Listing of static VLANs to show jumbo status per VLAN*

```
HP Switch(config)# show vlans
 Status and Counters - VLAN Information

  Maximum VLANs to support : 256
  Primary VLAN : DEFAULT_VLAN
  Management VLAN :

  VLAN ID Name                             Status       Voice  Jumbo
  ------- -------------------------------- ------------ -----  -----
  1       DEFAULT_VLAN                     Port-based   No     Yes
  5       VLAN5                            Port-based   No     No
  22      VLAN22                           Port-based   No     No
```

Indicates which static VLANs are configured to enable jumbo frames.

show vlans port <PORT-LIST>

If port 1 belongs to VLAN 1, port 2 belongs to VLAN 10, and port 3 belongs to VLAN 15, executing this command with a *PORT-LIST* of **1 - 3** results in a listing of all three VLANs, even though none of the ports belong to all three VLANS. (See **Figure 66: Listing the VLAN memberships for a range of ports** on page 231.)

Figure 66: *Listing the VLAN memberships for a range of ports*

```
HP Switch(config)# show vlans ports A1-A3

 Status and Counters - VLAN Information - for ports A1-A3

  VLAN ID Name                             Status       Voice  Jumbo
  ------- -------------------------------- ------------ -----  -----
  1       DEFAULT_VLAN                     Port-based   No     Yes
  10      VLAN10                           Port-based   No     No
  15      VLAN15                           Port-based   No     No
```

Indicates which static VLANs are configured to enable jumbo frames.

show vlans <VID>

Figure 67: *Listing the port membership and jumbo status for a VLAN*

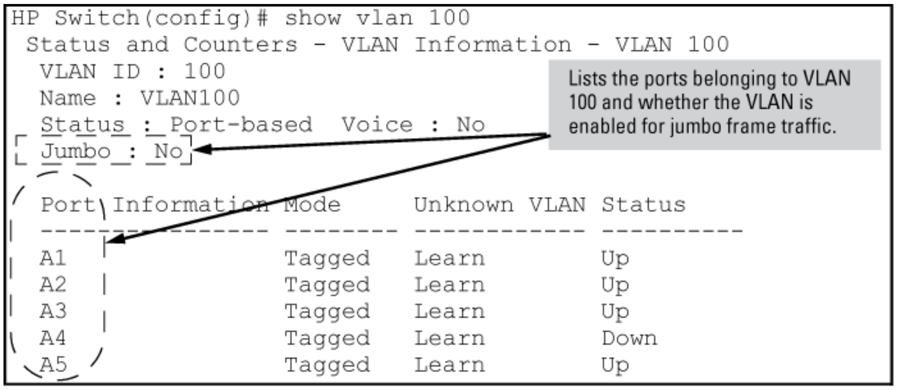

```
HP Switch(config)# show vlan 100
 Status and Counters - VLAN Information - VLAN 100
  VLAN ID : 100
  Name : VLAN100
  Status : Port-based  Voice : No
  Jumbo : No

  Port Information Mode     Unknown VLAN Status
  ---------------- -------- ------------ ----------
  A1               Tagged   Learn        Up
  A2               Tagged   Learn        Up
  A3               Tagged   Learn        Up
  A4               Tagged   Learn        Down
  A5               Tagged   Learn        Up
```


Enabling or disabling jumbo traffic on a VLAN

vlan jumbo

Syntax

```
vlan <VID>
       jumbo
```

Description

Configures the specified VLAN to allow jumbo frames on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, `vlan vid jumbo` also creates the VLAN.

A port belonging to one jumbo VLAN can receive jumbo frames through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo frames.Jumbos default to disabled on the specified VLAN.

Parameters

no

Disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are `jumbo` and `no jumbo`.

Options

<VID>

Shows port membership and jumbo configuration for the specified *<VID>* .

Configuring a maximum frame size

You can globally set a maximum frame size for jumbo frames that will support values from 1518 bytes to 9216 bytes for untagged frames.

jumbo max-frame-size

Syntax

```
jumbo max-frame-size <SIZE>
```

Description

Sets the maximum frame size for jumbo frames on a global level. The range is from 1518 bytes to 9216 bytes. (Default: 9216 bytes)

Configuring IP MTU

This feature is available on the switches covered in this guide. `jumbos` support is required for this feature. On switches that do not support this command, the IP MTU value is derived from the maximum frame size and is not configurable.

You can set the IP MTU globally by entering this command. The value of `max-frame-size` must be greater than or equal to 18 bytes more than the value selected for `ip-mtu`. For example, if `ip-mtu` is set to 8964, the `max-frame-size` is configured as 8982.

jump ip-mtu

Syntax

`jumbo ip-mtu <SIZE>`

Description

Globally sets the IP MTU size.

Options

<SIZE>

Values range between 1500 and 9198 bytes. This value must be 18 bytes less than the value of `max-frame-size`. Defaults to 9198 bytes

Viewing the maximum frame size

show jumbos

Syntax

`show jumbos`

Description

Displays the globally configured untagged maximum frame size for the switch

show jumbos

Use the `show jumbos` command to display the globally configured untagged maximum frame size for the switch.

```
switch# show jumbos

 Jumbos Global Values

  Configured : MaxFrameSize : 9216   Ip-MTU : 9198
  In Use     : MaxFrameSize : 9216   Ip-MTU : 9198
```

Operating notes for maximum frame size

- When you set a maximum frame size for jumbo frames, it must be on a global level. You cannot use the `jumbo max-frame-size` command on a per-port or per-VLAN basis.
- The original way to configure jumbo frames remains the same, which is per-VLAN, but you cannot set a maximum frame size per-VLAN.
- Jumbo support must be enabled for a VLAN from the CLI or through SNMP.
- Setting the maximum frame size does not require a reboot.

- When you upgrade to a version of software that supports setting the maximum frame size from a version that did not, the `max-frame-size` value is set automatically to 9216 bytes.
- Configuring a jumbo maximum frame size on a VLAN allows frames up to `max-frame-size` even though other VLANs of which the port is a member are not enabled for jumbo support.

Troubleshooting Jumbo frames

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames

Symptom

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames

Cause

The port may not be operating at a minimum of 10 Mbps on the 3500 switches or 1 Gbps on the other switches covered in this guide. Regardless of a port's configuration, if it is actually operating at a speed lower than 10 Mbps for 3500 switches or 1 Gbps for the other switches, it drops inbound jumbo frames. For example, if a port is configured for `Auto` mode (`speed-duplex auto`), but has negotiated a 7 Mbps speed with the device at the other end of the link, the port cannot receive inbound jumbo frames.

Action

1. To determine the actual operating speed of one or more ports, view the `Mode` field in the output for the following command:

   ```
   show interfaces brief  <PORT-LIST>
   ```

"Excessive undersize/giant frames" messages in the Event Log

Symptom

A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log.

Cause

The switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo frames received on the jumbo VLAN to non-jumbo ports.

Overview

Detection of link-flap and taking action on the port is done via fault-finder command at 3 different sensitivity levels (low, medium and high). The configuration in fault-finder for link-flap is a global configuration affecting all ports on the switch/stack. To provide further granularity to link-flap detection and action which provides different link-flap detection and action configuration for each port rather than the same configuration for all ports on the switch/ stack. The per-port configuration will supersede the global configuration for fault-finder link-flap.

A configurable option to re-enable ports disabled by link-flap after a waiting period is also been added. The waiting period time is expressed in unit of seconds in the range 0 to 604800. Maximum allowed waiting period is one week. Zero is the default value, meaning that the port will not be re-enabled automatically.

NOTE

A very important point is the wording of "link-flap" itself – i.e. the word "link". This condition should be at the link/port-level granular, allowing alerts and actions only on those certain links/ports where the functionality is needed.

fault-finder link-flap

Syntax

```
fault-finder link-flap [ethernet <PORT-LIST>] action [warn|warn-and-disable
<SECONDS>] sensitivity [low|medium|high]
```

Description

In the Config context, configures the link-flap on a port. Defaults to `warn`.

Parameters and options

no

Remove the current configuration of link-flap on a port.

link-flap

Configure link-flap control.

ethernet *<PORT-LIST>*

Enable link-flap control on a list of ports.

action

Configure the action taken when a fault is detected.

warn

Warn about faults found and log the event.

warn-and-disable

Warn and disable faulty component. Log the event.

<SECONDS>

Use with warn-and-disable to configure the number of seconds for which the port remains disabled. A value of 0 means that the port will remain disabled until manually re-enabled.

sensitivity

Indicates the sensitivity of the link-flap control threshold within a 10-second interval.

Low

10 link-flaps

Medium

6 link-flaps

High

3 link-flaps

Usage

Enable a linkFault-Finder check and set parameters for it. These commands may be repeated to enable additional checks. The default sensitivity is medium and the default action is warn.

```
[no] fault-finder [all | fault] sensitivity [low | medium | high] action [warn |
warn-and-disable]
```

```
[no] fault-finder broadcast-storm sensitivity [low | medium | high] action [warn |
warn-and-disable <SECONDS>]
```

```
[no] fault-finder link-flap sensitivity [low | medium | high] action [warn | warn-
and-disable]
```

```
[no] fault-finder link-flap PORT-LIST action [warn | warn-and-disable <SECONDS>]
sensitivity [low | medium | high]
```

Configure ports for link-flap detection with high sensitivity

Configure ports A1 to A5 for link-flap detection with sensitivity of high (3 flaps over 10s) and to log and disable port for 65535s if the link-flap threshold is exceeded.

```
switch(config)# fault-finder link-flap ethernet A1-A5 action warn-and-disable
65535
sensitivity high
```

Configure ports for link-flap detection with medium sensitivity

Configure ports A8 for link-flap detection with sensitivity of medium (6 flaps over 10s) and to log and disable port if the link-flap threshold is exceeded. User will need to re-enable the port if disabled.

```
switch(config)# fault-finder link-flap ethernet A8 action warn-and-disable 0
sensitivity medium
```

Configure ports for link-flap detection with low sensitivity

Configure ports A22 for link-flap detection with sensitivity of low (10 flaps over 10s) and to log if the link-flap threshold is exceeded

```
switch(config)# fault-finder link-flap ethernet A22 action warn sensitivity low
```

Disable link-flap detection

Disable link-flap detection for port A5

```
switch(config)# no fault-finder link-flap ethernet A5
```

Show fault-finder link-flap

Syntax

`show fault-finder link-flap ethernet PORT-LIST`

Description

Display the link-flap control configuration.

Show fault-finder link-flap

```
switch# show fault-finder link-flap A1

        Link  | Port                                   Disable    Disable Time
  Port   Flap  | Status Sensitivity Action               Timer      Left
  ------ ----- + ------ ----------- ------------------ ---------- ------------
  A1     Yes     Down      Low        warn-and-disable 65535      45303

switch# show fault-finder link-flap
        Link  | Port                                     Disable    Disable Time
  Port   Flap  | Status Sensitivity Action                 Timer      Left
  ------ ----- + ------ ----------- ------------------- ---------- ------------
  A1     Yes     Down      Low         warn-and-disable  65535      45303
  A5     No      Up        None        None                -          -
  A22    Yes     Down      Low         warn-and-disable    -          -
  A23    Yes     Down      High        warn-and-disable   100         -
```

This example displays only the list of ports configured via the above per-port config commands, does not include the global configuration ports.

Event Log

Cause

Message	Cause
FFI: port <ID>- Excessive link state transitions.	Link-flap is detected by fault-finder per the sensitivity configured.
FFI: port <ID>- Excessive link state transitions. FFI: port <ID>-Port disabled by Fault-finder. FFI: port <ID>-Administrator action is required to re-enable. ports: Fault-finder (71) has disabled port <ID>. ports: port <ID> is now offline. vlan: VLAN<VLAN-ID> virtual LAN is disabled.	Link-flap is detected and the action is to disable the port with no disable timer.
FFI: port <ID>- Excessive link state transitions. FFI: port <ID>-Port disabled by Fault-finder. ports: Fault-finder(71) has disabled port <ID> for <SECONDS> seconds. ports: port <ID> is now off-line. vlan: VLAN<VLAN-ID> virtual LAN is disabled.	Link-flap is detected and the action is to disable the port with disable timer.
port <ID> timer (71) has expired. ports: port <ID> is now on-line. vlan: VLAN<VLAN-ID> virtual LAN is enabled.	The port is enabled when the disable timer expires.

Restrictions

- Per port configuration for options – link-flap only. Global settings for other options.
- No support for menu interface.
- No support for Web UI.
- No support for trunks.

Configuring the switch to filter untagged traffic

Enter this command to configure the switch not to learn CDP, LLDP, or EAPOL traffic for a set of interfaces.

ignore-untagged-mac

Syntax

`ignore-untagged-mac <PORT-LIST>`

Description

Prevents MAC addresses from being learned on the specified ports when the VLAN is untagged and the destination MAC address is one of the following:

- 01000C-CCCCCC (CDP)
- 0180c2- 00000e (LLDP)
- 0180c2-000003 (EAPOL)

ignore packet MAC

Configuring the switch to ignore packet MAC address learns for an untagged VLAN.

```
switch(config) ignore-untagged-mac 1-2
```

Viewing configuration file change information

show running-config

Syntax

`show running-config [changes-history <1-32>] [detail]`

Description

Displays the history of changes made to the running-configuration file.

Parameters

changes-history

Shows up to 32 events and displays changes in descending order (the most recent change at the top of the list). You can specify from 1 to 32 entries for display.

detail

Displays a more detailed amount of information for the configuration changes.

Configuration change output

Figure 68: *Output for running configuration changes history for all ports*

```
HP Switch(config)# show running-config changes-history

 Running Config Last Changed    : 02/19/10 16:30:09
 Number of Changes Since Reboot : 150086


             Config
  Event ID   Method     Date     Time
  ---------- ---------- -------- --------
  150086     CLI        02/19/10 16:30:09
  150085     SNMP       02/03/10 14:50:12
  150084     SNMP       02/03/10 14:50:12
  150083     SNMP       02/03/10 14:45:59
  150082     SNMP       02/03/10 14:27:15
  150081     SNMP       02/03/10 13:11:00
  150080     SNMP       02/03/10 13:11:00
  150079     CLI        01/18/10 09:09:17
```

Figure 69: *Example of output for running config changes history with detail*

```
HP Switch(config)# show running-config changes-history detail

  Running Config Last Changed: 01/01/90 00:35:44
  Number of changes since last boot : 6

Event ID           : 6
User               :
Remote IP Address  :
Config Method      : CLI
Date               : 01/01/90
Time               : 00:35:44

Event ID           : 5
User               :
Remote IP Address  :
Config Method      : CLI
Date               : 01/01/90
Time               : 00:35:39

Event ID           : 4
User               :
Remote IP Address  :
Config Method      : CLI
Date               : 01/01/90
Time               : 00:35:33

Event ID           : 3
User               :
Remote IP Address  :
Config Method      : CLI
Date               : 01/01/90
Time               : 00:35:27
```

Current status of SNMP trap type

Figure 70: *SNMP trap configuration status information*

```
HP Switch(config)# show snmp-server traps

  Trap Receivers

   Link-Change Traps Enabled on Ports [All] : All

   Traps Category                  Current Status
   _____________________________  ______________
   SNMP Authentication           : Extended
   Password change               : Enabled
   Login failures                : Enabled
   Port-Security                 : Enabled
   Authorization Server Contact  : Enabled
   DHCP-Snooping                 : Enabled
   Dynamic ARP Protection        : Enabled
   Dynamic IP Lockdown           : Enabled
   Running Configuration Changes : Enabled

   Address                Community            Events   Type  Retry   Timeout
   ---------------------- -------------------- -------  ----- ------- -------
   173.33.25.201          public               None     trap  3       15


  Excluded MIBs
```

SNMP trap status for running-config changes is enabled.

Minimal interval for successive data change notifications

setmib

Syntax

setmib lldpnotificationinterval.0 -i *1 - 3600*

Description

Change the minimum interval for successive data change notifications for the same neighbor. Globally changes the interval between successive traps generated by the switch. If multiple traps are generated in the specified interval, only the first trap is sent. The remaining traps are suppressed. (A network management application can periodically check the switch MIB to detect any missed change notification traps. See IEEE P802.1AB or later for more information.)

(Default: 5 seconds)

Limiting change notification traps

The following command limits change notification traps from a particular switch to one per minute.

```
switch# setmib lldpnotificationinterval.0 -i 60 lldpNotificationInterval.0=60
```

Viewing the current port speed and duplex configuration on a switch port

show interfaces

Syntax

```
show interfaces [brief | config | custom | display | port-utilization |
transceiver | status |
tunnel | ethernet <PORT-LIST>]
```

Description

Show port configuration and status information.

Parameters

brief

Show port operational parameters.

config

Show port configuration information.

custom

Show port parameters in a customized table.

display

Show summary of network traffic handled by the ports.

internal-use

Show reserved or eligible internal ports.

[ethernet] *<PORT-LIST>*

Show summary of network traffic handled by the ports.

port-utilization

Show port bandwidth utilization.

status

Show interfaces tagged or untagged VLAN information.

transceiver

Show the transceiver information.

tunnel

Show tunnel configuration and status information.

Show interfaces

```
switch# show interfaces
Status and Counters - Port Counters
                                                                  Flow Bcast
 Port   Total Bytes    Total Frames   Errors Rx    Drops Tx     Ctrl Limit
 ------ -------------- -------------- ------------ ------------ ---- -----
A1     419,179        1555            0            0            off  0
A2     4217           24              0            0            off  0
A3     0              0               0            0            off  0
A4     0              0               0            0            off  0
A5     0              0               0            0            off  0
A6     0              0               0            0            off  0
A7     0              0               0            0            off  0
```

```
A8      0              0              0             0             off   0
A9      0              0              0             0             off   0
A10     0              0              0             0             off   0
A11     0              0              0             0             off   0
A12     0              0              0             0             off   0
A13     0              0              0             0             off   0
A14     0              0              0             0             off   0
A15     0              0              0             0             off   0
A16     0              0              0             0             off   0
A17     0              0              0             0             off   0
A18     0              0              0             0             off   0
A19     0              0              0             0             off   0
A20     0              0              0             0             off   0
A21     3846           21             0             0             off   0
A22     3855           19             0             0             off   0


MACsec Port Counters:
Port   Errors Rx    Drops Tx
------ ------------ ------------
A2     0            0
```

Viewing the configuration

show running-config

Syntax

`show running-config`

Description

Displays information about the configuration.

Example

show running-config

Configuration showing interfaces to ignore packet MAC address learns.

```
switch(config) show running-config
Running configuration:
; J9627 Configuration Editor; Created on release K.15.XX
; Ver #03:03.1f.ef:f0
hostname "HP Switch"
interface 1
ignore-untagged-mac
exit
interface 2
ignore-untagged-mac
exit
...
vlan 1
name "DEFAULT_VLAN"
untagged 1-24
ip address dhcp-bootp
exit
```

RMON advanced management

The switch supports RMON (remote monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the HPE Switch Manager network management software.

The CLI supports the configuration of RMON alarm threshold settings. The settings can be saved in the configuration file.

rmon alarm

Syntax

```
[no] rmon alarm  entry number alarm-variable sampling-interval  absolute | delta
rising-threshold threshold-value1 falling-threshold2 threshold-value2  owner string
```

Description

This command configures RMON sampling periods and threshold parameters.

Parameters

no

Removes the alarm entry.

entry number <1–65535>

An alarm number that uniquely identifies the alarm threshold entry.

alarm-variable <object-string>

Object identifier of the particular variable to be sampled. Variables must be of type Integer in order to be sampled.

sampling-interval <5-65535>

Time interval in seconds over which data is sampled and compared with the rising-threshold and the falling-threshold.

absolute

The value of the selected variable is compared directly with the thresholds at the end of the sampling interval.

If the absolute option is used for alarm variables of counter-type, an RMON trap is generated only once, when the threshold limit is reached. The RMON trap is never generated again. Hewlett Packard Enterprise recommends using the delta option instead when using a counter-type alarm variable.

delta

The value of the selected variable at the last sample is subtracted from the current value, and the difference is compared with the thresholds.

rising-threshold <threshold-value>

An integer value for the upper threshold for the sampled statistic. A single event is generated when the current sampled value of the specified statistic becomes greater than or equal to this threshold, and if the value at the last sampling intervals was less than this threshold.

The value of the rising-threshold must be greater than the value of the falling-threshold.

falling-threshold <threshold-value2>

An integer value for the lower threshold for the sampled statistic. A single event is generated when the current sampled value of the specified statistic becomes less than or equal to this threshold, and if the value at the last sampling interval was greater than this threshold.

owner <string>

The name of the owner of this alarm.

Using RMON alarm parameters

Figure 71: *Configuring the RMON Alarm Parameters in the CLI*

```
HP Switch(config)# rmon alarm 1 1.3.6.1.2.1.16.1.1.1.4.1 100 absolute rising-
threshold 500 falling-threshold 300 owner Joe
```

Figure 72: *Removing an RMON Alarm*

```
HP Switch(config)# no rmon alarm 1
```

Figure 73: *Show Command Output for a Specific Alarm*

```
HP Switch(config)# show rmon alarm 1

Alarm table 1 owned by Joe is VALID.
    Sample Type           : absolute
    Object Variable       : 1.3.6.1.2.1.16.1.1.1.4.1 <etherStatsOctets.1>

    Sampling Interval     : 100(sec)
    Rising Threshold      : 500
    Falling Threshold     : 300
    Startup Alarm         : risingOrFallingAlarm
    Last Sampled Value    : 2550
    Last Threshold Time   : Fri Jun 19 10:03:31 2012
```

Figure 74: *Display Command Output for a Specific Alarm*

```
HP Switch(config)# display rmon alarm 1

Alarm table 1 owned by Joe is VALID.
    Samples type          : absolute
    Variable formula      : 1.3.6.1.2.1.16.1.1.1.4.1 <etherStatsOctets.1>

    Sampling interval     : 100(sec)
    Rising threshold      : 500
    Falling threshold     : 300
    Startup alarm         : risingOrFallingAlarm
    Latest value          : 2550
    Last threshold time   : Fri Jun 19 10:03:31 2012
```

Figure 75: *Output of the running-config File Displaying the Configured RMON Alarm Parameters*

```
HP Switch(config)# show running-config

Running configuration:

; J9470A Configuration Editor; Created on release #K.15.13.0000x
; Ver #04:2f.ff.3f.ef:04
hostname "HP-3500-24"
module 1 type j94dda
snmp-server community "public" unrestricted
snmp-server host 15.255.133.156 community "public"
snmp-server host 15.255.133.146 community "public"
vlan 1
   name "DEFAULT_VLAN"
   untagged 1-24
   ip address dhcp-bootp
   exit
vlan 2
   name "VLAN2"
   ip address 10.10.10.100 255.255.255.0
   exit
spanning-tree
rmon alarm 1 "etherStatsOctets.1" 100 absolute rising-threshold 500
 falling-threshold 300 owner "Joe"
```

Configuring UDLD verify before forwarding

When an UDLD enabled port transitions to link-up, the port will begin with a UDLD blocking state. UDLD will probe via protocol packet exchange to determine the bidirectional state of the link. Until UDLD has completed the probe, all data traffic will be blocked. If the link is found to be bidirectional, UDLD will unblock the port for data traffic to pass. Once UDLD unblocks the port, other protocols will see the port as up and data traffic can be safely forwarded.

The default mode of a switch is "forward first then verify''. Enabling UDLD link-up will default to "forward first then verify". To change the mode to "verify then forward", you need to configure using the commands found in section 6.72.

Link-UP data traffic will resumed after probing the link partner completes. All other protocols running will see the port as down.

UDLD time delay

UDLD protocol informs the link partner simultaneously as it detects a state change from unidirectional to bidirectional traffic. Additional packet exchanges will be carried out by UDLD in addition to the existing UDLD exchanges whenever state changes from unidirectional to bidirectional.

Interval 1

5 seconds

With triggered updates: State=blockedPeer; State=blocked

Without triggered updates: State=blockedPeer; State=blocked

Interval 1 + delta

5+(<5) seconds (delta is the time when the unblock event occurs on local side)

With triggered updates: Inform PeerState=unblockedPeer; State=unblocked

Without triggered updates: State=unblockedPeer; State=blocked

Interval 2

10 seconds

With triggered updates: Regular UDLD TX

Without triggered updates: Inform PeerState=unblocked; Peer State=unblocked

Interval 3

15 seconds

With triggered updates: Regular UDLD TX

Without triggered updates: Regular UDLD TX

Restrictions

- There is no support available when configuring this mode from the web and menu interface.
- There are no new packet types are introduced with UDLD.
- There are no new UDLD timers being introduced.

UDLD configuration commands

link-keepalive mode

Syntax

`link-keepalive mode [verify-then-forward | forward-then-verify]`

Description

This command configures the link-keepalive mode.

Parameters

verify-then-forward

The port is in a blocking state until the link configured for UDLD establishes bidirectional communication.

forward-then-verify

The port forwards the data then verifies the status of the link-in state. When a unidirectional state is detected, the port is moved to a blocked state. When a bidirectional state is detected, the data is forwarded without interruption.

interval <DECISECONDS>

Configure the interval for `link-keepalive`. The `link-keepalive` interval is the time between sending two UDLD packets. The time interval is entered in deciseconds (1/10 sec). For example, a value of 10 is 1 second; 11 is 1.1 seconds. The default keepalive interval is 50 deciseconds.

retries <NUMBER>

Maximum number of sending attempts for UDLD packets before declaring the link as faulty.

Default keepalive attempt is 4.

show link-keepalive

Syntax

`show link-keepalive`

Description

Sample output

```
Total link-keepalive enabled ports: 8
Keepalive Retries : 4
Keepalive Interval: 5 sec
Keepalive Mode : verify-then-forward
Physical Keepalive Adjacent UDLD

Port  Enabled Status     Status   Switch    VLAN
----- ------- ---------------- --------- ------------- ----------
1    Yes    down      off-line 000000-000000 untagged
2    Yes    down      off-line 000000-000000 untagged
3    Yes    down      off-line 000000-000000 untagged
4    Yes    down      off-line 000000-000000 untagged
5    Yes    down      off-line 000000-000000 untagged
6    Yes    down      off-line 000000-000000 untagged
7    Yes    down      off-line 000000-000000 untagged
8    Yes    down      off-line 000000-000000 untagged
```

RMON generated when user changes UDLD mode

RMON events are generated when UDLD is configured. The first time you configure the mode, the UDLD states will be re-initialized. An event log entry is initiated to include the reason for the initial UDLD blocking state during link up.

UDLD mode `[verify-then-forward | forward-then-verify]` is configured

Severity: - Info.

MAC configurations

Configuring the MAC address count option

The MAC Address Count feature provides a way to notify the switch management system when the number of MAC addresses learned on a switch port exceeds the permitted configurable number.

To enable the mac-count-notify option, enter this command in global config context.

snmp-server mac-count-notify

Syntax

```
[no] snmp-server enable traps mac-count-notify
```

Description

Sends a trap when the number of MAC addresses learned on the specified ports exceeds the configured `<learned-count>` value.

Parameters

no

Disables `mac-count-notify traps`.

mac-count-notify

To configure the mac-count-notify option on a port or ports, enter this command. When the configured number of MAC addresses is exceeded (the learned-count), a trap is sent.

traps <PORT-LIST>

Configures `mac-count-notify traps` on the specified ports for the entire switch. With no <PORT-LIST> specified, confiures all ports.

<LEARNED-COUNT>

The number of MAC addresses learned before sending a trap. Values range between 1-128. Defaults to 32.

Usage

```
[no] mac-count-notify traps <PORT-LIST> [<learned-count>]
```

Configuring mac-count notify traps on ports 5–7

```
switch(config)# mac-count-notify traps 5-7 50
```

Configuring the MAC address table change option

When enabled, this feature allows the generation of SNMP traps for each MAC address table change. Notifications can be generated for each device that connects to a port and for devices that are connected through another device (daisy-chained.)

The `snmp-server enable traps mac-notify` command globally enables the generation of SNMP trap notifications upon MAC address table changes.

snmp-server mac-notify

Syntax

`[no] snmp-server enable traps mac-notify [mac-move | trap-interval <0- 120>]`

Description

Globally enables or disables generation of SNMP trap notifications.

Parameters

trap-interval

The time interval (in seconds) that trap notifications are sent. A value of zero disables the interval and traps are sent as events occur. If the switch is busy, notifications can be sent prior to the configured interval. Notifications may be dropped in extreme instances and a system warning is logged.The range is 0-120 seconds. Default: 30 seconds.

mac-move

Configures the switch to capture data for MAC addresses that are moved from one port to another port. The `snmp-server enable traps mac-notify` command must have been enabled in order for this information to be sent as an SNMP notification.

trap-interval

```
switch(config)# snmp-server enable traps mac-notify trap-interval 60
```

mac-move

```
switch(config)# snmp-server enable traps mac-notify mac-move
```

mac-notify at the interface context level

You can also execute the `mac-notify traps` command from the interface context.

```
switch# int 11
switch(int-11)# mac-notify traps learned
```

Per-port MAC change options for mac-notify

mac-notify traps

Syntax

`[no] mac-notify traps <PORT-LIST>[learned | removed]`

Description

Configure SNMP traps for learned or removed MAC addresses on a per-port basis.

The switch captures learned or removed events on the selected ports, but does not send an SNMP trap unless you have enabled mac-notify with the `snmp-server enable traps mac-notify` command.

When this command is executed without the learned or removed option, it enables or disables the capture of both learned and removed MAC address table changes for the selected ports in *<PORT-LIST>*.

Parameters

learned

Enables the capture of learned MAC address table changes on the selected ports.

removed

Enables the capture of removed MAC address table changes table on the selected ports.

Options

<PORT-LIST>

Configures MAC address table changes capture on the specified ports. Use all to capture changes for all ports on the switch.

Configuring traps on a per-port basis for learned MAC addresses

```
switch# mac-notify traps 5-6 learned
switch# show mac-notify traps 5-6
Mac Notify Trap Information
Mac-notify Enabled : Yes
Mac-move Enabled : Yes
Trap-interval : 60
Port    MAC Addresses    trap learned/removed
------ ----------------------------------
5       Learned
6       Learned
```

Configuring traps on a port-bases for removed MAC addresses

```
switch# mac-notify traps 3-4 removed
switch(config)# show mac-notify traps
Mac Notify Trap Information
Mac-notify Enabled : Yes
Mac-move Enabled : Yes
Trap-interval : 60
Port    MAC Addresses    trap learned/removed
------ ----------------------------------
1         None
2         None
3         Removed
4         Removed
```

Viewing the mac-count-notify option

show mac-count-notify

Syntax

```
show mac-count-notify traps [<PORT-LIST>]
```

Description

Displays information about the configured value for sending a trap, the current count, and if a trap has been sent.

Command output

```
switch(config)# show mac-count-notify traps

Mac-count-notify Enabled: Yes

Port                Count for          Count            Trap Sent
                    sending Trap
-----------------------------------------------------------------
1
2
3
4
5                    50                 0                No
6                    50                 2                No
7                    50                 0                No
8
9
...
```

Configuring mac-count-notify traps from the interface context

The interface context can be used to configure the value for sending a trap.

```
switch(config)# interface 5
```

```
switch(eth-5)# mac-count-notify traps 35
```

View information about SNMP traps, including MAC address count being Enabled/Disabled

The `show snmp-server traps` command displays whether the MAC Address Count feature is enabled or disabled.

```
switch# show snmp-server traps
Trap Receivers
Link-Change Traps Enabled on Ports [All] : All
Traps Category                Current Status
_____________________________ __________________
SNMP Authentication :             Extended
Password change :                 Enabled
Login failures :                  Enabled
Port-Security :                   Enabled
Authorization Server Contact :    Enabled
DHCP-Snooping :                   Enabled
Dynamic ARP Protection :          Enabled
Dynamic IP Lockdown :             Enabled
MAC address table changes :       Disabled
MAC Address Count :               Enabled

Address                   Community              Events   Type   Retry  Timeout
---------------------- ---------------------- -------- ------ ------- -------
15.146.194.77             public                  None    trap     3      15
15.255.134.252            public                  None    trap     3      15
16.181.49.167             public                  None    trap     3      15
16.181.51.14              public                  None    trap     3      15
Excluded MIBs
```

Viewing mac-notify traps configuration

show mac-notify traps

Syntax

`show mac-notify traps <PORT-LIST>`

Description

Displays information about SNMP trap configuration for MAC Address Table changes.

Output of SNMP trap configuration

Displays SNMP trap information for all ports, or each port in the *<PORT-LIST>*.

```
switch# show mac-notify traps
Mac Notify Trap Information
Mac-notify Enabled : Yes
Mac-move Enabled : Yes
Trap-interval : 60
Port     MAC Addresses    trap learned/removed
------ ----------------------------------
1        None
2        None
3        Removed
4        Removed
5        Learned
6        Learned
```

Running config file with mac-notify parameters configured

The configured `mac-notify` commands are displayed in the `show running-configuration` output.

```
switch# show running-config
Running configuration:
; J9087A Configuration Editor; Created on release #R.11.XX
hostname "Switch"
snmp-server community "public" Unrestricted
snmp-server host 15.255.133.236 "public"
snmp-server host 15.255.133.222 "public"
snmp-server host 15.255.133.70 "public"
snmp-server host 15.255.134.235 "public"
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-28
  ip address dhcp-bootp
  exit
snmp-server enable traps mac-notify mac-move
snmp-server enable traps mac-notify trap-interval 60
snmp-server enable traps mac-notify
mac-notify traps 5-6 learned
mac-notify traps 3-4 removed
```

Configuring sFlow

Under the multiple instance implementation, sFlow can be configured via the CLI or via SNMP. However, CLI-owned sFlow configurations cannot be modified via SNMP, whereas SNMP-owned instances can be disabled via the CLI using the `no sflow <receiver-instance>` command.

sflow

Syntax

```
[no] sflow <RECEIVER-INSTANCE> destination [<UDP-PORT-NUM> <IP-ADDRESS> [ipv4 |
ipv6 <UDP-PORT-NUM> oobm] sampling [<PORT-LIST> <SAMPLING RATE>] polling [<PORT-
LIST> <POLLING INTERVAL>]
```

Description

sFlow commands allow you to configure sFlow instances using the CLI.

Parameters and options

no

To disable an sFlow receiver/destination, enter `no sflow <RECEIVER-INSTANCE>` .

<RECEIVER-INSTANCE> destination

Enables an sFlow receiver/destination. The receiver-instance number must be a 1, 2, or 3.

oobm

A configurable option for sending sFlow packets to a destination through the OOBM port on the switch. Use the OOBM port to reach the specified sFlow receiver.

ipv4 | ipv6

Supports both IPv4 and IPv6 addresses.

<UDP-PORT-NUM>

The sFlow collector collects sample packets through the OOBM port, allowing the monitoring of network traffic. By default, the udp destination port number is 6343.

<IP-ADDRESS>

The IP address of a single destination.

<RECEIVER-INSTANCE> sampling

Once an sFlow receiver/destination has been enabled, this command enables flow sampling for that instance. The receiver-instance number is 1, 2, or 3.

<PORT-LIST>

Port or list of ports on which to enable flow-sampling. To disable flow-sampling for the specified *<PORT-LIST>* use a sampling rate of `0`.

<SAMPLING-RATE>

The allowable non-zero skipcount for the specified port or ports.

<RECEIVER-INSTANCE> polling

Once an sFlow receiver/destination has been enabled, this command enables counter polling for that instance. The receiver-instance number is 1, 2, or 3.

<PORT-LIST>

Port or list of ports on which to enable polling. To disable counter-polling for the specified *<PORT-LIST>* use a polling interval of `0`.

<POLLING INTERVAL>

An allowable non-zero value to enable polling on the specified port or ports.

Usage

```
[no] sflow <RECEIVER-INSTANCE> destination <IP-ADDRESS> <UDP-PORT-NUM>

sflow <RECEIVER-INSTANCE> sampling <PORT-LIST> <SAMPLING RATE>

sflow <RECEIVER-INSTANCE> polling <PORT-LIST> <POLLING INTERVAL>

[no] sflow <RECEIVER-INSTANCE> destination [ipv4 | ipv6] <UDP-PORT-NUM> oobm
```

sFlow Destination is OOBM port

```
HP_Switch (Config#) sflow 1 destination 192.168.2.3 6000 oobm
```

Figure 76: *Output showing OOBM Support Enabled*

```
HP Switch# show sflow 1 destination

Destination Instance        : 1
  sflow                     : Enabled
  Datagrams Sent            : 0
  Destination Address       : 192.168.2.3
  Receiver Port             : 6343
  Owner                     : Administrator, CLI-Owned, Instance 1
  Timeout (seconds)         : 2147479533
  Max Datagram Size         : 1400
  Datagram Version Support  : 5
  OOBM Support              : Enabled
```

Figure 77: *Output of the running-config File showing the sFlow Destination is the OOBM Port*

```
HP Switch# show running-config

Running configuration:

; J9091A Configuration Editor; Created on release #K.15.06.0006
; Ver #01:0d:0c

hostname "HP Switch"
module 1 type J8702A
module 7 type J8708A
module 12 type J8702A
sflow 1 destination 192.168.2.3 oobm
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,G1-G4,L1-L24
   ip address dhcp-bootp
   exit
snmp-server community "public" unrestricted
```

sFlow Configuring multiple instances

In earlier software releases, sFlow was configured on the switch via SNMP using a single sFlow instance. Beginning with software release K.11.34, sFlow can also be configured via the CLI for up to three distinct sFlow instances: once enabled, an sFlow receiver/destination can be independently configured for full flow-sampling and counter-polling. CLI-configured sFlow instances may be saved to the startup configuration to persist across a switch reboot.

Viewing sFlow Configuration and Status

show sflow agent

The following sFlow commands allow you to display sFlow configuration and status via the CLI. **Figure 79: Viewing sFlow destination information** on page 256 is an example of `sflow agent` information.

Syntax

`show sflow agent`

Description

Displays sFlow agent information. The agent address is normally the IP address of the first VLAN configured.

The `show sflow agent` command displays read-only switch agent information. The version information shows the sFlow version, MIB support, and software versions; the agent address is typically the IP address of the first VLAN configured on the switch.

sflow agent

Figure 78: *Viewing sflow agent information*

```
HP Switch# show sflow agent

  Version             1.3;HP;K.11.40
  Agent Address       10.0.10.228
```

show sflow destination

Syntax

`show sflow <RECEIVER INSTANCE> destination`

Description

Displays information about the management station to which the sFlow sampling-polling data is sent. Includes management station information such as destination address, receiver port, and owner, as shown in **Figure 79: Viewing sFlow destination information** on page 256

sflow destination

Figure 79: *Viewing* `sFlow destination` *information*

```
HP Switch# show sflow 2 destination

  Destination Instance          2
  sflow                         Enabled
  Datagrams Sent                221
  Destination Address           10.0.10.41
  Receiver Port                 6343
  Owner                         Administrator, CLI-owned, Instance 2
  Timeout (seconds)             99995530
  Max Datagram Size             1400
  Datagram Version Support      5
```

Note the following details:

- Destination Address remains blank unless it has been configured.
- Datagrams Sent shows the number of datagrams sent by the switch agent to the management station since the switch agent was last enabled.
- Timeout displays the number of seconds remaining before the switch agent will automatically disable sFlow (this is set by the management station and decrements with time.)
- Max Datagram Size shows the currently set value (typically a default value, but this can also be set by the management station.)

show sflow sampling-polling

Syntax

```
show sflow <RECEIVER INSTANCE> sampling-polling <PORT-LIST/RANGE>
```

Description

Displays status information about sFlow sampling and polling on the switch as shown in **Figure 80: Viewing sFlow sampling and polling information** on page 257.

Options

<RECEIVER INSTANCE>

The receiver-instance number is 1, 2, or 3.

<PORT-LIST/RANGE>

You can specify a list or range of ports for which to view sampling information.

sflow sampling-polling

Figure 80: *Viewing sFlow sampling and polling information*

```
HP Switch# show sflow 2 sampling-polling A1-A4
                 Number denotes the sampling/polling instance to which the receiver is coupled.

 Port  | Sampling                      Dropped     | Polling
       | Enabled  Rate     Header      Samples     | Enabled    Interval
 ----- + -------- -------- ------ ----------      + --------   --------
  A1     Yes(2)   40       128    1234567890         ---        ---
  A2      ---     ---      ---             0         Yes(1)     60
  A3      No(1)    0       100        898703         No         30
  A4     Yes(3)   50       128             0         No(3)       0
```

The sampling and polling instances (noted in parentheses) coupled to a specific receiver instance are assigned dynamically, and so the instance numbers may not always match. The key thing to note is whether sampling or polling is enabled on a port, and the sampling rates or polling intervals for the receiver instance configured on each port.

show snmpv3 user

Syntax

```
show snmpv3 user
```

Description

Displays information about the management stations configured for SNMPv3

Management stations configured on VLAN 1 to access the switch

```
switch# configure terminal
switch# vlan 1
switch(vlan-1)# show snmpv3 user

Status and Counters - SNMPv3 Global Configuration Information

 User Name       Auth. Protocol    Privacy Protocol
```

```
-----------    --------------    -----------------
initial        MD5               CFB AES-128
NetworkAdmin   MD5               CBC-DES
```

Configuring SNMP

Network security notifications

By default, a switch is enabled to send the SNMP notifications listed in **Supported Notifications** on page 261 when a network security event (for example, authentication failure) occurs. However, before security notifications can be sent, you must first configure one or more trap receivers or SNMPv3 management stations as described in:

- **SNMP trap receiver configuration** on page 273
- **Configuring SNMPv3 notifications** on page 275

You can manage the default configuration of the switch to disable and re-enable notifications to be sent for the following types of security events:

- ARP protection events
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- Dynamic IP Lockdown hardware resources consumed
- Link change notification
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Manager password changes
- Port-security (web, MAC, or 802.1X) authentication failure
- SNMP authentication failure
- Running configuration changes

SNMP traps on running configuration changes

You can send a specific SNMP trap for any configuration change made in the switch's running configuration file. The trap will be generated for changes made from any of these interfaces:

- CLI
- Menu
- SNMP (remote SNMP set requests.)

The SNMP trap contains the following information.

Information	Description
Event ID	An assigned number that identifies a specific running configuration change event.
Method	Method by which the change was made—CLI, Menu, or remote SNMP.For configuration changes triggered by internal events, the term "Internal-Event" is used as the source of the change.
IP Address Type	Indicates the source address type of the network agent that made a change. This is set to an address type of "unknown" when not applicable.

Table Continued

Information	Description
IP address	IP address of the remote system from which a user accessed the switch. If not applicable, this is an empty string and nothing is displayed, for example, if access is through a management console port.
User Name	User name of the person who made the change. Null if not applicable.
Date and Time	Date and time the change was made.

The SNMP trap alerts any interested parties that someone has changed the switch's configuration and provides information about the source for that change. It does not specify what has been changed.

Source IP address for SNMP notifications

The switch uses an interface IP address as the source IP address in IP headers when sending SNMP notifications (traps and informs) or responses to SNMP requests.

For multi-netted interfaces, the source IP address is the IP address of the outbound interface of the SNMP reply, which may differ from the destination IP address in the IP header of the received request. For security reasons, it may be desirable to send an SNMP reply with the IP address of the destination interface (or a specified IP address) on which the corresponding SNMP request was received.

To configure the switch to use the source IP address on which an SNMP request was received in SNMP notification/traps and replies, enter the `snmp-server response-source` and `snmp-server trap-source` commands.

Listening mode

For switches that have a separate out-of-band management port, you can specify whether a configured SNMP server listens for SNMP queries over the OOBM interface, the data interface, or both. By default, the switch listens over both interfaces.

This option is not available for switches that do not have a separate OOBM port.

The listening mode is set with parameters to the `snmp-server` command.

Group access levels

The switch supports eight predefined group access levels. There are four levels for use by version 3 users and four are used for access by version 2c or version 1 management applications.

Table 17: *Predefined group access levels*

Group name	Group access type	Group read view	Group write view
managerpriv	Ver3 Must have Authentication and Privacy	ManagerReadView	ManagerWriteView
managerauth	Ver3 Must have Authentication	ManagerReadView	ManagerWriteView
operatorauth	Ver3 Must have Authentication	OperatorReadView	DiscoveryView

Table Continued

Group name	Group access type	Group read view	Group write view
operatornoauth	Ver3 No Authentication	OperatorReadView	DiscoveryView
commanagerrw	Ver2c or Ver1	ManagerReadView	ManagerWriteView
commanagerr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorr	Ver2c or Ver1	OperatorReadView	DiscoveryView

Each view allows you to view or modify a different set of MIBs:

- Manager Read View – access to all managed objects
- Manager Write View – access to all managed objects except the following:
    - vacmContextTable
    - vacmAccessTable
    - vacmViewTreeFamilyTable
- OperatorReadView – no access to the following:
    - icfSecurityMIB
    - hpSwitchIpTftpMode
    - vacmContextTable
    - vacmAccessTable
    - vacmViewTreeFamilyTable
    - usmUserTable
    - snmpCommunityTable
- Discovery View – Access limited to samplingProbe MIB.

NOTE

All access groups and views are predefined on the switch. There is no method to modify or add groups or views to those that are predefined on the switch.

SNMPv3 communities

SNMP commuities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch. The communities are mapped to Group Access Levels that are used for version 2c or version 1 support. This mapping happens automatically based on the communities access privileges, but special mappings can be added with the `snmpv3 community` command.

SNMP community features

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.

SNMPv2c informs

On a switch enabled for SNMPv2c, you can use the `snmp-server host inform` command (**SNMPv2c inform option** on page 274) to send inform requests when certain events occur. When an SNMP Manager receives an inform request, it can send an SNMP response back to the sending agent on the switch to let the agent know that the inform request reached its destination.

If the sending agent on the switch does not receive an SNMP response back from the SNMP Manager within the timeout period, the inform request may be resent, based on the retry count value.

When you enable SNMPv2c inform requests to be sent, you must specify the IP address and community name of the management station that will receive the inform notification.

SNMP notifications

The switches:

- Default Traps: A switch automatically sends default traps to trap receivers using the configured community name. You have to configure and supply the community name to use in the trap-receiver config, there is no default. Use the `snmp-server host <IP_ADDRESS> community "<COMMUNITY_NAME>"` command to configure a community name and the `snmp-server host <IP_ADDRESS> community "<COMMUNITY_NAME>" trap-level [all | critical | not-info | debug | none]` command to set the level of traps to send to the community.
- SNMPv2c informs
- SNMP v3 notification process, including traps

This section describes how to configure a switch to send network security and link-change notifications to configured trap receivers.

Supported Notifications

By default, the following notifications are enabled on a switch:

- Manager password changes
- SNMP authentication failure
- Link-change traps: when the link on a port changes from up to down (linkDown) or down to up (linkUp)
- Port-security (web, MAC, or 802.1X) authentication failure
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- ARP protection events

Configuring SNMP notifications

Procedure

1. Determine the versions of SNMP notifications that you want to use in your network.

   If you want to use SNMPv1 and SNMPv2c traps, you must also configure a trap receiver.

   If you want to use SNMPv3 notifications (including traps), you must also configure an SNMPv3 management station.
2. To reconfigure any of the SNMP notifications that are enabled by default to be sent to a management station (trap receiver.)
3. (Optional) See the following sections to configure optional SNMP notification features and verify the current configuration:
   a. **Source IP address for SNMP notifications** on page 280
   b. **SNMP notification configuration** on page 282

SNMPv1 and SNMPv2c Traps

The switches support the following functionality from earlier SNMP versions (SNMPv1 and SNMPv2c):

- Trap receivers: A **trap receiver** is a management station to which the switch sends SNMP traps and (optionally) event log messages sent from the switch. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch.
- Default Traps: A switch automatically sends default traps to trap receivers using the configured community name. You have to configure and supply the community name to use in the trap-receiver config, there is no default. Use the snmp-server host <IP_ADDRESS> community "<COMMUNITY_NAME>" command to configure a community name and the snmp-server host <IP_ADDRESS> community "<COMMUNITY_NAME>" trap-level [all | critical | not-info | debug | none] command to set the level of traps to send to the community.
- Thresholds: A switch automatically sends all messages created when a system threshold is reached to the network management station that configured the threshold, regardless of the trap receiver configuration.

SNMP trap receivers

Use the snmp-server host command to configure a trap receiver that can receive SNMPv1 and SNMPv2c traps, and (optionally) Event Log messages. When you configure a trap receiver, you specify its community membership, management station IP address, and (optionally) the type of Event Log messages to be sent.

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps are sent to that trap receiver until the community to which it belongs has been configured on the switch.

NOTE

To replace one community name with another for the same IP address, you must first enter the

no snmp-server host *<COMMUNITY-NAME> <IP-ADDRESS>*

command to delete the unwanted community name. Otherwise, if you add a new community name with an IP address that is already used with a different community name, two valid community name entries are created for the same management station.

If you do not specify the event level ([none | all | not-info | critical | debug]), the switch does not send Event Log messages as traps. However, "well-known" traps and threshold traps (if configured) are still sent.

SNMP trap when MAC address table changes

An SNMP trap is generated when a laptop/PC is removed from the back of an IP phone and the laptop/PC MAC address ages out of the MAC table for the Switch 2920 and 5400 series switch.

The mac-notify trap feature globally enables the generation of SNMP trap notifications on MAC address table changes (learns/moves/removes/ages.)

The following command enables trap for aged MAC addresses:

Show mac-notify traps

Syntax

show mac-notify traps

Description

Displays the different mac-notify traps configured on an interface.

Output of show mac-notify traps

```
Mac Notify Trap Information
Mac-notify Enabled : No
Mac-move Enabled : No
Trap-interval : 30
Port   MAC Addresses trap learned/removed/aged
------ ---------------------------------------
```

```
1        Learned, Removed & Aged
2      Removed & Aged
3      Learned & Aged
4      Learned & Removed
5      Aged
6      Learned
7      Removed
```

show mac-notify for port 1

```
show mac-notify traps 1
```

```

1 Aged
```

Physical hardware removal/insertion trap generation

The specific events related to a physical module insertion or removal are being added to the default trap list.

Aruba 3810M Switch Series (JL071A, JL072A, JL073A, JL074A, JL075A, JL076A)
Aruba 5400Rzl2 Switch Series (J8698A, J8700A, J9823A-J9824A, J9825A, J9826A, J9868A, J9447A, J9448A)
Aruba 5406R Switch Series (JL002A, JL003A, JL095A,J9850A)
Aruba 5412R Switch Series (J9851A, JL001A)
Aruba 3800 Switch Series (J9573A, J9574A, J9575A, -J9576A, J9584A, J9585A, J9586A, J9587A, J9588A)

Current default traps

The default event scenarios for currently generated traps on ArubaOS-Switches are:

- Device cold start notifications
- Device warm start notifications
- Port down notifications
- Port up notifications
- Authentication failure notifications
- Enterprise change notifications
- Intrusion alarm notifications

Event scenario matrix

Different event scenarios for which traps are generated:

Event Id	Severity	Action	Message
68	Info	Slot Insertion	I 06/20/16 09:18:43 00068 chassis: AM1: Slot C Inserted
67	Info	Slot Removal	I 06/20/16 09:18:50 00067 chassis: AM1: Slot C Removed
405	Info	Transceiver Insertion	I 06/20/16 09:18:56 00405 ports: AM1: port A23 xcvr hot-swap insert
406	Info	Transceiver Removal	I 06/20/16 09:19:04 00406 ports: AM1: port A23 xcvr hot-swap remove

Table Continued

Event Id	Severity	Action	Message
552	Warning	Stacking module Insertion	W 04/20/16 09:20:43 00552 chassis: ST1-CMDR: Stacking Module insertion detected: Reboot required
552	Warning	Stacking module Removal	W 06/20/16 09:19:43 00552 chassis: ST1-CMDR: Stacking Module removal detected: Reboot required

Enabling and disabling traps

Action	Command
Disable both the log and trap	`setMib eventType.<event_Id> -i 1 – to disable both log & Trap`
Enable log only	`setMib eventType.<event_Id> -i 2 – to allow only log`
Enable both the log and trap (Default)	`setMib eventType.<event_Id> -i 4 – to allow both log & Trap`
Enable trap only	`setMib eventType.<event_Id> -i 3 – to allow only trap`

NOTE

If the event is configured to disable a trap, then the trap will not be sent for that particular event. In all other scenarios, a trap is generated for the listed events.

SNMP trap captures examples

Inserting a slot module

Event Id: 68

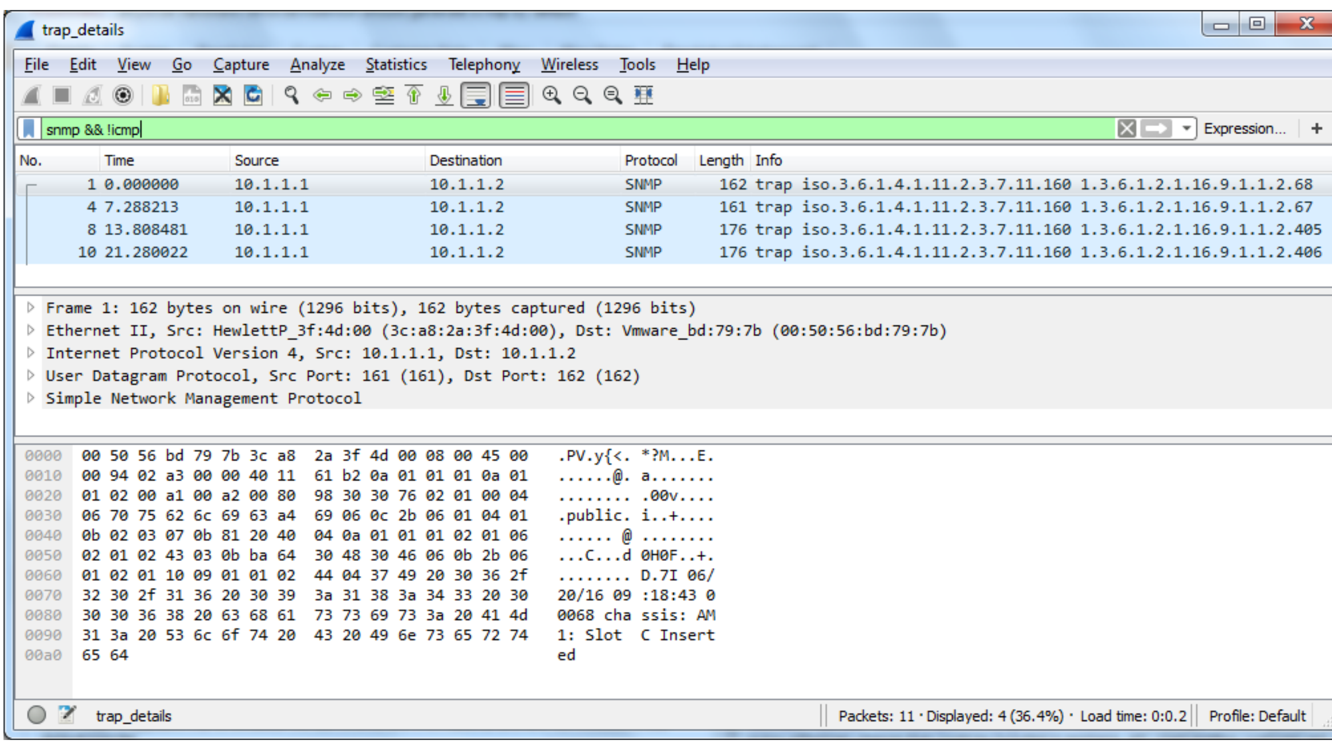

Removing a slot module

Event Id: 67

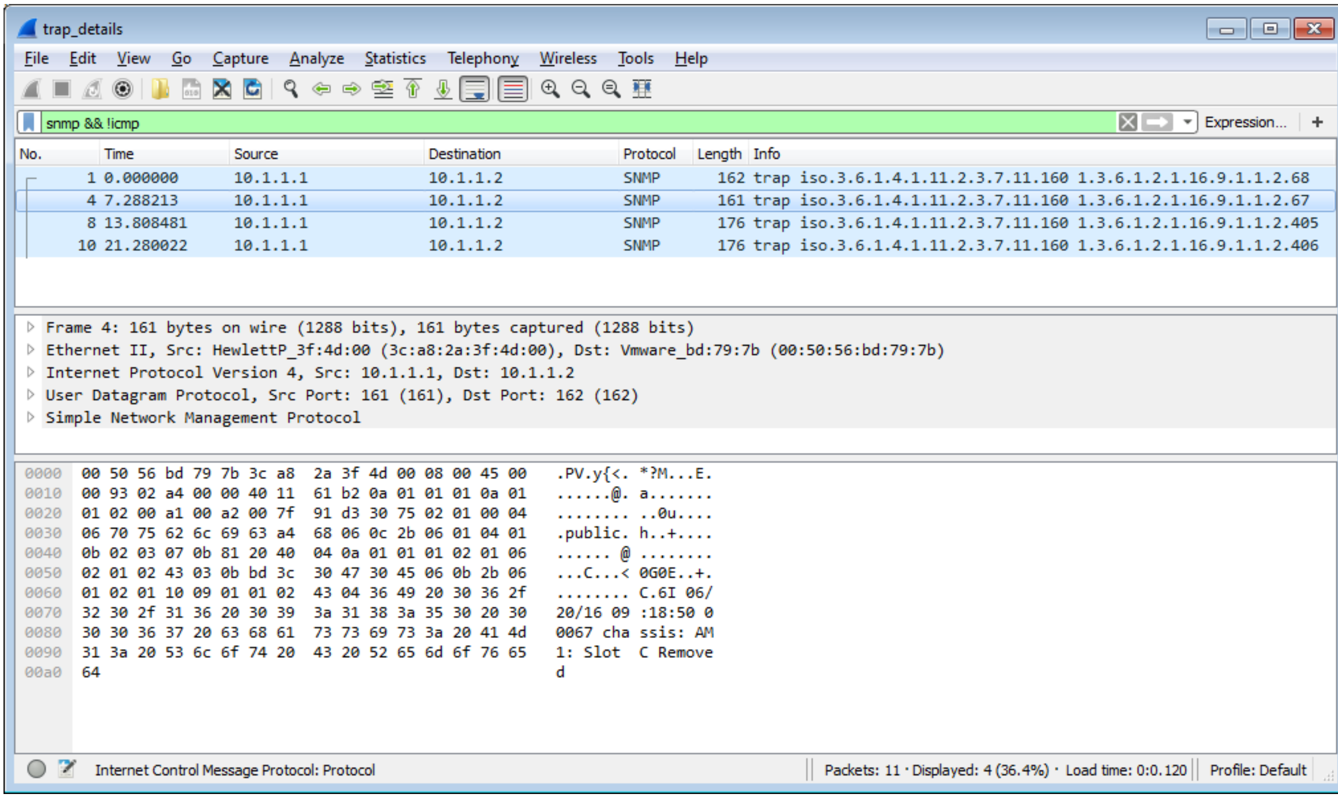

Inserting a transceiver

Event Id: 405

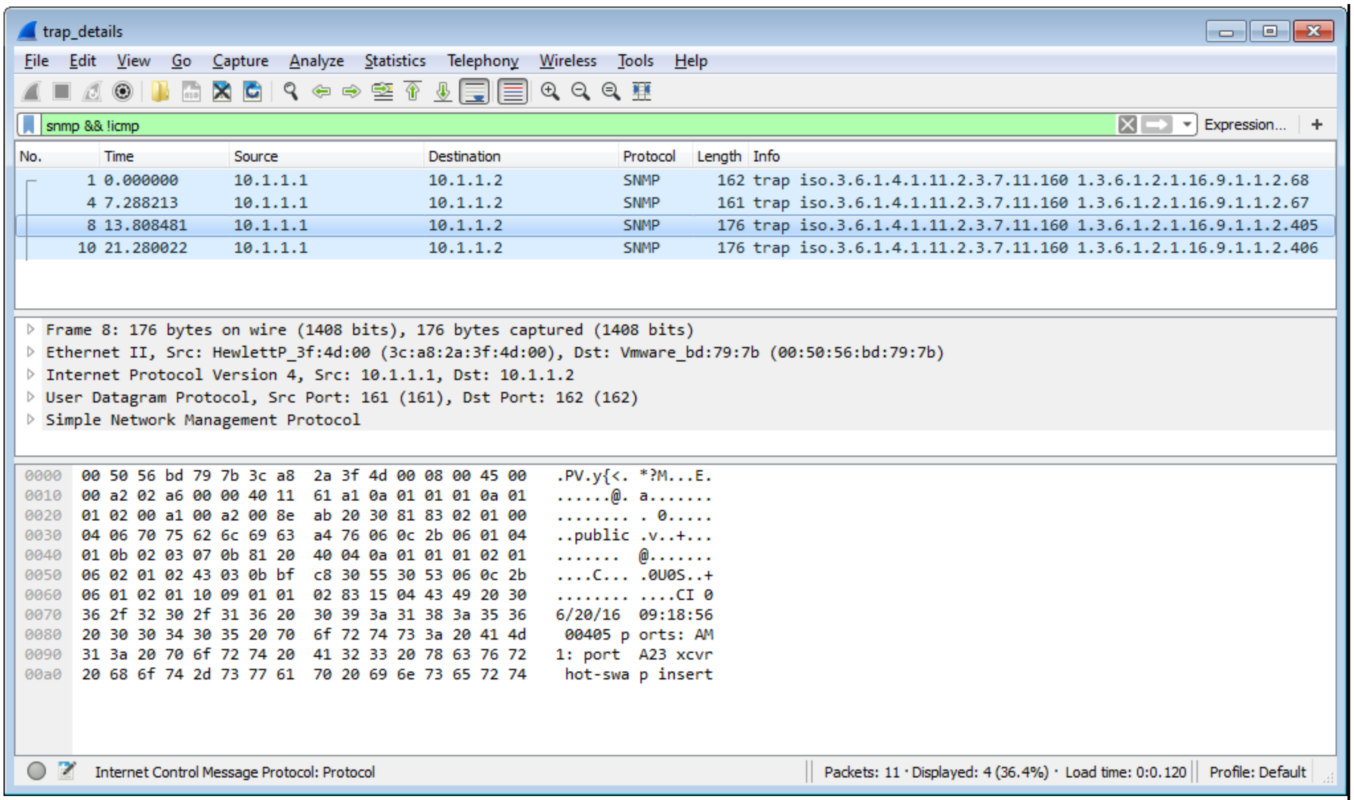

Removing a transceiver

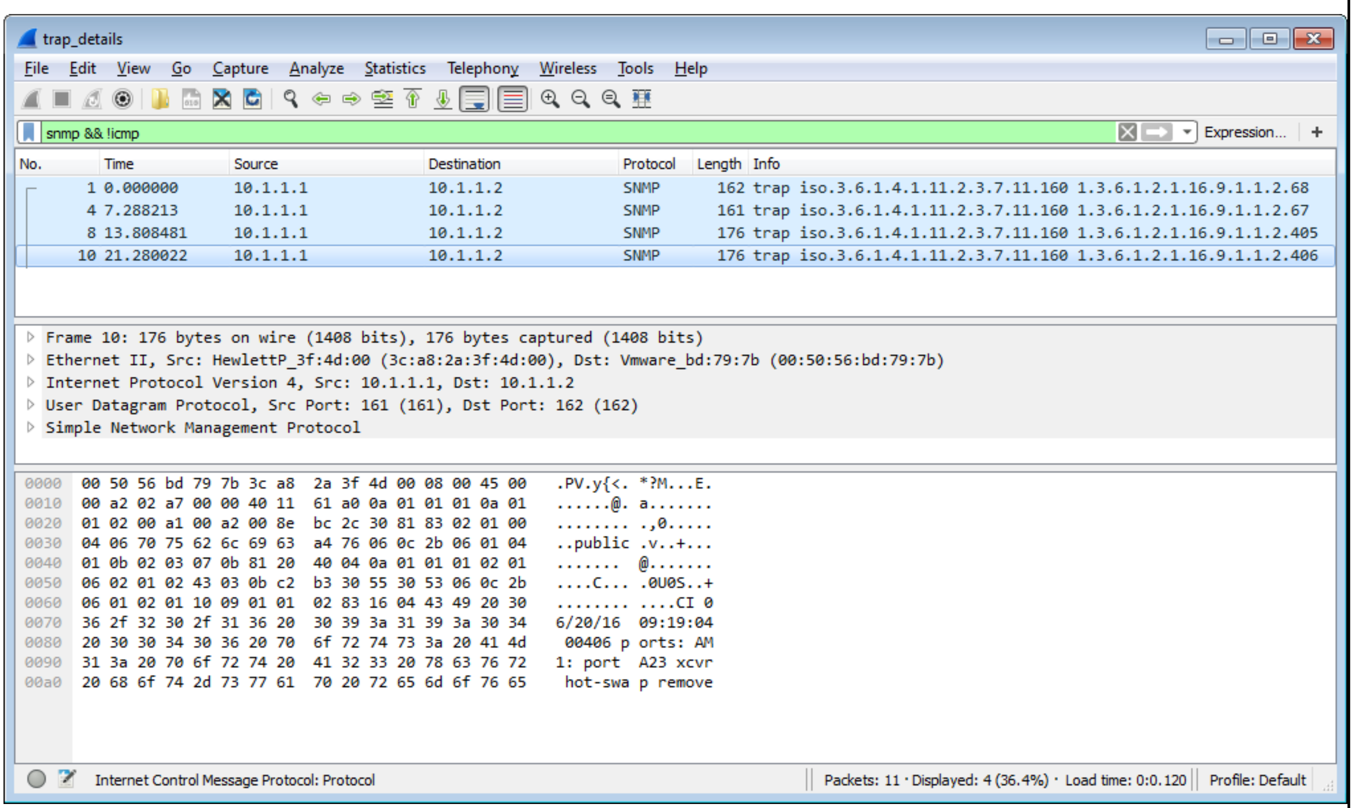

SNMP trap when power supply is inserted or removed

SNMP traps generate while inserting or removing a powered up Power Supply Unit (PSU) without pulling out the power cable and also when removing a powered down PSU from the Switch 5406 Series. RMON log events are used to generate SNMP traps for PSU insertion and removal in both powered up and powered down states.

Log event

```
Chassis: Power Supply 1 inserted
Chassis: Power Supply 1 removed while powered
Chassis: Power Supply 2 removed while not powered
```

Power supply inserted while powered off

```
W 09/13/13 09:10:18 03834 chassis: AM1: Power Supply 1 inserted
W 09/13/13 09:10:19 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures:
4
```

Power supply removed while powered off

```
W 09/13/13 09:08:57 03835 chassis: AM1: Power Supply 1 removed while not powered
W 09/13/13 09:08:57 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures:
3
```

Power supply removed while powered on

WARNING

PSUs should never be inserted or removed while the PSU has power applied.

```
W 09/13/13 09:03:36 03835 chassis: AM1: Power Supply 1 removed while powered
W 09/13/13 09:03:36 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures:
2
```

SNMP notification support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices, and control the interval between successive notifications of data changes on the same neighbor.

SNMPv3 users

To create new users, most SNMPv3 management software requires an initial user record to clone. The initial user record can be downgraded and provided with fewer features, but not upgraded by adding new features. For this reason, Hewlett Packard Enterprise recommends that when you enable SNMPv3, you also create a second user with SHA authentication and DES privacy.

To use SNMPv3 on the switch, you must configure the users that will be assigned to different groups:

Procedure

1. Configure users in the User Table with the `snmpv3 user` command.

To view the list of configured users, enter the `show snmpv3 user` command.

2. Assign users to Security Groups based on their security model with the `snmpv3 group` command.

If you add an SNMPv3 user without authentication, privacy, or both, to a group that requires either feature, the user will not be able to access the switch. Ensure that you add a user with the appropriate security level to an existing security group.

Add users

To configure an SNMPv3 user, you must first add the user name to the list of known users with the `snmpv3 user` command, as shown in **Figure 81: Adding SNMPv3 users and displaying SNMPv3 configuration** on page 268.

Figure 81: *Adding SNMPv3 users and displaying SNMPv3 configuration*

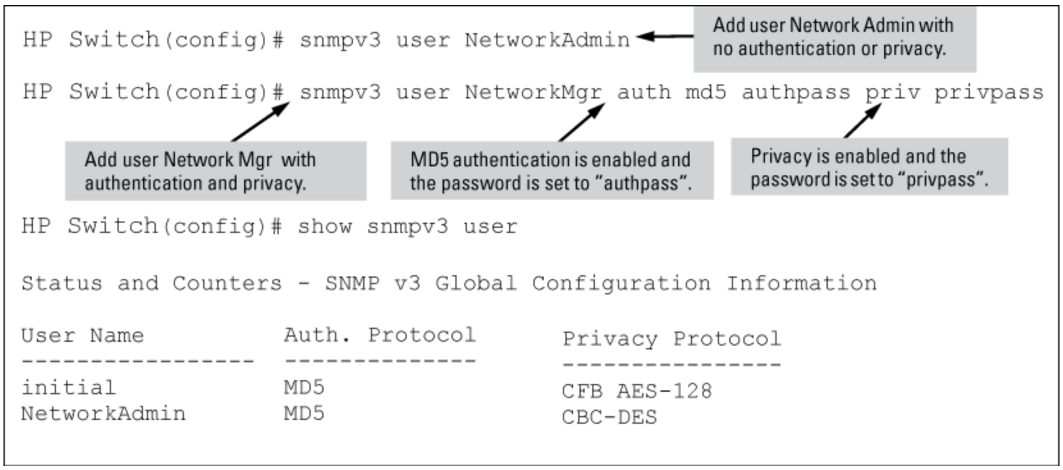

SNMP tools for switch management

SNMP is a management protocol that allows an SNMP client application to retrieve device configuration and status information and to configure the device (**get** and **set**.) You can manage the switch via SNMP from a network management station.

To implement SNMP management, the switch must have an IP address configured either manually or dynamically (using DHCP or Bootp.) If multiple VLANs are configured, each VLAN interface should have its own IP address.

If you use the switch's Authorized IP Managers and Management VLAN features, ensure that the SNMP management station, the choice of switch port used for SNMP access to the switch, or both, are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked.

SNMP management features

SNMP management features on the switch include:

- SNMP version 1, version 2c, or version 3 over IP
- Security via configuration of SNMP communities (**SNMPv3 communities** on page 260)
- Security via authentication and privacy for SNMPv3 access
- Event reporting via SNMP
    - Version 1 traps
    - RMON: groups 1, 2, 3, and 9
- PCM/PCM+

- Flow sampling using sFlow
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in a Hewlett Packard Enterprise proprietary MIB (management information base) file.

Downloading the latest MIB file

Procedure

1. Go to the Networking website at:

   http://www.hpe.com/Networking/support
2. Enter the model number of your switch (for example, 8212) or the product number in the **Auto Search** text box.
3. Select an appropriate product from the drop down list.
4. Click **Display selected**.
5. From the options that appear, select **Software downloads**.
6. Locate the list of MIBs with the switch software in the Other category.
7. Click **software updates**, and then click **MIBs**.

SNMPv1 and v2c access to the switch

SNMP access requires an IP address and subnet mask configured on the switch. If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address.

Once an IP address is configured, the main steps for configuring SNMPv1 and v2c access management features are:

Procedure

1. Configure the appropriate SNMP communities. (See **SNMPv3 communities** on page 260.)
2. Configure the appropriate trap receivers. (See **SNMP notifications** on page 261.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (See the access security guide.)

SNMPv3 access to the switch

SNMPv3 access requires an IP address and subnet mask configured on the switch. If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address.

Once you have configured an IP address, the main steps for configuring SNMPv3 access management features are the following:

Procedure

1. Enable SNMPv3 for operation on the switch.
2. Configure the appropriate SNMP users.
3. Configure the appropriate SNMP communities.
4. Configure the appropriate trap receivers.

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct User and community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the IP Authorized Manager feature for the switch. (See the access security guide.)

SNMP version 3 (SNMPv3) adds some new commands to the CLI for configuring SNMPv3 functions. To enable SNMMPv3 operation on the switch, use the `snmpv3 enable` command. An initial user entry will be generated with MD5 authentication and DES privacy.

You may (optionally) restrict access to only SNMPv3 agents by using the `snmpv3 only` command. To restrict write-access to only SNMPv3 agents, use the `snmpv3 restricted-access` command.

Restricting access to only version 3 messages will make the community named “public” inaccessible to network management applications (such as autodiscovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

Enabling SNMPv3

The `snmpv3 enable` command allows the switch to:

- Receive SNMPv3 messages.
- Configure initial users.
- Restrict non-version 3 messages to "read only" (optional.)

Restricting access to only version 3 messages makes the community named "public" inaccessible to network management applications (such as autodiscovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

SNMP version 3 `enable` command

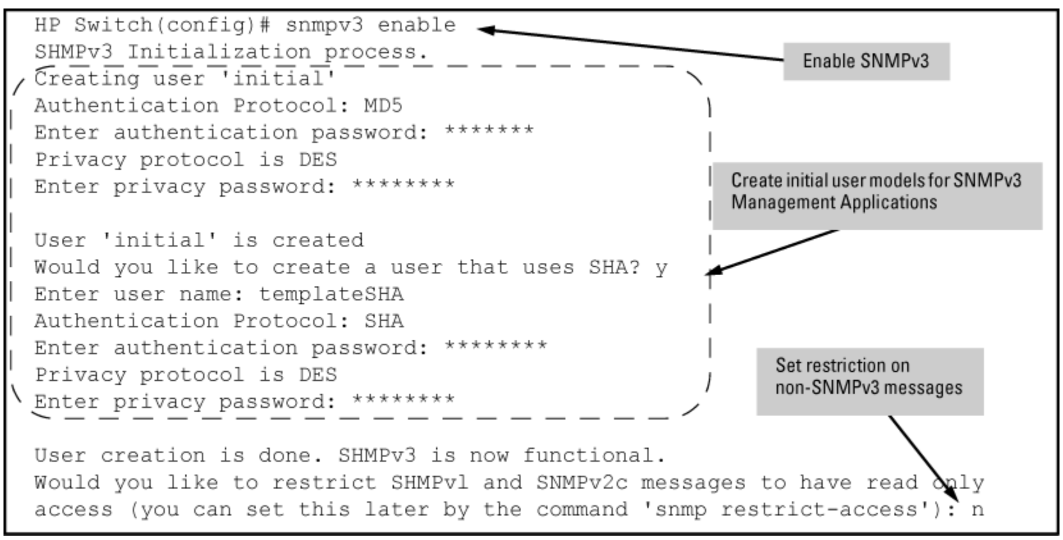

Configuring users in SNMPv3

snmpv3 user

Syntax

```
[no] snmpv3 user <USER_NAME> [auth md5|sha] <AUTH_PASS> [priv des|aes] <PRIV_PASS>
```

Description

Adds or deletes a user entry for SNMPv3. Authorization and privacy are optional, but to use privacy, you must use authorization.

Parameters and options

no

Used to delete a user entry. When you delete a user, only the user name is required.

<USER_NAME>

<AUTH_PASS>

With authorization, you can set either MD5 or SHA authentication. The authentication password *auth_pass* must be 6 to 32 characters and is mandatory when you configure authentication.

priv des|aes

With privacy, the switch supports DES (56-bit) and AES (128-bit) encryption. Defaults to DES. Only AES 128-bit and DES 56-bit encryption are supported as privacy protocols. Other non-standard encryption algorithms, such as AES-172, AES-256, and 3-DES are not supported.

<PRIV_PASS>

The privacy password *priv_pass* must be 6 to 32 characters and is mandatory when you configure privacy.

For the 3800 switches, when the switch is in enhanced secure mode, commands that take a password as a parameter have the echo of the password typing replaced with asterisks. The input for the password is prompted for interactively. Additionally, the DES option is not available. For more information, see the access security guide.

Switch access from SNMPv3 agents

snmpv3 enable

Syntax

```
snmpv3 enable
```

Description

Enables switch access from SNMPv3 agents, including the creation of the initial user record.

Restrict access from SNMPv3 agents

snmpv3 only

Syntax

```
[no] snmpv3 only
```

Description

When enabled, the switch rejects all non-SNMPv3 messages.

Restrict non-SNMPv3 agents to read-only access

snmpv3 restricted-access

Syntax

```
[no] snmpv3 restricted-access
```

Description

Enable or disable restrictions from all non-SNMPv3 agents (read-only access).

Operating status of SNMPv3

show snmpv3

Syntax

`show snmpv3 enable`

Description

View the operating status of SNMPv3 (enabled or disabled).

Non-SNMPv3 message reception status

show snmpv3 only

Syntax

`show snmpv3 only`

Description

Shows the message reception status of non-SNMPv3 messages.

Non-SNMPv3 write message status

show snmpv3 restricted-access

Syntax

`show snmpv3 restricted-access`

Description

Shows the status of non-SNMPv3 write messages.

Viewing and configuring non-version-3 SNMP communities (Menu)

Procedure

1. From the Main Menu, select:

   2. Switch Configuration…

6. SNMP Community Names

Figure 82: *SNMP Communities screen (default values)*

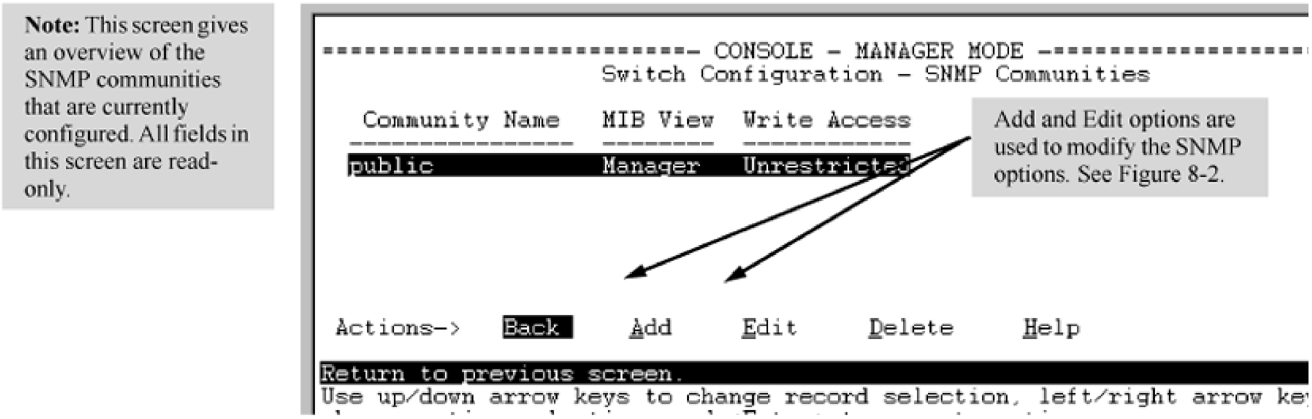

2. Press **[A]** (for **Add**) to display the following screen:

Figure 83: *SNMP add or edit screen*

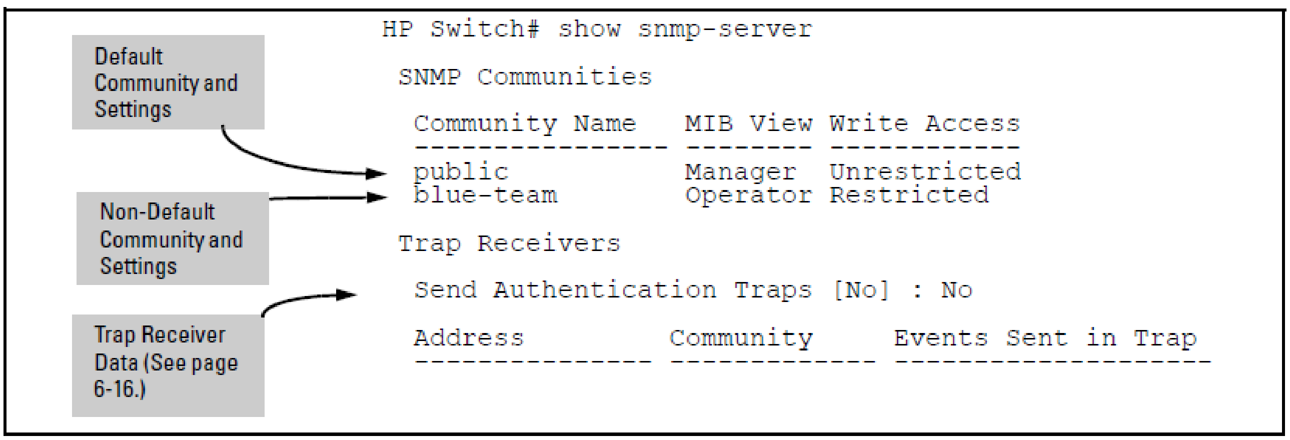

If you need information on the options in each field, press **[Enter]** to move the cursor to the Actions line, then select the Help option. When you are finished with Help, press **[E]** (for Edit) to return the cursor to the parameter fields.

3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the **[Tab]** key to move from one field to the next.)
4. Press **[Enter]**, then **[S]** (for **Save**.)

SNMP trap receiver configuration

snmp-server host

Syntax

`snmp-server host [ipv4-addr|ipv6-addr] <community name>`

Description

Configures a destination network management station to receive SNMPv1/v2c traps and (optionally) Event Log messages sent as traps from the switch, using the specified community name and destination IPv4 or IPv6 address. You can specify up to ten trap receivers (network management stations.) <COMMUNITY NAME> defaults to **public**.

Parameters and options

Event log messages

Optional: Configures the security level of the Event Log messages you want to send as traps to a trap receiver.

- The type of Event Log message that you specify applies only to Event Log messages, not to threshold traps.
- For each configured event level, the switch continues to send threshold traps to all network management stations that have the appropriate threshold level configured.
- If you do not specify an event level, the switch uses the default value (none) and sends no Event Log messages as traps.

none

Sends no Event Log messages.

all

Sends all Event Log messages.

not info

Sends all Event Log messages that are not for information only.

critical

Sends only Event Log messages for critical error conditions.

debug

Sends only Event Log messages needed to troubleshoot network- and switch-level problems.

[inform]

Optional: Configures the switch to send SNMPv2 inform requests when certain events occur.

Configure a trap receiver

To configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" event log messages, you can enter the following command:

```
switch# snmp-server host 10.28.227.130 red-team critical
```

SNMPv2c inform option

snmp-server host

Syntax

```
[no] snmp-server host [ipv4-addr|ipv6-addr] <COMMUNITY NAME> inform [retries
<COUNT>][timeout <INTERVAL>]
```

Description

Enables (or disables) the `inform` option for SNMPv2c on the switch and allows you to configure options for sending SNMP inform requests.

Parameters and options

The `retries` and `timeout` values are not used to send trap requests.

retries

Maximum number of times to resend an `inform` request if no SNMP response is received. Defaults to 3.

timeout

Number of seconds to wait for an acknowledgement before resending the `inform` request. Defaults to 15 seconds.

Verify SNMPv2c inform configuration

```
switch# show snmp-server

 SNMP Communities

  Community Name   MIB View Write Access
  ---------------- -------- ------------ public          Manager  Unrestricted

 Trap Receivers

  Link-Change Traps Enabled on Ports [All] : All
  ...
  Address                Community        Events Sent Notify Type Retry Timeout
  ---------------------- ---------------- ----------- ----------- ----- --------
  15.28.333.456          guest            All         inform      3     15

 Excluded MIBs

 Snmp Response Pdu Source-IP Information

  Selection Policy   : Default rfc1517

 Trap Pdu Source-IP Information
  Selection Policy   : Configured IP
  Ip Address        : 10.10.10.10
```

Configuring SNMPv3 notifications

The SNMPv3 notification process allows messages that are passed via SNMP between the switch and a network management station to be authenticated and encrypted.

Procedure

1. Enable SNMPv3 operation on the switch by entering the `snmpv3 enable` command.

   When SNMPv3 is enabled, the switch supports:

   - Reception of SNMPv3 notification messages (traps and informs)
   - Configuration of initial users
   - (Optional) Restriction of non-SNMPv3 messages to "read only"
2. Configure SNMPv3 users by entering the `snmpv3 user` command. Each SNMPv3 user configuration is entered in the User Table.
3. Assign SNMPv3 users to security groups according to their level of access privilege by entering the `snmpv3 group` command.
4. Define the name of an SNMPv3 notification configuration by entering the `snmpv3 notify` command (see **snmpv3 notify** on page 276).

5. Configure the target address of the SNMPv3 management station to which SNMPv3 informs and traps are sent by entering the `snmpv3 targetaddress` command (see **snmpv3 targetaddress** on page 276).
6. Create a configuration record for the target address with the `snmpv3 params` command (see **snmpv3 params** on page 277).

Configuring SNMPv2 notification

Figure 84: *SNMPv3 notification configuration*

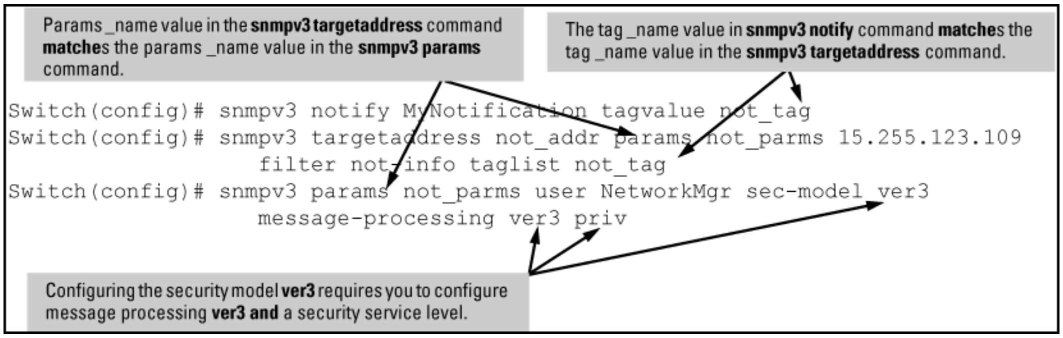

snmpv3 notify

Syntax

`[no] snmpv3 notify` *`notify_name`* `tagvalue` *`tag_name`*

Description

Associates the name of an SNMPv3 notification configuration with a tag name used (internally) in SNMPv3 commands. To delete a notification-to-tag mapping, enter `no snmpv3 notify` *`notify_name`* .

Options

notify *notify_name*

Specifies the name of an SNMPv3 notification configuration.

tagvalue *tag_name*

Specifies the name of a tag value used in other SNMPv3 commands, such as `snmpv3 targetaddress params taglist` *`tag_name`* in Step 5.

snmpv3 targetaddress

Syntax

`[no] snmpv3 targetaddress [ipv4-addr|ipv6-addr]` *`<NAME>`*

Description

Configures the IPv4 or IPv6 address, name, and configuration filename of the SNMPv3 management station to which notification messages are sent.

Parameters and options

params *parms_name*

Name of the SNMPv3 station's parameters file. The parameters filename configured with `params` *`params_name`* must match the `params` *`params_name`* value entered with the `snmpv3 params` command in Step 6.

taglist *tag_name* [tag_name] …

Specifies the SNMPv3 notifications (identified by one or more *tag_name* values) to be sent to the IP address of the SNMPv3 management station.You can enter more than one *tag_name* value.

Each *tag_name* value must be already associated with the name of an SNMPv3 notification configuration entered with the `snmpv3 notify` command in Step 4.

Use a blank space to separate *tag_name* values. You can enter up to 103 characters in *tag_name* entries following the `taglist` keyword.

[filter <none|debug|all|not-info|critical>]

(Optional) Configures the type of messages sent to a management station. Defaults to none.

udp-port <PORT>

(Optional) Specifies the UDP port to use. Defaults to 162.

port-mask <MASK>

(Optional) Specifies a range of UDP ports. (Default: 0.)

[addr-mask <MASK>

(Optional) Specifies a range of IP addresses as destinations for notification messages. Defaults to 0.

retries <VALUE>

(Optional) Number of times a notification is retransmitted if no response is received. Range: 1-255. Defaults to 3.

timeout <VALUE>

(Optional) Time (in millisecond increments) allowed to receive a response from the target before notification packets are retransmitted. Range: 0-2147483647. Defaults to 15 seconds.

max-msg-size <SIZE>

Optional) Maximum number of bytes supported in a notification message to the specified target. Defaults to 1472.

snmpv3 params

Syntax

```
[no] snmpv3 params <PARAMS_NAME> user <USER_NAME>
```

Description

Applies the configuration parameters and IP address of an SNMPv3 management station (from the `params` *`params_name`* value configured with the `snmpv3 targetaddress` command in Step 5) to a specified SNMPv3 user (from the `user` *`user_name`* value configured with the `snmpv3 user` command in Step 2.)

If you enter the `snmpv3 params user` command, you must also configure a security model (`sec-model`) and message processing algorithm (`msg-processing`.)

Parameters and options

[sec-model <ver1|ver2c|ver3>]

Configures the security model used for SNMPv3 notification messages sent to the management station configured with the `snmpv3 targetaddress` command in Step 5. If you configure the security model as `ver3`, you must also configure the message processing value as `ver3`.

[msg-processing <ver1|ver2c|ver3|noaut|auth|priv>]

Configures the algorithm used to process messages sent to the SNMPv3 target address. If you configure the message processing value as `ver3` and the security model as `ver3`, you must also configure a security services level (`noauth`, `auth`, or `priv`.)

SNMPv3 community mapping

SNMP communities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch.

snmpv3 community

Syntax

`[no] snmpv3 community`

Description

Maps or removes a mapping of a community name to a group access level. To remove a mapping you need to specify only the `index_name` parameter.

Parameters and options

index <INDEX_NAME>

An index number or title for the mapping. The values of 1 to 5 are reserved and can not be mapped.

name <COMMUNITY_NAME>

The community name that is being mapped to a group access level.

sec-name <SECURITY_NAME>

The group level to which the community is being mapped.

tag <TAG_VALUE>

This is used to specify which target address may have access by way of this index reference.

Assign a community to a group

Figure 85: Assigning a community to a group access level on page 278 shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator has an access level of CommunityOperatorReadOnly.

Figure 85: *Assigning a community to a group access level*

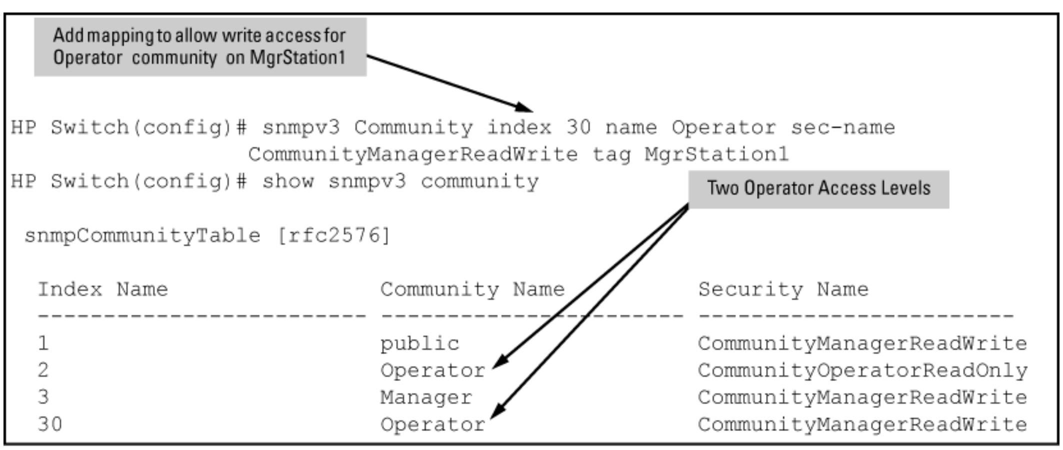

Running configuration changes and SNMP traps

Syntax

```
[no] snmp-server enable trapsfig-change transmission-interval <0-4294967295>
```

Description

Enables SNMP traps on running configurations.

Parameters and options

running-con

Enables SNMP traps being sent when changes to the running configuration file are made. Defaults to disabled.

transmission-interval <0-2147483647>

Controls the egress rate for generating SNMP traps for the running configuration file. The value configured specifies the time interval in seconds that is allowed between the transmission of two consecutive traps.

None of the running configuration change events that occur within the specified interval generate SNMP traps, although they are logged in the Configuration Changes History Table.

A value of 0 (zero) means there is no limit; traps can be sent for every running configuration change event. Defaults to 0.

Startup configuration changes and SNMP traps

You can send a specific SNMP trap for any configuration change made in the switch's startup configuration file when the change is written to flash. Changes to the configuration file can occur when executing a CLI write command, executing an SNMP set command directly using SNMP, or when using the WebAgent

NOTE

A log message is always generated when a startup configuration change occurs. An example log entry is:

```
I 07/06/10 18:21:39 02617 mgr: Startup configuration changed by SNMP. New
seq. number 8
```

The corresponding trap message is sent if the snmp-server enable traps startupconfig- change command is configured.

snmp-server enable traps startup-config-change

Syntax

```
snmp-server enable traps startup-config-change
```

Description

Enables notification of a change to the startup configuration. The change event is logged. Default: Disabled

Startup configuration changes

The number that displays when show config is executed is global for the switch and represents the startup configuration sequence number.

Figure 86: *Notification of changes to the Startup Configuration file*

```
Switch(config)# snmp-server enable traps startup-config-change

Switch(config)# show config

Startup configuration: 16

; J8697A Configuration Editor; Created on release #K.14.54

hostname "Switch"
module 1 type J8702A
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24, B1-B10
   ip address dhcp-bootp
   exit
snmp-server community "public" unrestricted
```

The number "16" is global for the switch and represents the startup configuration sequence number.

Fields in the trap when making a change

Fields in the trap when a change is made via SNMP (station ip=0xAC161251 (172.22.18.81), no username is set, and the new sequence number is 16.)

Figure 87: *Fields when the SNMP trap is set*

```
+ Internet Protocol, Src: 172.22.18.57 (172.22.18.57), Dst: 172.22.18.81 (172.22.18.81)
+ User Datagram Protocol, Src Port: snmp (161), Dst Port: snmptrap (162)
- Simple Network Management Protocol
    version: version-1 (0)
    community: public
  - data: trap (4)
    - trap
        enterprise: 1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1 (SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1)
        agent-addr: 172.22.18.57 (172.22.18.57)
        generic trap: enterpriseSpecific (6)
        specific-trap: 6
        time-stamp: 65437
      - variable-bindings: 6 items
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.9 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.9): 16
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.1 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.1): 2
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.2 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.2): 4
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.3 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.3): AC161251
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.4 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.4): <MISSING>
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.5 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.5): 1
```

Source IP address for SNMP notifications

When you use the `snmp-server response-source` and `snmp-server trap-source` commands, note the following behavior:

- The snmp-server response-source and snmp-server trap-source commands configure the source IP address for IPv4 interfaces only.
- You must manually configure the snmp-server response-source value if you wish to change the default user-defined interface IP address that is used as the source IP address in SNMP traps (RFC 1517.)
- The values configured with the snmp-server response-source and snmp-server trap-source commands are applied globally to all interfaces that are sending SNMP responses or SNMP trap PDUs.
- Only the source IP address field in the IP header of the SNMP response PDU can be changed.
- Only the source IP address field in the IP header and the SNMPv1 Agent Address field of the SNMP trap PDU can be changed.

snmp-server response-source

Syntax

```
[no] snmp-server response-source [dst-ip-of-request <ipv4-addr|ipv6-addr> loopback
<0-7>]
```

Description

Specifies the source IP address of the SNMP response PDU. The default SNMP response PDU uses the IP address of the active interface from which the SNMP response was sent as the source IP address. Defaults to Interface IP address.

Parameters and options

no

The no form of the command resets the switch to the default behavior (compliant with rfc-1517.)

dst-ip-of-request

Destination IP address of the SNMP request PDU that is used as the source IP address in an SNMP response PDU.

ipv4–addr|ipv6–addr

User-defined interface IP address that is used as the source IP address in an SNMP response PDU. Both IPv4 and IPv6 addresses are supported.

loopback 0-7

IP address configured for the specified loopback interface that is used as the source IP address in an SNMP response PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.

Destination interface IP as source IP

To use the IP address of the destination interface on which an SNMP request was received as the source IP address in the IP header of SNMP traps and replies, enter the following command:

```
switch# snmp-server response-source dst-ip-of-request
```

snmp-server trap-source

Syntax

```
[no] snmp-server trap-source ipv4-addr loopback0-7
```

Description

Specifies the source IP address to be used for a trap PDU. To configure the switch to use a specified source IP address in generated trap PDUs, enter the snmp-server trap-source command. Defaults to the interface IP address in generated trap PDUs.

Parameters and options

no

The `no` form of the command resets the switch to the default behavior (compliant with rfc-1517.)

dst-ip-of-request

Destination IP address of the SNMP request PDU that is used as the source IP address in an SNMP response PDU.

ipv4–addr

User-defined interface IPv4 address that is used as the source IP address in generated traps. IPv6 addresses are not supported.

loopback `0-7`

IP address configured for the specified loopback interface that is used as the source IP address in a generated trap PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.

SNMP replies and traps configuration

To verify the configuration of the interface IP address used as the source IP address in IP headers for SNMP replies and traps sent from the switch, enter the `show snmp-server` command to display the SNMP policy configuration, as shown in **Figure 88: Display of source IP address configuration** on page 282.

Figure 88: *Display of source IP address configuration*

```
HP Switch(config)# show snmp-server

 SNMP Communities

  Community Name   MIB View Write Access
  ---------------- -------- ------------
  public           Manager  Unrestricted

 Trap Receivers
  Link-Change Traps Enabled on Ports [All] : All

  ...

 Excluded MIBs
 Snmp Response Pdu Source-IP Information
  Selection Policy   : dstIpOfRequest

 Trap Pdu Source-IP Information
  Selection Policy   : Configured IP
  Ip Address         : 10.10.10.10
```

dstIpOfRequest: The destination IP address of the interface on which an SNMP request is received i s used as the source IP address in SNMP replies.

SNMP notification configuration

show snmp-server

Syntax

`show snmp-server`

Description

Displays the currently configured notification settings for versions SNMPv1 and SNMPv2c traps, including SNMP communities, trap receivers, link-change traps, and network security notifications.

show snmp-server output

In the following example, the `show snmp-server` command output shows that the switch has been configured to send SNMP traps and notifications to management stations that belong to the "public," "red-team," and "blue-team" communities.

Figure 89: *Display of SNMP notification configuration*

```
HP Switch(config)# show snmp-server

 SNMP Communities                                               SNMP
  Community Name   MIB View Write Access                        Community
  ---------------- -------- ------------
  public           Operator Restricted
  blue-team        Manager  Unrestricted
  red-team         Manager  Unrestricted
                                                                Link-change
 Trap Receivers                                                 trap setting

  Link-Change Traps Enabled on Ports [All] : All

  Trap Category                     Current Trap Configuration
  ------------------------------    --------------------------
  SNMP Authentication               extended
  Password change                   enabled
  Login failures                    enabled             Network security
  Port-Security                     enabled             notification
  Authorization Server Contact      enabled
  ARP Protection                    enabled
  DHCP Snooping                     enabled

  Address          Community   Events Sent  Notify Type  Retry  Timeout
  ---------------  ----------  -----------  -----------  -----  -------
  10.28.227.200    public      All          trap           3      15
  10.28.227.105    red-team    Critical     trap           3      15
  10.28.227.120    blue-team   Not-INFO     trap           3      15
...
```

Assign users to groups

snmpv3 group

Syntax

`snmpv3 group`

Description

Sets the group access level for the user by assigning the user to a group. Assigns or removes a user to a security group for access rights to the switch. To delete an entry, include all parameters in the command.

Parameters and options

group <GROUP_NAME>

Identifies the group that has the privileges that will be assigned to the user.

user <USER_NAME>

Identifies the user to be added to the access group. This must match the user name added with the `snmpv3 user` command.

sec-model <aver1|ver2|ver3>

Defines which security model to use for the added user. An SNMPv3 access group should use only the ver3 security model.

snmpv3 group

Figure 90: *Using snmpv3 group*

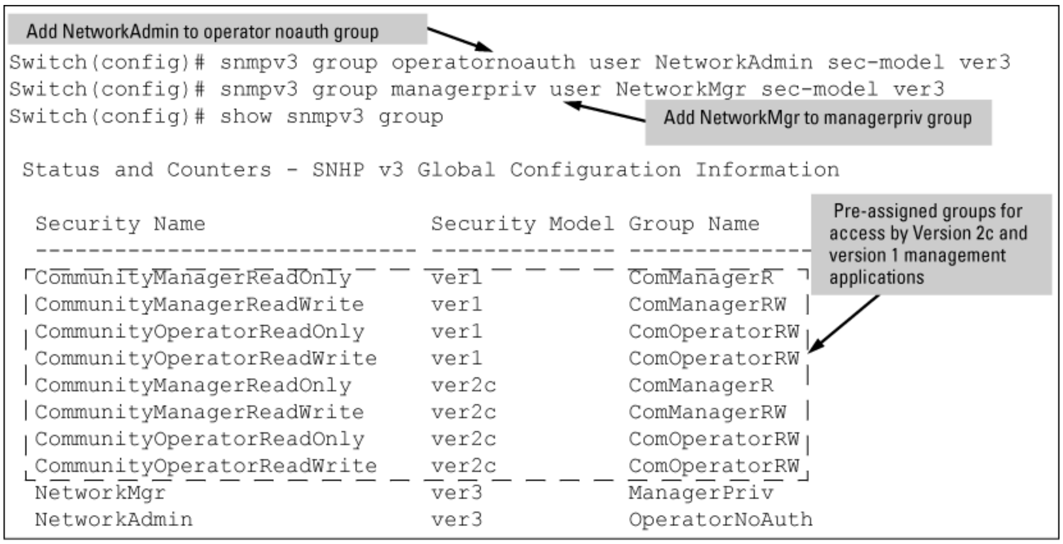

snmp-server community

Syntax

`[no] snmp-server community` *`community-name`*

Description

Configures a new community name.

- If you do not also specify `operator` or `manager`, the switch automatically assigns the community to the `operator` MIB view.
- If you do not specify `restricted` or `unrestricted`, the switch automatically assigns the community to `restricted` (read-only) access.

Parameters and options

no

The `no` form uses only the *`community-name`* variable and deletes the named community from the switch.

operator|manager

Optionally assigns an access level.

- At the `operator` level, the community can access all MIB objects except the CONFIG MIB.
- At the `manager` level, the community can access all MIB objects.

restricted|unrestricted

Optionally assigns MIB access type.

- Assigning the `restricted` type allows the community to read MIB variables, but not to set them.
- Assigning the `unrestricted` type allows the community to read and set MIB variables.

snmp-server community

This example adds the following communities and access level/types:

Community	Access Level	Type of Access
red-team	manager (**Access to all MIB objects.**)	unrestricted (**read/write**)
blue-team	operator (**Access to all MIB objects except the CONFIG MIB.**)	restricted (**read-only**)

```
switch# snmp-server community red-team
  manager unrestricted
switch# snmp-server community blue-team
  operator restricted
```

no snmp-server community

Eliminates a previously configured community named "gold-team."

```
switch(config) # no snmp-server community gold-team
```

Community names and values

The `snmp-server` command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax

`show snmp-server <COMMUNITY-STRING>`

Description

This command lists the data for currently configured SNMP community names along with trap receivers and the setting for authentication traps.

show snmp-server

Lists the data for all communities in a switch; that is, both the default "public" community name and another community named "blue-team."

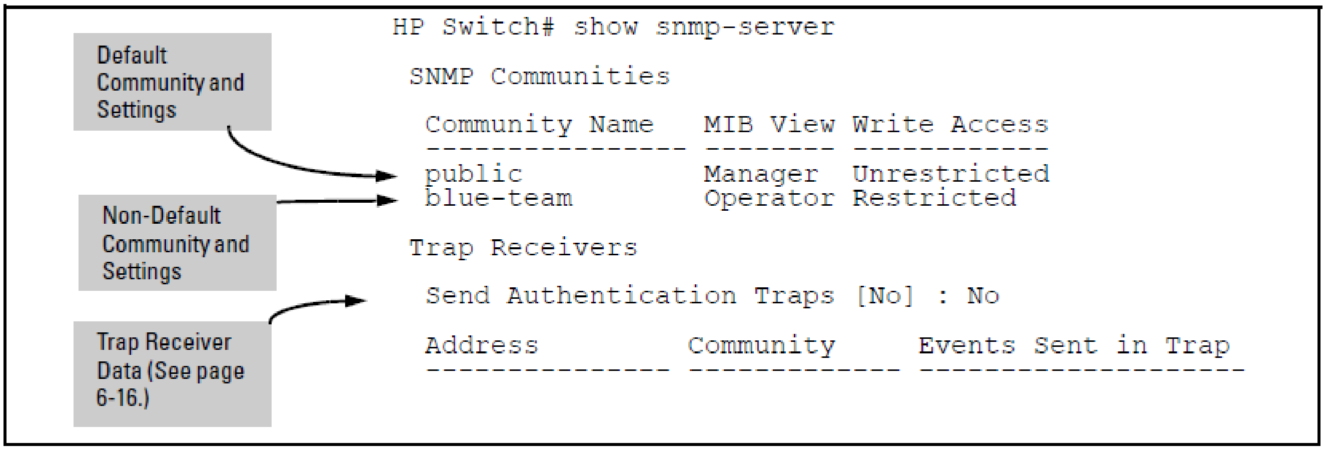

Figure 91: *SNMP community listing with two communities*

show snmp-server public

To list the data for only one community, such as the "public" community, use the above command with the community name included. For example:

```
switch# show snmp-server public
```

Notification/traps for network security failures and other security events

snmp-server enable traps

Syntax

`[no] snmp-server enable traps [snmp-auth|password-change-mgr|login-failure-mgr|port-security|auth-server-fail|dhcp-snooping|arp-protect|running-config-change|macsec failure`

Description

Enables or disables sending one of the security notification types listed below to configured trap receivers. (Unless otherwise stated, all of the following notifications are enabled in the default configuration.)

Parameters and options

arp-protect

If ARP packets are received with an invalid source or destination MAC address, an invalid IP address, or an invalid IP-to-MAC binding.

auth-server-fail

If the connection with a RADIUS or TACACS+ authentication server fails.

dhcp-snooping

If DHCP packets are received from an untrusted source or if DHCP packets contain an invalid IP-to-MAC binding.

dyn-ip-lockdown

If the switch is out of hardware resources needed to program a dynamic IP lockdown rule

link-change ***<PORT-LIST>***

When the link state on a port changes from up to down, or the reverse.

login-failure-mgr

For a failed login with a manager password.

password-change-mgr

When a manager password is reset.

mac-notify

Globally enables the generation of SNMP trap notifications upon MAC address table changes.

port-security

For a failed authentication attempt through a web, MAC, or 801.X authentication session.

running-config-change

When changes to the running configuration file are made.

snmp-authentication [extended|standard]

For a failed authentication attempt via SNMP. Defaults to extended.

startup-config-change

Sends a trap when changes to the startup configuration file are made.(Defaults to disabled.)

macsec failures

Set the trap for MACsec Connectivity Association (CA) failure. This trap is sent when establishing a MACsec CA fails or when a MACsec CA terminates due to MKA keep-alive timeout.

Show snmp-server traps

To determine the specific cause of a security event, check the Event Log in the console interface to see why a trap was sent.

```
Trap Receivers
Link-Change Traps Enabled on Ports [All] : All

Traps Category                         Current Status
------------------       ----------------
SNMP Authentication                    : Extended
Password change                         : Enabled
Login failures                          : Enabled
Port-Security                            : Enabled
Authorization Server Contact   : Enabled
DHCP-Snooping                            : Enabled
Dynamic ARP Protection                : Enabled
Dynamic IP Lockdown                    : Enabled
Startup Config change                 : Disabled
Running Config Change                 : Disabled
MAC address table changes           : Disabled
MAC Address Count                       : Disabled
```

```
MACsec Failures               : Enabled

Address                 Community              Events    Type     Retry    Timeout
---------- ---------------- -------- ------ ------- -------

Excluded MIBs
Snmp Response Pdu Source-IP Information
Selection Policy    : rfc1517
Trap Pdu Source-IP Information
Selection Policy    : rfc1517
```

Current network security notification configuration

show snmp-server traps

`show snmp-server traps`

Displays the current configuration for network security notifications. The command output is a subset of the information displayed with the `show snmp-server` command.

Figure 92: *show snmp-server traps*example output

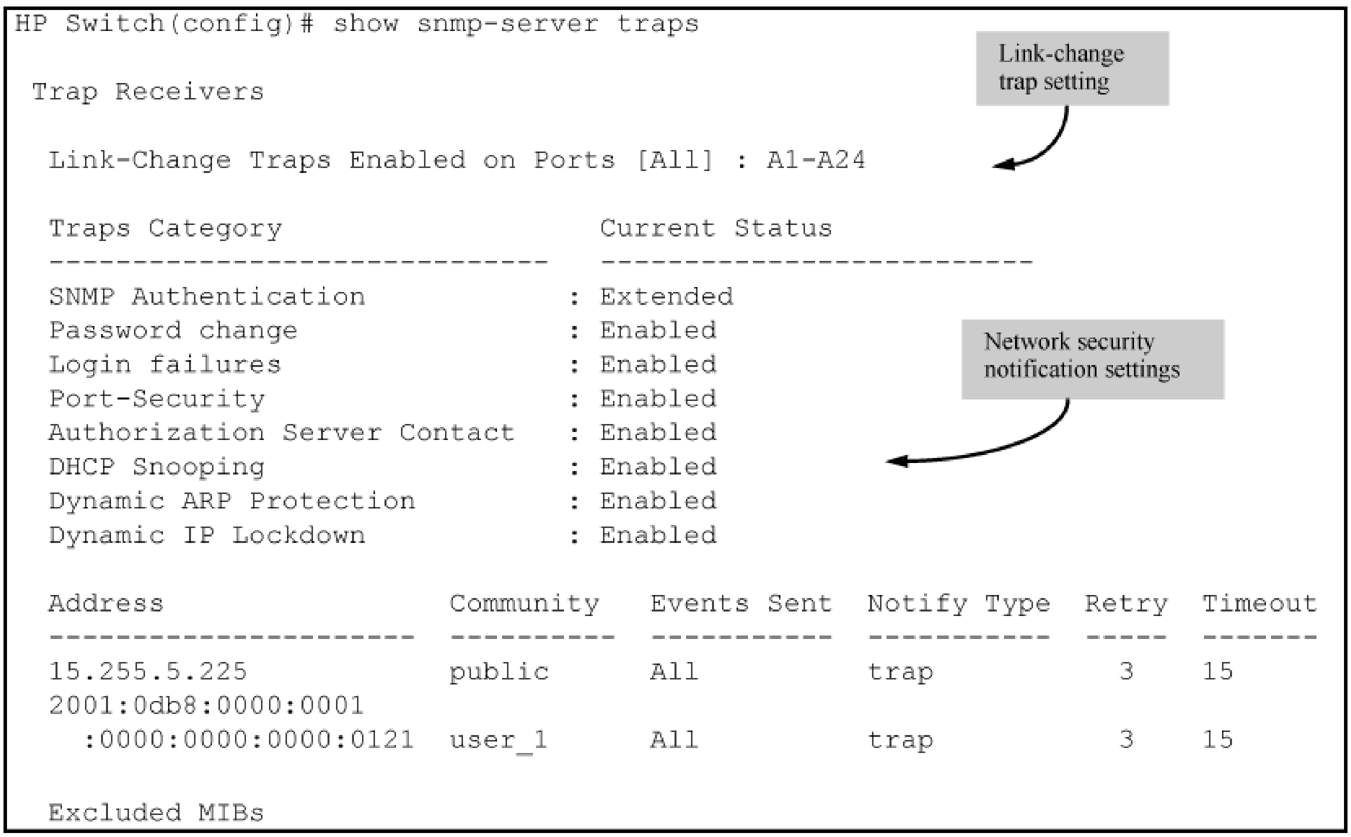


```
HP Switch(config)# show snmp-server traps

 Trap Receivers

  Link-Change Traps Enabled on Ports [All] : A1-A24

  Traps Category                   Current Status
  ------------------------------   --------------------------
  SNMP Authentication            : Extended
  Password change                : Enabled
  Login failures                 : Enabled
  Port-Security                  : Enabled
  Authorization Server Contact   : Enabled
  DHCP Snooping                  : Enabled
  Dynamic ARP Protection         : Enabled
  Dynamic IP Lockdown            : Enabled

  Address                  Community   Events Sent  Notify Type  Retry  Timeout
  ----------------------   ----------  -----------  -----------  -----  -------
  15.255.5.225             public      All          trap            3   15
  2001:0db8:0000:0001
    :0000:0000:0000:0121   user_1      All          trap            3   15

  Excluded MIBs
```

Link-Change Traps

snmp-server enable traps link-change

Syntax

[no] snmp-server enable traps link-change<*PORT-LIST*> [all]

Description

By default, a switch is enabled to send a trap when the link state on a port changes from up to down (linkDown) or down to up (linkUp.) This command allows you to reconfigure the switch to send link-change traps to configured trap receivers.

Parameters and options

`all`

Enables or disables link-change traps on all ports on the switch

Listening mode

snmp-server listen

Syntax

`snmp-server listen [oobm|data|both]`

Description

Enables or disables inbound SNMP access on a switch. The `listen` parameter is not available on switches that do not have a separate out-of-band management port.

Parameters and options

no

Disables inbound SNMP access.

listen

Available only on switches that have a separate out-of-band management port. Defaults to both.

oobm

Inbound SNMP access is enabled only on the out-of-band management port.

data

Inbound SNMP access is enabled only on the data ports.

both

Inbound SNMP access is enabled on both the out-of-band management port and on the data ports.

CDP configuration

CDP mode

cdp moden

Syntax

`[no] cdp moden[pass-through|rxonly]`

Description

Sets the selected mode of CDP processing. Use this command to set the CDP mode to pass-through or receive only.

CDPv2 for voice transmission

Legacy Cisco VOIP phones only support manual configuration or using CDPv2 for voice VLAN auto-configuration. LLDP-MED is not supported. CDPv2 exchanges information such as software version, device capabilities, and voice VLAN information between directly connected devices such as a VOIP phone and a switch.

When the Cisco VOIP phone boots up (or sometimes periodically), it queries the switch and advertises information about itself using CDPv2. The switch receives the VOIP VLAN Query TLV (type 0x0f) from the phone and then immediately sends the voice VLAN ID in a reply packet to the phone using the VLAN Reply TLV (type 0x0e.) The phone then begins tagging all packets with the advertised voice VLAN ID.

Configure a voice VLAN

A voice VLAN must be configured before the voice VLAN can be advertised. For example, to configure VLAN 10 as a voice VLAN tagged for ports 1 through 10, enter these commands:

```
switch# vlan 10

switch(vlan-10)# tagged 1-10

switch(vlan-10)# voice

switch(vlan-10)# exit
```

The switch CDP packet includes these TLVs:

- CDP Version: 2
- CDP TTL: 180 seconds
- Checksum
- Capabilities (type 0x04): 0x0008 (is a switch)
- Native VLAN: The PVID of the port
- VOIP VLAN Reply (type 0xe): voice VLAN ID (same as advertised by LLDPMED)
- Trust Bitmap (type 0x12): 0x00
- Untrusted port COS (type 0x13): 0x00

CDP should be enabled and running on the interfaces to which the phones are connected. Use the `cdp enable` and `cdp run` commands.

The `pre-standard-voice` option for the `cdp mode` command allows the configuration of CDP mode so that it responds to received CDP queries from a VoIP phone.

cdp mode pre-standard-voice

Syntax

```
[no] cdp mode pre-standard-voice [admin-status <PORT-LIST> [tx_rx | rxonly]]
```

Description

Enable CDP-compatible voice VLAN discovery with pre-standard VoIP phones. In this mode, when a CDP VoIP VLAN query is received on a port from pre-standard phones, the switch replies back with a CDP packet that contains the VID of the voice VLAN associated with that port.

Not recommended for phones that support LLDP-MED.

Parameters and options

pre-standard-voice

Enables CDP-compatible voice VLAN discovery with pre-standard VoIP phones.

admin-status

Sets the port in either transmit and receive mode, or receive mode only.

Default: tx-rx.

<PORT-LIST>

Sets this port in transmit and receive mode, or receive mode only.

rxonly

Enable receive-only mode of CDP processing.

tx_rx

Enable transmit and receive mode.

cdp mode pre-standard-voice

```
switch# cdp mode pre-standard-voice admin-status A5 rxonly
```

show cdp without cdp run

Show CDP output when `CDP Run` is disabled.

```
switch(config)# show cdp
Global CDP information
Enable CDP [yes] : no
```

show cdp with cdp run and sdp

`show cdp` output when `cdp run` and `sdp mode` are enabled.

```
switch# show cdp
Global CDP Information
Enable CDP [Yes] : Yes
CDP mode [rxonly] : pre-standard-voice
CDP Hold Time [180] : 180
CDP Transmit Interval [60] : 60
Port CDP admin-status
---- --------- ------------
A1   enabled   rxonly
A2   enabled   tx_rx
A3   enabled   tx_rx
```

show cdp with cdp run and cdp mode rxonly

`show cdp` output when `cdp run` and `cdp mode rxonly` are enabled. When CDP mode is not `pre-standard voice`, the admin-status column is note displayed.

```
switch# show cdp
Global CDP Information
Enable CDP [Yes} : Yes
CDP mode [rxonly] : rxonly
Port CDP
---- --------
A1   enabled
A2   enabled
A3   enabled
```

Example

show running-config

`show running-config` when admin-status is configured.

```
switch# show running-config
Running configuration:
; J9477A Configuration Editor; Created on release #K.16.09.0000x
; Ver #03:01:1f:ef:f2
hostname “HPSwitch”
module 1 type J9307A
cdp mode pre-standard-voice admin-status A5 RxOnly
```

CDP operation on individual ports

In the factory-default configuration, the switch has all ports enabled to receive CDP packets. Disabling CDP on a port causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table.

cdp enable

Syntax

`[no] cdp enable [e] <PORT-LIST>`

Description

Enable or disable ports to receive CDP packets.

Disable CDP on port A1

```
switch# no cdp enable a1
```

CDP Operation

cdp run

Syntax

`[no] cdp run`

Description

Enables or disables CDP read-only operation on the switch. Defaults to enabled.

Parameters and options

no

Disabling CDP operation clears the switch's CDP Neighbors table and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

run

Enabling CDP operation (the default) on the switch causes the switch to add entries to its CDP Neighbors table for any CDP packets it receives from other neighboring CDP devices.

Disable CDP read-only

```
switch# no cdp run
```

When CDP is disabled:

- `show cdp neighbors` displays an empty CDP Neighbors table
- `show cdp` displays

  Global CDP information

  Enable CDP [Yes]: No

CDP information filter

In some environments it is desirable to be able to configure a switch to handle CDP packets by filtering out the MAC address learns from untagged VLAN traffic from IP phones. This means that normal protocol processing occurs for the packets, but the addresses associated with these packets is not learned or reported by the software address management components. This enhancement also filters out the MAC address learns from LLDP and 802.1x EAPOL packets on untagged VLANs.

The feature is configured per-port.

CDP switch configuration view

show cdp

Syntax

`show cdp`

Description

Lists the global and per-port CDP configuration of the switch. CDP is shown as enabled/disabled both globally on the switch and on a per-port basis.

Show CDP with the default CDP configuration

This example shows the default CDP configuration.

```
switch# show cdp

 Global CDP information

  Enable CDP [Yes] : Yes (Receive Only)


  Port CDP
  ---- --------
  1    enabled
  2    enabled
  3    enabled
  .       .
  .       .
  .       .
```

CDP neighbors switch table view

show cdp neighors

Syntax

```
show cdp neighbors
```

Description

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device's CDP packet. Devices are listed by the port on which they were detected.

Parameters and options

[e] <PORT-NUM> [detail]

Lists the CDP device connected to the specified port (allows only one port at a time). Using `detail` provides a longer list of details on the CDP device the switch detects on the specified port.

[detail [e] <PORT-NUM>]

Provides a list of the details for all of the CDP devices the switch detects. Using `port-num` produces a list of details for the selected port.

CDP neighbors table listing

This example displays the CDP devices that the switch has detected by receiving their CDP packets.

```
switch# show cdp neighbors

 CDP neighbors information

  Port Device ID                     | Platform                     Capability
  ---- ----------------------------- + ---------------------------- -----------
  1    Accounting (0030c1-7fcc40)    | J4812A HP Switch. . .        S
  2    Resear¢1-1 (0060b0-889e43)    | J4121A HP Switch. . .        S
  4    Support (0060b0_761a45)       | J4121A HP Switch. . .        S
  7    Marketing (0030c5_33dc59)     | J4313A HP Switch. . .        S
  12   Mgmt NIC(099a05-09df9b        | NIC Model X666               H
  12   Mgmt NIC(099a05-09df11        | NIC Model X666               H
```

LLDP configuration

LLDP and CDP data management

This section describes points to note regarding LLDP and CDP (Cisco Discovery Protocol) data received by the switch from other devices. LLDP operation includes both transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices. CDP operation is limited to reading incoming CDP packets from neighbor devices. (switches do not generate CDP packets.)

Incoming CDP and LLDP packets tagged for VLAN 1 are processed even if VLAN 1 does not contain any ports. VLAN 1 must be present, but it is typically present as the default VLAN for the switch.

The switch may pick up CDP and LLDP multicast packets from VLAN 1 even when CDP- and /or LLDP-enabled ports are not members of VLAN 1.

LLDP and CDP neighbor data

With both LLDP and (read-only) CDP enabled on a switch port, the port can read both LLDP and CDP advertisements, and stores the data from both types of advertisements in its neighbor database. (The switch

stores only CDP data that has a corresponding field in the LLDP neighbor database.) The neighbor database itself can be read by either LLDP or CDP methods or by using the `show lldp` commands. Take note of the following rules and conditions:

- If the switch receives both LLDP and CDP advertisements on the same port from the same neighbor, the switch stores this information as two separate entries if the advertisements have different chassis ID and port ID information.
- If the chassis and port ID information are the same, the switch stores this information as a single entry. That is, LLDP data overwrites the corresponding CDP data in the neighbor database if the chassis and port ID information in the LLDP and CDP advertisements received from the same device is the same.
- Data read from a CDP packet does not support some LLDP fields, such as "System Descr," "SystemCapSupported," and "ChassisType." For such fields, LLDP assigns relevant default values. Also:
  - The LLDP "System Descr" field maps to CDP's "Version" and "Platform" fields.
  - The switch assigns "ChassisType" and "PortType" fields as "local" for both the LLDP and the CDP advertisements it receives.
  - Both LLDP and CDP support the "System Capability" TLV. However, LLDP differentiates between what a device is capable of supporting and what it is actually supporting, and separates the two types of information into subelements of the System Capability TLV. CDP has only a single field for this data. Thus, when CDP System Capability data is mapped to LLDP, the same value appears in both LLDP System Capability fields.
  - System Name and Port Descr are not communicated by CDP, and thus are not included in the switch's Neighbors database.

Because switches do not generate CDP packets, they are not represented in the CDP data collected by any neighbor devices running CDP.

A switch with CDP disabled forwards the CDP packets it receives from other devices, but does not store the CDP information from these packets in its own MIB.

LLDP data transmission/collection and CDP data collection are both enabled in the switch's default configuration. In this state, an SNMP network management application designed to discover devices running either CDP or LLDP can retrieve neighbor information from the switch regardless of whether LLDP or CDP is used to collect the device-specific information.

Protocol state	Packet generation	Inbound data management	Inbound packet forwarding
CDP Enabled	N/A	Store inbound CDP data.	No forwarding of inbound CDP packets.
CDP Disabled	N/A	No storage of CDP data from neighbor devices.	Floods inbound CDP packets from connected devices to outbound ports.
LLDP Enabled1	Generates and transmits LLDP packets out all ports on the switch.	Store inbound LLDP data.	No forwarding of inbound LLDP packets.
LLDP Disabled	No packet generation.	No storage of LLDP data from neighbor devices.	No forwarding of inbound LLDP packets.

CDP operations

By default the switches have CDP enabled on each port. This is a read-only capability, meaning that the switch can receive and store information about adjacent CDP devices but does not generate CDP packets.

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received—and does not forward the packet. The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet and cannot be controlled in the switch receiving the packet.) A switch reviews the list of CDP neighbor entries every three seconds and purges any expired entries.

NOTE

For details on how to use an SNMP utility to retrieve information from the switch's CDP Neighbors table maintained in the switch's MIB, see the documentation provided with the particular SNMP utility.

LLDP

To standardize device discovery on all switches, LLDP will be implemented while offering limited read-only support for CDP, as documented in this manual. For the latest information on your switch model, consult the Release Notes (available on the Networking website.) If LLDP has not yet been implemented (or if you are running an older version of software), consult a previous version of the *Management and Configuration Guide* for device discovery details.

- LLDP (Link Layer Discovery Protocol): Provides a standards-based method for enabling the switches covered in this guide to advertise themselves to adjacent devices and to learn about adjacent LLDP devices.
- LLDP-MED (LLDP Media Endpoint Discovery): Provides an extension to LLDP and is designed to support VoIP deployments.

NOTE

LLDP-MED is an extension for LLDP, and the switch requires that LLDP be enabled as a prerequisite to LLDP-MED operation.

An SNMP utility can progressively discover LLDP devices in a network by:

1. Reading a given device's Neighbors table (in the Management Information Base, or MIB) to learn about other, neighboring LLDP devices.
2. Using the information learned in step 1 to find and read the neighbor devices' Neighbors tables to learn about additional devices, and so on.

Also, by using show commands to access the switch's neighbor database for information collected by an individual switch, system administrators can learn about other devices connected to the switch, including device type (capability) and some configuration information. In VoIP deployments using LLDP-MED on the switches, additional support unique to VoIP applications is also available.

LLDP operations

An LLDP packet contains data about the transmitting switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled and by reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. (LLDP is a one-way protocol and does not include any acknowledgement mechanism.) An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a Neighbor database (MIB.)

LLDP-MED

This capability is an extension to LLDP and is available on the switches. See **LLDP-MED** on page 302.

Packet boundaries in a network topology

- Where multiple LLDP devices are directly connected, an outbound LLDP packet travels only to the next LLDP device. An LLDP-capable device does not forward LLDP packets to any other devices, regardless of whether they are LLDP-enabled.
- An intervening hub or repeater forwards the LLDP packets it receives in the same manner as any other multicast packets it receives. Thus, two LLDP switches joined by a hub or repeater handle LLDP traffic in the same way that they would if directly connected.
- Any intervening 802.1D device or Layer-3 device that is either LLDP-unaware or has disabled LLDP operation drops the packet.

LLDP operation configuration options

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings, which apply to all active ports on the switch, and per-port settings, which affect only the operation of the specified ports.

The commands in the LLDP sections affect both LLDP and LLDP-MED operation.

LLDP on the switch

In the default configuration, LLDP is globally enabled on the switch. To prevent transmission or receipt of LLDP traffic, you can disable LLDP operation.

LLDP-MED

In the default configuration for the switches, LLDP-MED is enabled by default which requires that LLDP is also enabled.

LLDP packet transmissions to neighbor devices

On a global basis, you can increase or decrease the frequency of outbound LLDP advertisements.

Time-To-Live for LLDP packets sent to neighbors

On a global basis, you can increase or decrease the time that the information in an LLDP packet outbound from the switch will be maintained in a neighbor LLDP device.

Transmit and receive mode

With LLDP enabled, the switch periodically transmits an LLDP advertisement (packet) out each active port enabled for outbound LLDP transmissions and receives LLDP advertisements on each active port enabled to receive LLDP traffic (**Configuring per-port transmit and receive modes** on page 304.) Per-port configuration options include four modes:

- Transmit and receive (`tx_rx`): This is the default setting on all ports. It enables a given port to both transmit and receive LLDP packets and to store the data from received (inbound) LLDP packets in the switch's MIB.
- Transmit only (`txonly`): This setting enables a port to transmit LLDP packets that can be read by LLDP neighbors. However, the port drops inbound LLDP packets from LLDP neighbors without reading them. This prevents the switch from learning about LLDP neighbors on that port.
- Receive only (`rxonly`): This setting enables a port to receive and read LLDP packets from LLDP neighbors and to store the packet data in the switch's MIB. However, the port does not transmit outbound LLDP packets. This prevents LLDP neighbors from learning about the switch through that port.
- Disable (`disable`): This setting disables LLDP packet transmissions and reception on a port. In this state, the switch does not use the port for either learning about LLDP neighbors or informing LLDP neighbors of its presence.

SNMP notification

You can enable the switch to send a notification to any configured SNMP trap receiver(s) when the switch detects a remote LLDP data change on an LLDP-enabled port (**SNMP notification support** on page 301.)

Per-port (outbound) data options

The following table lists the information the switch can include in the per-port, outbound LLDP packets it generates. In the default configuration, all outbound LLDP packets include this information in the TLVs transmitted to neighbor devices. However, you can configure LLDP advertisements on a per-port basis to omit some of this information (**Remote management address for outbound LLDP advertisements** on page 304.)

Table 18: *Data available for basic LLDP advertisements*

Data type	Configuration options	Default	Description
Time-to-Live	[1].	120 Seconds	The length of time an LLDP neighbor retains the advertised data before discarding it.
Chassis Type[2]	N/A	Always Enabled	Indicates the type of identifier used for Chassis ID.
Chassis ID[3]	N/A	Always Enabled	Uses base MAC address of the switch.
Port Type[43]	N/A	Always Enabled	Uses "Local," meaning assigned locally by LLDP.
Port Id[3]	N/A	Always Enabled	Uses port number of the physical port. This is an internal number reflecting the reserved slot/port position in the chassis.
Remote Management Address			
Type[3]	N/A	Always Enabled	Shows the network address type.
Address[5]	Default or Configured	Uses a default address selection method unless an optional address is configured.	
System Name[3]	Enable/Disable	Enabled	Uses the switch's assigned name.
System Description[3]	Enable/Disable	Enabled	Includes switch model name and running software version, and ROM version.
Port Description[3]	Enable/Disable	Enabled	Uses the physical port identifier.

Table Continued

Data type	Configuration options	Default	Description
System capabilities supported[3]	Enable/Disable	Enabled	Identifies the switch's primary capabilities (bridge, router.)
System capabilities enabled[6] [3]	Enable/Disable	Enabled	Identifies the primary switch functions that are enabled, such as routing.

[1] The packet time-to-live value is included in LLDP data packets. (See **Changing the time-to-live for transmitted advertisements** on page 309.)

[2] Subelement of the Chassis ID TLV.

[3] Populated with data captured internally by the switch. For more on these data types, refer to the IEEE P802.1AB Standard.

[4] Subelement of the Port ID TLV.

[5] Subelement of the Remote-Management-Address TLV.

[6] Subelement of the System Capability TLV.

Remote management address

The switch always includes an IP address in its LLDP advertisements. This can be either an address selected by a default process or an address configured for inclusion in advertisements.

Debug logging

You can enable LLDP debug logging to a configured debug destination (Syslog server, a terminal device, or both) by executing the `debug lldp` command. Note that the switch's Event Log does not record usual LLDP update messages.

Options for reading LLDP information collected by the switch

You can extract LLDP information from the switch to identify adjacent LLDP devices. Options include:

- Using the switch's `show lldp info` command options to display data collected on adjacent LLDP devices—as well as the local data the switch is transmitting to adjacent LLDP devices (**Global LLDP, port admin, and SNMP notification status** on page 315.)
- Using an SNMP application that is designed to query the Neighbors MIB for LLDP data to use in device discovery and topology mapping.
- Using the `walkmib` command to display a listing of the LLDP MIB objects

LLDP and LLDP-MED standards compatibility

The operation covered by this section is compatible with these standards:

- IEEE P802.1AB
- RFC 2922 (PTOPO, or Physical Topology MIB)
- RFC 2737 (Entity MIB)
- RFC 2863 (Interfaces MIB)
- ANSI/TIA-1057/D6 (LLDP-MED; refer to **LLDP-MED** on page 302.)

Port trunking

LLDP manages trunked ports individually. That is, trunked ports are configured individually for LLDP operation, in the same manner as non-trunked ports. Also, LLDP sends separate advertisements on each port in a trunk, and not on a per-trunk basis. Similarly, LLDP data received through trunked ports is stored individually, per-port.

IP address advertisements

In the default operation, the port uses the following preference order for IP advertisement:

IPv4:

1. IPV4 address of the Management VLAN
2. IPV4 address of first VLAN that the port is a member of (including tagged and untagged memberships)
3. IPV4 address configured on the primary VLAN
4. IPV4 address configured anywhere else in the device

IPv6:

1. Global unicast address (GUC), site local address (SL), and unique local address (UL) of the Management VLAN
2. GUC, SL, and UL of first VLAN that the port is a member of (including tagged and untagged memberships)
3. GUC, SL, and UL configured on the primary VLAN
4. GUC, SL, and UL configured anywhere else in the device
5. LinkLocal Address of a VLAN that the port is a member of

If no IP address is configured on any of the above, the port's MAC address will be advertised.

NOTE

This order is always overridden by the user configured lldp mgmt address on the port.

You can override the default operation by configuring the port to advertise any IP address that is manually configured on the switch, even if the port does not belong to the VLAN configured with the selected IP address (**Remote management address for outbound LLDP advertisements** on page 304.) (Note that LLDP cannot be configured through the CLI to advertise an addresses acquired through DHCP or Bootp. However, as mentioned above, in the default LLDP configuration, if the lowest-order IP address on the VLAN with the lowest VID for a given port is a DHCP or Bootp address, the switch includes this address in its LLDP advertisements unless another address is configured for advertisements on that port.) Also, although LLDP allows configuring multiple remote management addresses on a port, only the lowest-order address configured on the port will be included in outbound advertisements. Attempting to use the CLI to configure LLDP with an IP address that is either not configured on a VLAN or has been acquired by DHCP or Bootp results in the following error message.

```
xxx.xxx.xxx.xxx: This IP address is not configured or is a DHCP address.
```

Spanning-tree blocking

Spanning tree does not prevent LLDP packet transmission or receipt on STP-blocked links.

802.1X blocking

Ports blocked by 802.1X operation do not allow transmission or receipt of LLDP packets.

LLDP operation on the switch

Enabling LLDP operation (the default) causes the switch to:

- Use active, LLDP-enabled ports to transmit LLDP packets describing itself to neighbor devices.
- Add entries to its neighbors table based on data read from incoming LLDP advertisements.

Time-to-Live for transmitted advertisements

The Time-to-Live value (in seconds) for all LLDP advertisements transmitted from a switch is controlled by the switch that generates the advertisement and determines how long an LLDP neighbor retains the advertised data before discarding it. The Time-to-Live value is the result of multiplying the `refresh-interval` by the `holdtime-multiplier`.

Delay interval between advertisements

The switch uses a delay-interval setting to delay transmitting successive advertisements resulting from these LLDP MIB changes. If a switch is subject to frequent changes to its LLDP MIB, lengthening this interval can reduce the frequency of successive advertisements. You can change the delay-interval by using either an SNMP network management application or the CLI `setmib` command.

Re-initialize delay interval

In the default configuration, a port receiving a `disable` command followed immediately by a `txonly`, `rxonly`, or `tx_rx` command delays re-initializing for two seconds, during which LLDP operation remains disabled. If an active port is subjected to frequent toggling between the LLDP disabled and enabled states, LLDP advertisements are more frequently transmitted to the neighbor device. Also, the neighbor table in the adjacent device changes more frequently as it deletes, then replaces LLDP data for the affected port which, in turn, generates SNMP traps (if trap receivers and SNMP notification are configured.) All of this can unnecessarily increase network traffic. Extending the re-initialization-delay interval delays the ability of the port to re-initialize and generate LLDP traffic following an LLDP disable/enable cycle.

SNMP notification support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices and control the interval between successive notifications of data changes on the same neighbor.

Changing the minimum interval

If LLDP trap notification is enabled on a port, a rapid succession of changes in LLDP information received in advertisements from one or more neighbors can generate a high number of traps. To reduce this effect, you can globally change the interval between successive notifications of neighbor data change.

Basic LLDP per-port advertisement content

In the default LLDP configuration, outbound advertisements from each port on the switch include both mandatory and optional data.

Mandatory Data

An active LLDP port on the switch always includes the mandatory data in its outbound advertisements. LLDP collects the mandatory data, and, except for the Remote Management Address, you cannot use LLDP commands to configure the actual data.

- Chassis Type (TLV subelement)
- Chassis ID (TLV)
- Port Type (TLV subelement)
- Port ID (TLV)
- Remote Management Address (TLV; actual IP address is a subelement that can be a default address or a configured address)

Optional Data

You can configure an individual port or group of ports to exclude one or more of the following data types from outbound LLDP advertisements.

- Port description (TLV)
- System name (TLV)
- System description (TLV)
- System capabilities (TLV)

    - System capabilities Supported (TLV subelement)
    - System capabilities Enabled (TLV subelement)
- Port speed and duplex (TLV subelement)

Optional data types, when enabled, are populated with data internal to the switch; that is, you cannot use LLDP commands to configure their actual content.

Support for port speed and duplex advertisements

This feature is optional for LLDP operation, but is **required** for LLDP-MED operation.

Port speed and duplex advertisements are supported on the switches to inform an LLDP endpoint and the switch port of each other's port speed and duplex configuration and capabilities. Configuration mismatches between a switch port and an LLDP endpoint can result in excessive collisions and voice quality degradation. LLDP enables discovery of such mismatches by supporting SNMP access to the switch MIB for comparing the current switch port and endpoint settings. (Changing a current device configuration to eliminate a mismatch requires intervention by the system operator.)

An SNMP network management application can be used to compare the port speed and duplex data configured in the switch and advertised by the LLDP endpoint. You can also use the CLI to display this information.

Port VLAN ID TLV support on LLDP

The `port-vlan-id` option enables advertisement of the port VLAN ID TLV as part of the regularly advertised TLVs. This allows discovery of a mismatch in the configured native VLAN ID between LLDP peers. The information is visible using `show` commands and is logged to the Syslog server.

SNMP support

The LLDP-EXT-DOT1-MIB has the corresponding MIB variables for the Port VLAN ID TLV. The TLV advertisement can be enabled or disabled using the MIB object `lldpXdot1ConfigPortVlanTxEnable` in the `lldpXdot1ConfigPortVlanTable`.

The port VLAN ID TLV local information can be obtained from the MIB object `lldpXdot1LocPortVlanId` in the local information table `lldpXdot1LocTable`.

The port VLAN ID TLV information about all the connected peer devices can be obtained from the MIB object `lldpXdot1RemPortVlanId` in the remote information table `lldpXdot1RemTable`.

LLDP-MED

LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED in the switches uses the standard LLDP commands described earlier in this section, with some extensions, and also introduces new commands unique to LLDP-MED operation. The `show` commands described elsewhere in this section are applicable to both LLDP and LLDP-MED operation. LLDP-MED benefits include:

- Plug-and-play provisioning for MED-capable, VoIP endpoint devices
- Simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network
- Automatic deployment of convergence network policies (voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority)
- Configurable endpoint location data to support the Emergency Call Service (ECS) (such as Enhanced 911 service, 999, 112)
- Detailed VoIP endpoint data inventory readable via SNMP from the switch
- Power over Ethernet (PoE) status and troubleshooting support via SNMP
- support for IP telephony network troubleshooting of call quality issues via SNMP

This section describes how to configure and use LLDP-MED features in the switches to support VoIP network edge devices (media endpoint devices) such as:

- IP phones
- Voice/media gateways
- Media servers
- IP communications controllers
- Other VoIP devices or servers

LLDP-MED interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- Auto-negotiate speed and duplex configuration with the switch
- Use the following network policy elements configured on the client port
- Voice VLAN ID
- 802.1p (Layer 2) QoS
- Diffserv codepoint (DSCP) (Layer 3) QoS
- Discover and advertise device location data learned from the switch
- Support ECS (such as E911, 999, and 112)
- Advertise device information for the device data inventory collected by the switch, including:
  - Hardware revision
  - Firmware revision
  - Software revision
  - Serial number
  - Manufacturer name
  - Model name
  - Asset ID
- Provide information on network connectivity capabilities (for example, a multi-port VoIP phone with Layer 2 switch capability)
- Support the fast-start capability

LLDP-MED is intended for use with VoIP endpoints and is not designed to support links between network infrastructure devices, such as switch-to-switch or switch-to-router links.

LLDP-MED network elements

Figure 93: *LLDP-MED network elements*

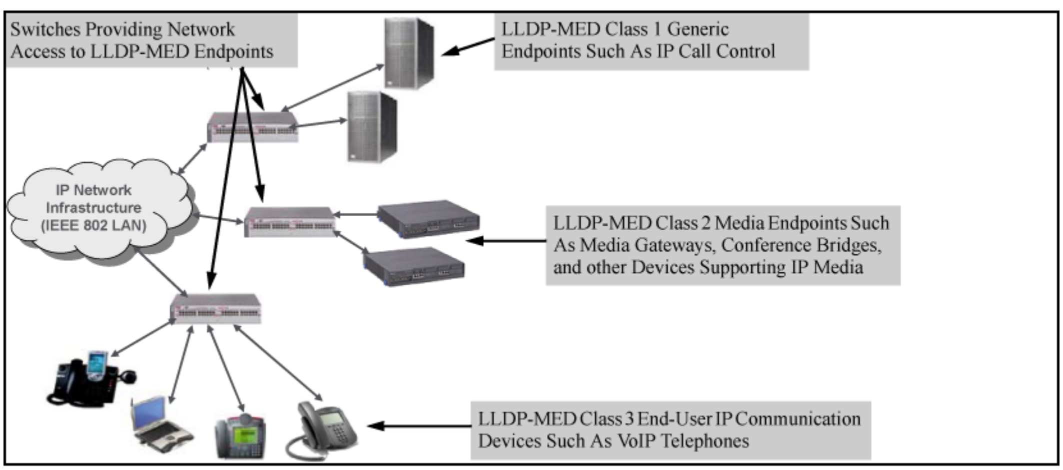

LLDP-MED classes

LLDP-MED endpoint devices are, by definition, located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (generic endpoint devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (media endpoint devices): These devices offer all Class 1 features plus media-streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.
- Class 3 (communication devices): These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

LLDP-MED operational support

The switches offer two configurable TLVs supporting MED-specific capabilities:

- medTlvEnable (for per-port enabling or disabling of LLDP-MED operation)
- medPortLocation (for configuring per-port location or emergency call data)

LLDP-MED operation also requires the port speed and duplex TLV (dot3TlvEnable; page 14-41), which is enabled in the default configuration.

Topology change notifications provide one method for monitoring system activity. However, because SNMP normally employs UDP, which does not guarantee datagram delivery, topology change notification should not be relied upon as the sole method for monitoring critical endpoint device connectivity.

Configuring per-port transmit and receive modes

lldp admin-status

Syntax

```
lldp admin-status <PORT-LIST> [txonly|rxonly|tx_rx|disable]
```

Description

With LLDP enabled on the switch in the default configuration, each port is configured to transmit and receive LLDP packets. The options allow you to control which ports participate in LLDP traffic and whether the participating ports allow LLDP traffic in only one direction or in both directions. Defaults to tx_rx.

Parameters and options

txonly

Configures the specified ports to transmit LLDP packets, but block inbound LLDP packets from neighbor devices.

rxonly

Configures the specified ports to receive LLDP packets from neighbors, but block outbound packets to neighbors.

tx_rx

Configures the specified ports to both transmit and receive LLDP packets. (This is the default setting.)

disable

Disables LLDP packet transmit and receive on the specified ports.

Remote management address for outbound LLDP advertisements

lldp config ipAddrEnable

Syntax

```
[no] lldp config  <PORT-LIST>  ipAddrEnable  ip-address
```

Description

This is an optional command you can use to include a specific IP address in the outbound LLDP advertisements for specific ports. Replaces the default IP address for the port with an IP address you specify. This can be any IP address configured in a static VLAN on the switch, even if the port does not belong to the VLAN configured with the selected IP address.

Default: The port advertises the IP address of the lowest-numbered VLAN (VID) to which it belongs. If there is no IP address configured on the VLANs to which the port belongs, and if the port is not configured to advertise an IP address from any other (static) VLAN on the switch, the port advertises an address of 127.0.0.1.)

This command does not accept either IP addresses acquired through DHCP or Bootp, or IP addresses that are not configured in a static VLAN on the switch.

Parameters and options

no

Deletes the specified IP address. If there are no IP addresses configured as management addresses, the IP address selection method returns to the default operation.

lldp config

If port 3 belongs to a subnetted VLAN that includes an IP address of 10.10.10.100 and you want port 3 to use this secondary address in LLDP advertisements, you need to execute the following command:

```
switch# lldp config 3 ipAddrEnable 10.10.10.100
```

lldp config basicTlvEnable

Syntax

```
lldp config  <PORT-LIST>  basicTlvEnable  <TLV-Type>
```

Description

Enable basic TLV for LLDP by type and *<PORT-LIST>* .

Parameters and options

<PORT_DESC>

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the port. Defaults to enabled.

<SYSTEM_NAME>

For outbound LLDP advertisements, this TLV includes an alphanumeric string showing the assigned name of the system. Defaults to enabled.

<SYSTEM_DESCR>

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the full name and version identification for the hardware type, software version, and networking application of the system. Defaults to enabled.

<SYSTEM_CAP>

For outbound advertisements, this TLV includes a bitmask of supported system capabilities (device functions.) Also includes information on whether the capabilities are enabled. Defaults to enabled.

no lldp config

To exclude the system name TLV from the outbound LLDP advertisements for all ports on a switch, use this command:

```
switch# no lldp config 1-24 basicTlvEnable system_name
```

lldp config

To reinstate the system name TLV on ports 1-5, use this command:

```
switch# lldp config 1-5 basicTlvEnable system_name
```

Port speed and duplex advertisement support

lldp config dot3TlvEnable

Syntax

`[no] lldp config <PORT-LIST> dot3TlvEnable macphy_config`

Description

For outbound advertisements, this TLV includes the (local) switch port's current speed and duplex settings, the range of speed and duplex settings the port supports, and the method required for reconfiguring the speed and duplex settings on the device (autonegotiation during link initialization, or manual configuration.)

Using SNMP to compare local and remote information can help in locating configuration mismatches. Defaults to enabled.

For LLDP operation, this TLV is optional. For LLDP-MED operation, this TLV is mandatory.

Location data for LLDP-MED devices

lldp config medPortLocation

Syntax

`[no] lldp config <PORT-LIST> medPortLocation Address-Type`

Description

Configures location of emergency call data the switch advertises per port in the `location_id` TLV. This TLV is for use by LLDP-MED endpoints employing location-based applications. Enables configuration of a physical address on a switch port and allows up to 75 characters of address information.

The switch allows one medPortLocation entry per port (without regard to type.) Configuring a new medPortLocation entry of any type on a port replaces any previously configured entry on that port.

Parameters and options

COUNTRY-STR

A two-character country code, as defined by ISO 3166. Some examples include `FR` (France), `DE` (Germany), and `IN` (India.) This field is required in a `civic-addr` command. (For a complete list of country codes, see **http://www.iso.org**.)

WHAT

A single-digit number specifying the type of device to which the location data applies: 0: Location of DHCP server 1: Location of switch 2: Location of LLDP-MED endpoint (recommended application) This field is required in a `civic-addr` command.

Type/Value Pairs [*CA-TYPE*|*CA-VALUE*]

A series of data pairs, each composed of a location data "type" specifier and the corresponding location data for that type. That is, the first value in a pair is expected to be the civic address "type" number (`CA-TYPE`), and the second value in a pair is expected to be the corresponding civic address data (`CA-VALUE`.)

For example, if the `CA-TYPE` for "city name" is "3," the type/value pair to define the city of Paris is "3 Paris."

Multiple type/value pairs can be entered in any order, although Hewlett Packard Enterprise recommends that multiple pairs be entered in ascending order of the `CA-TYPE`.

When an emergency call is placed from a properly configured class 3 endpoint device to an appropriate PSAP, the country code, device type, and type/value pairs configured on the switch port are included in the transmission. The "type" specifiers are used by the PSAP to identify and organize the location data components in an understandable format for response personnel to interpret.

A `civic-addr` command requires a minimum of one type/value pair, but typically includes multiple type/value pairs as needed to configure a complete set of data describing a given location.

`CA-TYPE:` This is the first entry in a type/value pair and is a number defining the type of data contained in the second entry in the type/value pair (`CA-VALUE`.) Some examples of `CA-TYPE` specifiers include:

- 3=city
- 6=street (name)
- 25=building name

(Range: 0 - 255)

`CA-VALUE:` This is the second entry in a type/value pair and is an alphanumeric string containing the location information corresponding to the immediately preceding `CA-TYPE` entry.

Strings are delimited by either blank spaces, single quotes (' … '), or double quotes ("… ".)

Each string should represent a specific data type in a set of unique type/value pairs comprising the description of a location, and each string must be preceded by a `CA-TYPE` number identifying the type of data in the string.

NOTE

A switch port allows one instance of any given `CA-TYPE`. For example, if a type/value pair of 6 Atlantic (to specify "Atlantic" as a street name) is configured on port A5 and later another type/value pair of 6 Pacific is configured on the same port, Pacific replaces Atlantic in the civic address location configured for port A5.

elin-addr emergency-number

This feature is intended for use in ECS applications to support class 3 LLDP-MED VoIP telephones connected to a switch in an MLTS infrastructure.

An ELIN is a valid NANP format telephone number assigned to MLTS operators in North America by the appropriate authority. The ELIN is used to route emergency (E911) calls to a PSAP.

(Range: 1-15 numeric characters)

Usage

`civic-addr <COUNTRY-STR> <WHAT> <CA-TYPE> <CA-VALUE> … <CA-TYPE> <CA-VALUE> … <CA-TYPE> <CA-VALUE>`

LLDP data change notification for SNMP trap receivers

lldp enable-notification

Syntax

`[no] lldp enable-notification <PORT-LIST>`

Description

Enables or disables each port in *<PORT-LIST>* for sending notification to configured SNMP trap receivers if an LLDP data change is detected in an advertisement received on the port from an LLDP neighbor. Defaults to disabled.

Enable SNMP notification on ports 1 - 5

```
switch# lldp enable-notification 1-5
```

LLDP operation on the switch

lldp run

Syntax

`[no] lldp run`

Description

Enables or disables LLDP operation on the switch.

Parameters and options

no

Regardless of individual LLDP port configurations, prevents the switch from transmitting outbound LLDP advertisements and causes the switch to drop all LLDP advertisements received from other devices.

The switch preserves the current LLDP configuration when LLDP is disabled. After LLDP is disabled, the information in the LLDP neighbors database remains until it times-out. Defaults to enabled.

Disable lldp on the switch

```
switch# no lldp run
```

LLDP-MED fast start control

lldp fast-start-count

Syntax

```
lldp fast-start-count  <1 - 10>
```

Description

An LLDP-MED device connecting to a switch port may use the data contained in the MED TLVs from the switch to configure itself. However, the `lldp refresh-interval` setting (default: 30 seconds) for transmitting advertisements can cause an unacceptable delay in MED device configuration.

To support rapid LLDP-MED device configuration, the `lldp fast-start-count` command temporarily overrides the `refresh-interval` setting for the `fast-start-count` advertisement interval. This results in the port initially advertising LLDP-MED at a faster rate for a limited time. Thus, when the switch detects a new LLDP-MED device on a port, it transmits one LLDP-MED advertisement per second out the port for the duration of the `fast-start-count` interval. In most cases, the default setting should provide an adequate `fast-start-count` interval. Defaults to 5 seconds.

This global command applies only to ports on which a new LLDP-MED device is detected. It does not override the `refresh-interval` setting on ports where non-MED devices are detected.

Changing the packet transmission interval

This interval controls how often active ports retransmit advertisements to their neighbors.

lldp refresh-interval

Syntax

```
lldp refresh-interval <5 - 32768>
```

Description

Changes the interval between consecutive transmissions of LLDP advertisements on any given port. Defaults to 30 seconds.

The `refresh-interval` must be greater than or equal to (4 x `delay-interval`.) (The default `delay-interval` is 2.) For example, with the default `delay-interval`, the lowest `refresh-interval`you can use is 8 seconds (4 x 2=8.) Thus, if you want a `refresh-interval` of 5 seconds, you must first change the delay interval to 1 (that is, 4 x 1 5.) If you want to change the `delay-interval`, use the `setmib` command.

Changing the time-to-live for transmitted advertisements

lldp holdtime-multiplier

Syntax

```
lldp holdtime-multiplier <2 - 10>
```

Description

Changes the multiplier an LLDP switch uses to calculate the Time-to-Live for the LLDP advertisements it generates and transmits to LLDP neighbors. When the Time-to-Live for a given advertisement expires, the advertised data is deleted from the neighbor switch's MIB. Defaults to 4.

If the refresh-interval on the switch is 15 seconds and the `holdtime-multiplier` is at the default, the Time-to-Live for advertisements transmitted from the switch is 60 seconds (4 x 15.)

Reduce time-to-live

To reduce the Time-to-Live, you could lower the `holdtime-interval` to 2, which would result in a Time-to-Live of 30 seconds.

```
switch# lldp holdtime-multiplier 2
```

Delay interval

To change the delay interval between advertisements generated by value or status changes to the LLDP MIB, use the following command.

set mib lldpTxDelay.0

Syntax

setmib lldpTxDelay.0 -i *<1 - 8192>*

Uses `setmib` to change the minimum time (delay-interval) any LLDP port will delay advertising successive LLDP advertisements because of a change in LLDP MIB content. Defaults to 2.

The LLDP refresh-interval (transmit interval) must be greater than or equal to (4 x delay-interval.) The switch does not allow increasing the delay interval to a value that conflicts with this relationship. That is, the switch displays `Inconsistent value` if (4 x delay-interval) exceeds the current transmit interval, and the command fails. Depending on the current refresh-interval setting, it may be necessary to increase the refresh-interval before using this command to increase the delay-interval.

For the 3800 switches, when the switch is in enhanced secure mode, the following prompt appears before the sensitive information for the setmib command is displayed:

```
The setmib command should not be used in enhanced secure mode.
```

For more information, see the access security guide.

Change the delay-interval

To change the delay-interval from 2 seconds to 8 seconds when the refresh-interval is at the default 30 seconds, you must first set the refresh-interval to a minimum of 32 seconds (32 = 4 x 8.) (See **Figure 94: Changing the transmit-delay interval** on page 310.)

Figure 94: *Changing the transmit-delay interval*

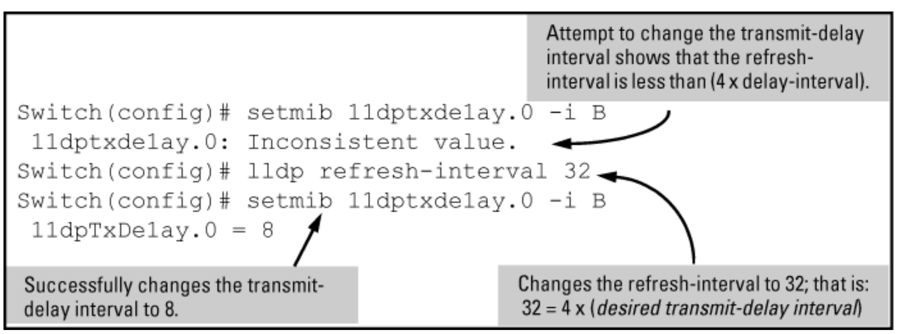

Changing the reinitialization delay interval

setmib lldpReinitDelay.0

Syntax

```
setmib lldpReinitDelay.0 -i <1-10>
```

Uses `setmib` to change the minimum time (reinitialization delay interval) an LLDP port will wait before reinitializing after receiving an LLDP disable command followed closely by a txonly or tx_rx command. The delay interval commences with execution of the `lldp admin-status <PORT-LIST> disable` command. Defaults to 2.

Change the reinitialization delay interval

The following command changes the reinitialization delay interval to five seconds:

```
switch# setmib lldpreinitdelay.0 -i 5
```

PVID mismatch log messages

PVID mismatches are logged when there is a difference in the PVID advertised by a neighboring switch and the PVID of the switch port which receives the LLDP advertisement. Logging is an LLDP feature that allows detection of possible vlan leakage between adjacent switches. However, if these events are logged too frequently, they can overwhelm the log buffer and push relevant logging data out of log memory, making it difficult to troubleshoot another issue.

Use the following command to enable or disable the logging of the PVID mismatch log messages:

logging filter

Syntax

```
logging filter [<filter-name> enable] [<filter-name><sub filter
id><regularexpression> deny]
```

Description

Filters out PVID mismatch log messages on a per-port basis, allowing you to disable or enable logging using the CLI. This includes displaying the Mac-Address in the PVID mismatch log message when the port ID is Mac-Address instead of displaying garbage characters in the peer device port ID field.

Parameters and options

Regular-expression

The regular expression should match the message which is to be filtered.

Viewing port configuration details

show lldp config

Syntax

```
show lldp config <PORT-LIST>
```

Description

Displays the LLDP port-specific configuration for all ports in *<PORT-LIST>*, including which optional TLVs and any non-default IP address that are included in the port's outbound advertisements.

show lldp config

Figure 95: *Per-port configuration display*

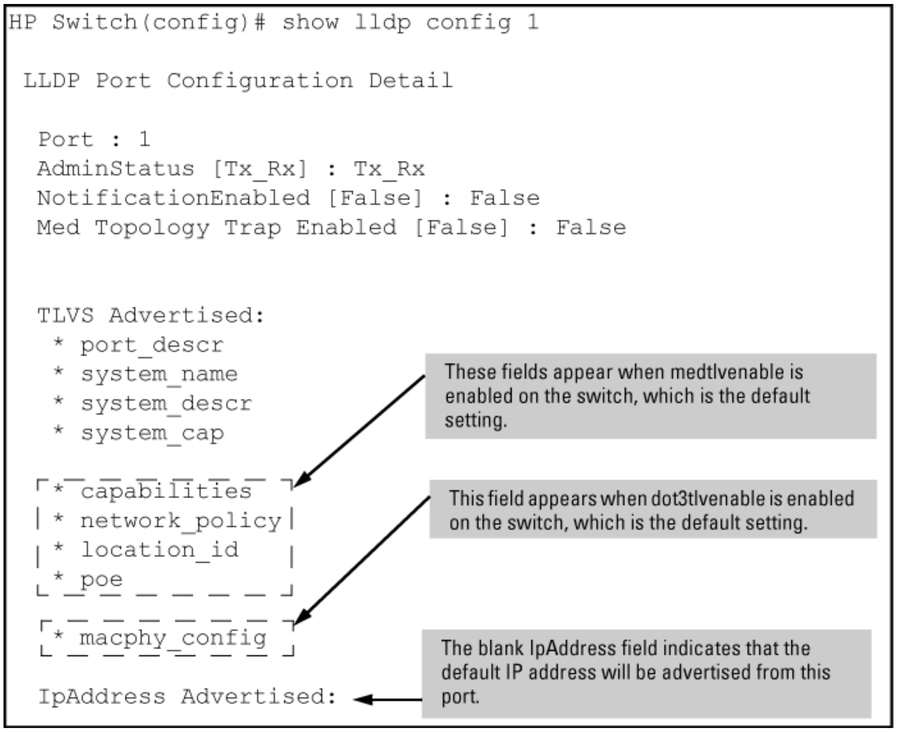

Available switch information available outbound advertisements

show lldp info local-device

Syntax

`show lldp info local-device<PORT-LIST>`

Description

Displays global switch information and per-port information currently available for populating outbound LLDP advertisements. This command displays the information available on the switch. Use the `lldp config <PORT-LIST>` command to change the selection of information that is included in actual outbound advertisements. In the default LLDP configuration, all information displayed by this command is transmitted in outbound advertisements.

Parameters and options

<PORT-LIST>

Without the *<PORT-LIST>* option, displays the global switch information and the per-port information currently available for populating outbound LLDP advertisements.

With the *<PORT-LIST>* option, displays only the following port-specific information that is currently available for outbound LLDP advertisements on the specified ports:

- `PortType`
- `PortId`
- `PortDesc`

show lldp info local-device output

In the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in **Figure 96: Displaying the global and per-port information available for outbound advertisements** on page 313.

Figure 96: *Displaying the global and per-port information available for outbound advertisements*

```
HP Switch(config)# show lldp info local-device

 LLDP Local Device Information

  Chassis Type : mac-address
  Chassis Id   : 00 23 47 4b 68 00
  System Name  : HP Switch1
  System Description : HP J9091A Switch 3500yl, revision K.15.06...
  System Capabilities Supported:bridge
  System Capabilities Enabled:bridge

  Management Address  :
     Type:ipv4
     Address:

 LLDP Port Information

  Port     | PortType PortId   PortDesc
  -------- + -------- -------- --------
  1        | local    1        1
  2        | local    2        2
  3        | local    3        3
  4        | local    4        4
  5        | local    5        5
```

The Management Address field displays only the LLDP-configurable IP addresses on the switch. (Only manually-configured IP addresses are LLDP-configurable.) If the switch has only an IP address from a DHCP or Bootp server, then the Management Address field is empty (because there are no LLDP-configurable IP addresses available). For more on this topic, refer to "Remote Management Address" on page 6-52.

Default per-port information content for ports 1 and 2

```
switch# show lldp info local 1-2

  LLDP Local Port Information Detail

   Port     : 1
   PortType : local
   PortId   : 1
   PortDesc : 1

----------------------------------------
   Port     : 2
   PortType : local
   PortId   : 2
   PortDesc : 2
```

LLDP statistics

show lldp stats

Syntax

`show lldp stats<PORT-LIST>`

Description

Displays (globally) an overview of neighbor detection activity on the switch, plus data on the number of frames sent, received, and discarded per-port. The **per-port LLDP** statistics command enhances the list of per-port statistics provided by the global statistics command with some additional per-port LLDP statistics.

Parameters and options

Neighbor Entries List Last Updated

The elapsed time since a neighbor was last added or deleted.

New Neighbor Entries Count

The total of new LLDP neighbors detected since the last switch reboot. Disconnecting, and then reconnecting a neighbor increments this counter.

Neighbor Entries Deleted Count

The number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports.

For example, if the admin status for port on a neighbor device changes from `tx_rx` or `txonly` to `disabled` or `rxonly`, the neighbor device sends a "shutdown" packet out the port and ceases transmitting LLDP frames out that port.

The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter.

This can occur, for example, when a new neighbor is detected when the switch is already supporting the maximum number of neighbors. See **Neighbor maximum** on page 325.

Neighbor Entries Dropped Count

The number of valid LLDP neighbors the switch detected, but could not add.

Neighbor Entries AgeOut Count

The number of LLDP neighbors dropped on all ports because of Time-to-Live expiring.

NumFramesRecvd

The total number of valid, inbound LLDP advertisements received from any neighbors on *<PORT-LIST>* .

Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources.

NumFramesSent

The total number of LLDP advertisements sent from *<PORT-LIST>* .

NumFramesDiscarded

The total number of inbound LLDP advertisements discarded by *<PORT-LIST>* .

This can occur, for example, when a new neighbor is detected on the port, but the switch is already supporting the maximum number of neighbors. See **Neighbor maximum** on page 325. This can also be an indication of advertisement formatting problems in the neighbor device.

Frames Invalid

The total number of invalid LLDP advertisements received on the port.

An invalid advertisement can be caused by header formatting problems in the neighbor device.

TLVs Unrecognized

The total number of LLDP TLVs received on a port with a type value in the reserved range.

This can be caused by a basic management TLV from a later LLDP version than the one currently running on the switch.

TLVs Discarded

The total number of LLDP TLVs discarded for any reason. In this case, the advertisement carrying the TLV may be accepted, but the individual TLV is not usable.

Neighbor Ageouts

The number of LLDP neighbors dropped on the port because of Time-to-Live expiring.

A global LLDP statistics display

```
switch# show lldp stats

 LLDP Device Statistics

  Neighbor Entries List Last Updated : 2 hours
  New Neighbor Entries Count : 20
  Neighbor Entries Deleted Count : 20
  Neighbor Entries Dropped Count : 0
  Neighbor Entries AgeOut Count : 20

 LLDP Port Statistics

  Port   | NumFramesRecvd NumFramesSent NumFramesDiscarded
  ------ + -------------- ------------- ------------------
  A1     | 97317          97843         0
  A2     | 21             12            0
  A3     | 0              0             0
  A4     | 446            252           0
  A5     | 0              0             0
  A6     | 0              0             0
  A7     | 0              0             0
  A8     | 0              0             0
```

A per-port LLDP statistics display

```
switch# show lldp stats 1

 LLDP Port Statistics Detail

  PortName : 1
  Frames Discarded   : 0
  Frames Invalid     : 0
  Frames Received    : 7309
  Frames Sent        : 7231
  TLVs Unrecognized : 0
  TLVs Discarded     : 0
  Neighbor Ageouts   : 0
```

Global LLDP, port admin, and SNMP notification status

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings that apply to all active ports on the switch, and per-port settings that affect only the operation of the specified ports.

The commands in this section affect both LLDP and LLDP-MED operation.

show lldp config

Syntax

`show lldp config`

Displays the LLDP global configuration, LLDP port status, and SNMP notification status.

View default LLDP

`show lldp config` produces the following display when the switch is in the default LLDP configuration. The values displayed in the LLDP column correspond to the `lldp refresh-interval` command.

```
switch# show lldp config

LLDP Global Configuration

LLDP Enabled [Yes] :              Yes
LLDP Transmit Interval [30] :     30
LLDP Hold time Multiplier [4] :   4
LLDP Delay Interval [2] :         2
LLDP Reinit Interval [2] :        2
LLDP Notification Interval [5] :  5
LLDP Fast Start Count [5] :       5

LLDP Port Configuration
Port | AdminStatus NotificationEnabled  Med Topology Trap Enabled
---- + ----------- -------------------  -------------------------
A1 |    Tx_Rx          False              False
A2 |    Tx_Rx          False              False
A3 |    Tx_Rx          False              False
A4 |    Tx_Rx          False              False
A5 |    Tx_Rx          False              False
A6 |    Tx_Rx          False              False
A7 |    Tx_Rx          False              False
A8 |    Tx_Rx          False              False
```

LLDP-MED connects and disconnects—topology change notification

This optional feature provides information an SNMP application can use to track LLDP-MED connects and disconnects.

lldp top-change-notify

Syntax

`lldp top-change-notify <PORT-LIST>`

Description

Defaults to disabled. When enabled on an LLDP port, topology change notification causes the switch to send an SNMP trap if it detects LLDP-MED endpoint connection or disconnection activity on the port, or an age-out of the LLDP-MED neighbor on the port.

The trap includes the following information:

- The port number (internal) on which the activity was detected.
- The LLDP-MED class of the device detected on the port.

To send traps, this feature requires access to at least one SNMP server.

If a detected LLDP-MED neighbor begins sending advertisements without LLDP-MED TLVs, the switch sends a top-change-notify trap.

View topology change notification status

You can use the `show running` command shows whether the topology change notification feature is enabled or disabled. For example, if ports A1 to A10 have topology change notification enabled, the following entry appears in the `show running` output:

```
lldp top-change-notify A1-A10
```

Device capability, network policy, PoE status and location data

The medTlvEnable option on the switch is enabled in the default configuration and supports the following LLDP-MED TLVs:

- LLDP-MED capabilities: This TLV enables the switch to determine:
    - Whether a connected endpoint device supports LLDP-MED
    - Which specific LLDP-MED TLVs the endpoint supports
    - The device class (1, 2, or 3) for the connected endpoint

  This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.
- Network policy operating on the port to which the endpoint is connected (VLAN, Layer 2 QoS, Layer 3 QoS.)
- PoE (MED Power-over-Ethernet.)
- Physical location data.

LLDP-MED operation requires the macphy_config TLV subelement (enabled by default) that is optional for IEEE 802.1AB LLDP operation. For more information, see the `dot3TlvEnable macphy_config` command.

Network policy advertisements

Network policy advertisements are intended for real-time voice and video applications, and include these TLV sub-elements:

- Layer 2 (802.1p) QoS
- Layer 3 DSCP (diffserv code point) QoS
- Voice VLAN ID (VID)

VLAN operating rules

These rules affect advertisements of VLANs in network policy TLVs:

- The VLAN ID TLV subelement applies only to a VLAN configured for voice operation (`vlan` *`vid`* `voice` .)
- If there are multiple voice VLANs configured on a port, LLDP-MED advertises the voice VLAN having the lowest VID.
- The voice VLAN port membership configured on the switch can be tagged or untagged. However, if the LLDP-MED endpoint expects a tagged membership when the switch port is configured for untagged, or the reverse, a configuration mismatch results. (Typically, the endpoint expects the switch port to have a tagged voice VLAN membership.)
- If a given port does not belong to a voice VLAN, the switch does not advertise the VLAN ID TLV through this port.

Policy elements

These policy elements may be statically configured on the switch or dynamically imposed during an authenticated session on the switch using a RADIUS server and 802.1X or MAC authentication. (Web authentication does not apply to VoIP telephones and other telecommunications devices that are not capable of accessing the switch through a Web browser.) The QoS and voice VLAN policy elements can be statically configured with the following CLI commands:

```
vlan <VID> voice

vlan <VID> [tagged|untagged] <PORT-LIST>

int <PORT-LIST> qos priority 0 - 7

vlan vid qos dscp codepoint
```

A codepoint must have an 802.1p priority before you can configure it for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows `No Override` in the `Priority` column of the DSCP policy table (display with `show qos-dscp map`, then use `qos-dscp map` *`codepoint`* `priority 0 - 7` to configure a priority before proceeding.

For more information on this topic, see the advanced traffic management guide.

PoE advertisements

These advertisements inform an LLDP-MED endpoint of the power (PoE) configuration on switch ports. Similar advertisements from an LLDP-MED endpoint inform the switch of the endpoint's power needs and provide information that can be used to identify power priority mismatches.

PoE TLVs include the following power data:

Power type

Indicates whether the device is a power-sourcing entity (PSE) or a PD. Ports on the J8702A PoE zl module are PSE devices. A MED-capable VoIP telephone is a PD.

Power source

Indicates the source of power in use by the device. Power sources for PDs include PSE, local (internal), and PSE/local. The switches advertise `unknown`.

Power priority

Indicates the power priority configured on the switch (PSE) port or the power priority configured on the MED-capable endpoint.

Power value

Indicates the total power in watts that a switch port (PSE) can deliver at a particular time, or the total power in watts that the MED endpoint (PD) requires to operate.

Location data for LLDP-MED devices

You can configure a switch port to advertise location data for the switch itself, the physical wall-jack location of the endpoint (recommended), or the location of a DHCP server supporting the switch, endpoint, or both. You also have the option of configuring these different address types:

Civic address

Physical address data such as city, street number, and building information.

ELIN (Emergency Location Identification Number)

An emergency number typically assigned to MLTS (Multiline Telephone System) Operators in North America.

Coordinate-based location

Attitude, longitude, and altitude information (Requires configuration via an SNMP application.)

Coordinate-based locations

Latitude, longitude, and altitude data can be configured per switch port using an SNMP management application. For more information, see the documentation provided with the application. A further source of information on this topic is the *RFC 3825: Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information*.

Endpoint use of data from a medPortLocation TLV sent by the switch is device-dependent. See the documentation provided with the endpoint device.

The code assignments in the following table are examples from a work-in-progress (the internet draft titled "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information draft-ietf-geopriv-dhcp-civil-06" dated May 30, 2005.) For the actual codes to use, contact the PSAP or other authority responsible for specifying the civic addressing data standard for your network.

Table 19: *Some location codes used in CA-TYPE fields*

Location element	Code	Location element	Code
national subdivision	1	street number	19
regional subdivision	2	additional location data	22
city or township	3	unit or apartment	26
city subdivision	4	floor	27
street	6	room number	28
street suffix	18		

Example

Suppose a system operator wants to configure the following information as the civic address for a telephone connected to her company's network through port A2 of a switch at the following location:

Description	CA-type	CA-VALUE
national subdivision	1	CA
city	3	Widgitville
street	6	Main
street number	19	1433
unit	26	Suite 4-N
floor	27	4
room number	28	N4-3

Example of a civic address configuration on page 320 shows the commands for configuring and displaying the above data.

Example of a civic address configuration

```
switch# lldp config 2 medportlocation civic-addr US 2 1 CA 3
Widgitville 6 Main 19 1433 26 Suite_4—N 27 4 28 N4—3

switch# show lldp config 2
 LLDP Port Configuration Detail
  Port : A2
  AdminStatus [Tx_Rx] : Tx_Rx
  NotificationEnabled [False] : False
  Med Topology Trap Enabled [False] : False
  Country Name          : US
  What                  : 2
  Ca-Type               : 1
  Ca-Length             : 2
  Ca-Value              : CA
  Ca-Type               : 3
  Ca-Length             : 11
  Ca-Value              : Widgitville
  Ca-Type               : 6
  Ca-Length             : 4
  Ca-Value              : Main
  Ca-Type               : 19
  Ca-Length             : 4
  Ca-Value              : 1433
  Ca-Type               : 26
  Ca-Length             : 9
  Ca-Value              : Suite_4-N
  Ca-Type               : 27
  Ca-Length             : 1
  Ca-Value              : 4
  Ca-Type               : 28
  Ca-Length             : 4
  Ca-Value              : N4-3
```

Viewing the current port speed and duplex configuration

You can compare port speed and duplex information for a switch port and a connected LLDP-MED endpoint for configuration mismatches by using an SNMP application. You can also use the switch CLI to display this information, if necessary. The `show interfaces brief` *<PORT-LIST>* and `show lldp info remote-device`*<PORT-LIST>* commands provide methods for displaying speed and duplex information for switch ports. For information on displaying the currently configured port speed and duplex on an LLDP-MED endpoint.

Viewing LLDP statistics

LLDP statistics are available on both a global and a per-port levels. Rebooting the switch resets the LLDP statistics counters to zero. Disabling the transmit and/or receive capability on a port "freezes" the related port counters at their current values.

LLDP over OOBM

The following commands enable the user to configure LLDP for OOBM ports.

lldp admin-status oobm

Syntax

`lldp admin-status oobm [ txonly | rxonly | tx_rx | disable ]`

Description

This command sets the OOBM port operational mode.

Parameters and options

`txonly`

Sets in transmit only mode.

`rxonly`

Sets in receive mode.

`tx_rx`

Sets in transmit and receive mode.

`disable`

Disables lldp on OOBM port.

lldp enable-notification oobm

Syntax

`[no] lldp enable-notification oobm`

Description

This command enables or disables notification on the OOBM port.

Parameters and options

`oobm`

Enables notification on the OOBM port.

`no`

Disables notification.

Enable-notification

```
switch(config)#lldp enable-notification ?
oobm                  Enable or disable notification on the OOBM port.
[ethernet] PORT-LIST  Enable notification on the specified ports.
```

show lldp config

Syntax

`show lldp config [[ethernet] PORT-LIST | oobm]`

Description

This command shows LLDP configuration information.

Parameters and options

`[ethernet] PORT-LIST`

Shows port-list configuration information.

`oobm`

Shows oobm LLDP configuration information.

show lldp config

```
switch(config)#show lldp config

LLDP Global Configuration
```

```
 LLDP Enabled [Yes] : Yes
 LLDP Transmit Interval    [30] : 30
 LLDP Hold time Multiplier  [4] : 4
 LLDP Delay Interval        [2] : 2
 LLDP Reinit Interval       [2] : 2
 LLDP Notification Interval [5] : 5
 LLDP Fast Start Count      [5] : 5


LLDP Port Configuration

 Port   | AdminStatus NotificationEnabled Med Topology Trap Enabled
 ------ + ----------- ------------------- -------------------------
 1      | Tx_Rx       False               False
 2      | Tx_Rx       False               False
 3      | Tx_Rx       False               False
 4      | Tx_Rx       False               False
 5      | Tx_Rx       False               False
 6      | Tx_Rx       False               False
 7      | Tx_Rx       False               False
 8      | Tx_Rx       False               False
 9      | Tx_Rx       False               False
OOBM    | Tx_Rx       False               False
```

show lldp config oobm

Syntax

`show lldp config oobm`

Description

This command shows oobm LLDP configuration information.

show lldp config oobm

```
switch(config)#show lldp config oobm

LLDP Port Configuration Detail

  Port : OOBM
  AdminStatus [Tx_Rx] : Tx_Rx
  NotificationEnabled [False] : False
  Med Topology Trap Enabled [False] : False


  TLVS Advertised:
   * port_descr
   * system_name
   * system_descr
   * system_cap

  IpAddress Advertised:
   * 10.0.0.1
```

show lldp info

Syntax

`show lldp info <local-device | remote-device> [[ethernet] PORT-LIST | oobm]`

Description

This command shows LLDP information about a local or remote device.

Parameters and options

`local-device`

Shows LLDP information about a local device.

`remote-device`

Shows LLDP information about a remote device.

Subcommands

The following are next level parameters of a local-or remote-device.

`[ethernet] PORT-LIST`

Shows port-list configuration information.

`oobm`

Shows oobm LLDP configuration information.

show lldp info local-device

Syntax

`show lldp info local-device`

Description

This command shows LLDP information about a local device.

show lldp info local-device

```
switch(config)# show lldp info local-device

LLDP Local Device Information

  Chassis Type : mac-address
  Chassis Id   : 08 2e 5f 69 8c 00
  System Name  : HPE Switch
  System Description : HPE Switch, revision XX.15.15.000...
  System Capabilities Supported: bridge, router
  System Capabilities Enabled: bridge

  Management Address  :
     Type: ipv4
     Address: 20.0.0.1

  OOBM Management Address:
     Type: ipv4
     Address: 100.0.0.1


 LLDP Port Information

  Port     PortType PortId   PortDesc
  -------- -------- -------- --------
  1        local    1        1
  2        local    2        2
  3        local    3        3
  4        local    4        4
  5        local    5        5
 OOBM      local    4000     OOBM
```

show lldp info local-device oobm

Syntax

```
show lldp info local-device oobm
```

Description

This command shows LLDP information about a local device for the specified oobm ports.

show lldp info local-device oobm

```
switch(config)# show lldp info local-device oobm
LLDP Local Port Information Detail

  Port     : OOBM
  PortType : local
  PortId   : 4000
  PortDesc : OOBM
  Pvid     : n/a
```

show lldp info remote-device oobm

Syntax

```
show lldp info remote-device oobm
```

Description

This command shows LLDP information about a remote device for the specified oobm ports.

show lldp info remote-device oobm

```
switch(config)# show lldp info remote-device oobm

 LLDP Remote Device Information Detail

  Local Port   : OOBM
  ChassisType  : mac-address
  ChassisId    : b4 b5 2f a8 84 00
  PortType     : local
  PortId       : 21
  SysName      : HPE Switch
  System Descr : HPE Switch, revision XX.15.15.000...
  PortDescr    : 21
  Pvid         :

  System Capabilities Supported  : bridge, router
  System Capabilities Enabled    : bridge

  Remote Management Address
     Type    : all802
     Address : b4 b5 2f a8 84 00
```

show lldp info remote-device 21

```
switch(config)# show lldp info remote-device 21

 LLDP Remote Device Information Detail

  Local Port   : 21
  ChassisType  : mac-address
```

```
 ChassisId    : b4 b5 2f a8 84 00
 PortType     : local
 PortId       : OOBM
 SysName      : HPE Switch
 System Descr : HPE Switch, revision XX.15.15.000...
 PortDescr    : OOBM
 Pvid         :

 System Capabilities Supported  : bridge, router
 System Capabilities Enabled    : bridge

 Remote Management Address
    Type    : all802
    Address : b4 b5 2f a8 84 00
```

show lldp stats

Syntax

`show lldp stats [[ethernet] PORT-LIST | oobm]`

Description

This command shows LLDP statistics.

Parameters and options

`oobm`

Shows statistics for the specified ports.

show lldp stats

```
switch(config)# show lldp stats

 LLDP Device Statistics

  Neighbor Entries List Last Updated : 45 mins
  New Neighbor Entries Count : 2
  Neighbor Entries Deleted Count : 0
  Neighbor Entries Dropped Count : 0
  Neighbor Entries AgeOut Count : 0

 LLDP Port Statistics

  Port    | NumFramesRecvd NumFramesSent NumFramesDiscarded
  ------ + -------------- ------------- ------------------
  1       | 91             96            0
  2       | 91             96            0
 OOBM     |  1              6            0
```

LLDP operating notes

Neighbor maximum

The neighbors table in the switch supports as many neighbors as there are ports on the switch. The switch can support multiple neighbors connected through a hub on a given port, but if the switch neighbor maximum is reached, advertisements from additional neighbors on the same or other ports will not be stored in the neighbors table unless some existing neighbors time-out or are removed.

LLDP packet forwarding

An 802.1D-compliant switch does not forward LLDP packets, regardless of whether LLDP is globally enabled or disabled on the switch.

One IP address advertisement per port

LLDP advertises only one IP address per port, even if multiple IP addresses are configured by `lldp config` *`<PORT-LIST>`* `ipAddrEnable` on a given port.

802.1Q VLAN information

LLDP packets do not include 802.1Q header information and are always handled as untagged packets.

Effect of 802.1X operation

If 802.1X port security is enabled on a port, and a connected device is not authorized, LLDP packets are not transmitted or received on that port. Any neighbor data stored in the neighbor MIB for that port prior to the unauthorized device connection remains in the MIB until it ages out. If an unauthorized device later becomes authorized, LLDP transmit and receive operation resumes.

Disconnecting a neighbor LLDP device

After disconnecting a neighbor LLDP device from the switch, the neighbor can continue to appear in the switch's neighbor database for an extended period if the neighbor's `holdtime-multiplier` is high; especially if the `refresh-interval` is large. See **Changing the time-to-live for transmitted advertisements** on page 309.

Mandatory TLVs

All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

Topology change notification

Enabling topology change notification on a switch port and then connecting or disconnecting an LLDP-MED endpoint on that port causes the switch to send an SNMP trap to notify the designated management stations. The port number included in the trap corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number appearing in an SNMP trap, use the `walkmib ifDescr` command, as shown in **Figure 97: Matching internal port numbers to external slot/port numbers** on page 327.

Figure 97: *Matching internal port numbers to external slot/port numbers*

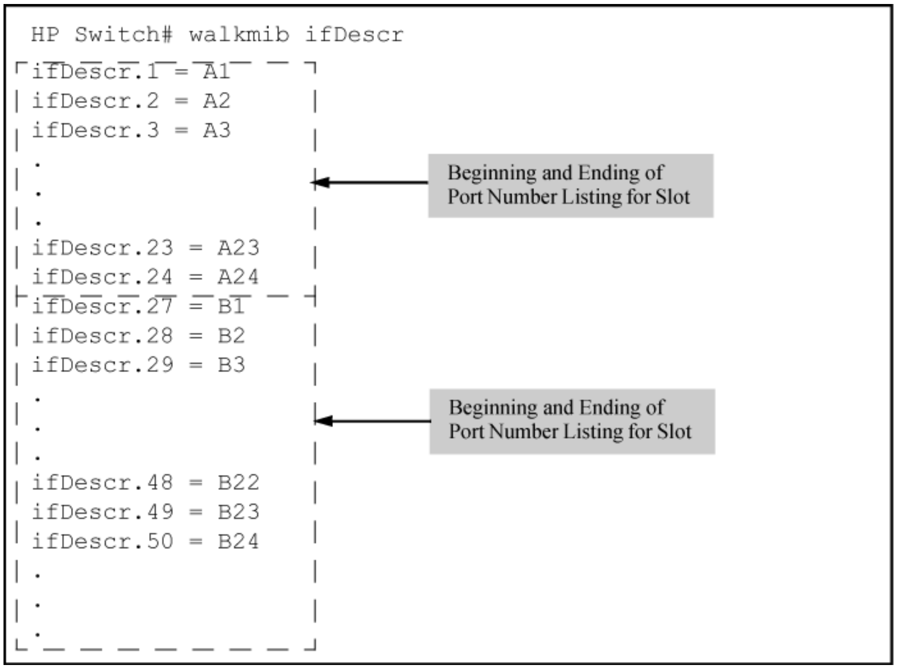

Advertisements currently in the neighbors MIB

show lldp info remote-device

Syntax

```
show lldp info remote-device<PORT-LIST>
```

Description

Without the *<PORT-LIST>* option, provides a global list of the individual devices it has detected by reading LLDP advertisements. Discovered devices are listed by the inbound port on which they were discovered.

Multiple devices listed for a single port indicates that such devices are connected to the switch through a hub.

Discovering the same device on multiple ports indicates that the remote device may be connected to the switch in one of the following ways:

- Through different VLANS using separate links. (This applies to switches that use the same MAC address for all configured VLANs.)
- Through different links in the same trunk.
- Through different links using the same VLAN. (In this case, spanning-tree should be invoked to prevent a network topology loop. Note that LLDP packets travel on links that spanning-tree blocks for other traffic types.)

With the *<PORT-LIST>* option, provides a listing of the LLDP data that the switch has detected in advertisements received on the specified ports.

See also **Data available for basic LLDP advertisements** on page 298.

A global listing of discovered devices

```
switch# show lldp info remote

 LLDP Remote Devices Information

  LocalPort | ChassisId                 PortId PortDescr SysName
  --------- + ------------------------- ------ --------- -------------
  1         | 00 11 85 35 3b 80         6      6         HP Switch 3500yl
  2         | 00 11 85 cf 66 60         8      8         HP Switch 3500yl
```

Figure 98: *An LLLDP-MED listing of an advertisement received from an LLDP-MED (VoIP telephone) source*

```
HP Switch(config)# show lldp info remote-device 1

 LLDP Remote Device Information Detail

  Local Port    : A2
  ChassisType   : network-address
  ChassisId     : 0f ff 7a 5c
  PortType      : mac-address
  PortId        : 08 00 0f 14 de f2
  SysName       : HP Switch
  System Descr  : HP Switch, revision K.15.06.0000x
  PortDescr     : LAN Port

  System Capabilities Supported  : bridge, telephone
  System Capabilities Enabled    : bridge, telephone

  Remote Management Address

  MED Information Detail
    EndpointClass          :Class3
    Media Policy Vlan id   :10
    Media Policy Priority  :7
    Media Policy Dscp      :44
    Media Policy Tagged    :False
    Poe Device Type        :PD
    Power Requested        :47
    Power Source           :Unknown
    Power Priority         :High
```

Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.

PoE advertisements

show lldp info remote-device

Syntax

`show lldp info remote-device <PORT-LIST>`

Description

Display the current power data for an LLDP-MED device connected to a port.

show power

`show power <PORT-LIST>`

Description

Display the current PoE configuration on the switch.

LLDP Management TLV Transmission disablement

Overview

A command has been written to suppress the IPv4 / IPv6 management address transmission in outgoing LLDP packets.

A local LAN device transmits organization-specific information in the form of type, length, and value (TLV). The organization-associated values are stored in the LLDP organizationally defined local device LLDP MIB extensions. Management address TLV in IPv4 and IPv6 environments is optional from the TLV basic management set.

TLV basic management set

- Port description TLV — Describes the port in an alpha-numeric format. The value equals the `ifDescr` object, if the LAN device supports RFC 2863.
- System name TLV — Assigns the system name in an alpha-numeric format. The value equals the `sysName` object, if the LAN device supports RFC 3418.
- System description TLV — Describes the network entity in an alpha-numeric format. The system description TLV includes the system name, versions of hardware, operating system, and networking software supported in the device. The value equals the `sysDescr` object, if the LAN device supports RFC 3418.
- System capabilities TLV — Indicates primary functions of the device and if they are enabled in the device.
- Management address TLV — Indicates the addresses of the local LLDP agent. Other remote managers can use this address to obtain information related to the local device.

The command `lldp config <all> ipAddrEnable <IP_ADDR>` is used to advertise specific IP address through the port.

The command `[no] lldp config <PORT NO> basicTlvEnable management_addr` suppresses the IP address to be advertised.

Commands

[no] lldp config basicTlvEnable management_addr

Syntax

In the configure context:

`[no] lldp config <PORT_NUM> basicTlvEnable management_addr`

Description

The feature suppresses the IPv4 or IPv6 address as well as suppresses the MAC address if the `[no] ip address` is configured. By default this management address TLV is enabled in switch. No other TLV (except management address TLV) suppression will occur when this command is used.

Parameters

Management_addr

Management TLV

Example

```
[no] lldp config all basicTlvEnable management_addr
```

lldp config

Syntax

`lldp config <port-number>`

Description

Configure the lldp for the desired port by number.

Parameter

`basicTlvEnable`

Enables the basic advertised TLV for each port by number.

Options

<management_addr>

Use the option *<management_addr>* to specify specific devices to enable TLV advertisement.

Usage

`lldp config <port_num> basicTlvEnable <management_addr>`

Show commands

Use the command `show running-config` to view the lldp configuration.

Example

```
switch# show running-config

Running configuration:
...
no lldp config 1 basicTlvEnable management_addr
```

Example

```
switch# show lldp config 1

LLDP Port Configuration Detail

  Port : 1
  AdminStatus [Tx_Rx] : Tx_Rx
  NotificationEnabled [False] : False
  Med Topology Trap Enabled [False] : False
  TLVS Advertised:

* management_addr


  IpAddress Advertised:
```

TVL configuration

VLAN ID TLV

lldp config dot1T1vEnable

Syntax

`[no] lldp config <PORT-LIST> dot1TlvEnable port-vlan-id`

Description

Use this command to enable or disable the VLAN ID TLV advertisement. Defaults to enabled.

Parameters and options

no

Disables the TLV advertisement.

Enabling the VLAN ID TLV

```
switch# lldp config a1 dot1TlvEnable port-vlan-id
```

Advertised TLVs

show lldp config

Syntax

`show lldp config <PORT_NAME>`

Description

The show commands display the configuration of the TLVs. The command `show lldp config` lists the TLVs advertised for each port.

Figure 99: *Displaying the TLVs for a port*

```
HP Switch(config)# show lldp config a1

 LLDP Port Configuration Detail

  Port : a1
  AdminStatus [Tx_Rx] : Tx_Rx
  NotificationEnabled [False] : False
  Med Topology Trap Enabled [False] : False


  TLVS Advertised:
   * port_descr
   * system_name
   * system_descr
   * system_cap

   * capabilities
   * network_policy
   * location_id
   * poe

   * macphy_config

   * port_vlan_id  <----------  The VLAN ID TLV is being advertised.

  IpAddress Advertised:
     :
     :
```

Figure 100: *Example of local device LLDP information*

```
HP Switch(config)# show lldp info local-device a1

LLDP Local Port Information Detail

  Port     : A1
  PortType : local
  PortId   : 1
  PortDesc : A1

  Port VLAN ID : 1 <---  The information that LLDP used in its
                         advertisement.
```

Figure 101: *Example of remote device LLDP information*

```
HP Switch(config)# show lldp info remote-device a1

LLDP Remote Device Information Detail

  Local Port   : A1
  ChassisType  : mac-address
  ChassisId    : 00 16 35 22 ca 40
  PortType     : local
  PortId       : 1
  SysName      : esp-dback
  System Descr : HP J8693A Switch 3500yl-48G, revision K.13.03, ROM ...
  PortDescr    : A1

  System Capabilities Supported  : bridge, router
  System Capabilities Enabled    : bridge, router

  Port VLAN ID : 200

  Remote Management Address
     Type    : ipv4
     Address : 192.168.1.1
```

TLVs controlled by medTLvEnable

lldp config medTlvEnable

Syntax

`[no] lldp config <PORT-LIST> medTlvEnable <MEDTLV>`

Description

TLVs controlled by medTlvEnable in the LLDP-MED configuration default to enabled.

This command enables or disables advertisement of the following TLVs on the specified ports:

- Device capability TLV
- Configured network policy TLV
- Configured location data TLV
- Current PoE status TLV

Helps to locate configuration mismatches by allowing use of an SNMP application to compare the LLDP-MED configuration on a port with the LLDP-MED TLVs advertised by a neighbor connected to that port.

Parameters and options

`capabilities`

This TLV enables the switch to determine:

- Which LLDP-MED TLVs a connected endpoint can discover
- The device class (1, 2, or 3) for the connected endpoint

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

Defaults to enabled. Cannot be disabled unless the `network_policy`, `poe`, and `location_id` TLVs are already disabled.

`network-policy`

This TLV enables the switch port to advertise its configured network policies (voice VLAN, Layer 2 QoS, Layer 3 QoS), and allows LLDP-MED endpoint devices to autoconfigure the voice network policy advertised by the switch. This also enables the use of SNMP applications to troubleshoot statically configured endpoint network policy mismatches.

Network policy is advertised only for ports that are configured as members of the voice VLAN. If the port belongs to more than one voice VLAN, the voice VLAN with the lowest-numbered VID is selected as the VLAN for voice traffic.

Defaults to enabled. If disabled, this TLV cannot be enabled unless the `capability` TLV is already enabled.

`location_id`

This TLV enables the switch port to advertise its configured location data (if any.)

Defaults to enabled. If disabled, this TLV cannot be enabled unless the capability TLV is already enabled.

`poe`

This TLV enables the switch port to advertise its current PoE state and to read the PoE requirements advertised by the LLDP-MED endpoint device connected to the port.

Defaults to enabled. If disabled, this TLV cannot be enabled unless the capability TLV is already enabled.

Generic header ID in configuration file

DHCP auto deployment

Auto deployment relies on DHCP options and the current DHCP auto-configuration function. Auto deployment is platform independent, avoiding the J-number validation of the downloaded configuration file when downloaded using DHCP option 66/67. The downloaded configuration file has an `IGNORE` tag immediately after the J-number in its header.

An option to add an `add-ignore-tag` to an existing copy command will insert an `ignore` tag into the configuration header. This insertion happens while transferring the configurations, (`startup configuration files`and `running configuration files`) from the switch to a configuration file setup on a remote server. The process uses TFTP/SFTP or can be accomplished with a serially connected workstation using XMODEM.

Add-Ignore-Tag option

The `add-ignore-tag` option is used in conjunction with the `copy` command to transfer the `startup configuration` or `running configuration` files from the switch to a remote server with `IGNORE` tag inserted into it.

The `IGNORE` tag is inserted into the first line of the configuration file directly after the J-number.

Configuration file

```
; J9782A IGNORE Configuration Editor; Created on release #YB.15.14.0000x
; Ver #04:63.ff.37.27:88
hostname "HP-2530-24"
snmp-server community "public" unrestricted
vlan 1
name "DEFAULT_VLAN"
no untagged 2,20-25
untagged 1,3-19,26-28
ip address dhcp-bootp
```

The J-number validation is ignored only when configuration file that contains the `IGNORE` tag is downloaded to a switch via DHCP option 66/67. When a `configuration` file containing the `IGNORE` tag is downloaded to a switch using CLI, SNMP or WebUI, the downloaded `configuration` file is only accepted if the J-number in it matches the J-number on the switch.

There is no change to the current switch configuration when executing the copy command with the `add-ignore-tag` option. The `IGNORE` tag is only added to the `configuration` file being exported to the external server. The `configuration` file stored on an external server is then downloaded to the switch using DHCP option 66 during bootup. If the `IGNORE` tag is available in the downloaded `configuration` file then the switch will avoid the J-number validation of the `configuration` file. The downloaded `configuration` file will then go through a line by line validation. Once the `configuration` file passes this validation, it gets updated in the flash. Once the `configuration` file has been updated, the switch will reboot automatically.

The J-number in the downloaded `configuration` file is replaced with that of the switch. The IGNORE tag is removed from the downloaded `configuration` file before updating it to flash. The `show running-configuration` command will not display the `IGNORE` tag but displays the switch's J-number as part of the output.

Copy with add-ignore-tag

```
switch(config)# copy startup-config tftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy running-config tftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy startup-config sftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy running-config sftp <ip-addr> <filename> add-ignore-tag
```

```
switch(config)# copy startup-config xmodem add-ignore-tag
switch(config)# copy running-config xmodem add-ignore-tag
```

Configuration commands for the add-ignore-tag option

Configuration files can be transferred to the switch from a server using the following `copy` commands:

- `copy tftp`
- `copy xmodem`
- `copy sftp`

Copy commands

```
copy tftp < startup-config | running-config > < ip-address > < remote-file >[ pc |
unix ]
copy xmodem startup-config < pc | unix >
copy sftp < startup-config | running-config > < ip-address > < remote-file >
```

`Configuration` files that are downloaded using the `copy` commands as described in the **example** will be accepted by the switch if they pass J-number validations and line by line validations after download. The downloaded configuration file will be discarded by the switch if the validations fail. If the validations fail, the switch will work with it's previous configuration.

Show logging commands for the add-ignore-tag option

The `show logging` command is used to locate errors during a configuration validation process. The event log catalogs entries with the ID#00158 and updates for each invalid entry found in the `configuration` file.

Show logging

```
-- Reverse event Log listing: Events Since Boot ----
W 01/07/14 00:29:31 00158 update: line 13. Module command missing for port or
invalid port: 36
I 01/07/14 00:29:30 00131 tftp: Transfer completed
I 01/07/14 00:29:29 00090 dhcp: Trying to download Config File (using TFTP)
received in DHCP from 192.168.1.1
```

NOTE

Downloading manually edited configuration file is not encouraged.

Exclusions

The `IGNORE` tag is not an available option when using external SCP, SFTP or TFTP clients such as PuTTY, Open SSH, WinSCP and SSH Secure Shell to transfer `configuration` files out of the switch.

Overview

The Dynamic Host Configuration Protocol (DHCP) is a network protocol that enables a server to automate assignment of IP addresses to hosts. A DHCP server can be configured to provide other network information like IP addresses of TFTP servers, DNS server, boot file name and vendor specific options. Commonly there are two types of address assignments, dynamic and manual. The lease of dynamic addresses is renewed periodically; manual leases are permanently assigned to hosts. With this feature, you can configure multiple pools of IP addresses for IP address assignment and tracking.

IP pools

A DHCP server is configured with IP pools. The server is then instructed to use IP addresses falling into the specified range of IP while offering leases. Multiple IP pools are configured to not have duplicate or overlapping IP subnets. You can also configure a DHCP server with multiple IP ranges within an IP subnet; this confines the allocatable IP addresses within the configured IP pool.

An IP pool will be claimed valid only if it is either:

- Dynamic pool – Has a network address, subnet mask and IP range(s)
- Static pool – Should have a static IP-to-MAC binding.

The DHCP server will discard the invalid and incomplete pools and will only operate on the valid IP pools. The DHCP server will require at least one valid pool to start.

DHCP options

On a DHCP server, an IP pool is configured with various options. These options signify additional information about the network. Options are supported with explicit commands such as `boot-file`. Option codes that correspond to explicit commands can not be configured with a generic option command; the generic option command requires an option code and TLV.

NOTE

RFC 2132 defines various network information that a client may request when trying to get the lease.

BootP support

The DHCP server also functions as BootP server. A manual binding configured in a static IP Pool may either service a BootP client request or a DHCP client request.

Authoritative server and support for DHCP inform packets

The server message `DHCPinform`may be received when the server is already configured for static IPv4 addresses so that the server can to get configuration parameters dynamically.

RFC 2131 states that if a client has obtained a network address through some other means (e.g., manual configuration), it may use a `DHCPinform` request message to obtain other local configuration parameters. Servers receiving a `DHCPinform` message construct a `DHCPACK` message with any local configuration parameters appropriate for the client without: allocating a new address, checking for an existing binding, filling in `yiaddr` or including lease time parameters.

Authoritative pools

To process the `DHCPINFORM` packets received from a client within the given IP pool, a DHCP server has to be configured as `authoritative` for that IP pool. The server is the sole authority for this IP pool so when a client requests an IP address lease where the server is authoritative, and the server has no record of that IP address, the server will respond with `DHCPNAK` message which indicates that the client should no longer use that IP address. Any `DHCPINFORM` packet received for a non-authoritative pool will be ignored by the DHCP server.

The `authoritative` command has no effect when configured on a static pool or an incomplete pool without a network statement. In such cases, the server intentionally not send an error message.

A CLI toggle is provided under the **pool** context that will allow the `authoritative` configuration.

The `authoritative` command requires a network statement to be configured on a pool.

Authoritative dummy pools

A dummy pool, without the range statement, can be configured and made authoritative. A dummy pool allows static-bind entries which do not have matching dynamic pools with network statements to be configured. By creating a dummy pool on a DHCP server, the support for `DHCPinform` packets will not be actively serving the client on this pool. No active leases or resource consumption will be sent to the DHCP server when this option is used.

Dummy pools help the DHCP server learn the network topology.

Example

```
dhcp-server pool dummy192
network 192.168.10.0 255.255.255.255
option 1…
option 2…
:
option n…
authoritative
exit
```

Change in server behavior

Making the server authoritative for an IP pool changes how the server processes `DHCP REQUEST` packets.

The following table exhibits the behavior on the receiving `DHCP REQUEST` and `DHCP inform` packets from DHCP clients residing on either authoritative and non-authoritative pools.

Table 20: *Authoritative and non-authoritative pools*

	Authoritative Pool			Non-authoritative pool		
When a DHCP client sending..	For Own IP	For IP belonging to different client	Unknown IP falling outside the range	For Own IP	For IP belonging to different client	Unknown IP falling outside the range
DHCP INFORM	send ACK	send ACK	send ACK	DROP	DROP	DROP
DHCP REQUEST	send ACK	send NACK	send NACK	send ACK	DROP	DROP

DHCPv4 configuration commands

DHCPv4 server

dhcp-server

Syntax

`[no] dhcp-server [enable | disable]`

Description

Use this command to nable/disable the DHCPv4 server in a switch. Defaults to disabled.

Parameters and options

no

Removes all DHCPv4 server configurations.

enable

Enables the DHCPv4 server on the device. The `no` form of this command

disable

Disables the DHCPv4 server on the device.

DHCP address pool name

dhcp-server pool

Syntax

`[no] dhcp-server pool <POOL-NAME>`

Description

Configure the DHCPv4 server IP address pool with either a static IP or a network IP range.

Parameters and options

pool

DHCPv4 server IP address pool.

ASCII-STR

Enter an ASCII string.

authoritative

Configure the DHCP server authoritative for a pool.

bootfile-name

Specify the boot file name which is used as a boot image.

default-router

List of IP addresses of the default routers.

dns-server

List of IP addresses of the DNS servers.

domain-name

Configure the DNS (Domain Name System) domain name for translation of hostnames to IP addresses.

lease

Lease period of an IP address.

netbios-name-server

List of IP addresses of the NetBIOS (WINS) name servers.

netbios-node-type

NetBIOS node type for a Microsoft DHCPv4 client.

network

Subnet IP and mask of the DHCPv4 server address pool.

option

Raw DHCPv4 server options.

range

Range of IP addresses for the DHCPv4 server address pool.

static-bind

Static binding information for the DHCPv4 server address pool.

tftp-server

Configure a TFTP server for the DHCPv4 server address pool.

Validations

Validation	Error/Warning/Prompt
Configuring pool when maximum Number of pools already configured.	Maximum number of pools (128) has already been reached
Configuring Pool with a name that exceeds the maximum length requirement.	String %s too long. Allowed length is 32 characters.

Table Continued

Trying to delete non existing pool	The specified address pool does not exist.
Only alphanumeric characters, numerals and underscore is allowed in the pool name. Violating this would throw the following error message.	Invalid name. Only alphanumeric characters and hyphen are allowed.
Trying to delete existing pool or adding new pool when DHCP server enabled.	DHCP server should be disabled before changing the configuration.

Authoritative

Syntax

`[no] authoritative`

Description

The DHCP server is the sole authority for the network configured under this pool. When the DHCP server is configured as authoritative, the server will respond with DHCP ACK or NACK as appropriate for all the received `DHCP REQUEST` and `DHCP INFORM` packets belonging to the subnet.

Non-authoritative `DHCP INFORM` packets received from the clients on a non-authoritative pool will be ignored.

Parameters and options

authoritative

Configure the DHCP server authoritative for a pool.

DHCP client boot file

bootfile-name

Syntax

`[no] bootfile-name <FILENAME>`

Description

Specify the boot file name to be used as the boot image.

DHCP client default router

default-router

Syntax

`[no] default-router <IP-ADDR-STR> [IP-ADDR2 IP-ADDR8]`

Description

Configure the DHCP pool context to the default router for a DHCP client. List all of the IP addresses of the default routers.

Two IP addresses must be separated by a comma.

Maximum of eight default routers can be configured.

DNS IP servers

dns-server

Syntax

`[no] dns-server <IP-ADDR> [IP-ADDR2 IP-ADDR8]`

Description

Configure the DHCP pool context to the DNS IP servers that are available to a DHCP client. List of IP addresses of the DNS servers.

Two IP addresses must be separated by comma.

Maximum of eight DNS servers can be configured.

Configure a domain name

domain-name

Syntax

`[no] domain-name <NAME>`

Description

Configure the DNS domain name for translation of hostnames to IP addresses.

Configure lease time

lease

Syntax

`[no] lease [DD:HH:MM | infinite]`

Description

Configure the lease time for an IP address in the DHCP pool. Lease time is infinite for static pools.

The default lease period is one day.

Parameters and options

DD:HH:MM

Enter lease period.

Lease

Lease period of an IP address.

NetBIOS WINS servers

Syntax

`[no] netbios-name-server <IP-ADDR-STR> [IP-ADDR2  IP-ADDR8]`

Description

Configure the DHCP pool for the NetBIOS WINS servers that are available to a Microsoft DHCP client. List all IP addresses of the NetBIOS(WINS) name servers. The Windows Internet Naming Service (WINS) is a name resolution service that Microsoft DHCP clients use to correlate host names to IP addresses within a general grouping of networks.

Two IP addresses must be separated by a comma.

Maximum of 8 NetBIOS (WINS) name servers can be configured.

NetBIOS node type

net bios-ode-type

Syntax

```
[no] netbios-node-type [ broadcast | hybrid | mixed | peer-to-peer ]
```

Description

Configure the DHCP pool mode to the NetBIOS node type for a Microsoft DHCP. The NetBIOS node type for Microsoft DHCP clients can be one of four settings: broadcast, peer-to-peer, mixed, or hybrid.

Parameters and options

broadcast

Broadcast node.

hybrid

Hybrid node.

mixed

Mixed node.

peer-to-peer

Peer to peer node.

Subnet and mask

network

Syntax

```
[no] network <ip-addr/mask-lenght>
```

Description

Configure the DHCPv4 server pool subnet and mask for the DHCP server address pool.

Range is configured to enable pool.

Parameters and options

ip-addr/mask-lenght

Interface IP address/mask.

DHCP server options

option

Syntax

```
[no] option <CODE> ascii <ASCII-STRING>|hex <HES-STRING>|ip <IP-ADDR-STR>[IP-ADDR2
… IP-ADDR8]
```

Description

Configure the raw DHCP server options.

Parameters and options

ascii

Specify ASCII string as option code value.

hex

Specify hexadecimal string as option code value.

ip

Specify one or more IP addresses as option code value.

ip-addr-str

Specify IP address.

ascii-str

Enter an ASCII string.

hex-str

Specify Hexadecimal string.

IP address range

range

Syntax

```
[no] range <IP-ADDR>[<IP-ADDR>]
```

Description

Configure the DHCP pool to the range of IP address for the DHCP address pool.

Parameters and options

range

Range of IP addresses for the DHCPv4 server address pool.

ip-addr

Low IP address.

High IP address.

Static bindings

static-bind

Syntax

```
static-bind ip <IP-ADDR/MASK-LENGTH> mac <MAC-ADDR>
```

Description

Configure static binding information for the DHCPv4 server address pool. Manual bindings are IP addresses that have been manually mapped to the MAC addresses of hosts that are found in the DHCP database. Manual bindings are just special address pools. There is no limit on the number of manual bindings but you can only configure one manual binding per host pool.

Parameters and options

ip

Specify client IP address.

static-bind

Static binding information for the DHCPv4 server address pool.

ip-addr / mask-length

Interface IP address or mask.

mac

Specify client MAC address.

mac-addr

Enter a MAC address.

TFTP server domain name

tftp-server

Syntax

[no] `tftp-server [server-name <server-name> | server-ip < ip-address >]`

Description

Configure the TFTP server domain name for the DHCP address pool.

Parameters and options

tftp-server

Configure a TFTP server for the DHCPv4 server address pool.

server-name

TFTP server name for the DHCPv4 server address pool.

Configure the TFTP server address

tftp-server

Syntax

`tftp-server server-ip <IP-ADDRESS>`

Description

Configure the TFTP server address for the DHCP address pool.

Parameters and options

server-ip

TFTP server IP addresses for the DHCPv4 server address pool.

ip-addr

Specify TFTP server IP address.

Number of ping packets

dhcp-server ping

Syntax

```
[no] dhcp-server ping [packets <0-10>|timeout <0-10>]
```

Description

Specify, in the global configuration context, the number of ping packets the DHCP server will send to the pool address before assigning the address. The default is two packets.

Parameters and options

ping

Specify DHCPv4 ping parameters.

packets <0-10>

Specify number of ping packets in the range of 0 to 10. 0 disables ping.

timeout <1-10

Ping timeout in the range of 1–10 seconds. Indicates the amount of time the DHCPv4 server must wait before timing out a ping packet. Defaults to one second.

Save DHCP server automatic bindings

dhcp-server database

Syntax

```
[no] dhcp-server database [file ASCII-STR] [delay<15-86400>][timeout <0-86400>]
```

Description

Specifies DHCPv4 database agent and the interval between database updates and database transfers.

Parameters and options

delay

Seconds to delay writing to the lease database file.

file

URL Format: "tftp://<ip-address>/<filename>".

database

Specifies DHCPv4 database agent and the interval between database updates and database transfers.

timeout

Seconds to wait for the transfer before failing.

ascii-str

Database URL.

<15-86400>

Delay in seconds.

<0-86400>

Timeout in seconds.

DHCP server and SNMP notifications

snmp-server enable traps

Syntax

```
[no] snmp-server enable traps dhcp-server
```

Description

Configure a DHCP server to send SNMP notifications to the SNMP entity. This command enables or disables event traps sent by the switch.

Parameters and options

dhcp-server

Traps for DHCP-Server.

Conflict logging on a DHCP server

dhcp-server conflict-logging

Syntax

```
[no] dhcp-server conflict-logging
```

Description

Enable conflict logging on a DHCP server. Default is disabled.

Parameters and options

conflict-logging

Enable DHCPv4 server address conflict logging.

Enable the DHCP server on a VLAN

dhcp-server

Syntax

```
dhcp-server
```

Description

Enable DHCPv4 server on a VLAN. DHCPv4 client or DHCPv4 relay cannot co-exist with DHCPv4 server on a VLAN.

Parameters and options

dhcp-server

Enable DHCPv4 server on a VLAN.

Clear commands

clear dhcp-server conflicts

Syntax

```
clear dhcp-server conflicts <IP-ADDR>
```

Description

Reset DHCPv4 server conflicts database. If IP address is specified, reset only that conflict.

Parameters and options

dhcp-server

Clears theDHCPv4 server information.

ip-addr

Specify the IP address whose conflict is to be cleared.

Reset all DHCP server and BOOTP counters

clear dhcp-server statistics

Syntax

```
clear dhcp-server statistics
```

Description

Reset all DHCP server and BOOTP counters

Parameters and options

statistics

Reset DHCPv4 server and BOOTP counters.

Delete an automatic address binding

clear dhcp-server statistics

Syntax

```
clear dhcp-server statistics
```

Description

Delete an automatic address binding from the DHCP server database.

Parameters and options

binding

Reset DHCPv4 server automatic address bindings.

ip-addr

Specify IP address of the binding is to be cleared.

Show commands

show dhcp-server

Syntax

```
show dhcp-server [binding|conflicts|database|statistics|pool <POOL-NAME>]
```

Description

Show DHCPv4 server global configuration information for the device.

Parameters and options

binding

Display the DHCPv4 server address bindings on the device..

conflicts

Display address conflicts found by a DHCPv4 server when addresses are offered by a client.

database

Display DHCPv4 server database agent information.

statistics

Display DHCPv4 server statistics.

pool <POOL-NAME>

Display the DHCPv4 server IP pool information.

Event log

Event Log Messages

Cause

Table 21: *Event Log Messages*

Events	Debug messages
DHCP server is enabled globally.	`DHCP server is enabled globally.`
DHCP server is enabled globally. Warnings - One or more incomplete pool configurations are found during the server startup.	`DHCP server is enabled globally.Warning -One or more incomplete pool configurations are found during the server startup.`
A dynamic pool is considered invalid, if network IP or subnet mask is not configured. A static pool is considered incomplete, if network IP, subnet mask or MAC address is not configured.	
DHCP server failed to start. The reason for failure is printed as the argument.	`DHCP server failed to start: %s "with a manual binding.`
DHCP server is disabled globally.	`DHCP server is disabled globally.`
The DHCP server configurations are deleted.	`The DHCP server configurations are deleted`
Decline from client when server assigns an illegal Ipv6 address.	`%s: Decline offer from %x (server) of %x because the address is illegal.`
DHCP server is enabled on a specific VLAN.	`DHCP server is enabled on VLAN %d`

Table Continued

Events	Debug messages
DHCP server is disabled on a specific VLAN.	`DHCP server is disabled on VLAN %d`
Ping check is enabled and configured with specified retry count and timeout values	`Ping-check configured with retry count = %d, timeout = %d`
Ping check is disabled	`Ping-check is disabled`
Conflict-logging is enabled	`Conflict-logging is enabled`
Conflict-logging is disabled.	`Conflict-logging is disabled.`
A specific IP address is removed from the conflict logging database.	`IP address %s is removed from the conflict-logging database.`
All IP addresses are removed from the conflict-logging database.	`"All IP addresses are removed from the conflict-logging database`
Dynamic binding for a specific IP address is freed.	`Dynamic binding for IP address %s is freed`
All the dynamic IP bindings are freed.	`All the dynamic IP bindings are freed`
Remote binding database is configured for a specific URL.	`Remote binding database is configured at %s`
Remote biding database is disabled.	`Remote binding database is disabled`
Binding database is read from the specified URL at the specified time	`Binding database read from %s at %s`
Failed to read the remote binding from the specified URL.	`Failed to read the remote binding database at %s`
Binding database is written to the specified URL at the specified time.	`Binding database written to %s at %s`
Failed to write the binding database to the specified URL. The reason for failure is printed as argument.	`Failed to write the binding database to %s. Error: %s`
Invalid bindings are found in the database at the specified URL.	`Invalid binding database at %s`
The specified VLAN does not have a matching IP pool configured.This occurs when the DHCP-server is enabled on the specified VLAN, but no IP pool is configured with a network IP matching the VLAN network IP.	`VLAN %d does not have a matching IP pool`

Table Continued

Events	Debug messages
Binding database is replicated to standby management module.	`Binding database is replicated to standby management module`
DHCP server is listening for DHCP packetsThis message is displayed when DHCP server is enabled globally and DHCP server is enabled on at-least one VLAN.	`DHCP server is listening for DHCP packets`
DHCP server is disabled on all the VLANs. Server is no longer listening for DHCP packets.	`DHCP server is disabled on all the VLANs. Server is no longer listening for DHCP packets`
The specified IP is not offered to the DHCP client, as it is already in use.	`IP address %s is not offered, as it is already in use`
No IP addresses available on the specified pool.	`No IP addresses to offer from pool %s`
High threshold reached for the specified pool. Count of Active bindings and Free bindings are printed as arguments.	`High threshold reached for pool %s. Active bindings: %d, Free bindings: %d`
Low threshold reached for the specified pool. Count of Active bindings and Free bindings are printed as arguments.	`Low threshold reached for pool %s. Active bindings: %d, Free bindings: %d`
No active VLAN with an IP address is available to read binding database from the configured URL.	`No active Vlan with an IP address available to read binding database`

DHCPv6 hardware address

The incremental deployment of IPv6 to existing IPv4 networks results in dual-stacking network environments. Some devices will act as both DHCPv4 and DHCPv6 clients. For these dual-stack situation, here is a need to associate DHCPv4 and DHCPv6 messages with the same client interface. A DHCPv4 server uses the client link-layer address as the customer identifier and a key for lookup in the client database. The DHCPv6 Relay-Forward message carries the client link-layer address to the DHCPv6 server allowing the association of both DHCPv4 and DHCPv6 messages with the same client interface.

As defined in RFC-6939, DHCPv6 relay agents receiving solicit and request messages that originate from DHCPv6 clients include the link-layer source address of the received DHCPv6 message. This is accomplished in the Client Link-Layer Address option within DHCPv6 Relay-Forward messages. The Client Link-Layer Address enables the server to recognize and service specific clients. DHCPv6 relay agent behavior (as set by the configuration) decides whether the Client Link-Layer Address option is included for each client.

DHCPv6 relays agents include Option–79 for all message types when enabled. The message types are: solicit, request, confirm, decline, renew, rebind, release and information-request. DHCPv6 provides additional information for event debugging and logging related to the client at the server.

NOTE

All cascading relay-agents simply encapsulate the message received and relay-forward to the server. The service function does not receive any message-types directly from the client even when the feature is enabled.

DHCPv6 snooping and relay

dhcpv6–snooping

Syntax

`[no] dhcpv6-snooping [vlan <VLAN-ID-RANGE>]`

Description

Enable or disable the global administrative status of DHCPv6 snooping. You must enable DHCP snooping globally (`dhcpv6-snooping`) to enable snooping on any VLAN.

Parameters and options

no

Disabling global administrative status (`no dhcpvt-snooping`) disables snooping on all VLANs.

vlan *<VLAN-ID-RANGE>]*

Disables snooping on a VLAN or a range of VLANs. Requires enabling DHCP global snooping (`dhcpv6-snooping`)

Validation rules for DHCPv6 global snooping

Validation	Error/Warning/Prompt
Verify whether entered ipv6 address is valid	Invalid Ipv6 address:< ipv6-address>
If an invalid server address is configured	Invalid IP address. Only IPv6 unicast or link-local addresses are supported.
If the limit on configuring the authorized servers had reached.	Cannot configure the authorized server as only 20 authorized servers can be configured.

Validation rules for DHCPv6–snooping VLAN

Validation	Error/Warning/Prompt
if the VLAN is a SVLAN and the bridge mode is mixed mode	DHCPv6-snooping is not supported on SVLANs and SVLAN ports in QinQ mixed VLAN mode
If number of snooped VLAN count is greater than max_vlans_with_dipv6ld and also the max binding limit has reached.	DHCPv6 snooping cannot be enabled on %s VLANs. The switch will support only 8 DHCPv6 snooping enabled VLANs when Dynamic IPv6 Lockdown feature is enabled.
If the VLAN which is being configured for DHCPv6 Snooping has a Smart Link enabled port.	Cannot configure DHCPv6 Snooping on a VLAN containing Smart Link ports.
If a VLAN is being configured as a Smart Link protected VLAN and DHCPv6 Snooping is enabled on it.	Cannot configure a VLAN as a protected VLAN when DHCPv6 Snooping is enabled on it .
If Smart Link is being configured on a port which is a part of DHCPv6 Snooping VLAN..	Canot configure the Smart Link feature on a port when DHCPv6 Snooping is enabled on that port.

dhcpv6 snooping trust

Syntax

```
[no] dhcpv6-snooping trust ethernet <PORT-LIST>
```

Description

Configure trusted interfaces. The system forwards server packets received on trusted interfaces only.

Parameters and options

no

Marks the specified interfaces as untrusted. Port state defaults to untrusted.

Validation rules

Validation	Error/Warning/Prompt
Verify whether the port exist in the device.	Module not present for port or invalid port: <PORT-LIST>
If the port is a part of a SVLAN and the bridge mode is mixed mode.	Port %s cannot be configured as trusted port as it is part of a SVLAN in QinQ mixed VLAN mode.
If the port is not a part of a dsnoopv6 enabled VLAN	Port %s is not a part of a DHCPv6-snooping VLAN.
If trusted attribute is being configured on a port on which max-binding has been already configured.	Disable max-binding feature configured on the port before configuring it as a trusted port.
If a Dynamic trunk is configured as a trusted port.	Cannot configure a port as a DHCPv6 Snooping trusted port when Dynamic Trunking is enabled on that port.
If a Smart Link port is being configured as a trusted port	Cannot configure a Smart Link port as a DHCPv6 Snooping trusted port.
If a trusted port is being configured as a Smart Link port	Cannot configure a DHCPv6 Snooping trusted port as a Smart Link port

dhcpv6–snooping authorized-server

Syntax

```
[no] dhcpv6-snooping authorized-server <IPV6-ADDRESS>
```

Description

Configure authorized DHCPv6 servers. For DHCPv6 snooping to allow a server to client packet to be forwarded, it must be received on a trusted port from an authorized server. If no authorized servers are configured, all server addresses are valid.

ddhcpv6–snooping database file

Syntax

```
[no] dhcpv6-snooping database file [ASCII-STR|delay <15-86400>| timeout<0-86400>]
```

Description

Configure a lease entry file and its options for storing DHCPv6 snooping binding database.

Parameters and options

ASCII-STR

Copies the DHCPv6 snooping lease file to a TFTP server. The parameter ASCII-STR is a URL and is in the format `tftp://<IP-ADDR>/<FILENAME>`. The TFTP address can be up to 255 characters. IP-ADDR can be an IPv4 address or an IPv6 address. The IPv6 address must be enclosed in square brackets [].

timeout *seconds*

Configures the number of seconds to wait for the DSNOOPv6 lease file transfer to complete. An error message is displayed if the file transfer is not completed within the timeout value. A value of zero indicates that the attempt to transfer the DHCPv6 lease file retries indefinitely. The default timeout value is 300 seconds.

database

Configure the parameters to copy the DHCPv6 Snooping lease file to a TFTP server.

delay

Configure the number of seconds to wait before copying the DSNOOPv6 lease file to a TFTP server.

file

Copy the DHCPv6 Snooping lease file to a TFTP server.

timeout

Configure the number of seconds to wait for the DSNOOPv6 lease file transfer to complete.

Validation rules

Validation	Error/Warning/Prompt
Verify whether file name entered is in URL format	database: Bad URL format.
Verify whether the timeout value is within the limit	Invalid input: <value>
Verify whether the delay value is within the limit.	Invalid input: <value>
If the URL format is not proper	Bad URL format.
If the entered URL does not have a valid transfer mode.	URL Transport mode is not supported.

dhcpv6–snooping max-bindings

Syntax

`[no] dhcpv6-snooping max-bindings <PORT-LIST 1-8192>`

Description

Configure the maximum number of binding addresses allowed per binding anchor. A binding anchor is a unique attribute that can be associated with a client address.

Parameters and options

max-bindings

Configuring maximum number of binding addresses allowed per port.

- If the max-bindings value is configured **before** enabling `dhcpv6-snooping` the limit is immediately applied and the bindings are not allowed to exceed the max-bindings value.
- The max-bindings value is **setafter** enabling `dhcpv6-snooping`.
- The current bindings are greater than the max-binding value, the configuration will be applied as and when clients release their IPv6 addresses.
- Current bindings are lesser than that of the value entered, the configuration will be immediately applied.

<PORT-LIST 1-8192>

Specify the ports on which max-bindings need to be applied in the range of 1–8192.

Validation rules

Validation	Error/Warning/Prompt
Verify max-bindings value entered is in the range	Invalid input: <value>
If DHCPv6-Snooping is already configured before entering the command and current bindings are greater than the value being set.	Existing bindings %d are more than the max-bindings being configured, and the maximum limit will be applied once the number of existing bindings fall below this limit
If the value is being configured for a trusted port	Cannot configure maximum binding for DHCPv6 snooping feature on a trusted port
If the value is being configured for a port which is not a part of a dhcpv6-snooped vlan	Port %s is not a part of a DHCPv6-snooping VLAN.
If the max-binding value is being set for a Dynamic trunk.	Cannot configure DHCPv6 Snooping on a port when Dynamic Trunking is enabled on that port.
If the number of static bindings is greater than the max-binding value being set.	Cannot configure the maximum binding value because the number of static bindings on the port exceeds the maximum binding value.
If a port on which max-binding is enabled is being put into a trunk.	Cannot add a port to a trunk group when DHCPv6 Snooping Maxbinding is configured on that port.
If a trunk has max-bindings configured on it. And the trunk is being removed.	Cannot remove the port %s from the trunk group because DHCPv6 Snooping max-binding is configured on the trunk and removing the port will delete the trunk.
If DT trunk is being configured on a max-binding enabled port.	Cannot configure Distributed Trunking on a port when DHCPv6 Snooping max-binding is configured on that port.

DT trunks can use jumbo VLAN as usual, but user needs to ensure that jumbo is configured on both the DT pairs, otherwise packet drops/fragmentations can be seen.

dhcpv6–relay option 79

Syntax

```
[no] dhcpv6-relay option 79
```

Description

Enabling option 79 will force the DHCPv6 Relay agent to forward the client Link-layer address. Defaults to disabled.

snmp-server enable traps dhcpv6-snooping

Syntax

```
[no] snmp-server enable traps dhcpv6-snooping [out-of-resources|errant-reply]
```

Description

Configure the traps for DHCPv6 snooping.

Parameters and options

out-of-resources

This trap is sent when the number of bindings exceed the maximum limit of 8192 bindings.

errant-reply

This trap is sent when a DHCPv6 reply packet is received on an untrusted port or from an un-authorized server.

clear dhcpv6–snooping stats

Syntax

`clear dhcpv6-snooping stats`

Description

Clears dhcpv6 snooping statistics.

Validation rules

Validation	Error/Warning/Prompt
If dhcp-snooping not enabled globally	DHCPv6 snooping is disabled.

debug security dhcpv6–snooping

Syntax

`debug security dhcpv6-snooping [config|event|packet]`

Description

Enable debug for DHCPv6 snooping.

Parameters and options

config

Debug DHCPv6 snooping configuration.

event

Debug a DHCPv6 snooping event.

packet

Debug DHCPv6 snooping by packet.

ipv6 source-lockdown ethernet

Syntax

`[no] ipv6 source-lockdown ethernet <PORT-LIST>`

Description

Used to configure DIPv6LD lockdown globally and on specific ports which can be configured on per-port basis using the PORT-LIST option.

Parameters and options

[ethernet] PORT-LIST

Specify the ports being configured for Ipv6 source-lockdown.

source-lockdown

Enable IPv6 source lockdown for a specific port.

Validation rules

Validation	Error/Warning/Prompt
Verify whether dhcpv6-snooping is enabled globally	DHCPv6 snooping is disabled.
Verify whether port configured is in the VLAN which is dhcpv6-snooping enabled.	Ports <PORT-LIST> are not in a DHCPv6 Snooping VLAN.
If lockdown is being configured on a trusted port	Port %s is a trusted port.
If the HW resources are not available for changing dipv6ld global or a port characteristic	Cannot enable DIPLDv6 as required resources are unavailable.
If global GVRP is enabled	DIPLDv6 cannot be enabled when GVRP is enabled
If no of snooped VLAN count is greater than max_vlans_with_dipv6ld	DHCPv6 snooping cannot be enabled on %s VLANs. The switch support only 8 DhCPv6 snooping enabled VLANs when Dynamic Ipv6 lockdown is enabled.
If Binding limits are exceeded	Cannot enable Dynamic Ipv6 Lockdown on ports %s as manual binding limits are exceeded.
If lockdown is being enabled on an interface which is part of a dynamic trunk (LACP)	Cannot configure Dynamic Ipv6 Lockdown on interface %s, it is a Dynamic trunk.
If lock down is being configured on a mesh port	Cannot configure Dynamic Ipv6 Lockdown on a logical mesh port.
If trunk is being formed using a port which has DIPLDv6 enabled on it.	Cannot add a port to a trunk group when Dynamic IPv6 Lockdown is enabled on that port.
If DIPLDv6 is configured on a trunk and the trunk is being removed.	Cannot remove the port %s from the trunk group because Dynamic IPv6 Lockdown is configured on the trunk and removing the port will delete the trunk.
If DIPLDv6 is being is configured on a Smart Link port	Cannot enable Dynamic IPv6 Lockdown feature on a Smart Link port.
If Smart Link is being enabled on a DIPLDv6 enabled port	Cannot configure the Smart Link feature on a port when the Dynamic IPv6 Lockdown feature is enabled on that port.

ipv6 source-binding

Syntax

`[no] ipv6 source-binding` *`VLAN-ID IPV6-ADDR MAC-ADDR PORT-NUM IPV6-ADDR`*

Description

Add a DHCPv6 static binding entry into the binding table. Static binding entries will have infinite lifetime.

Parameters and options

VLAN-ID

The VLAN ID of the static binding entry.

Ipv6-ADDRESS

The Ipv6 address of the static binding entry.

MAC-ADDRESS

The MAC address of the static binding entry.

[ethernet] PORT-NUM

Port number of the static binding entry.

IPV6-ADDR

The Ipv6 link-local address of the static binding entry.

Validation rules

Validation	Error/Warning/Prompt
Verify whether the vlan id is proper	Invalid input:%s
Verify whether the mac-address is valid	Invalid input:%s
Verify whether the ipv6 address is valid	Invalid input:%s
Verify whether the port number is valid on the device	Module not present for port or invalid port: <port-num>
If any other addresses other than global unicast address are entered	Invalid Ipv6 address
If the ipv6 address entered is not a unicast.	Only Ipv6 unicast addresses are supported.
If a multicast ipv6 address is entered to configure a binding.	Cannot configure a binding using a multicast IPV6 address.
If an invalid MAC address is being added into the binding table.	Cannot add a %s MAC address to the table.
If an invalid port is used for configuring a static binding	Port %s is invalid.
If DSNOOPV6 is globally disabled when configuring a static binding.	Cannot configure static binding whenDHCPv6 Snooping is disabled.

Table Continued

Validation	Error/Warning/Prompt
While configuring a static binding if the Ipv6 address is already present in the Binding table but the entered vlanid and MAC address doesnot match with the one present in the binding table.	%s has already been assigned to a VID/MAC. Delete the existing binding first.
If a binding which does not exist in the binding table is tried to be removed.	Binding for %s not found.
If DIPLDv6 limits are exceeded on the switch.	Cannot add the IPv6 source binding because the number of source bindings exceeds the maximum limit of "STR(DSNOOPV6_MAX_STATIC_LEASES)".
If more than 4 IPv6 addresses are being assigned to a VID/MAC pair	Cannot add the IPv6 source binding because only "STR(DHCPV6_MAX_IAADDRS)"IPv6 addresses can be bound to a VID-MAC pair.
If a VID-MAC pair is bound to a link-local address and the same VID-MAC pair is being assigned another link-local address.	%s is already bound to a link-local address. To bind another link-local address, delete the existing binding .
If a binding exists for a particular client in the BST and the same binding is being configured for another port.	The IPv6 source binding already exists for another port.
If the switch total limit for bindings is exceeded.	Cannot add the IPv6 source binding because the number of source bindings exceeds the maximum limit of STR(DSNOOPV6_MAX_STATIC_LEASES).
If a trunk is being configured for a port which has static binding configured on it.	Cannot add a port to a trunk group when IPv6 source binding is configured on that port.
If static binding is being configured on a Smart Link enabled port	Cannot configure IPv6 source binding on a Smart Link port.
If Smart Link is being configured on a port with static binding.	Cannot configure Smart Link feature on a port when IPv6 source binding is configured on that port.

snmp-server enable traps dyn-ipv6-lockdown

Syntax

```
[no] snmp-server enable traps dyn-ipv6-lockdown [out-of-resources | violations]
```

Description

The Dynamic IPv6 Lockdown trap is sent when resources are unavailable for configuring. This trap is sent when a source lockdown violation takes place.

Parameters and options

out-of-resources

Dynamic IPv6 Lockdown out of resources.

violations

Dynamic IPv6 lockdown violations.

debug security dynamic-ipv6-lockdown

Syntax

`debug security dynamic-ipv6-lockdown`

Description

Enable debug for DIPLDv6

Show commands for DHCPv6–snooping

show dhcpv6-snooping

Syntax

`show dhcpv6-snooping`

Description

Show dhcpv6 snooping configuration.

Validaton rules

Validation	Error/Warning/Prompt
If dhcpv6-snooping not enabled	DHCPv6 snooping is disabled

show dhcpv6 snooping bindings

Syntax

`show dhcpv6-snooping bindings`

Description

Show dhcpv6 snooping binding entries. This would show both dynamic and static binding entries.

Validation rules

Validation	Error/Warning/Prompt
If dhcpv6-snooping not enabled	DHCPv6 snooping is disabled

show dhcpv6 snooping statistics

Syntax

`show dhcpv6-snooping stats`

Description

Show dhcpv6-snooping statistics.

show ipv6 source-lockdown

Syntax

`show ipv6 source-lockdown [bindings | status]`

Description

Shows IPv6 source bindings that are configured using the command `IPv6 source-bindings`.

Parameters and options

bindings

Show source bindings for Dynamic IPv6 Lockdown ports.

status

Show source bindings for Dynamic IPv6 Lockdown status.

Show source bindings Dynamic IPv6 Lockdown status

```
Dynamic IPv6 Lockdown Bindings

Port    IPv6 Address                           Vlan   MAC          Not
Address                                                            in HW
---- ---------------------------------------- ----- ------------- -------
A1   3000:abbb:1234:3456:1234: 1234:1234:1234  1     123456-789101  Yes
F23  300:ab::2                                 4092  abcdef-123455  No
```

show ipv6 source-lockdown status

Syntax

`show ipv6 source-lockdown status`

Description

Used to show IPV6 source-lockdown status per port.

Parameters and options

source-lockdown

Show dynamic IPv6 Lockdown.

Show dynamic IPv6 Lockdown configuration

```
Dynamic IPv6 Lockdown information
Global State: Enabled
Port   Operational State
------ --------------------------
1    Active
2    Active
IPv6 Source Lockdown is disabled on Ports 3-24.
```

show snmp-server traps

Syntax

```
show snmp-server traps <COMMUNITY-STR>
```

Description

Shows traps controlled. Shows all information on SNMP communities, trap receivers and SNMP response or trap source-ip policy configured on the switch.

Parameters and options

traps

Show all configured traps.

<COMMUNITY-STR>

Displays information for the specified community only.

Show snmp-server traps

```
switch(config)# sh snmp-server traps

Trap Receivers
Link-Change Traps Enabled on Ports [All] : All

Traps Category                                           Current Status
_____________________________       __________________
SNMP Authentication                                    : Extended
Password change                                             : Enabled
Login failures                                               : Enabled
Port-Security                                                 : Enabled
Authorization Server Contact                  : Enabled
DHCP-Snooping                                                 : Enabled
DHCPv6-Snooping Out of Resource             : Enabled
DHCPv6-Snooping Errant Replies              : Enabled
Dynamic ARP Protection                              : Enabled
Dynamic IP Lockdown                                    : Enabled
Dynamic IPv6 Lockdown Out of Resource : Enabled
Dynamic IPv6 Lockdown Violations            : Enabled
Startup Config change                                 : Disabled
Running Config Change                                 : Disabled
MAC address table changes                         : Disabled
MAC Address Count                                         : Disabled

Address                Community            Events  Type  Retry  Timeout

Excluded MIBs
switch(config)#
Alignment change - right shifted
```

show distributed-trunking consistency-parameters

Syntax

```
show distributed-trunking consistency-parameters global feature
```

Description

Parameters and options

dhcp-snooping

Display DHCP snooping peer consistency details.

IGMP

Display IGMP peer consistency details.

loop-protect

Display Loop protect peer consistency details.

MLD

Display MLD peer consistency details.

pim-dm

Display PIM-DM peer consistency details.

pim-sm

Display PIM-SM peer consistency details.

Display PIM-SM peer consistency details.

```
show distributed-trunking consistency-parameters global feature pim-sm

PIM-SM Enabled VLANs on Local : 20,30
PIM-SM Enabled VLANs on Peer : 20,30
```

show distributed-trunking consistency-parameters

Syntax

`show distributed-trunking consistency-parameters global <PIM-SM>`

Description

Display global peer consistency details. If the platforms do not match an error message similar to `inconsistant criteria` is returned.

Parameters and options

global

Display global peer consistency details.

<PIM-SM>

Display PIM-SM peer consistency details.

Show distributed-trunking consistency-parameters global

```
switch# show distributed-trunking consistency-parameters global
Local Peer
----- ----
Peer config unavailable.
Image Version KB.15.18.0000x

IP Routing Disabled Disabled
Peer-keepalive interval 1000 0
PIM-DM Support Disabled Disabled
PIM-SM Support Disabled Disabled
```

```
IGMP enabled VLANs on Local :
IGMP enabled VLANs on Peer :
PIM-DM Enabled VLANs on Local : <List of Vlans>
PIM-DM Enabled VLANs on Peer : <List of Vlans>
PIM-SM enabled VLANs on Local : <List of Vlans>
PIM-SM enabled VLANs on Peer : <List of Vlans>
DHCP-snooping Enabled on Local :
DHCP-Snooping Enabled on Peer                 : Yes
DHCP-Snooping Enabled VLANs on Local          : 1
DHCP-Snooping Enabled VLANs on Peer           : 1
DHCP-Snooping Max-Binding Configured on Local  : Yes

Ports  Max-Bindings
------- ------------
Trk2   6

DHCP-Snooping Max-Binding Configured on Peer   : No
```

Feature pim-sm

```
show distributed-trunking consistency-parameters global feature pim-sm

PIM-SM Enabled VLANs on Local : 20,30
PIM-SM Enabled VLANs on Peer : 20,30
```

show dhcpv6 relay

Syntax

`show dhcpv6-relay`

Description

Displays the DHCPv6 relay configuration. Cannot be configured from the WebUI or Menu.

Sample output

```
show dhcpv6-relay

DHCPV6 Relay Agent : Enabled
Option 79 : Disabled
```

DHCPv6 event log

Cause

Event	Message
RMON_DSNOOPV6_UNTRUSTED_PORT_SERVER_RELAY	%s: %s message received on the untrusted port %s from %s.
RMON_DSNOOPV6_UNTRUSTED_PORT_SERVER_SUSP	%s: Ceasing the log messages for the server packets received on an untrusted port for %s.
RMON_DSNOOPV6_UNTRUSTED_PORT_CLIENT_DEST	%s: Client packet destined to the untrusted port %s is dropped.
RMON_DSNOOPV6_UNTRUSTED_PORT_CLIENT_DEST_SUSP	%s: Ceasing the log messages for the client packets destined to an untrusted port for %s.
RMON_DSNOOPV6_UNAUTHORIZED_SERVER	%s: Unauthorized server %s detected on port %s
RMON_DSNOOPV6_UNAUTHORIZED_SERVER_SUSP	%s: Ceasing unauthorized server logs for %s
RMON_DSNOOPV6_BAD_RELEASE	%s: Illegal IPv6 release from %02X%02X%02X-%02X%02X %02X on port %s; Address leased to other client or not leased. %s.
RMON_DSNOOPV6_BAD_RELEASE_SUSP	%s: Ceasing the log messages for the illegal IPv6 release messages received from the clients for %s
RMON_DSNOOPV6_TABLE_FULL	%s: Unable to add the DHCPv6 lease because the lease table is full.
RMON_DSNOOPV6_TABLE_FULL_SUSP	%s: Ceasing the log messages for the failed lease table updates for %s.
RMON_DSNOOPV6_MAX_BINDING_CROSSED	%s: Droppped IPv6 request from %02X%02X%02X-%02X %02X%02X. The max-binding limit has reached on the port %s. %s

Table Continued

Event	Message
RMON_DSNOOPV6_MAX_BINDING_CROSSED_SUSP	%s: Ceasing max-binding limit crossed packet information logs for %s.
RMON_DSNOOPV6_EVENT_MAXBINDING_REMOVED	%s: The DHCPv6-Snooping max-binding configured on port %s is removed.
RMON_DSNOOPV6_EVENT_MAXBINDING_REMOVED_SUSP	%s: Ceasing the log messages for the removal of DHCPv6-Snooping max-binding from the ports for %s
RMON_DSNOOPV6_EVENT_BINDINGS_EQUALS_MAXBIND	%s: The number of bindings on the port %s equals the maximum binding configured on that port.
RMON_DSNOOPV6_EVENT_BINDINGS_EQUALS_MAXBIND_SUSP	%s: Ceasing the log messages for bindings on port that equals max-binding value for %s.
RMON_DSNOOPV6_EVENT_MAXBIND_BELOW_BINDINGS	%s: The number of bindings on the port %s exceeds the maximum binding configured on that port.
RMON_DSNOOPV6_EVENT_MAXBIND_BELOW_BINDINGS_SUSP	%s: Ceasing the log messages for bindings on port that exceeds max-binding for %s.
RMON_DSNOOPV6_READ_LEASES_ERROR	%s: Reading %s/%s %s
RMON_DSNOOPV6_READ_LEASES_SUSP	%s: Ceasing remote server lease file read status logs for %s
RMON_DSNOOPV6_WRITE_LEASES_ERROR	%s: Writing %s/%s %s
RMON_DSNOOPV6_WRITE_LEASES_SUSP	%s: Ceasing remote server lease file write status logs for %s
RMON_DSNOOPV6_TABLE_FULL_REM_LEASE	%s: The dynamic binding for %s on port %s was replaced with a manual binding.
RMON_DSNOOPV6_TABLE_FULL_REM_LEASE_SUSP	%s: Ceasing removed lease logs for %s.

Table Continued

Event	Message
RMON_DSNOOPV6_BAD_IP_REQ	%s: Illegal IPv6 request from %02X%02X%02X-%02X%02X %02X on port %s; %s.
RMON_DSNOOPV6_BAD_IP_REQ_SUSP	%s: s: Ceasing the log messages for illegal IPv6 requests for %s
RMON_DSNOOPV6_BAD_IP_OFFER	%s: Offered lease from %s conflicts other leases in BST. %s.
RMON_DSNOOPV6_BAD_IP_OFFER_SUSP	%s: Ceasing the log messages for duplicate IPv6 offers for %s.
RMON_DSNOOPV6_ILLEGAL_LEASE	%s: Dropped the IPv6 offer from %s because the offered address is illegal. %s.
RMON_DSNOOPV6_ILLEGAL_LEASE_SUSP	%s: Ceasing the log messages for illegal IPv6 offers for %s
RMON_DSNOOPV6_INVALID_PACKET	%s: Invalid DHCPv6 packet %s. %s.
RMON_DSNOOPV6_INVALID_PACKET_SUSP	%s: Ceasing the log messages for invalid DHCPv6 packets for %s
RMON_DSNOOPV6_DIPLDV6_PORT_REMOVED_VLAN	Port %s removed from dhcpv6-snooping enabled vlan %d
RMON_DIPLDV6_DSNOOPV6_DISABLED_GLOBAL	Dhcpv6-snooping disabled globally, dynamic Ipv6 lockdown also disabled.
RMON_DIPLDV6_DSNOOPv6_DISABLED_VLAN	Dhcpv6-snooping disabled on vlan %d, dynamic Ipv6 lockdown also disabled.
RMON_DSNOOPV6_PORT_TRUSTED_TO_VALIDATING	The port %s is configured as an untrusted port.
RMON_DSNOOPV6_PORT_ADD_TO_TRUNK_ERROR	Unable to add port %s to trunk, insufficient HW resources.
RMON_DIPLDV6_PORT_ADD_HW_RESOURCE_ERROR	Unable to apply dynamic Ipv6 lockdown to port %s, insufficient HW resources.

Table Continued

Event	Message
RMON_DIPLDV6_ADD_BINDING_OUT_OF_RESOURCES	Unable to add binding for %x, %02x%02x%02x-%02x%02x %02x on port %s.
RMON_DIPLDV6_VLAN_DENY_OUT_OF_RESOURCES	Unable to ipv6 lock-down VLAN %d on port %s, not enough HW resources.
RMON_DIPLDV6_VIOLATION	Access denied %s -> %s port %s, %d packets received since last log
RMON_DIPLDV6_DHCPV6_REQUEST_DROPPED	DHCPV6 REQUEST dropped for %02x%02x%02x-%02x %02x%02x port %s, unable to add the binding; a port or switch limit was reached.
RMON_DIPLDV6_VIOLATION_ON_VLAN	Access was denied on VLAN %d, %d packets received since last log.
RMON_DSNOOPV6_CONFLICT_IN_BST	%s: The IPv6 address %s provided by the DHCPv6 server to the client %s is already assigned to another client %s.
RMON_DSNOOPV6_CONFLICT_IN_BST_SUSP	%s: Ceasing status logs for Conflicts in BST for %s

DHCPv6 event messages

Cause

Events	Debug messages
When the BST becomes full, to indicate that lease bindings are being dropped.	Unable to add binding for %x, %02x%02x%02x-%02x %02x%02x on port %s. BST is full.
When DHCPv6 packet validation fails (packets are received on which they are not expected to).	Dropping packet as validation failed, reason %s
When a Dynamic binding is replaced with a static binding on a particular port.	The dynamic binding for %s on port %s was replaced with a manual binding.
While an attempt to release an Ipv6 address from port which is leased to another port.	Unable to release. %02x%02x%02x-%02x%02x%02x is bound not bound to port %u.

Table Continued

Events	Debug messages
Decline from client to server when client finds the address issued by server is already in use in the link where the client is connected.	%s: Decline offer from %x (server)of %x because the address is already assigned to another client.
Decline from client when server assigns an illegal Ipv6 address.	%s: Decline offer from %x (server) of %x because the address is illegal.
When TFTP transfer of binding state table is a success or failure.	TFTP of BST from the dsnoopv6 device is successful / failed.
When a DIPLDv6 enabled port is removed from a DsnoopV6 enabled vlan.	Port %u is removed from a dhcpv6-snooped VLAN
When DsnoopV6 is disabled globally which makes DIPLDv6 no longer configured?	Dhcpv6-snooping disabled globally, dynamic Ipv6 lockdown also disabled.
When DsnoopV6 is disabled on a particular VLAN which makes DIPLDv6 also disabled	Dhcpv6-snooping disabled on VLAN %s, dynamic Ipv6 lockdown also disabled.
When a port moved from SAVI-Trust to validating port.	Port %u is validating.
While adding a port to a trunk for which DIPLDv6 is already enabled.	Unable to add port %u to trunk, dynamic ipv6 lockdown is enabled on it.
While enabling DIPLDv6 on a port which is added to a trunk.	Unable to configure dynamic Ipv6 lockdown on port %u which is a part of a trunk.
While enabling DIPLDv6 on a port for which ACL is configured.	Unable to configure dynamic Ipv6 lockdown on port %u, ACL is configured on port.
When it is unable to add a lock on a particular VLAN for a particular port due to vlan deny rule.	Access was denied on VLAN %d, deny rule exists on the VLAN
When DIPLDv6 violations are detected on a VLAN	Access was denied on VLAN %d, %d packets received since last log.
When max-binding limit is reached on a Port	Max-binding limit reached on Port %s.

Chapter 13 Captive portal for ClearPass

The Captive Portal feature allows the support of the ClearPass Policy Manager (CPPM) into the ArubaOS-Switch product line. The switch provides configuration to allow you to enable or disable the Captive Portal feature. By default, Captive Portal is disabled to avoid impacting existing installations as this feature is mutually exclusive with the following web-based authentication mechanisms: Web Authentication, EWA, MAFR, and BYOD Redirect.

Captive Portal is user-based, rather than port or VLAN-based, therefore the configuration is on a switch global basis. ArubaOS-Switch supports the following authentication types on the switch with RADIUS for Captive Portal:

- Media Access Control (MAC)
- 802.1X

Once you enable Captive Portal, the redirect functionality is triggered only if a redirect URL attribute is provided as part of the RADIUS Access-Accept response from an authentication request of type 802.1X or MAC. The redirect enables the client to self-register or directly login with valid credentials via the CPPM. Upon subsequent re-authentication, it provides access to the network per the CPPM configured policies that are communicated via the RADIUS attributes.

The redirect feature offers:

- Client self-registration
- Client direct login with valid credentials via CPPM Captive Portal
- On-boarding
- Ability to quarantine devices to remedy their status

Requirements

- HTTPS support requires a certificate to be configured on the switch with a usage type of `all` or `captive-portal`.
- If you are running HPE 5400 Series v2 modules, you must turn off the compatibility mode with the following command:

  ```
  switch(config)# no allow-v1-modules
  ```

  This will ensure that the switch will only power up with the v2 modules.

Best Practices

- Use the Port Bounce VSA via a CoA message, instead of the Disconnect message, to cause the second RADIUS authentication to occur during the Captive Portal exchange. This is the more reliable method for forcing a re-DHCP for the client.
- Configure Captive Portal such that the first `ACCESS_ACCEPT` returns a rate limit VSA to reduce the risk of DoS attacks. This configuration enables rate limiting for the HTTP/HTTPS ACL for traffic sent to CPPM.
- Do not use the keyword `cpy` in any other `NAS-Filter-Rules`. The keyword `cpy` in the enforcement profile attributes is specific to CPPM use. It is only supported with the `deny` attribute. If you configure the `cpy` keyword to `permit`, no ACL will be applied.

Limitations

- Captive Portal will not work with RADIUS configured on a loopback port or on the Out-of-Band Management (OOBM) port.
- Captive Portal is supported in CPPM versions 6.5.5 and later. However, by manually modifying the RADIUS dictionary files, any CPPM version 6.5.* can be used.
- Captive Portal does not support v1 modules, and will not work unless compatibility mode is turned off.
- Captive Portal does not support IPv6.
- Simultaneous Captive Portal client connections: maximum of 512
- Captive Portal does not support web proxy. The permit CPPM ACLs and the steal ACLs only use port 80 and 443. Non-standard ports for HTTP and HTTPS are not supported.
- Captive Portal is mutually exclusive with the following web-based authentication mechanisms: Web Authentication, EWA, MAFR, and BYOD.
- URL-string limitation of 253 characters.

Features

High Availability

Captive Portal includes support for High Availability (HA). The Captive Portal configurations (such as enablement, authenticated clients, and redirect URLs) are replicated to standby or other members.

If the feature is enabled and a failover occurs, clients in the process of onboarding are still redirected to Captive Portal, and authenticated clients continue to have the same access to the network.

Clients that are in the process of authenticating via MAC or 802.1X authentication will not be replicated to the standby. Replication of client data is only done when MAC or 802.1X authentication has resulted in a successful authentication.

Load balancing and redundancy

The following options are available to create load balancing and provide redundancy for CPPM:

- Virtual IP use for a CPPM server member
- CPPM servers configured in the switch RADIUS server group
- External load balancer

Captive Portal when disabled

By default, Captive Portal is disabled. If the Captive Portal feature is disabled and the switch receives a redirect URL attribute from the RADIUS server as part of the Access-Accept, it will view the redirect as an error. The authentication success will be overridden, the session will be flushed, and the switch will send the Accounting Start and Accounting Stop messages to indicate the client is no longer authenticated.

The Captive Portal feature may be disabled while there are in flight authentication requests. These are authentication sessions that have not finished the final authentication with the switch. The switch flushes all sessions with a redirect URL associated with them when Captive Portal is disabled.

Fully authenticated sessions are not impacted when Captive Portal is disabled. If CPPM deems these sessions to be invalid, a RADIUS Disconnect can be sent to flush all these sessions.

Disabling Captive Portal

Use one of the following commands:

- **`aaa authentication captive-portal disable`**
- **`no aaa authentication captive-portal enable`**

Configuring Captive Portal on CPPM

Procedure

1. **Import the HP RADIUS dictionary**
2. **Create enforcement profiles**
3. **Create a ClearPass guest self-registration**
4. **Configure the login delay**

Importing the HP RADIUS dictionary

For CPPM versions 6.5.*, you must update the HP RADIUS dictionary. To import the dictionary in CPPM, follow these steps:

Procedure

1. Go to **Administration** -> **Dictionaries** -> **RADIUS** and click **Import**.
2. Select the XML HP RADIUS Dictionary from your Hard Drive.
3. Click **Import**.

Creating enforcement profiles

For the HPE Bounce Host-Port profile, configure Captive Portal so that the RADIUS CoA message that includes the Port Bounce VSA is sent to force the second RADIUS re-authentication after the user registers their device and makes it known.

Procedure

1. Create the HPE Bounce Host-Port profile and the Guest Login profile only if they do not already exist.
2. In CPPM, go to **Configuration** -> **Enforcement** -> **Profiles**
3. Click **Add**.
4. Enter the Profile Name: **HPE Bounce Host-Port**
5. Enter the Description: **Custom-defined profile to bounce host port (HPE)**.
6. Select the type **RADIUS_CoA**.
7. Select the action **CoA**.
8. Add all of the attributes required for a CoA message, and specify the port bounce duration (valid values are between 0 and 60). This is the amount of time in seconds the port will be held in the down state. The recommended setting is 12 seconds.

Summary | Profile | Attributes

Profile:

Name:	HPE Bounce Host-Port
Description:	Custom-defined profile to bounce host port (HPE)
Type:	RADIUS_CoA
Action:	CoA
Device Group List:	-

Attributes:

	Type	Name		Value
1.	Radius:IETF	User-Name	=	%{Radius:IETF:User-Name}
2.	Radius:IETF	Calling-Station-Id	=	%{Radius:IETF:Calling-Station-Id}
3.	Radius:IETF	NAS-Port	=	%{Radius:IETF:NAS-Port}
4.	Radius:IETF	NAS-IP-Address	=	%{Radius:IETF:NAS-IP-Address}
5.	Radius:IETF	Event-Timestamp	=	%{Radius:IETF:Event-Timestamp}
6.	Radius:HPE	HPE-Port-Bounce-Host	=	12

9. Repeat **Step 2** to **Step 6** to configure the Guest Login profile that will be sent as part of the first RADIUS Access-Accept and enforce the redirect to the Captive Portal on CPPM. For this profile, select **RADIUS** as the type and **Accept** as the action.
10. Add all of the NAS-Filter-Rule attributes specified below, replacing the IP address in the first two NAS-Filter-Rule attributes with your CPPM address. Add the HPE-Captive-Portal-URL attribute to specify the redirect URL, replacing the IP address with your CPPM address. This will cause the client to be redirected to the Captive Portal on CPPM. You can add other attributes, such as a VLAN to isolate onboarding clients, or a rate limit to help prevent DoS attacks.

The `HPE-Captive-Portal-URL` value must be a URL normalized string. The scheme and host must be in lower case, for example `http://www.example.com/`.

Summary | Profile | Attributes

Profile:

Name:	HPE Wired Guest Login
Description:	
Type:	RADIUS
Action:	Accept
Device Group List:	-

Attributes:

	Type	Name		Value
1.	Radius:IETF	Tunnel-Type	=	VLAN (13)
2.	Radius:IETF	Tunnel-Medium-Type	=	IEEE-802 (6)
3.	Radius:IETF	Tunnel-Private-Group-Id	=	100
4.	Radius:HPE	HPE-Captive-Portal-URL	=	http://10.73.4.136/guest/aruba_guest.php
5.	Radius:IETF	NAS-Filter-Rule	=	permit in tcp from any to 10.73.4.136 80
6.	Radius:IETF	NAS-Filter-Rule	=	permit in tcp from any to 10.73.4.136 443
7.	Radius:IETF	NAS-Filter-Rule	=	deny in tcp from any to any 80 cpy
8.	Radius:IETF	NAS-Filter-Rule	=	deny in tcp from any to any 443 cpy
9.	Radius:IETF	NAS-Filter-Rule	=	permit in udp from any to any 53
10.	Radius:IETF	NAS-Filter-Rule	=	permit in udp from any to any 67

Creating a ClearPass guest self-registration

Procedure

1. From the Customize Guest Registration window, select **Server-initiated** as the Login Method.
2. Optionally, under Security Hash, select the level of checking to apply to the redirect URL.

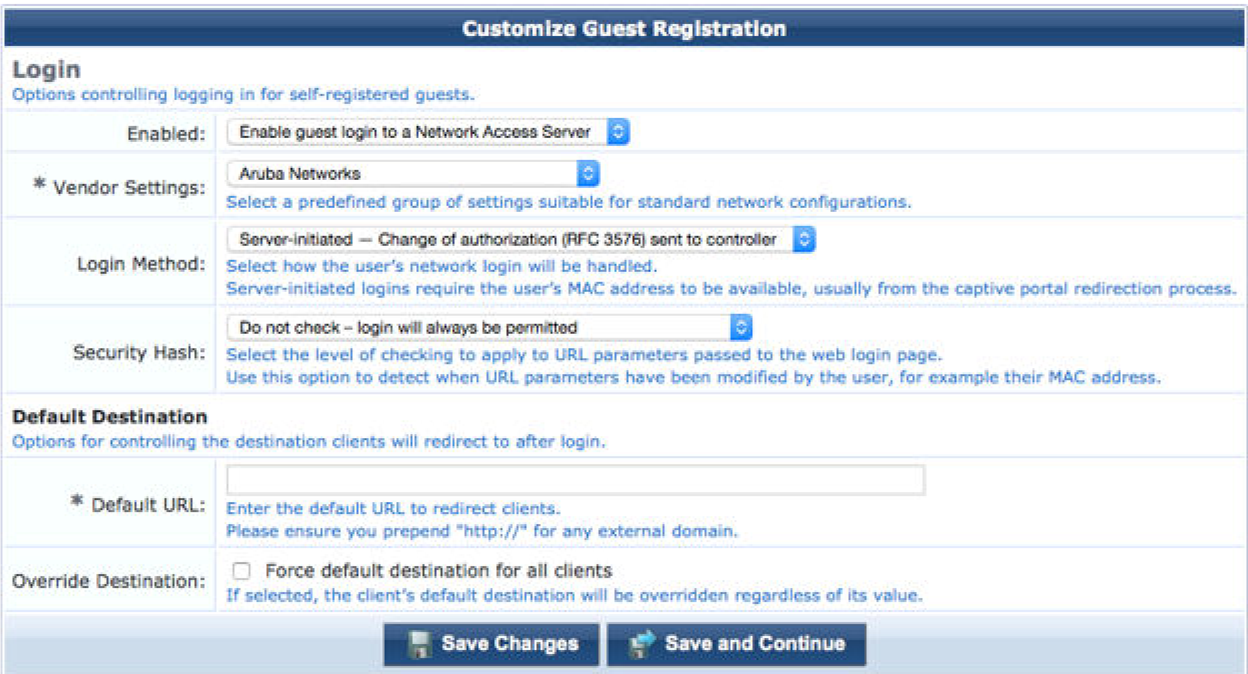

Configuring the login delay

Enter the **Login Delay** value. The value must be greater than the `HPE-Port-Bounce-Host` attribute. In this example, we set the login delay value to 20 seconds.

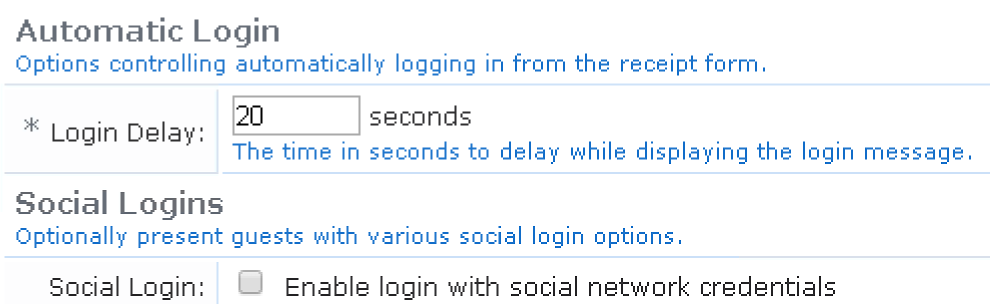

Configuring the switch

Once you have configured Captive Portal, you can configure the switch. To configure the switch, you must first configure the switch as a RADIUS client, then configure the ports that will be used for Captive Portal, as follows:

Procedure

1. Configure the switch as a RADIUS client. In this example, the CPPM IP address is `10.73.4.136` and `secret` is the secret key shared with the RADIUS server:
   a. `switch(config)# radius-server host 10.73.4.136 key "secret"`
   b. `switch(config)# radius-server host 10.73.4.136 dyn-authorization`
   c. `switch(config)# radius-server host 10.73.4.136 time-window 0`

   Make sure to set your time-window to `0`. See **Event Timestamp not working** on page 376.

2. Configure the ports that will be used for Captive Portal. In this example, the commands enable ports `B3-B5` for MAC Authentication:

a. `switch(config)# aaa authentication port-access chap-radius`
b. `switch(config)# aaa port-access mac-based B3-B5`

3. If you configured the `Security Hash` to `Deny login on validation error` in **Creating a ClearPass guest self-registration** on page 373, configure the URL key.

   See **The URL key** on page 375.
4. Configure the certificate. See **Configuring a certificate for Captive Portal usage** on page 375
5. Enable Captive portal:

   ```
   switch(config)# aaa authentication captive-portal enable
   ```

   Captive Portal defaults to disabled. Once enabled, you are redirected to the URL supplied via the HPE-Captive-Portal-URL VSA. Captive Portal is enabled on a global/switch wide basis.

The URL key

You can optionally configure a URL hash key to provide some security for the Captive Portal exchange with CPPM. The key is a shared secret between CPPM and the switch. When configured, the switch generates a HMAC-SHA1 hash of the entire redirect URL, and appends the hash to the URL to be sent to CPPM as part of the HTTP redirect. If CPPM is configured to check the hash, it will generate the hash of the URL using its version of the URL hash key and compare against the value provided by the switch. The action taken by CPPM upon a match or mismatch is determined by what is configured on CPPM.

CPPM provides the following options:

- Do not check - login will always be permitted
- Deny login on validation error - login will not be permitted

The URL hash key is globally configured and will be used for all redirects to Captive Portal. This key is not configured on a per CPPM or RADIUS server basis. If the key is not specified, the hash is not added to the URL. The URL hash key is an ASCII string with a maximum length of 64 characters.

The URL key supports the FIPS certification feature encrypt-credentials and can optionally be encrypted for more robust security. This option is only available when the global encrypt-credentials is enabled.

Configuring the captive portal URL key

Procedure

1. To configure a plain text captive-portal URL key:

   ```
   switch(config)# aaa authentication captive-portal url-hash-key plaintext <KEY>
   ```

2. To configure an encrypted captive-portal URL key when encrypt-credentials is enabled:

   ```
   switch(config)# aaa authentication captive-portal url-hash-key encrypted
   <ENCRYPTED-KEY>
   ```

3. To clear a captive-portal URL key:

   ```
   switch(config)# no aaa authentication captive-portal url-hash-key
   ```

Configuring a certificate for Captive Portal usage

HTTPS support requires the use of a certificate. If a certificate for Captive Portal does not exist, the certificate designated for all use is used instead.

Procedure

1. To create a certificate signing request for Captive Portal, enter:

```
switch(config)# crypto pki create-csr certificate-name <cert-name> usage captive-
portal
```

2. To create a self-signed certificate for Captive Portal, enter:

```
switch(config)# crypto pki enroll-self-signed certificate-name
```

Displaying the Captive Portal configuration

To display the Captive Portal configuration settings, enter the `show captive-portal` command:

```
switch(config)# show captive-portal

Captive Portal Configuration
 Redirection Enabled    : Yes
 URL Hash Key Configured : No
```

Showing certificate information

To view the certificate information, enter:

```
switch(config)# show crypto pki local-certificate
```

Name	Usage	Expiration	Parent / Profile
cp	Captive Portal	2016/08/14	default

Troubleshooting

Event Timestamp not working

Symptom

The client gets a credentials request on the web browser even though the valid credentials were already provided, or the client is not redirected to the Captive Portal.

Cause

- ClearPass 6.5.x does not support the sending of `Event Timestamp` in automated workflows (manual via Access Tracker works).
- The switch will reject CoA requests when the time on CPPM is ahead of the switch time by even a second.

Action

1. Set the time-window security feature in PVOS to 0:

```
radius-server host<CLEARPASS-IP> time-window 0
```

Cannot enable Captive Portal

Symptom

When running the `aaa authentication captive-portal enable` command, getting the following error message:

```
Captive portal cannot be enabled when BYOD redirect, MAC authentication failure
redirect, or web-based authentication are enabled.
```

Cause

The failure is due to a mutual exclusion restriction.

Action

1. Check which one of the following are enabled: BYOD redirect, MAC authentication failure redirect, or web-based authentication.
2. Disabled the enabled authentication method found in step 1.
3. Run the `aaa authentication captive-portal enable` command.

Unable to enable feature

Symptom

One of the following messages is displayed:

- `BYOD redirect cannot be enabled when captive portal is enabled.`
- `MAC authentication failure redirect cannot be enabled when captive portal is enabled.`
- `Web-based authentication cannot be enabled when captive portal is enabled.`
- `V1 compatibility mode cannot be enabled when captive portal is enabled.`

Cause

You cannot enable these features when Captive Portal is already enabled. They are mutually exclusive.

Action

1. You can either disable Captive Portal or avoid enabling these features.

Authenticated user redirected to login page

Symptom

User is redirected back to the login page to submit credentials even after getting fully authenticated.

Solution 1

Cause

The status is not changed to Known.

Action

1. After the client submits the credentials, the CPPM service must change the Endpoint Status to `Known`.

Solution 2

Cause

The cache value is set.

Action

1. Clear the CPPM Cache Timeout of the Endpoint Repository.

Unable to configure a URL hash key

Symptom

The following message is displayed:

```
Key exceeds the maximum length of 64 characters.
```

Cause

The URL hash key is not valid.

Action

1. Select a key that is 64 or less ASCII text. For example:

   ```
   switch(config)# aaa authentication captive-portal url-hash-key plaintext
   "8011A89FEAE0234BCCA"
   ```

ClearPass captive portal authentication commands

aaa authentication captive-portal

Syntax

`aaa authentication captive-portal`

Description

Configure ClearPass Captive Portal.

Parameters and options

enable

Enables redirection to a Captive Portal server for additional client authentication.

disable

Disables redirection to a Captive Portal server for additional client authentication (aaa authentication captive-portal disable).

no

Disables redirection to a Captive Portal server for additional client authentication (no aaa authentication captive-portal enable).

url-hash-key

Configures a hash key used to verify the integrity of the portal URL.

Clearpass captive portal show commands

show config

Syntax

`show [config|running-config|ip|captive-portal]`

Description

Show configuration information.

Parameters and options

config

Displays the saved configuration

running-config

Displays the running configuration

ip

Displays the switch IP addresses.

captive-portal

Displays the captive portal configuration.

show port-access clients

Syntax

`show port-access clients [port][detailed]`

Description

Displays the consolidated client view.;

For the summary view (without the detailed option), only the untagged VLAN is displayed.

Parameters and options

detailed

Displays the applied access policy. Only displays the IP address `dhcp-snooping` is enabled.

show radius

Syntax

show radius [authentication|dyn-authorization|accounting]

Description

Parameters and options

authentication

Displays NAS identifier and data on the configured RADIUS server and switch interactions with this server.

dyn-authorization

Statistics for Radius CoA and Disconnect.

accounting

Statistics for Radius accounting.

show crypto pki local-certificate

Syntax

`show crypto pki local-certificate [summary]`

Description

Installed certificates.

Parameters and options

Clearpass captive portal debug command

debug security

Syntax

`debug security [captive-portal] [port-access <mac-based|authenticator>] [radius-server]`

Description

Debug security issues.

Parameters and options

captive-portal

Enables debug logging for the Captive Portal sub-system.

port-access

MAC-based

Enables debug logging for the MAC-based sub-system.

authenticator

Enables debug logging for the 802.1X authenticator sub-system.

radius-server

Enables debug logging for the Radius sub-system.

debug destination

Syntax

`debug destination [session|logging|buffer]`

Description

Debug destination issues.

Parameters and options

session

Prints debug messages to terminal.

logging

Sends debug messages to the syslog server.

buffer

Prints debug messages to a buffer in memory.

Chapter 14 ZTP with AirWave Network Management

Beginning with switch software release 16.01, ZTP with AirWave Network Management is supported on the following switch models covered in this guide:

- 3800 (KA software)
- 3810 (KB software)
- 5400R (KB software)

AirWave is a Network Management Solution (NMS) tool. Once connected to AirWave using the WebUI and CLI interfaces, you can:

- Configure your switches using Zero Touch Provisioning (ZTP)
- Configure your switches using the CLI
- Troubleshoot your switches
- Monitor your switches
- Upgrade your firmware for your switches

Once you have configured your switch, you can monitor, manage, and upgrade your hardware using the AirWave Management Platform.

More information

Switch configuration options on page 382
Stacking and chassis switches on page 396
Troubleshooting on page 396
Aruba Networks and AirWave Switch Configuration Guide

Requirements

- DHCP server
- AirWave NMS
- HPE Aruba switches

Best Practices

- Implement ZTP in a secure and private environment. Any public access may compromise the security of the switch, as follows:
  - Since ZTP is enabled only on the factory default configuration of the switch, DHCP snooping is not enabled. You must manage the Rogue DHCP server.
  - The DHCP offer is in plain data without encryption. Therefore, the offer can be listened by any device on the network and they can in turn obtain the AirWave information.
  - The TLS certificate of the server is not validated by the switch during the HTTPs check-in to AirWave. The AirWave server is in the private environment of the switch.

Limitations

- ZTP is not supported through OOBM.
- The DNS/hostname in option 66 is not supported, only the IPv4 address.
- The switch does not validate peer certificate of the AirWave server as part of the TLS handshake.

- The HTTPS check-in to AirWave does not support HTTPS proxy.
- For non-ZTP cases, the AirWave check-in starts by validating the following condition: Primary or Management VLAN must be configured with the IP address and one of the interface must be `UP`. By default, `VLAN 1` is the primary VLAN.

Switch configuration options

To configure a switch, use one of the following methods:

- **Configure AirWave details in DHCP (preferred method)** on page 382

  .
- If you are using existing HPE switches and using the DHCP server for the configuration or firmware management, you can configure the AirWave details in DHCP using this method: **Configure AirWave details in DHCP (alternate method)** on page 387.
- If you are configuring the switch using a CLI, see **CLI switch configuration** on page 395.
- If you are using ZTP, the configuration is automatic and does not require any user interaction, see **Zero Touch Provisioning** on page 394.

The switch contacts the AirWave server that is configured on the switch and initiates the check-in process.

Once you have configured the DHCP server, the AirWave details received from the DHCP options are stored in the switch configuration. This assures that the configuration is retained across reboots.

Once AirWave completes the switch check-in, it lists the first switch as `New Devices`. The first switch is used to create a new configuration template for the specific group and device type. With this new template, the required configuration is generated for the group. Subsequent switch of the specific type and joining the same group as the first device are added directly to the group and the configuration is pushed using the configuration template via a SSH connection.

Configure AirWave details in DHCP (preferred method)

To configure a DHCP server for AirWave, from a Windows Server 2008, do the following steps:

Procedure

1. From the **Start** menu, select **Server Manager**.

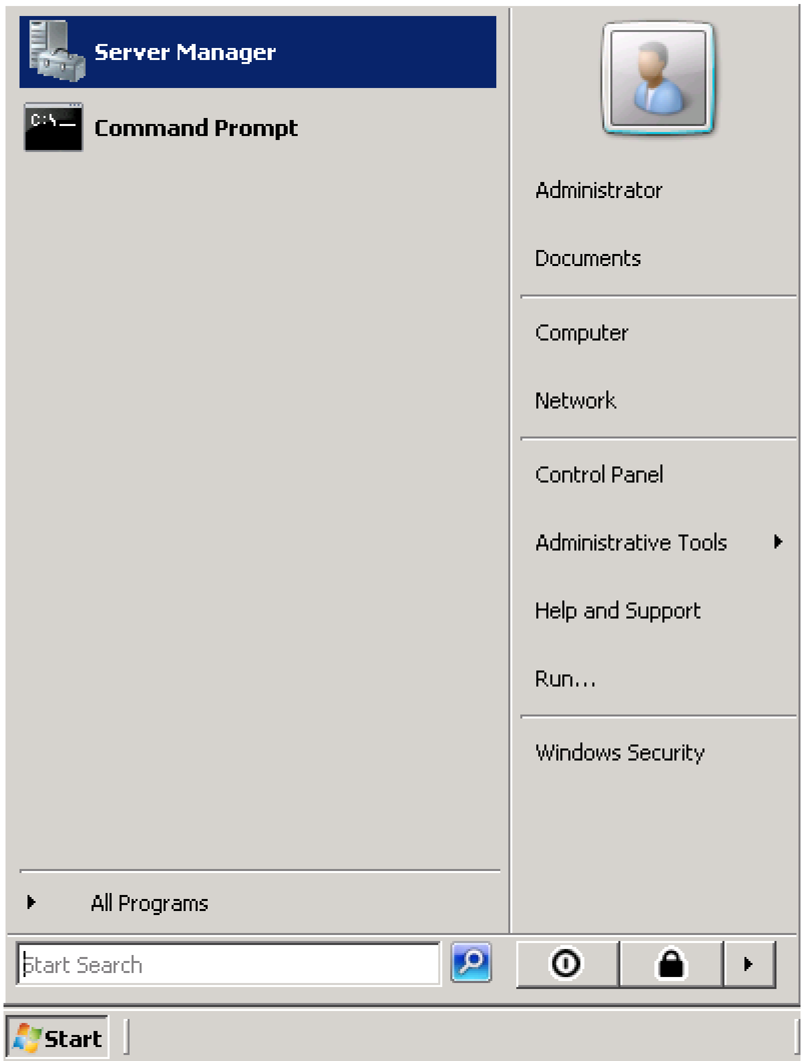

2. Select **Roles -> DHCP -> Server -> w2k8 -> IPv4**.

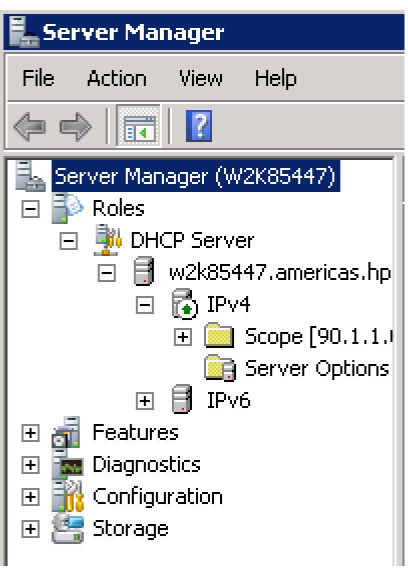

3. Right click on **IPv4** and select **Set Predefined Options...**

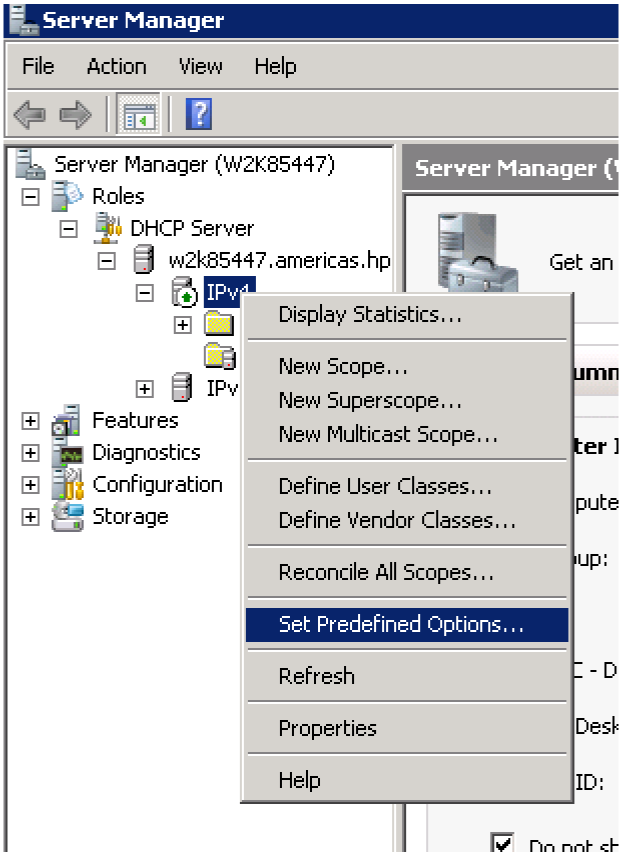

4. The Predefined Options and Values screen is displayed. Click **Add...**.

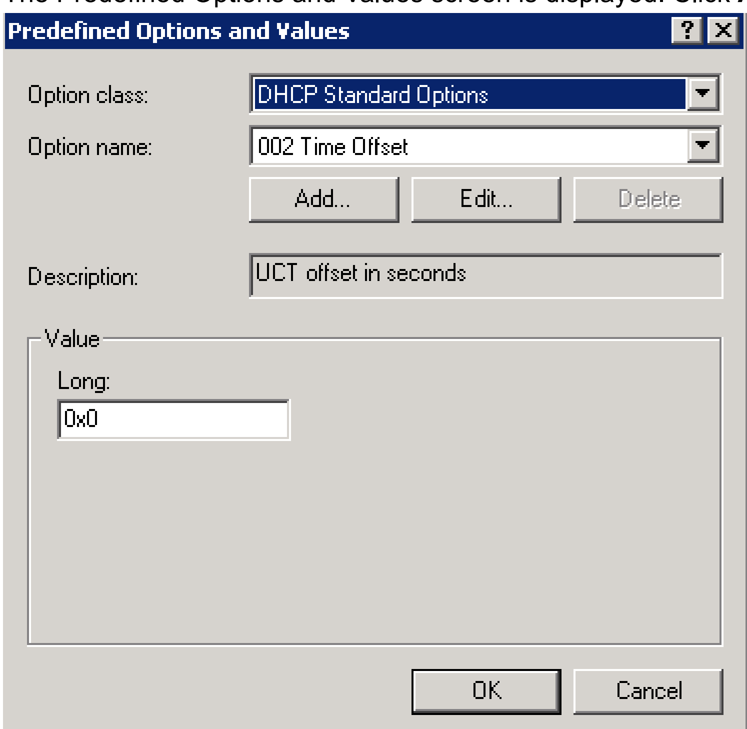

5. Enter the desired **Name** (any), **Data type** (select **String**), **Code** (enter **60**), and **Description** (any).

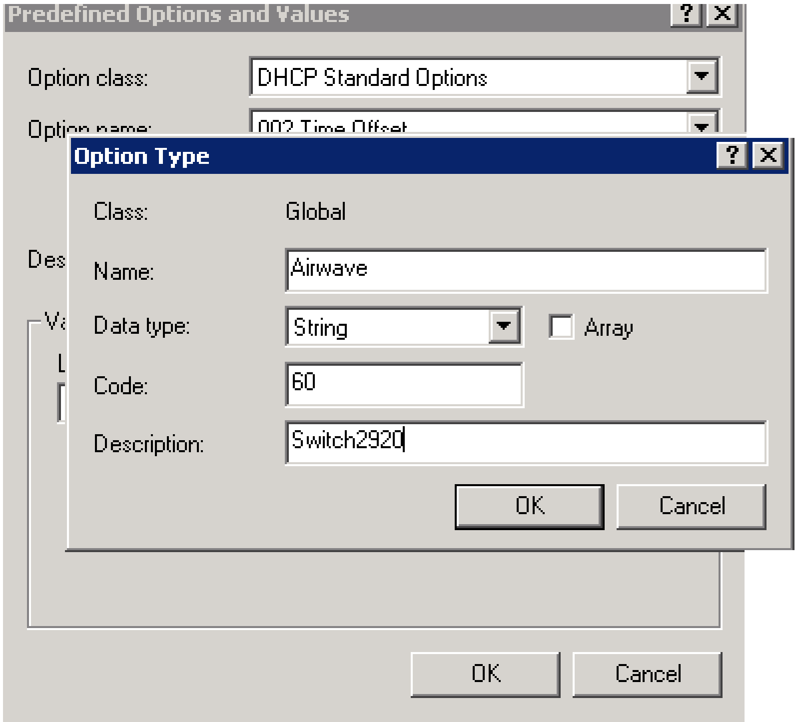

6. Click **OK**.
7. From the Predefined Options and Values screen, under Value, enter the String **ArubaInstantAP**. The string is case sensitive and must be `ArubaInstantAP`.

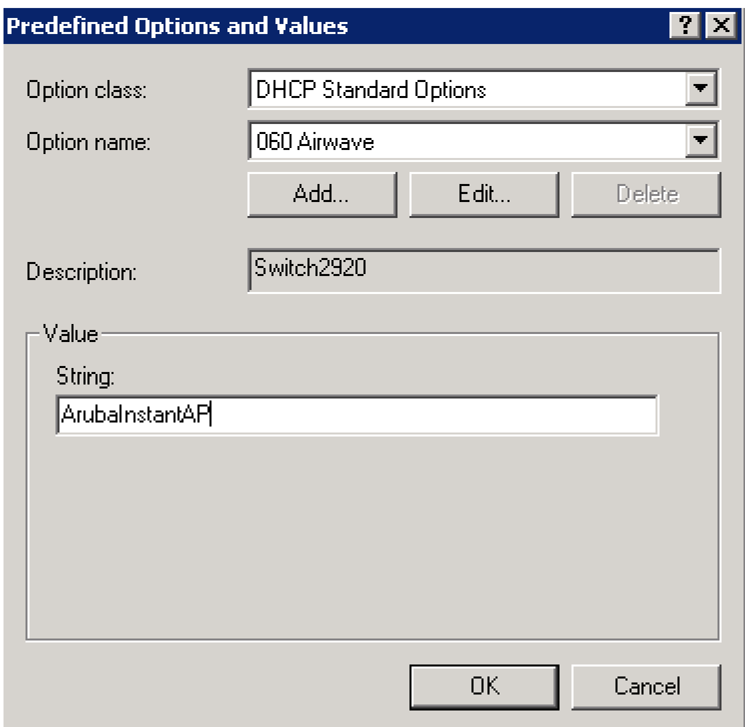

8. Click **OK**.
9. Under IPv4, expand **Scope**. Right click on **Scope Options** and select **Configure Options...**

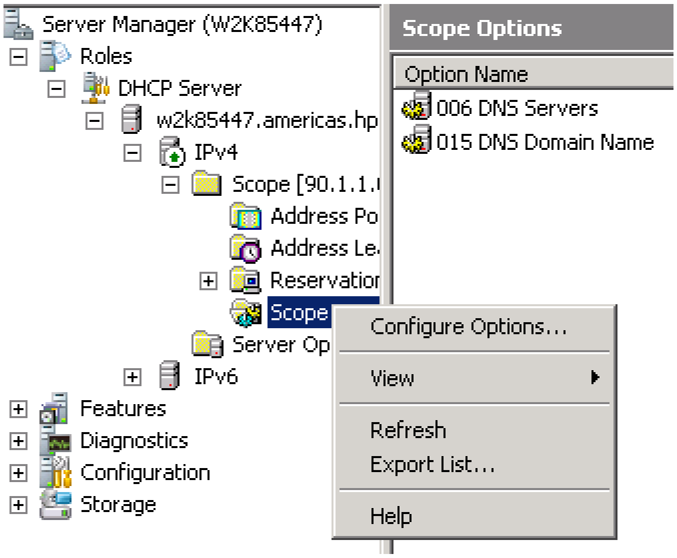

10. Under the General tab, select **043 Vendor Specific Info**. The Data entry data appears. Under ASCII, enter **hpeSwitch:hp2920,90.1.1.10,admin**. The ASCII value has the following format:
`<Group>:<Topfolder>,<AMP IP>,<shared secret>`
11. If you need to add sub-folders, use the following format:`<Group>:<Topfolder>:<folder1>,<AMP IP>,<shared secret>`

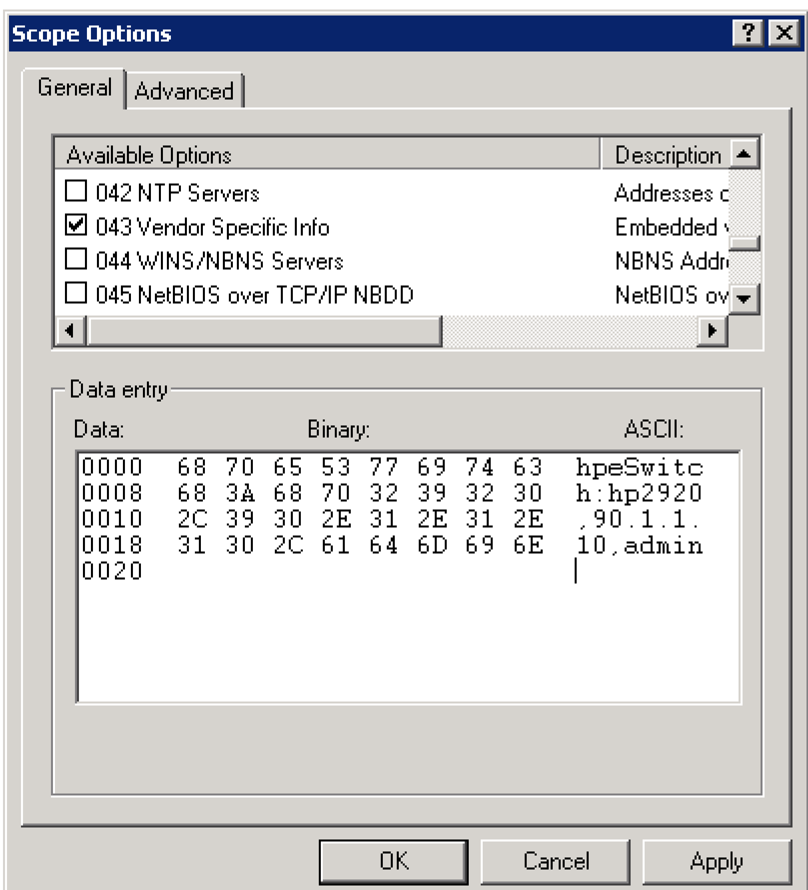

12. Under the General tab, select **060 AirWave**. Click **OK**.

NOTE

No changes are required to the 060 option.

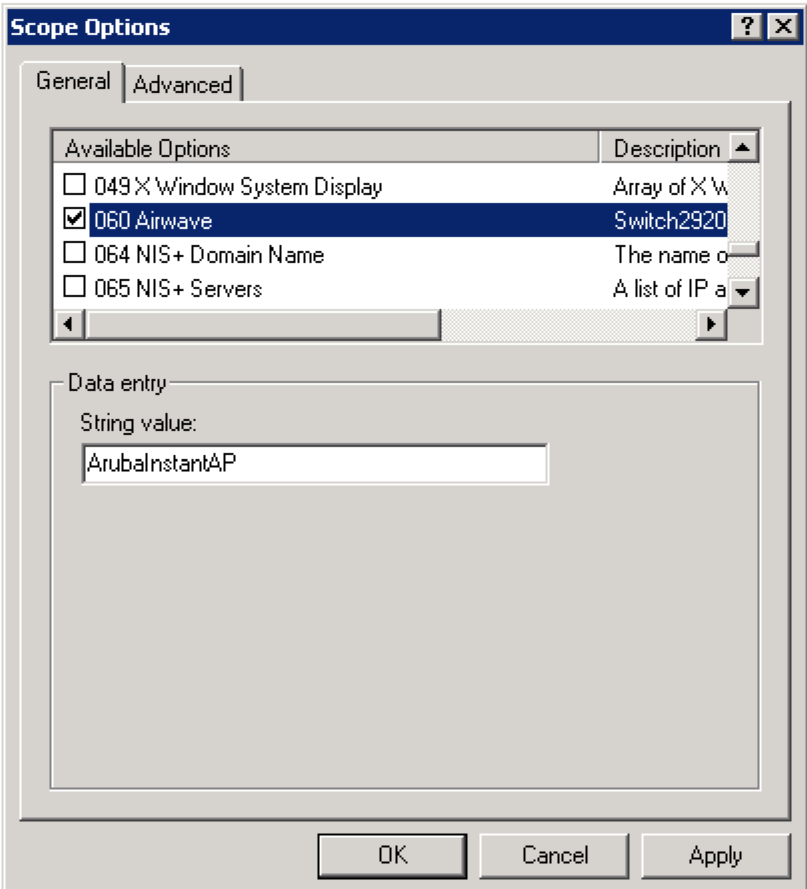

13. You can verify the AirWave details as follows:

```
switch# show amp-server
switch# show run
```

Configure AirWave details in DHCP (alternate method)

To configure a DHCP server for ZTP and AirWave, from a Windows Server 2008, do the following steps:

NOTE

You must repeat these steps for every type of switch that needs to be configured for ZTP, selecting a different Vendor Class for each type of switch.

Procedure

1. From the **Start** menu, select **Server Manager**.

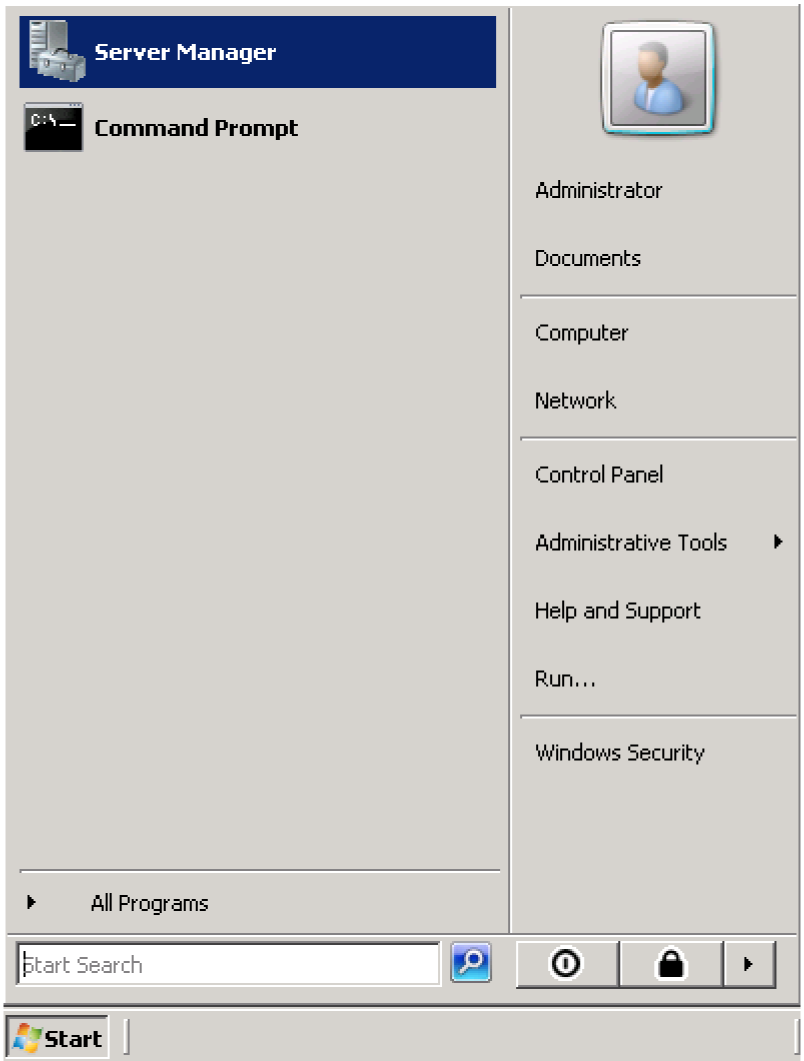

2. Select **Roles -> DHCP -> Server -> w2k8 -> IPv4**.

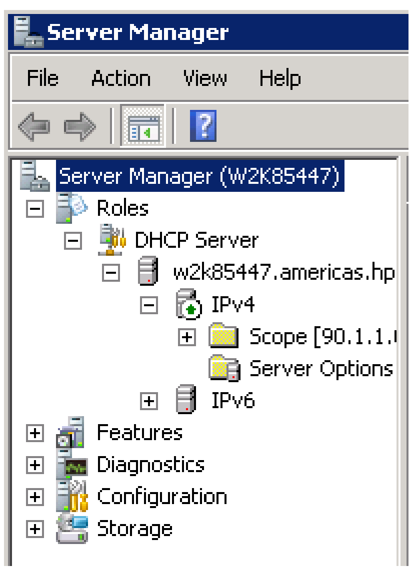

3. Right click on **IPv4** and select **Define Vendor Classes...**

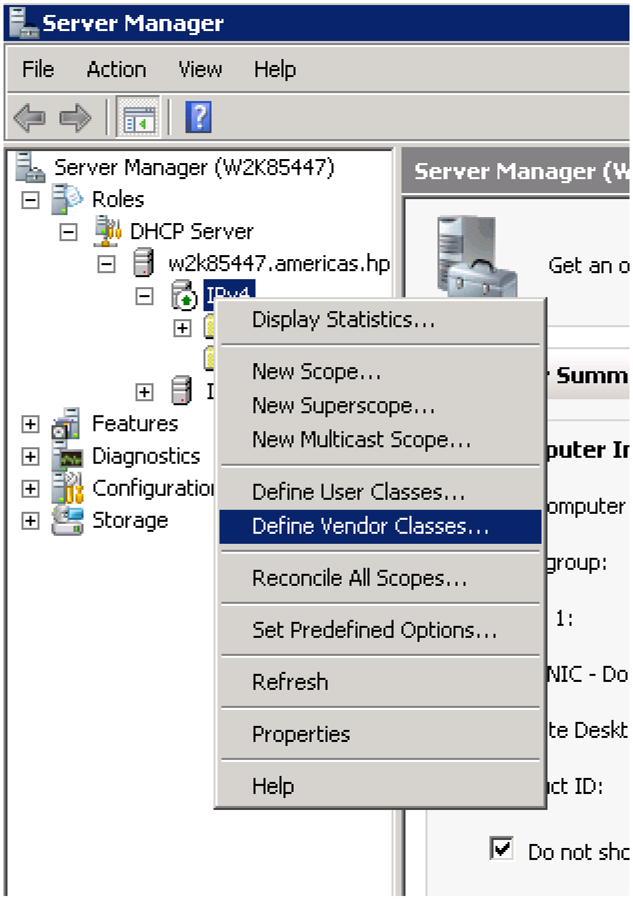

4. The DHCP Vendor Classes window is displayed. Click **Add...**.

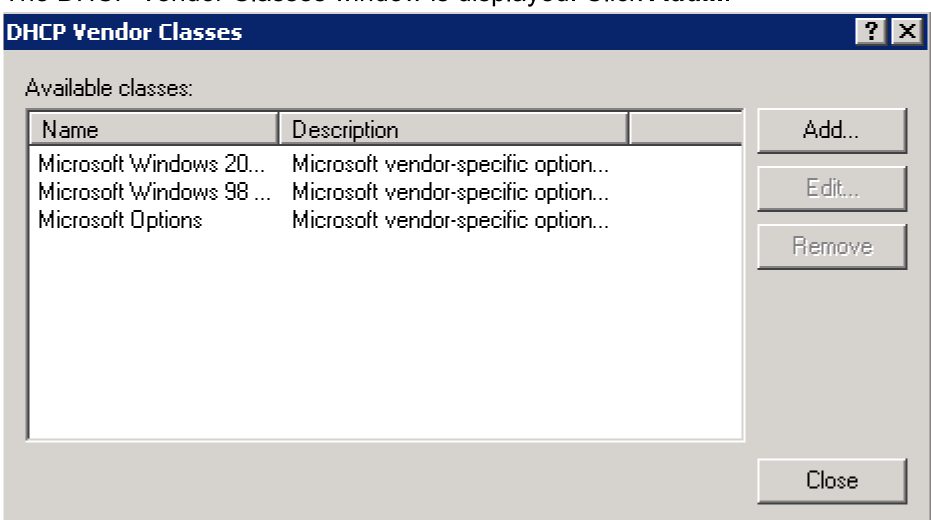

5. To get the vendor-specific value of a switch, go to the switch console and enter:

```
switch# show dhcp client vendor-specific
```

6. In our example, the command returns the following value:Processing of Vendor Specific Configuration is enabled

```
    Vendor Class Id = HP J9729A 2920-24G-PoE+ Switch dslforum.org
```

7. From the New Class window, enter the desired **Display name** (any) and the **Description** (any). For the **ASCII** field, enter the exact value that you got by executing the `show` command performed in the previous step. In this example, **HP J9729A 2920-24G-PoE+ Switch dslforum.org**.

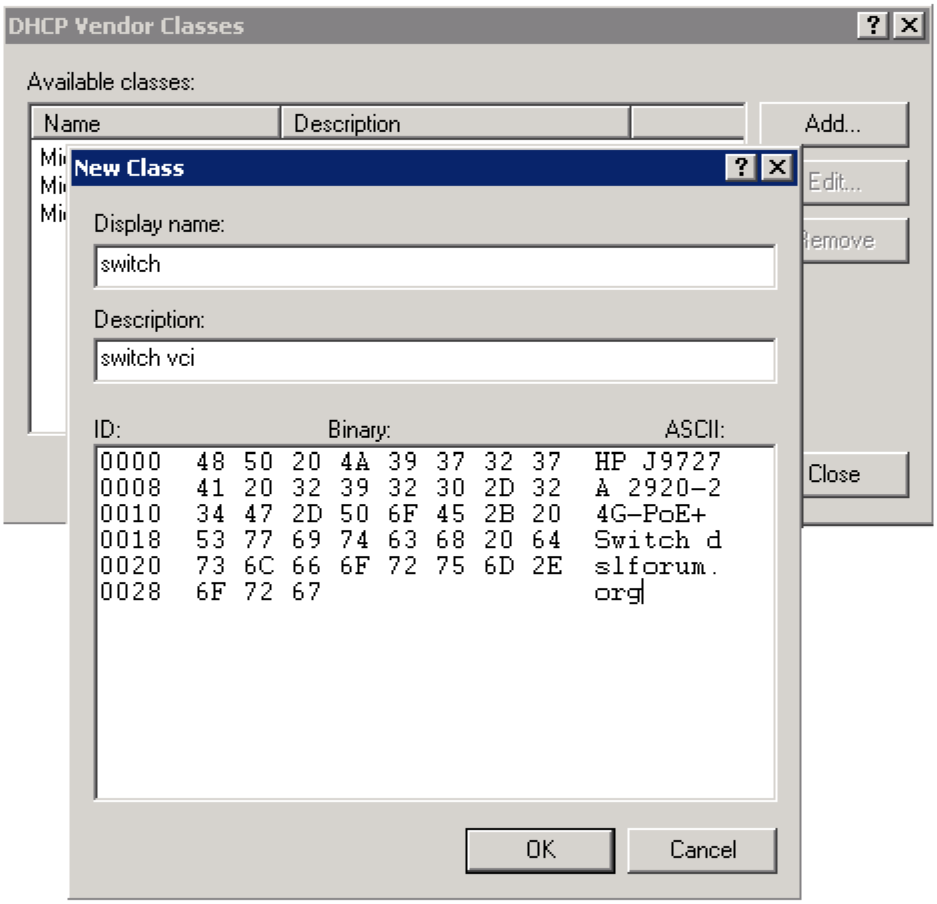

8. Click **OK**.
9. Right click on **IPv4** and select **Set Predefined Options...**.

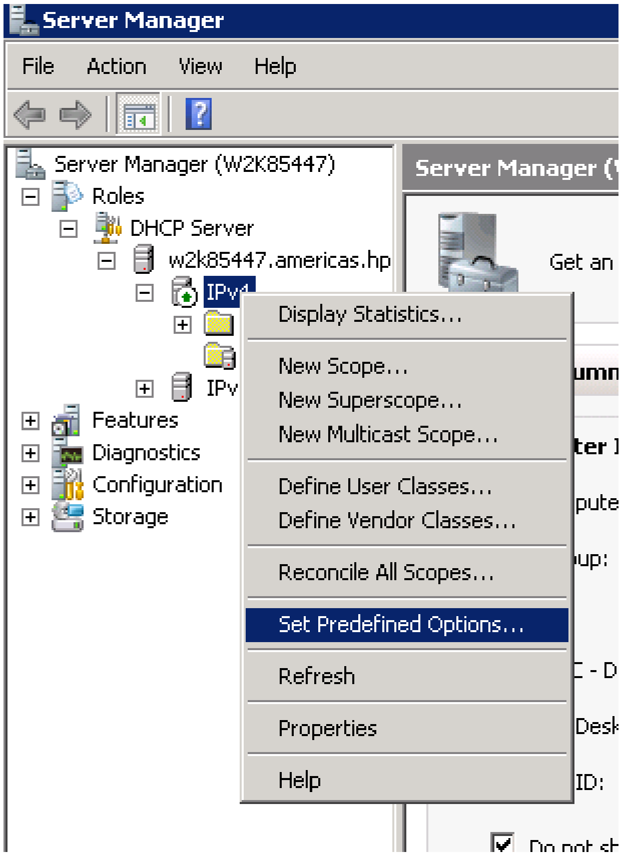

10. From the Predefined Options and Values window, select **Option class**. The Option Class displayed is the one that you configured under **DHCP Vendor Class**. In this example, the Option Class is **switch**.

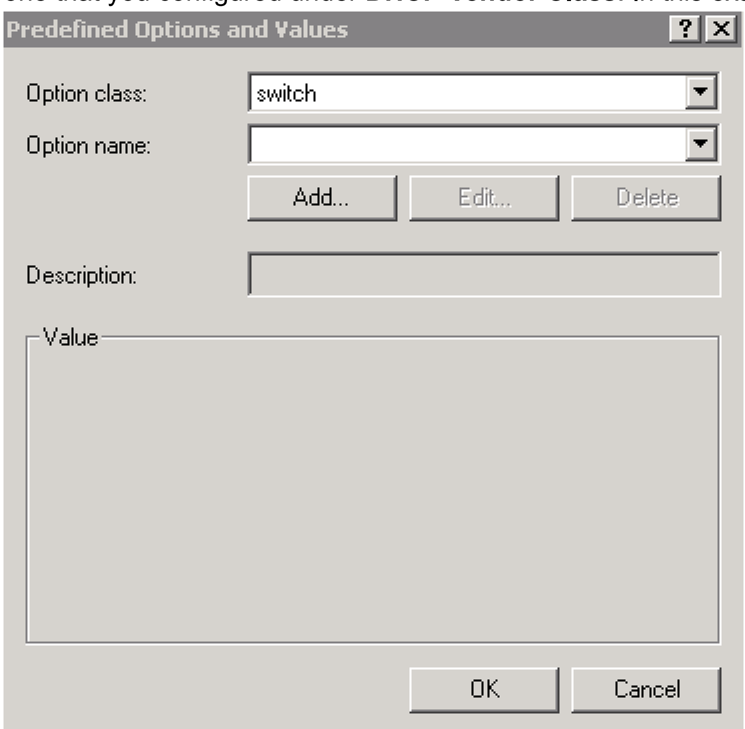

11. Click **Add...**.
12. From the Option Type window, enter the desired **Class** (any), the **Data type** (select **string**), the **Code** (enter **146**), and the **Description** (any).

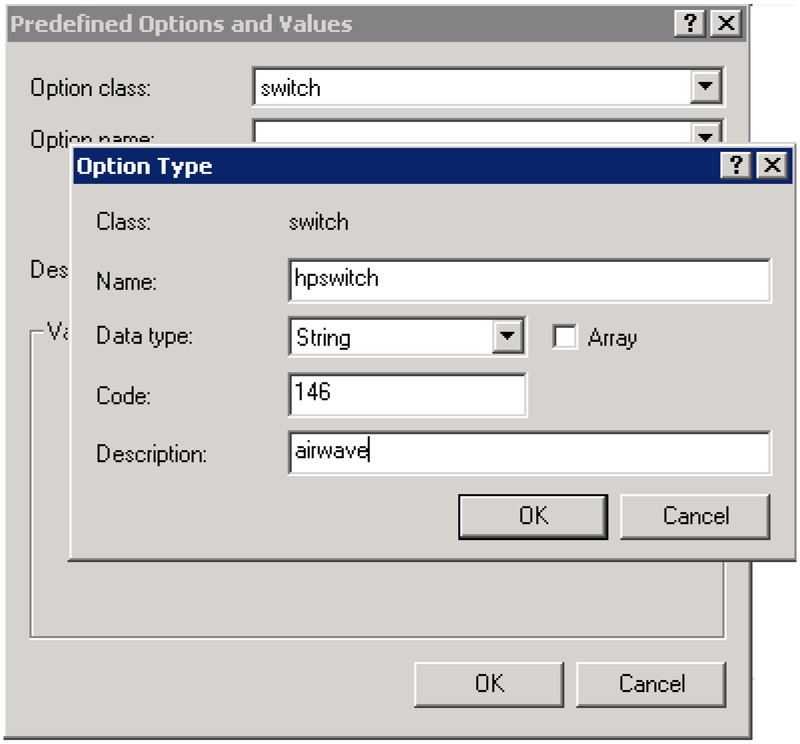

13. Click **OK**.
14. Under the Predefined Options and Values window, enter the Value String. In this example, we enter **hpeSwitch:hp2920,90.1.1.10,admin**. The String has the following format: `<Group>:<Topfolder>,<AMP IP>,<shared secret>`
15. If you need to add sub-folders, use the following format:`<Group>:<Topfolder>:<folder1>,<AMP IP>,<shared secret>`

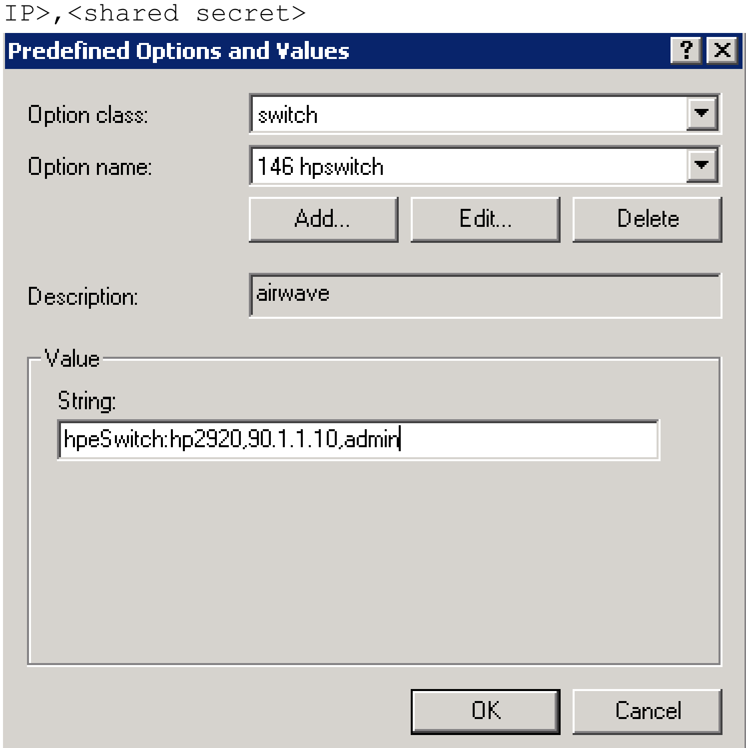

16. Click **OK**.
17. Under **IPv4**, expand **Scope**. Right click on **Scope Options** and select **Configure Options...**

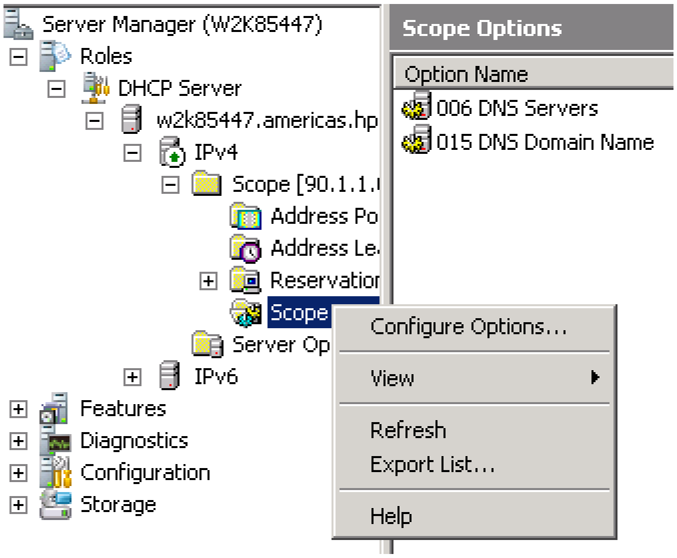

18. From the Scope Options window:

a. Select the **Advanced** tab.

b. Under Vendor class, select the desired switch. In this example, **switch**.

c. Select the **146 hpswitch** option.

d. Click **OK**.

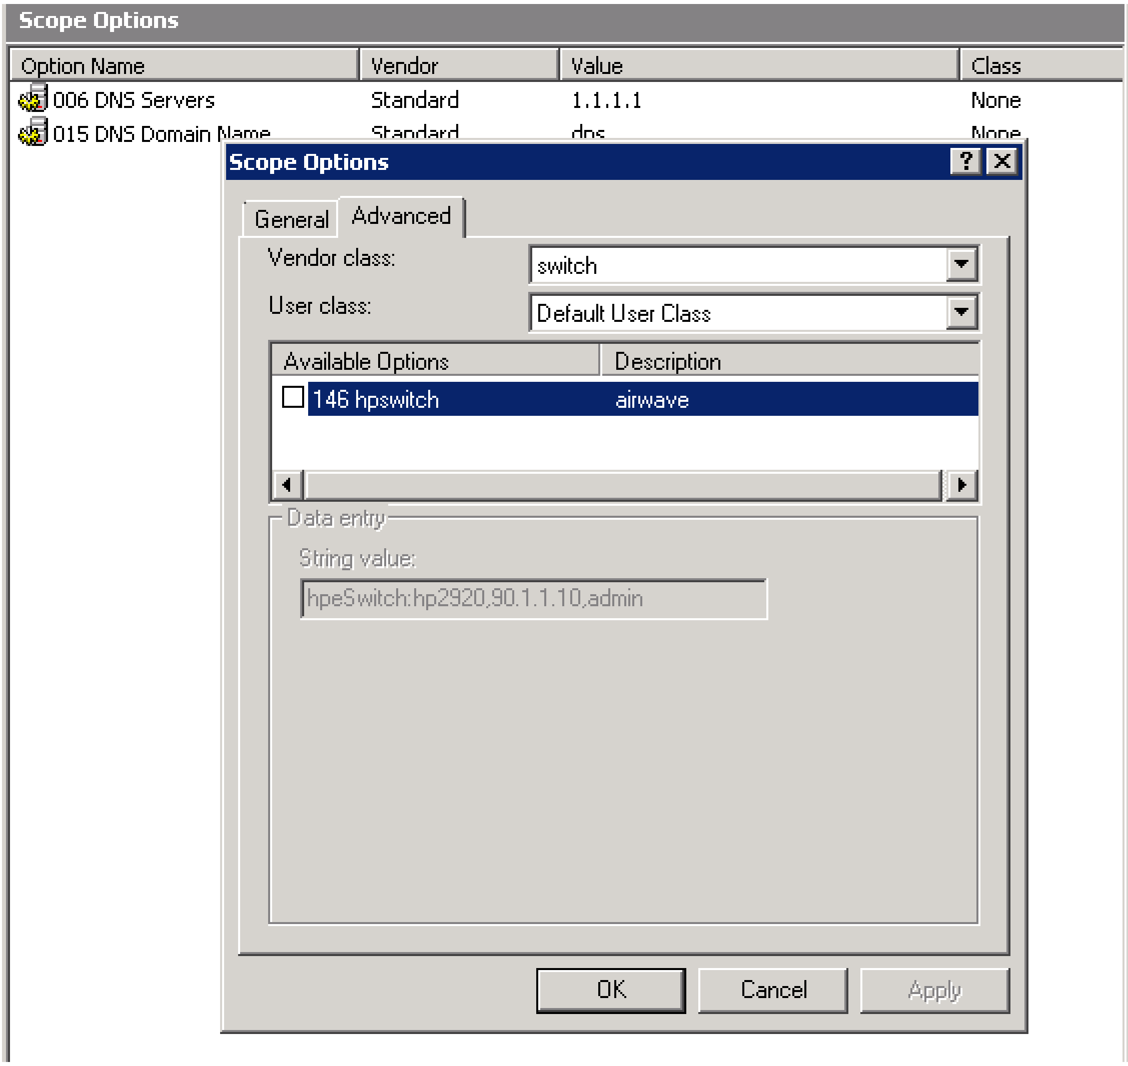

19. You can verify the AirWave details as follows:

```
switch# show amp-server
switch# show run
```

Zero Touch Provisioning

The Zero Touch Provisioning (ZTP) solution enables the auto-configuration of your switches on the first boot without requiring any administrator's intervention at the switch. The switches use DHCP server option configurations to support ZTP.

NOTE

If the switch does not contain the minimal configuration set, ZTP will get disabled. See **Image Upgrade**.

Auto-configuration using ZTP

ZTP auto-configures your switches as follows:

Procedure

1. The switch boots up with the factory default configuration.

2. The switch sends out a DHCP discovery from the primary VLAN interface.
   a. The preferred configuration method uses DHCP option 43 value as a string to parse AirWave configuration. Switch would expect a DHCP option 60 with value `ArubaInstantAP` along with DHCP option 43 to parse AirWave details
   b. The alternate configuration method supports both encapsulated values from option 43 and direct value from option 43. Encapsulated vendor-specific sub options, with sub-option code 146 is for AirWave details.
3. After the AirWave details are verified and configured, the switch initiates the check-in into the AirWave server using the HTTPS communication.

NOTE

The AirWave configuration must be in the following format:

<Group>:<Topfolder>:<folder1>,<AMP IP >,<shared secret>

4. After a successful registration, AirWave can monitor, configure, and troubleshoot the switches. Refer to *Aruba Networks and AirWave Switch Configuration Guide*.
5. Check-in failure retry is done every 60 seconds for 10 retries.
6. If the DHCP options are not configured for AirWave, the switch is left in its default state for manual configuration.

Disabling ZTP

Zero touch provisioning is disabled if you make any of the following changes to the switch's configuration:

- Enter the switch configuration mode using the `configure terminal` command.
- Enter into Menu and exit without doing any configuration
- Make any successful configuration that changes the running-configuration of the switch using a CLI, SNMP, REST APIs, menu interface, or the web GUI.
- If you upgrade with non-minimal configuration set from any 15.xx version to version 16.01, see **Image Upgrade**

Image Upgrade

If you upgrade from any 15.xx version to version 16.01, the following minimal set of configuration is validated to enable or disable the ZTP process:

- If the switch has any other VLAN apart from the default VLAN, ZTP gets disabled.
- In default VLAN, if the IPv4 address is not set as DHCP (default option is DHCP), ZTP gets disabled.
- In default VLAN, if IPv6 is enabled or configured, ZTP gets disabled.

If you have any other configuration during the upgrade, ZTP will be in the enabled state only.

CLI switch configuration

Use the `amp-server` command to configure the AirWave IP address, group, folder, and shared secret. You must have the `manager` role to execute this command.

For example:

```
switch(config)# amp-server ip 172.16.185.23 group 2530 folder 2530 secret secret
```

The `show amp-server` command shows the configuration details:

```
switch# show amp-server
AirWave Configuration details
  AMP Server IP            : 172.16.185.23
  AMP Server Group         : 2530
```

```
AMP Server Folder        : 2530
AMP Server Secret        : secret
AMP Server Config status: Configured
```

Stacking and chassis switches

The ZTP and AirWave interaction for stacked switches is similar to the one for the standalone switch, with the exception that only the commander in the stack processes the ZTP and AirWave interaction.

Stacking supports the following features:

- Backplane Stacking (BPS) running on:
  - HPE 3800 Switch Series
  - HPE Aruba 2920 Switch Series
  - HPE Aruba 3810M Series
- Virtual Switching Framework (VSF) running on HPE Aruba 5400R Switch Series v3 modules
- Chassis running on HPE Aruba 5400R Switch Series v3 modules

Troubleshooting

Cause

You can troubleshoot switches by using the SSH connection and the device logs available in AirWave. For a list of all RMON message, refer to *HPE ArubaOS-Switch Event Log Message Reference Guide*.

You can enable the debug logging with the `debug ztp` command, see **debug ztp**.

AMP server messages

To display the AMP server debug messages, enter:

```
switch# debug ztp
```

To print the debug messages to the terminal, enter:

```
switch# debug destination session
```

Validation Rules

Validation	Error/Warning
Invalid AirWave IP address	Invalid input: 300.300.300.300
Group name exceeds max length	String %s too long. Allowed length is 32 characters.
Folder name exceeds max length	String %s too long. Allowed length is 128 characters.
Secret name exceeds max length	String %s too long. Allowed length is 32 characters.
AirWave IP address or Group or folder or secret is not configured.	Incomplete input: amp-server

AirWave configuration details

amp-server

Syntax

```
[no] amp-server ip <IP ADDRESS> group <GROUP> folder <FOLDER> secret <SECRET>
```

Description

The `amp-server` command configures the AirWave Management Platform (AMP) IP address, group, folder, and shared secret and triggers the device registration with AMP.

Only the `manager` role can execute this command.

Parameters and options

ip

AMP server IP address.

group

AMP server group name.

folder

AMP server folder name.

secret

AMP server shared secret string.

no

The `no amp-server` command removes the configuration for the AMP server.

show amp-server

To view the AirWave configuration details, use the `show amp-server` command, for example:

```
AirWave Configuration details

  AMP Server IP           : 192.168.1.1
  AMP Server Group        : HP_GROUP
  AMP Server Folder       : folder
  AMP Server Secret       : secret123
  AMP Server Config Status: Configured
```

show running-configuration

```
switch# show running-config
```

```
hostname "HP-2920-24G"
module 1 type j9726a
snmp-server community "public" unrestricted
oobm
   ip address dhcp-bootp
   exit
vlan 1
```

```
   name "DEFAULT_VLAN"
   untagged 1-24
   ip address dhcp-bootp
   exit
amp-server ip 192.168.1.1 group "group" folder "folder" secret "secret123"
```

debug ztp

Syntax

`[no] debug ztp`

Description

Enables or disables ZTP debug logging.

Parameters and options

`ztp`

Zero Touch Provisioning.

`no`

The `no debug ztp` command disables the ZTP debug logging.

Beginning with switch software release 16.01, Auto configuration upon Aruba AP detection is supported on the following switch models covered in this guide:

- 3500 switch series (K software)
- 3800 switch series (KA software)
- 3810 switch series (KB software)
- 5400R switch series (KB software)
- 5400 v1/v2 switch series (K software)

Auto device detection and configuration

The auto device detection and configuration detects a directly connected Aruba AP dynamically and applies predefined configurations to ports on which the Aruba AP is detected.

You can create port configuration profiles, associate them to a device type, and enable or disable a device type. The only device type supported is `aruba-ap` and it is used to identify all the Aruba APs.

When a configured device type is connected on a port, the system automatically applies the corresponding port profile. Connected devices are identified using LLDP. When the LLDP information on the port ages out, the device profile is removed.

By default, the device profile feature is disabled. When you enable the device profile support for a device type, if no other device profile is mapped to the device type, the default device profile `default-ap-profile` is associated with the device type. You can modify the AP default device profile configuration but you cannot delete it. The `default-ap-profile` command supports only the AP device type.

Requirements

- Only APs directly connected to the switch will be detected.

Limitations

- Only one device type is supported, `aruba-ap`, and it is used to identify all the Aruba APs.
- You can modify the configuration parameters of the default profile, `default-ap-profile`, but you cannot delete it or change its name.
- For HPE 5400 Series v1 & v2 modules devices, the maximum value for `poe-max-power` is 30 W. For all other devices, the maximum value for `poe-max-power` is 33 W.
- If the port was part of any protocol VLANs prior to the device profile application, those VLANs will not be removed while applying the device profile.
- Egress rate limiting is not supported for devices running on:
    - HPE Aruba 2530 Switch Series
    - HPE Switch 2530G Series
    - HPE Switch 2620 Series
- The `egress-bandwidth` is only supported for devices running on:
    - HPE Aruba 2920 Switch Series
    - HPE Aruba 5400R Switch Series v2 & v3 modules
    - HPE 3800 Switch Series
- The `egress-bandwidth` option is not supported and not displayed in the CLI running on:

- HPE Switch 2530G Series
- HPE Aruba 2530 Switch Series
- HPE Switch 2620 Series
- 40G is not supported in egress rate-limit.

Feature Interactions

Profile Manager and 802.1X

Profile Manager interoperates with RADIUS when it is working in the client mode. When a port is blocked due to 802.1X authentication failure, the LLDP packets cannot come in on that port. Therefore, the Aruba AP cannot be detected and the device profile cannot be applied. When the port gets authenticated, the LLDP packets comes in, the AP is detected, and the device profile is applied.

You must ensure that the RADIUS server will not supply additional configuration such as VLAN or CoS during the 802.1X authentication as they will conflict with the configuration applied by the Profile Manager. If the RADIUS server supplies any such configurations to a port, the device profile will not be applied on such ports.

Profile Manager and LMA/WMA/MAC-AUTH

If either LMA, WMA, or MAC-AUTH is enabled on an interface, all the MAC addresses reaching the port must be authenticated. If LMA, WMA, or MAC-AUTH is configured on an interface, the user can have more granular control and does not need the device profile configuration. Therefore, the device profile will not be applied on such interface.

Profile manager and Private VLANs

When the device profile is applied, a check is performed to verify if the VLAN addition violates any PVLAN requirements. The following PVLAN related checks are done before applying the VLANs configured in the device profile to an interface:

- A port can be a member of only one VLAN from a given PVLAN instance.
- A promiscuous port cannot be a member of a secondary VLAN.

Procedure for creating a device identity and associating a device type

Procedure

1. Create a device identity using the command:

```
switch# device-identity name <DEVICE-NAME>
```

2. Specify the OUI used in LLDP's organization using specific TLV, (type =127). OUI should be in XXXXXX format. The default OUI "000000" indicates that device-identity will not use LLDP to identify device:

```
switch(config)# device-identity name <DEVICE-NAME> lldp oui <MAC_OUI>
 sub-type <SUBTYPE>
```

   To add new device on switch:

```
switch(config)# device-identity name abc lldp oui a1b2c3 sub 2
```

   To remove device from switch:

```
switch(config)# no device-identity name abc
```

3. Show device identity configuration:

```
switch(config)# show device-identity lldp

Device Identity Configuration

 Index  Device name             Oui        Subtype
 ------ ----------------------- ---------- -------
 1      abc                     a1b2c3     2
```

device-profile name

Syntax

```
[no] device-profile name <PROFILE-NAME> [untagged-vlan <VLAN-ID>
             tagged-vlan <VLAN-LIST> |
             cos <COS-VALUE> |
             ingress-bandwidth <Percentage> |
             egress-bandwidth <Percentage> |
             {poe-priority {critical | high | low} |
             speed-duplex {auto | auto-10 | auto-100 | ...} |
             poe-max-power <Watts>]
```

Description

Use this command to create an user-defined profile. A profile is a named collection of port settings applied as a group. You can modify the default profile, `default-ap-profile`, but you cannot delete it. You can create four additional profiles.

The `default-ap-profile` has the following values:

- `untagged-vlan`

  : 1
- `tagged-vlan`

  : None
- `ingress-bandwidth`

  : 100
- `egress-bandwidth`

  : 100
- `cos`

  : 0
- `speed-duplex`

  : auto
- `poe-max-power`

  : 33
- `poe-priority`

  : critical

You can modify these parameters. For example, you can execute `no untagged-vlan` to create a device profile with tagged only ports.

Parameters and options

`name`

Specifies the name of the profile to be configured. The profile names can be at most 32 characters long.

cos

The Class of Service (CoS) priority for traffic from the device.

untagged-vlan

The port is an untagged member of specified VLAN.

tagged-vlan

The port is a tagged member of the specified VLANs.

ingress-bandwidth

The ingress maximum bandwidth for the device port.

egress-bandwidth

The egress maximum bandwidth for the device port.

poe-priority

The PoE priority for the device port.

speed-duplex

The speed and duplex for the device port.

poe-max-power

The maximum PoE power for the device port.

no

Removes the user-defined profiles.

Restrictions

- You can modify the configuration parameters of the default profile, `default-ap-profile`, but you cannot delete it or change its name.
- For HPE Aruba 5400R Switch Series devices, the maximum value for `poe-max-power` is 30 W. For all other devices, the maximum value for `poe-max-power` is 33 W.
- Egress rate limiting is not supported for devices running on:
    - HPE Aruba 2530 Switch Series
    - HPE Switch 2530G Series
    - HPE Switch 2620 Series
- The `egress-bandwidth` is only supported for HP Switch 2920 Series, HP Switch 5400R Series v2 & v3 modules, and HP Switch 3800 Series.
- The `egress-bandwidth` option is not supported and not displayed in the CLI for devices on: HPE Switch 2530G Series, HPE Aruba 2530 Switch Series, and HPE Switch 2620 Series.
- The profile configuration is only applicable to access points.

device-profile type

Syntax

```
device-profile type <DEVICE> [associate <PROFILE-NAME> | enable | disable ]
```

Description

This command specifies an approved device type in order to configure and attach a profile to it. The profile's configuration is applied to any port where a device of this type is connected.

Only one device type is supported, `aruba-ap`, and it is used to identify all the Aruba access points.

Parameters

type

An approved device type in order to configure and attach a profile to it. The only device type supported is `aruba-ap` and it is used to identify all the Aruba APs.

associate

Associates a profile with a device type.

enable

Enables automatic profile association.

disable

Disables automatic profile association.

no

Removes the device type association and disables the feature for the device type. By default, this feature is disabled.

Rogue AP Isolation

The Rogue AP Isolation feature detects and blocks any unauthorized APs in the network. You can either log or block the rogue device. If the action requested is to log the rogue device, the MAC address of the rogue device is logged in the system logs (RMON). If the action is to block the rogue device, the traffic to and from the MAC address of the rogue device is blocked. The MAC is also logged in the system log.

When an Aruba AP detects a rogue AP on the network, it sends out the MAC address of the AP as well as the MAC of the clients connected to the AP to the switch using the ArubaOS-Switch proprietary LLDP TLV protocol. The switch then adds a rule in its hardware table to block all the traffic originating from the rogue AP's MAC address.

The `rogue-ap-isolation` command configures the rogue AP isolation for the switch and gives the option to enable or disable the rogue AP isolation feature. The `rogue-ap-isolation action` command gives you the ability to block the traffic to or from the rogue device or log the MAC of the rogue device. When the action is set to block, the rogue MAC is logged as well. By default, the action is set to block.

The `rogue-ap-isolation whitelist` command lets you add devices detected as possible rogue APs to the whitelist. A maximum of 128 MAC addresses are supported for the whitelist.

The `clear rogue-aps` command clears the detected rogue AP device MAC address.

NOTE

Rogue AP Containment feature in ArubaOS-Switch only works with Instant AP.

Limitations

- You can add a maximum of 128 MAC addresses to the whitelist.
- When a MAC is already authorized by any of the port security features such as LMA, WMA, or 802.1X, the MAC is logged but you cannot block it using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- When a MAC is already configured as an IP received MAC of a VLAN interface, the MAC is logged but you cannot block it by using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- When a MAC is already locked out via `lockout-mac` or locked down using the `static-mac` configuration, the MAC is logged but you cannot block it using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- The number of rogue MACs supported on a switch is a function of the value of `max-vlans` at boot time. Since the resources are shared with the `lockout-mac` feature, the scale is dependent on how many lockout

addresses have been configured on the switch using the `lockout-mac` feature. The following table lists the scale when there are no lockout addresses configured on the switch:

Max VLAN	Supported MACs
0 < VLAN <= 8	200
8 < VLAN <= 16	100
16 < VLAN <= 256	64
256 < VLAN <= 1024	16
1024 < VLAN <= 2048	8
2048 < VLAN <= 4094	4

The switch will create an RMON log entry and the rogue MAC will be ignored when the limit is reached.

NOTE

If the `max-vlans` value is changed to a different value, the scale of rogue MACs supported will not change until the next reboot.

Feature Interactions

MAC lockout and lockdown

The Rogue AP isolation feature uses the MAC lockout feature to block MACs in hardware. Therefore, any MAC blocked with the Rogue AP isolation feature cannot be added with the `lockout-mac` or `static-mac` command if the action type is set to `block`.

For example:

```
switch# lockout-mac 247703-7a8950
Cannot add the entry for the MAC address 247703-7a8950 because it is already
blocked by rogue-ap-isolation.
```

```
switch# static-mac 247703-7a8950 vlan 1 interface 1
Cannot add the entry for the MAC address 247703-7a8950 because it is already
blocked by rogue-ap-isolation.
```

Similarly, any MAC that was added with the `lockout-mac` or `static-mac` command and that is being detected as rogue will be logged, but not blocked in hardware as it already is set to block. If the MAC is removed from `lockout-mac` or `static-mac` but is still in the rogue device list, it will be blocked back in hardware if the action type is `block`.

LMA/WMA/802.1X/Port-Security

Any configuration using LMA, WMA, 802.1X, or Port-Security will not be blocked if the Rogue AP isolation feature is enabled. All these features act only when a packet with the said MAC is received on a port.

If `rogue-ap-isolation` blocks a MAC before it is configured to be authorized, packets from such MACs will be dropped until one of the following happens:

- Rogue action is changed to LOG.
- Rogue-AP isolation feature is disabled.

- The MAC is not detected as rogue anymore.
- LLDP is disabled on the port (or globally).

Once a MAC has been authorized by one of these features, it will not be blocked by Rogue AP isolation. A RMON will be logged to indicate the failure to block.

The Rogue AP module will retry to block any such MACs periodically. In the event of the MAC no longer being authorized, Rogue AP isolation will block the MAC again. No RMON is logged to indicate this event.

L3 MAC

The Rogue AP isolation feature will not block a MAC configured as an IP receive MAC address on a VLAN interface. This event will be logged in RMON if such MACs are detected as rogue.

Conversely, any MAC already blocked by Rogue AP isolation will not be allowed to be configured as an IP receive MAC address of a VLAN interface.

For example:

```
switch# vlan 1 ip-recv-mac-address 247703-3effbb
Cannot add an entry for the MAC address 247703-3effbb because it is already
blocked by rogue-ap-isolation.
```

Using the Rogue AP Isolation feature

Procedure

1. Check the feature state:

```
switch# show rogue-ap-isolation

 Rogue AP Isolation

  Rogue AP Status : Disabled
  Rogue AP Action : Block

  Rogue MAC Address Neighbour MAC Address
  ----------------- ---------------------
```

2. Enable the feature:

```
switch# rogue-ap-isolation enable
switch# show rogue-ap-isolation

 Rogue AP Isolation

  Rogue AP Status : Enabled
  Rogue AP Action : Block

  Rogue MAC Address Neighbour MAC Address
  ----------------- ---------------------
```

3. Change the action type from block to log:

```
switch# rogue-ap-isolation action log
switch# show rogue-ap-isolation

 Rogue AP Isolation

  Rogue AP Status : Enabled
  Rogue AP Action : Log
```

```
 Rogue MAC Address Neighbour MAC Address
 ----------------- ---------------------
```

4. List the current whitelist entries:

```
switch# show rogue-ap-isolation whitelist

Rogue AP Whitelist Configuration

 Rogue AP MAC
 ------------------
```

5. Add a new whitelist entry:

```
switch# rogue-ap-isolation whitelist 005056-00326a
switch# show rogue-ap-isolation whitelist

Rogue AP Whitelist Configuration

 Rogue AP MAC
 ------------------
 00:50:56:00:32:6a
```

rogue-ap-isolation

Syntax

`rogue-ap-isolation {enable | disable}`

Description

Configures the rogue AP isolation for the switch.

Parameters

`enable`

Enables the rogue AP isolation.

`disable`

Disables the rogue AP isolation.

rogue-ap-isolation action

syntax

`rogue-ap-isolation action [log | block]`

Description

Configures the action to take for the rogue AP packets. This function is disabled by default.

Parameters and options

`action`

Configure the action to take for rogue AP packets. By default, the rogue AP packets are blocked.

`log`

Logs traffic to or from any rogue access points.

block

Blocks and logs traffic to or from any rogue access points.

rogue-ap-isolation whitelist

Syntax

```
[no] rogue-ap-isolation whitelist <MAC-ADDRESS>
```

Description

Configures the rogue AP Whitelist MAC addresses for the switch. Use this command to add to the whitelist the MAC addresses of approved access points or MAC addresses of clients connected to the rogue access points. These approved access points will not be added to the rogue AP list even if they are reported as rogue devices.

Parameters and options

<MAC-ADDRESS>

Specifies the MAC address of the device to be moved from the rogue AP list to the whitelist. You can add a maximum of 128 MAC addresses to the whitelist.

no

Removes the MAC address individually by specifying the MAC.

clear rogue-ap-isolation

syntax

```
clear rogue-ap-isolation [<MAC-ADDRESS> | all]
```

Description

Removes the MAC addresses from the rogue AP list. The MAC addresses cleared using this option will be added back to the rogue list under the following circumstances:

- The LLDP administrator status of the port on which the AP that reported the MAC is disabled and enabled back.
- The data that is in the rogue AP TLV sent from the AP that informed the rogue MAC has changed.
- To permanently ignore a MAC from being detected as rogue, add it to the whitelist.

Parameters and options

<MAC-ADDRESS>

Specifies the MAC address of the device to be moved from the rogue AP list.

all

Clears all MAC addresses from the rogue AP list.

Troubleshooting

Dynamic configuration not displayed when using “show running-config”

Symptom

The `show running-config` command does not display the dynamic configuration applied through the device profile.

Cause

The `show running-config` command shows only the permanent user configuration and parameters configured through device profile.

Action

1. Use the specific `show device-profile` command to display the parameters dynamically configured through the device profile.

Switch does not detect the rogue AP TLVs

Symptom

The switch does not detect the rogue AP TLVs that could be sent from the neighboring device.

Cause

The LLDP administrator status of a port is moved from `txOnly` to `tx_rx` or `rx_only` within 120 seconds of the previous state change to `txOnly`.

Action

1. Wait for 120 seconds before moving from the state `txOnly` to the state `tx_rx` or `rx_only`.
2. Move the administrator status to `disable` and then back to `tx_rx` or `rx_only`.

The `show run` command displays non-numerical value for untagged-vlan

Symptom

The `show run` command displays one of the following values for `untagged-vlan`:

- `no untagged-vlan`
- `untagged-vlan : None`

Cause

The `no device-profile` or the `no rogue-ap-isolation whitelist` command is executed to configure `untagged-vlan` to `0`.

Action

1. No action is required.

Show commands

Use the following show commands to view the various configurations and status.

show device-profile

Syntax

`show device-profile [config|status]`

Description

Shows the device profile configuration and status.

Parameters and options

config

Shows the device profile configuration details for a single profile or all profiles.

status

Shows currently applied device profiles.

show rogue-ap-isolation

Syntax

```
show rogue-ap-isolation [whitelist]
```

Description

Shows the following information:

- The status of the feature: enabled or disabled.
- The current action type for the rogue MACs detected.
- The list of MAC addresses detected as rogue and the MAC address of the AP that reported them.

Parameters and options

whitelist

Shows the rogue AP whitelist configuration.

show run

Syntax

```
show run
```

Description

Shows the running configuration.

Validation Rules

Validation	Error/Warning/Prompt
`device-profile profile-name default-ap-profile`	Maximum tagged VLANs that can be associated with a device-profile is 256.
`device-profile profile-name` creation.	String too long. Allowed length is 32 characters.
`device-profile profile-name` creation.	Device profile <> already exists.
`device-profile profile-name` creation.	The maximum number of device profiles allowed is 5.
`device-profile profile-name` deletion.	Device profile <> does not exist.
`device-profile profile-name` deletion.	Cannot delete profile <> when associated with a device type.
`device-profile profile-name` deletion.	Default profile cannot be deleted.

Table Continued

Validation	Error/Warning/Prompt
`device-profile profile-name` modification via SNMP.	Default profile name cannot be changed.
`device-profileprofile-name` creation/ modification via SNMP.	Device profile index cannot be greater than 5.
`untagged-vlan`	Invalid VLAN.
`untagged-vlan`	Cannot configure the VLAN <> as an untagged VLAN because this is already used as a tagged VLAN.
`tagged-vlan 1-1000`	The maximum number of tagged VLANs in a profile is less than 512 or the maximum VLANs, `MAX_VLANs`, configurable in the system.
`tagged-vlan`	Cannot configure the VLAN <> as a tagged VLAN because this is already used as an untagged VLAN.
`ingress-bandwidth`	SNMP should return `WRONG_VALUE_ERROR`.
`egress-bandwidth`	SNMP should return `WRONG_VALUE_ERROR`.
`cos`	SNMP should return `WRONG_VALUE_ERROR`.
`speed-duplex`	SNMP should return `WRONG_VALUE_ERROR`.
`poe-max-power`	SNMP should return `WRONG_VALUE_ERROR`.
`poe-priority`	SNMP should return `WRONG_VALUE_ERROR`.
`device-profile type aruba-ap profile-name`	String <> too long. Allowed length is 32 characters.
`device-profile type aruba-ap profile-name`	Device profile <> does not exist.
`device-profile type aruba-switch-router`	Device type is not supported.
`rogue-ap-whitelist`	Whitelist MAC address already exists in the list.
`rogue-ap-whitelist`	Whitelist MAC address does not exist in the list.
`rogue-ap-whitelist`	The maximum number of whitelist MACs allowed is 128.
`rogue-ap-whitelist <MAC>`	Cannot add the whitelist entry because the specified MAC address is already configured as a lock-out MAC.

Table Continued

Validation	Error/Warning/Prompt
`lock-out <MAC>`	Cannot add the lock-out entry because the specified MAC address is already configured as a whitelist MAC.
`lockout-mac <MAC-ADDRESS>`OR`static-mac <MAC-ADDRESS> vlan <vlan-id> interface <interface>`OR`vlan <vlan-id> ip-recv-mac-address <MAC-ADDRESS`	Cannot add an entry for the MAC address <`MAC-ADDRESS`> because it is already blocked by `rogue-ap-isolation`.

LACP configuration

interface <PORT-LIST> lacp

Syntax

```
[no] interface <PORT-LIST> lacp [mad-passthrough [enable|disable]|active|passive|
key <KEY>]
```

Description

Defines whether LACP is enabled on a port, and whether it is in active or passive mode when enabled.

Parameters and options

mad-passthrough

Applies only to trunks and not to physical ports.

enable

Allows the port to send LACP packets.

disable

When LACP is disabled, the port ignores LACP packets.

active

When LACP is enabled and active, the port sends LACP packets and listens to them. Defaults to active.

passive

When LACP is enabled and passive, the port sends LACP packets only if it is spoken to.

key *<KEY>*

During dynamic link aggregation using LACP, ports with the same key are aggregated as a single trunk.

show lacp

Syntax

```
show lacp [counters [<PORT-LIST>] | local [<PORT-LIST>] |peer [<PORT-LIST>] |
distributed | mad-passthrough [counters [<PORT-LIST>]]]
```

Description

Show LACP-MAD passthrough configuration on LACP trunks, or show LACP-MAD passthough counters on ports.

Usage

```
show lacp mad-passthrough counters [<PORT-LIST>]
```

clear lacp statistics

Syntax

```
clear lacp statistics
```

Description

Clear all LACP statistics including MAD passthrough counters. Resets LACP packets sent and received on all ports.

LACP-MAD Operations

Link Aggregation Control Protocol-Multi-Active Detection (LACP-MAD) is a detection mechanism deployed by switches to recover from a breakup of the Virtual Switching Foundation (VSF) stack due to link or other failure.

LACP-MAD is implemented by sending extended LACP data units (LACPDUs) with a type length value (TLV) that conveys the active ID of an VSF virtual device. The active ID is identical to the member ID of the master and is thus unique to the VSF virtual device. When LACP MAD detection is enabled, the members exchange their active IDs by sending extended LACPDUs.

- When the VSF virtual device operates normally, the active IDs in the extended LACPDUs sent by all members are the same, indicating that there is no multi-active collision.
- When there is a breakup in the VSF virtual chassis, the active IDs in the extended LACPDUs sent by the members in different VSF virtual devices are different, indicating that there are multi-active collisions.

LACP-MAD passthrough helps VSF-capable devices detect multi-access and take corrective action. These devices do not initiate transmission of LACP-MAD frames or participate in any MAD decision making process. These devices simply forward LACP-MAD TLVs received on one interface to the other interfaces on the trunk. LACP-MAD passthrough can be enabled for 24 LACP trunks. By default, LACP-MAD passthrough is disabled.

File transfer methods

The switches support several methods for transferring files to and from a physically connected device or via the network, including TFTP, Xmodem, and USB. This chapter explains how to download new switch software, upload or download switch configuration files and software images, and upload command files for configuring ACLs.

TFTP

TFTP at the switch is allows for extensive use of scripts on various customer environments. Such environs, like FW, configurations, backups, and restores all use the TFTP network service.

- SSH/SFTP is needed to secure access to network components.
- Users are allowed to re-enable TFTP and make both TFTP and SFTP work in parallel.
- SFTP support for database of DSNOOPv4, v6 and DHCP Server are also available. To provide a secure way to transfer the database, the SFTP option has been added where the respective database can also be transferred to a SFTP Server.

Prerequisites

Use of the commands in this section assumes that:

- A software version for the switch has been stored on a TFTP server accessible to the switch. (The software file is typically available from the Switch Networking website at **http://www.hpe.com/networking/support**.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you proceed, complete the following:

- Obtain the IP address of the TFTP server in which the software file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the software file stored in the TFTP server for the switch (for example, E0820.swi.)

NOTE

If your TFTP server is a UNIX workstation, ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.

Downloading switch software

To download a switch software file named K.0800.swi from a TFTP server with the IP address of 10.28.227.103 to primary flash:

Procedure

1. Execute **copy tftp flash** on page 415 as shown below:

   Figure 102: *Download command for an OS (switch software)*

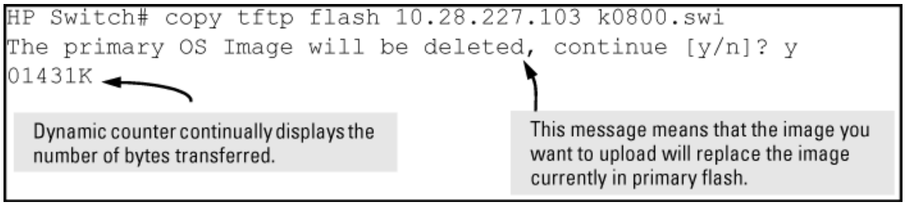

2. When the switch finishes downloading the software file from the server, it displays this progress message:
3. **Validating and Writing System Software to FLASH ...**
4. When the download finishes, you must reboot the switch to implement the newly downloaded software image. To do so, use either **boot system flash** on page 415 or **reload** on page 415.
5. To confirm that the software downloaded correctly, execute `show system` and check the **Firmware revision** line.

For information on primary and secondary flash memory and the boot commands, see the basic operation guide.

If you use

`auto-tftp`

to download a new image in a redundant management system, the active management module downloads the new image to both the active and standby modules. Rebooting after the

`auto-tftp`

process completes reboots the entire system.

---

copy tftp flash

Syntax

`copy tftp flash <IP-ADDRESS>  <REMOTE-FILE> [primary|secondary] [oobm]`

Automatically downloads a switch software file to primary or secondary flash. If you do not specify the flash destination, the TFTP download defaults to primary flash.

boot system flash

Syntax

`boot system flash [primary|secondary]`

Description

Boots from the selected flash.

reload

Syntax

`reload`

Description

Boots from the flash image and startup-config file. A switch covered in this guide (with multiple configuration files), also uses the current startup-config file.

Enabling tftp

Procedure

1. TFTP defaults to enabled on the switch. If TFTP operation has been disabled, you can re-enable it by specifying TFTP client or server functionality with the following command.

   > For information on how to configure TFTP file transfers on an IPv6 network, see the IPv6 configuration guide
   >
   > .

tftp

Syntax

`[no] tftp [client|server] listen [oobm|data|both]`

Description

Disables/re-enables TFTP for client or server functionality.

Parameters and options

`no`

Disables all TFTP client or server operation on the switch except for the auto-TFTP feature. When you disable TFTP, instances of TFTP in the CLI `copy` command and the Menu interface "Download OS" screen become unavailable.

> The `no tftp [client|server]` command does not disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the `no auto-tftp` command to remove the command entry from the switch's configuration.

`client`

Use TFTP client functionality to access TFTP servers in the network to receive downloaded files.

`server`

Use TFTP server functionality to upload files to other devices on the network.

`listen`

For switches that have a separate out-of-band management port, the `listen` parameter in a server configuration allows you to specify whether transfers take place through the out-of-band management (`oobm`) interface, the `data` interface, or `both`.

> Using IP SSH file transfer to enable SCP and SFTP functionality on the switch disables TFTP client and server functionality. After enabling `ip ssh` file transfer, you cannot re-enable TFTP and auto-TFTP from the CLI.

show running-configuration

```
switch(config)# show running-config

Running configuration:
; J8693A Configuration Editor; Created on release #K.15.15.0000x
; Ver #04:7f.ff.3f.ef:54
hostname "HP-3500yl-48G"
```

```
no tftp client
no tftp server
```

Enable TFTP client

```
switch(config)# tftp client
```

ip ssh filetransfer

The command `ip ssh filetransfer` disables the TFTP Client and TFTP Server, and the user can re-enable them. The command displays the following message.

```
ip ssh filetransfer
tftp and auto-tftp have been disabled.
```

Automatic software download from a TFTP server

The `auto-tftp` command lets you configure the switch to download software automatically from a TFTP server.

At switch startup, the auto-TFTP feature automatically downloads a specified software image to the switch from a specified TFTP server and then reboots the switch. To implement the process, you must first reboot the switch using one of the following methods:

- Enter the `boot system flash primary` command in the CLI.
- With the default flash boot image set to primary flash (the default), enter the `boot` or the `reload` command, or use the reset button on the switch. (To reset the boot image to primary flash, use `boot set-default flash primary`.)

auto-tftp

Syntax

`auto-tftp <IP-ADDR> filename`

Description

Configures the switch to automatically download the specified software file from the TFTP server at the specified IP address. The file is downloaded into primary flash memory at switch startup, and then the switch automatically reboots from primary flash. Defaults to disabled.

Parameters and options

no

Disables auto-TFTP operation by deleting the `auto-tftp` entry from the startup configuration. Does not affect the current TFTP-enabled configuration on the switch. However, entering the `ip ssh filetransfer` command automatically disables both `auto-tftp` and `tftp` operation.

To enable auto-TFTP to copy a software image to primary flash memory, the version number of the downloaded software file (for example, K_14_01.swi) must be different from the version number currently in the primary flash image.

The current TFTP client status (enabled or disabled) does not affect auto-TFTP operation. (See **Enabling tftp** on page 416.)

Completion of the auto-TFTP process may require several minutes while the switch executes the TFTP transfer to primary flash and then reboots again.

Downloading to primary flash using TFTP

The menu interface accesses only the primary flash.

Procedure

1. In the console Main Menu, select **Download OS** to display the screen in **Figure 103: Download OS (software) screen (default values)** on page 418. (The term "OS" or "operating system" refers to the switch software):

   Figure 103: *Download OS (software) screen (default values)*

```
==========================- CONSOLE - MANAGER MODE -============================
                              Download OS

  Current Firmware revision : K.11.00

  Method [TFTP]  : TFTP
  TFTP Server :

  Remote File Name :



 Actions->   Cancel     Edit     eXecute     Help

Select the file transfer method (TFTP and XMODEM are currently supported).
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

2. Press **[E]** (for **Edit** .)
3. Ensure that the **Method** field is set to **TFTP** (the default.)
4. In the **TFTP Server** field, enter the IP address of the TFTP server in which the software file has been stored.
5. In the **Remote File Name** field, enter the name of the software file (if you are using a UNIX system, remember that the filename is case-sensitive.)
6. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.
7. The screen shown in **Figure 104: Download OS (software) screen during a download** on page 418 appears:

   Figure 104: *Download OS (software) screen during a download*

```
==========================- CONSOLE - MANAGER MODE -============================
                              Download OS
  Current Firmware revision : E.08.00
  Method [TFTP]  : TFTP
  TFTP Server : 10.28.227.105

  Remote File Name : K.11.00.swi

             Received 370,000 bytes of OS download.
  +------------------------------------------------------------------------+
  |*******************                                                     |
  +------------------------------------------------------------------------+
```

8. A "progress" bar indicates the progress of the download. When the entire software file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**
9. After the primary flash memory is updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch** .)
10. You will see this prompt:

```
Continue reboot of system? : No
```

11. Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.

> When you use the menu interface to download a switch software, the new image is always stored in primary flash. Also, using the
>
> `Reboot Switch`
>
> command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI provides more options. See the
>
> basic operation guide
>
> .

12. After you reboot the switch, confirm that the software downloaded correctly:
    a. From the Main Menu, select **2. Switch Configuration...**, and then select **2. Port/Trunk Settings**
    b. Check the **Firmware revision** line.

Disabling TFTP and auto-TFTP for enhanced security

Disabling TFTP and auto-TFTP

Using the `ip ssh filetransfer` command to enable SFTP automatically disables TFTP and auto-TFTP (if either or both are enabled), as shown in **Figure 105: Example of switch configuration with SFTP enabled** on page 419.

Figure 105: *Example of switch configuration with SFTP enabled*

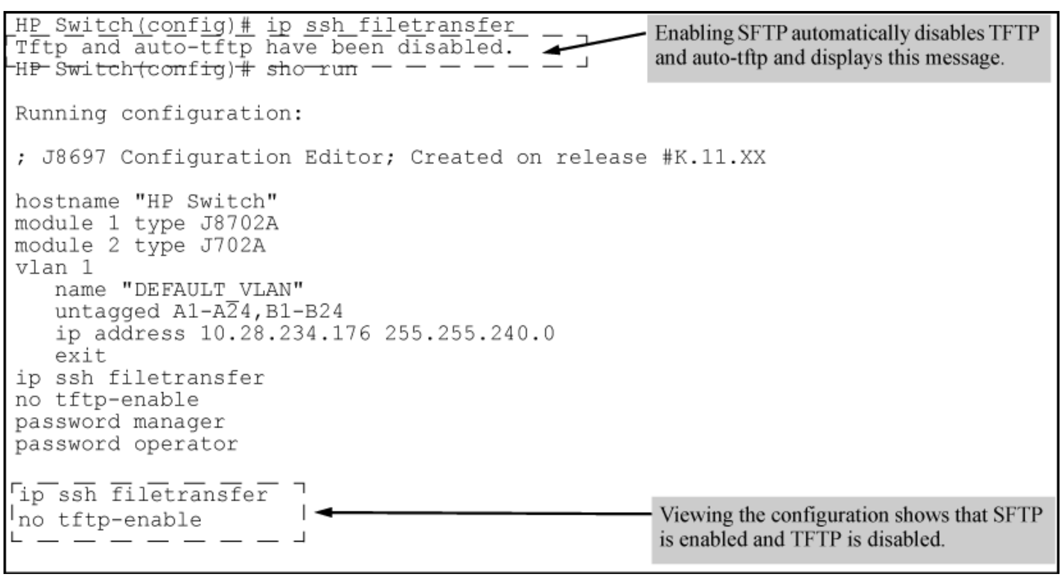

If you enable SFTP and then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating rules

Prerequisites

To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but you must use the CLI to disable auto-TFTP. The following CLI commands disable TFTP and auto-TFTP on the switch.

Procedure

1. The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface (see

   Figure 106: Using the Menu interface to disable TFTP on page 420

   ), or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

   Figure 106: *Using the Menu interface to disable TFTP*

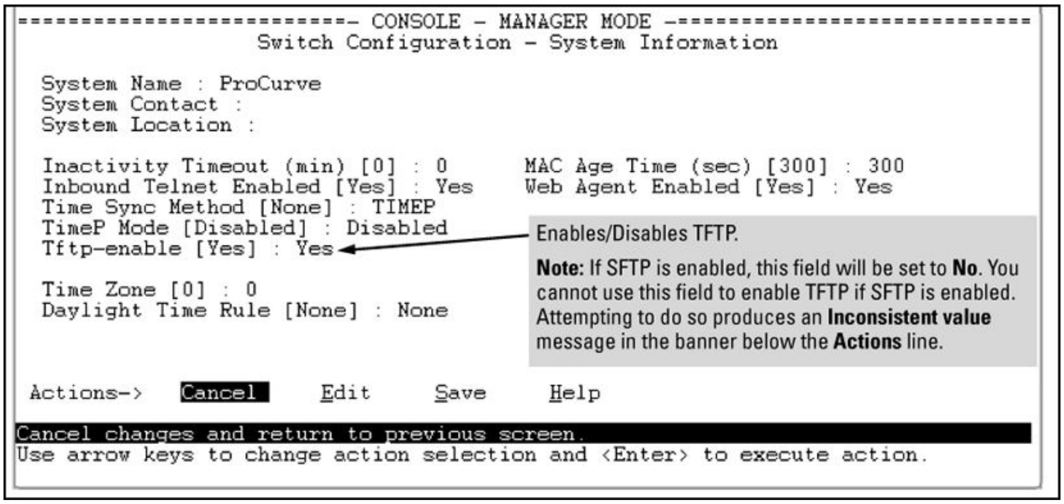

2. While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

```
SFTP must be disabled before enabling tftp.
SFTP must be disabled before enabling auto-tftp.
```

   Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an "inconsistent value" message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

Enabling SSH V2 (required for SFTP)

As a matter of policy, administrators should **not** enable the SSH V1-only or the SSH V1-or-V2 advertisement modes. SSHv1 is supported on only some legacy switches (such as the HPE Switch Series 2500 switches.)

Any attempts to use SCP or SFTP without using `ip ssh filetransfer` cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for example:

```
IP file transfer not enabled on the switch
```

1. Enter the following command to enable SSH v2:

   `switch(config)# ip ssh version 2`
2. Enter the `show ip ssh` command to confirm that you have enabled an SSH session:

   `switch(config)# show ip ssh`
3. Enter the `ip ssh filetransfer` command so that SCP and/or SFTP can run.
4. Open your third-party software client application to being using the SCP or SFTP commands to safely transfer files or issue commands to the switch.

Disabling secure file transfer

1. To disable SSH, enter the following command:

   `switch(config)# no ip ssh filetransfer`

Authentication

Switch memory allows up to ten public keys. This means the authentication and encryption keys you use for your third-party client SCP/SFTP software can differ from the keys you use for the SSH session, even though both SCP and SFTP use a secure SSH tunnel.

NOTE

SSH authentication is mutually exclusive with RADIUS servers.

Some clients, such as PSCP (PuTTY SCP), automatically compare switch host keys for you. Other clients require you to manually copy and paste keys to the **`$HOME/.ssh/known_hosts`** file. Whatever SCP/SFTP software tool you use, after installing the client software you must verify that the switch host keys are available to the client.

Because the third-party software utilities you may use for SCP/SFTP vary, you should refer to the documentation provided with the utility you select before performing this process.

SCP/SFTP operating notes

- When an SFTP client connects, the switch provides a file system displaying all of its available files and folders. No file or directory creation is permitted by the user. Files may be only uploaded or downloaded, according to the permissions mask. All of the necessary files the switch needs are already in place on the switch. You do not need to (nor can you) create new files.
- The switch supports one SFTP session or one SCP session at a time.
- All files have read-write permission. Several SFTP commands, such as `create` or `remove`, are not allowed and return an error message. The switch displays the following files:

```
/
+---cfg
|   running-config
|   startup-config
+---log
|   crash-data
|   crash-data-a
|   crash-data-b
|   crash-data-c
|   crash-data-d 8212zl only
|   crash-data-e            "      "
|   crash-data-f  ""
|   crash-data-g  8212zl only
|   crash-data-h            "      "
|   crash-data-I  ""
|   crash-data-J  ""
|   crash-data-K  ""
|   crash-data-L  "    "
```

```
|   crash-log
|   crash-log-a
|   crash-log-b
|   crash-log-c
|   crash-log-d   8212zl only
|   crash-log-e  ""
|   crash-log-f  ""
|   crash-log-g  8212zl only
|   crash-log-h  " "
|   crash-log-I  " "
|   crash-log-J  " "
|   crash-log-K  " "
|   crash-log-L  " "
|   event log
+---os
|   primary
|   secondary
\---ssh
    +---mgr_keys
    |   authorized_keys
    \---oper_keys
    |   authorized_keys
\---core   (this directory is not available on the 8212zl)
|   mm1.cor      management module or management function
|   im_a.cor     interface module (chassis switches only)
|   im_b.cor     interface module (chassis switches only)
|   im_1.cor     interface module (chassis switches only)
|   port_1-24.cor          core-dump for ports 1-24 (stackable switches only)
|   port_25-48.cor         core-dump for ports 25-48 (stackable switches only)
```

- When using SFTP to copy a software image onto the switch, the command return takes only a few seconds. However, this does not mean that the transfer is complete, because the switch requires additional time (typically more than one minute) to write the image to flash in the background. To verify the file transfer has been completed, you can use the `show flash` command or look for a confirmation message in the log, as in the following example:

```
I 01/09/09 16:17:07 00150 update: Primary Image updated.
```

Troubleshooting SSH, SFTP, and SCP operations

Cause

NOTE

You can verify secure file transfer operations by checking the switch's event log, or by viewing the error messages sent by the switch that most SCP and SFTP clients print out on their console.

Messages that are sent by the switch to the client depend on the client software in use to display them on the user console.

Broken SSH connection

Symptom

Any of the following types of error messages are logged, according to the type of session that is running (SSH, SCP, or SFTP):

```
ssh: read error Bad file number, session aborted I 01/01/90
00:06:11 00636 ssh: sftp session from ::ffff:10.0.12.35 W
01/01/90 00:06:26 00641 ssh:

sftp read error Bad file number, session aborted I 01/01/90
00:09:54 00637 ssh: scp session from ::ffff:10.0.12.35 W 01/01/90

ssh: scp read error Bad file number, session aborted
```

The Bad file number is from the system error value and may differ depending on the cause of the failure. In the third example, the device file to read was closed as the device read was about to occur.

Cause

If an ssh connection is broken at the wrong moment (for instance, the link goes away or spanning tree brings down the link), a fatal exception occurs on the switch. If this happens, the switch gracefully exits the session and produces an Event Log message indicating the cause of failure.

Attempt to start a session during a flash write

Symptom

With some client software, an error message similar to this appears on the client console:

```
Received disconnect from 10.0.12.31: 2: Flash access in progress
lost connection
```

Cause

If you attempt to start an SCP (or SFTP) session while a flash write is in progress, the switch does not allow the SCP or SFTP session to start.

Action

- Wait a few seconds for the flash writing to finish and then try again.

Failure to exit from a previous session

Symptom

```
Received disconnect from 10.0.12.31: 2: Wait for previous
session to complete

lost connection
```

Cause

The error message might appear on the client console if a new SCP (or SFTP) session is started from a client before the previous client session has been closed (the switch requires approximately ten seconds to timeout the previous session).

Action

Attempt to start a second session

Symptom

```
Received disconnect from 10.0.12.31: 2: Other SCP/SFTP
session running

lost connection
```

Cause

The switch supports only one SFTP session or one SCP session at a time. If a second session is initiated (for example, an SFTP session is running and then an SCP session is attempted), the error message might appear on the client console.

Action

Use USB to transfer files to and from the switch

The switch's USB port (labeled as **Aux Port**) allows the use of a USB flash drive for copying configuration files to and from the switch.

Operating rules and restrictions:

- Unformatted USB flash drives must first be formatted on a PC (Windows FAT format.) For devices with multiple partitions, only the first partition is supported. Devices with secure partitions are not supported.
- If they already exist on the device, subdirectories are supported. When specifying a ***filename*** , you must enter either the individual file name (if at the root) or the full path name (for example, /subdir/filename.)
- To view the contents of a USB flash drive, use the `dir` command. This lists all files and directories at the root. To view the contents of a directory, you must specify the subdirectory name (that is, `dir subdirectory`.)
- The USB port supports a single USB device. USB hubs to add more ports are not supported.

NOTE

Some USB flash drives may not be supported on your switch. Consult the latest *Release Notes* for information on supported devices.

SCP and SFTP

Enabling SCP and SFTP

Procedure

1. Open an SSH session as you normally would to establish a secure encrypted tunnel between your computer and the switch. Please note that this is a one-time procedure for new switches or connections. If you have already done it once you should not need to do it a second time.
2. For more detailed directions on how to open an SSH session, see the access security guide.
3. To enable secure file transfer on the switch (once you have an SSH session established between the switch and your computer), open a terminal window and enter the following command:

```
switch(config)# ip ssh filetransfer
```

Using SCP and SFTP

Using IP SSH file transfer to enable SCP and SFTP functionality on the switch disables TFTP client and server functionality. After enabling

```
ip ssh
```

file transfer, you cannot re-enable TFTP and auto-TFTP from the CLI.

The general process for using SCP and SFTP involves three steps:

Procedure

1. Open an SSH tunnel between your computer and the switch if you have not already done so.
2. (This step assumes that you have already set up SSH on the switch.)
3. Execute `ip ssh filetransfer` to enable secure file transfer.
4. Use a third-party client application for SCP and SFTP commands.

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session and enabling ip ssh file transfer, you can then use a third-party software application to take advantage of SCP and SFTP. SCP and SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially, you are creating a secure SSH tunnel as a way to transfer files with SFTP and SCP channels.

Once you have configured your switch to enable secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

To use these commands, you must install on the administrator workstation a third-party application software client that supports the SFTP and/or SCP functions. Some examples of software that supports SFTP and SCP are PuTTY, Open SSH, WinSCP, and SSH Secure Shell. Most of these are freeware and may be downloaded without cost or licensing from the internet. There are differences in the way these clients work, so be sure you also download the documentation.

As described earlier in this chapter you can use a TFTP client on the administrator workstation to update software images. This is a plain-text mechanism that connects to a standalone TFTP server or another switch acting as a TFTP server to obtain the software image files. Using SCP and SFTP allows you to maintain your switches with greater security. You can also roll out new software images with automated scripts that make it easier to upgrade multiple switches simultaneously and securely.

SFTP is unrelated to FTP, although there are some functional similarities. Once you set up an SFTP session through an SSH tunnel, some of the commands are the same as FTP commands. Certain commands are not allowed by the SFTP server on the switch, such as those that create files or folders. If you try to issue commands such as `create` or `remove` using SFTP, the switch server returns an error message.

You can use SFTP just as you would TFTP to transfer files to and from the switch, but with SFTP, your file transfers are encrypted and require authentication, so they are more secure than they would be using TFTP. SFTP works only with SSH version 2 (SSH v2.)

NOTE

SFTP over SSH version 1 (SSH v1) is not supported. A request from either the client or the switch (or both) using SSH v1 generates an error message. The actual text of the error message differs, depending on the client software in use. Some examples are:

```
Protocol major versions differ: 2 vs. 1
Connection closed

Protocol major versions differ: 1 vs. 2
Connection closed

Received disconnect from  ip-addr : /usr/local/libexec/
sftp-server: command not supported
Connection closed
```

SCP is an implementation of the BSD `rcp` (Berkeley UNIX remote copy) command tunneled through an SSH connection.

SCP is used to copy files to and from the switch when security is required. SCP works with both SSH v1 and SSH v2. Be aware that the most third-party software application clients that support SCP use SSHv1.

Xmodem

Downloading software using Xmodem

Prerequisites

Procedure

1. Connect the switch via the Console RS-232 port to a PC operating as a terminal. (For information on connecting a PC as a terminal and running the switch console interface, see the

   installation and getting started guide

   you received with the switch.)
2. Verify that the switch software is stored on a disk drive in the PC.

3. Verify that the terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the

   Send File

   option in the

   Transfer

   drop-down menu.)

Downloading to Flash

The following procedure downloads a switch software file named `E0822.swi` from a PC (running a terminal emulator program such as HyperTerminal) to primary flash.

Procedure

1. Execute the following command in the CLI:

```
HP Switch# copy xmodem flash
Press 'Enter' and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:
   a. Click on **Transfer**, then **Send File**.
   b. Type the file path and name in the Filename field.
   c. In the Protocol field, select **Xmodem**.
   d. Click on the **[Send]** button.The download can take several minutes, depending on the baud rate used in the transfer.
3. When the download finishes, you must reboot the switch to implement the newly downloaded software.
4. Use either **boot system flash** on page 426 or **reload** on page 426 commands.
5. To confirm that the software downloaded correctly:

```
HP Switch show system
```

6. Check the **Firmware revision** line. It should show the software version that you downloaded in the preceding steps.

boot system flash

Syntax

`boot system flash [primary|secondary]`

Description

Reboots from the selected flash

reload

Syntax

`reload`

Description

Reboots from the flash image currently in use

copy xmodem flash

Syntax

`copy xmodem flash [primary|secondary]`

Description

Downloads a software file to primary or secondary flash. If you do not specify the flash destination, the Xmodem download defaults to primary flash.

Downloading to primary flash using Xmodem (Menu)

The menu interface accesses only the primary flash.

Procedure

1. From the console Main Menu, select **7. Download OS**
2. Press **[E]** (for **Edit**) on the keyboard.
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.
5. The following message appears:
6. **Press enter and then initiate Xmodem transfer from the attached computer.....**
7. Press **[Enter]** and then execute the terminal emulator commands to begin Xmodem binary transfer.
8. For example, using HyperTerminal:
   a. Click on **Transfer**, then **Send File**.
   b. Enter the file path and name in the Filename field.
   c. In the Protocol field, select **Xmodem**.
   d. Click on the **[Send]** button.
9. The download then commences. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.
10. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**.) You then see the following prompt:
11. **Continue reboot of system? : No**
12. Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.
13. To confirm that the software downloaded correctly:
   a. From the Main Menu, select **1. Status and Counters**, and the select **1. General System Information**
   b. Check the **Firmware revision** line.

USB

Downloading switch software using USB

Enable or disable the USB port

This feature allows configuration of the USB port using either the CLI or SNMP.

Prerequisites

Procedure

1. Store a software version for the switch on a USB flash drive. (The latest software file is typically available from the Switch Networking website at **http://www.hpe.com/networking/support**.)
2. Insert the USB device into the switch's USB port.
3. Determine the name of the software file stored on the USB flash drive (for example, `K.0800.swi`).
4. Decide whether the image will be installed in the primary or secondary flash.

USB port status

show usb-port

Syntax

`show usb-port`

Description

Displays the status of the USB port. It can be enabled, disabled, or not present.

Command context

`operator`

Usage

One of the following messages indicates the presence or absence of a USB device:

- `Not able to sense device in USB port`
- `USB device detected in port`
- `no USB device detected in port`

Example

Display USB port status.

```
switch# show usb-port

  USB port status: enabled
  USB port power status: power on    (USB device detected in port)
```

Copy a switch software file named `K.0800.swi` from a USB device to primary flash.

Procedure

1. Issue the `copy usb flash` command as shown below:

   Figure 107: *The command to copy switch software from USB*

   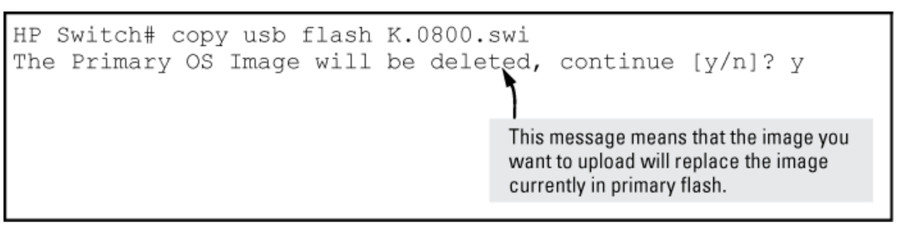

2. When the switch finishes copying the software file from the USB device, it displays this progress message: `Validating and Writing System Software to the Filesystem....`
3. When the save finishes, you must reboot the switch to load the newly loaded software.
4. To confirm that the software downloaded correctly, execute `show system` and check the **Software revision** line.

copy usb flash

Syntax

`copy usb flash <FILENAME> [primary|secondary]`

Description

This command automatically downloads a switch software file to primary or secondary flash. If you do not specify the flash destination, the USB download defaults to primary flash.

Switch to Switch

Switch-to-switch download

You can use TFTP to transfer a software image between two switches of the same series. The CLI enables all combinations of flash location options. The menu interface enables you to transfer primary-to-primary or secondary-to-primary.

OS download from another switch

Where two switches in your network belong to the same series, you can download a software image between them by initiating a `copy tftp` command from the destination switch. The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

copy tftp flash

Syntax

`copy tftp flash <IP-ADDR>  flash [primary|secondary][oobm]`

Description

When executed in the destination switch, downloads the software flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

Parameters and options

`primary`

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

`secondary`

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

`oobm`

For switches that have a separate OOBM port, the `oobm` parameter specifies that the TFTP traffic must come in through the OOBM interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Download from primary flash

To download a software file from primary flash in a switch with an IP address of 10.29.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:

Figure 108: *Switch-to-switch, from primary in source to either flash in destination*

Running Total of Bytes Downloaded

```
HP Switch# copy tftp flash 10.29.227.13 flash
Device will be rebooted, do you want to continue [y/n]? y
00107K
```

copy tftp flash os

Syntax

```
copy tftp flash <IP-ADDR>  [/os/primary|/os/secondary][primary|secondary] [oobm]
```

Description

This command (executed in the destination switch) gives you the most options for downloading between switches. If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

Parameters and options

oobm

For switches that have a separate out-of-band management port, the `oobm` parameter specifies that the TFTP traffic must come in through the out-of-band management interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The `oobm` parameter is not available on switches that do not have a separate out-of-band management port.

Switch-to-switch, from either flash in source to either flash in destination

To download a software file from secondary flash in a switch with an IP address of 10.28.227.103 to the secondary flash in a destination switch, you would execute the following command in the destination switch's CLI:

```
switch# copy tftp flash 10.29.227.13 flash /os/secondary secondary
Device will be rebooted, do you want to continue [y/n]? y
00184K
```

Switch-to-switch download to primary flash (Menu)

Using the menu interface, you can download a switch software file from either the primary or secondary flash of one switch to the primary flash of another switch of the same series.

Procedure

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default.)
3. In the **TFTP Server** field, enter the IP address of the remote switch containing the software file you want to download.
4. For the **Remote File Name** , enter one of the following:
   a. To download the software in the primary flash of the source switch, enter

      `flash`

      in lowercase characters.
   b. To download the software in the secondary flash of the source switch, enter

      `/os/secondary`

      .
5. Press **[Enter]**, and then **[X]** (for **eXecute**) to begin the software download.
6. A "progress" bar indicates the progress of the download. When the entire switch software download has been received, all activity on the switch halts and the following messages appear:
7. **Validating and writing system software to FLASH...**
8. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch** .) You then see this prompt:
9. **Continue reboot of system? : No**

10. Press the space bar once to change `No` to `Yes`, then press **[Enter]** to begin the reboot.
11. To confirm that the software downloaded correctly:
    a. From the Main Menu, select**Status and Counters General System Information**
    b. Check the **Firmware revision** line.

Copying

Software images

copy flash tftp

Syntax

copy flash tftp *<IP-ADDR>* *<FILENAME>* [oobm]

Description

Copies the primary flash image to a TFTP server.

Parameters and options

oobm

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Copy primary flash to TFTP

To copy the primary flash to a TFTP server having an IP address of 10.28.227.105:

```
switch# copy flash tftp 10.28.227.105 K.0800.swi
```

where `K.0800.swi` is the filename given to the flash image being copied.

copy flash xmodem

Syntax

copy flash xmodem [pc|unix]

Description

Uses Xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation.To use this method, the switch must be connected via the serial port to a PC or UNIX workstation.

Copy primary flash

To copy the primary flash image to a serially connected PC, execute the copy xmodem flash command:

```
switch# copy xmodem flash
Press 'Enter' and start XMODEM on your host...
```

At the prompt, press **Enter** on the keyboard, and then execute the terminal emulator commands to begin the file transfer.

Copying using USB

To copy the primary image to a USB flash drive:

Procedure

1. Insert a USB device into the switch's USB port.
2. Execute the command: `switch# copy flash usb K.0800.swi primary/secondary` where `K.0800.swi` is the name given to the primary flash image that is copied from the switch to the USB device.

copy flash usb

Syntax

```
copy flash usb <FILENAME>
```

Description

Uses the USB port to copy the specified flash image from the switch to a USB flash memory device. The default setting will use the primary image.

Copying diagnostic data to a remote host, USB device, PC, or UNIX workstation

copy command-output

Syntax

```
copy command-output <CLI-COMMAND> tftp <IP-ADDRESS> <FILEPATH-FILENAME> [oobm]
```

```
copy command-output <CLI-COMMAND> usb <FILENAME>
```

```
copy command-output <CLI-COMMAND> xmodem
```

Description

These commands direct the displayed output of a CLI command to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Use Xmodem to copy the output of show config

The command you specify must be enclosed in double quotation marks.

Figure 109: *Sending command output to a file on an attached PC*

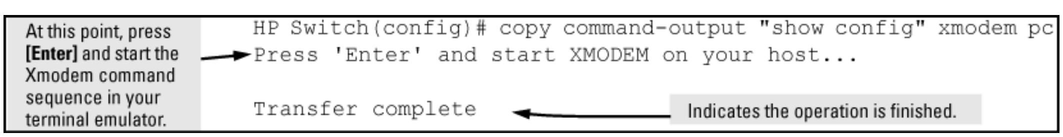

copy command-log

Syntax

```
copy command-log [sftp | smm | tftp | usb | xmodem]
```

Description

This command copies the Command Log content to a remote host or to a serially-connected PC or UNIX workstation.

- Use the `sftp` option to copy data to an SFTP server.
- Use the `smm` option to copy AMM and SMM log events.
- Use the `tftp` option to copy data to a TFTP server.
- Use the `usb` option to copy data to a USB flash drive.
- Use the `xmodem` option to copy data to the console using XMODEM.

copy event-log

Syntax

```
copy event-log [smm] [tftp <IP-ADDRESS> <FILEPATH_FILENAME> [oobm]] [usb
<FILENAME>] [xmodem <FILENAME>]
```

Description

These commands copy the Event Log content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

Parameters and options

smm

Copies the entire Event Log, both active management module events and standby management module events, to the selected host, USB device, or serially connected PC or UNIX workstation.

oobm

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Copy the event log to a PC connected to the switch

Figure 110: *Sending event log content to a file on an attached PC*

At this point, press **[Enter]** and start the Xmodem command sequence in your terminal

```
HP Switch(config)# copy event-log xmodem pc
Press 'Enter' and start XMODEM on your host...

Transfer complete
```

copy crash-data

Syntax

copy crash-data [<*SLOT-ID*>|master] tftp <*IP-ADDRESS*> <*FILENAME*> [oobm]

copy crash-data [<*SLOT-ID*>|mm] usb <*FILENAME*>

copy crash-data [<*SLOT-ID*>|mm] xmodem

Description

These commands copy the crash data content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation using TFTP, USB, or Xmodem. You can copy individual slot information or the management module's switch information. If you do not specify either, the command defaults to the management function's data. You can copy individual slot information or the management module (mm) switch information. If you do not specify either, the command defaults to the mm data.

Parameters and options

<*SLOT-ID*>

a - h—Retrieves the crash log or crash data from the processor on the module in the specified slot

mm

Retrieves crash log or crash data from the switch's chassis processor. When "mm" is specified, crash files from both management modules are copied.

oobm

For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

Copy crash data file to a PC

Figure 111: *Copying switch crash data content to a PC*

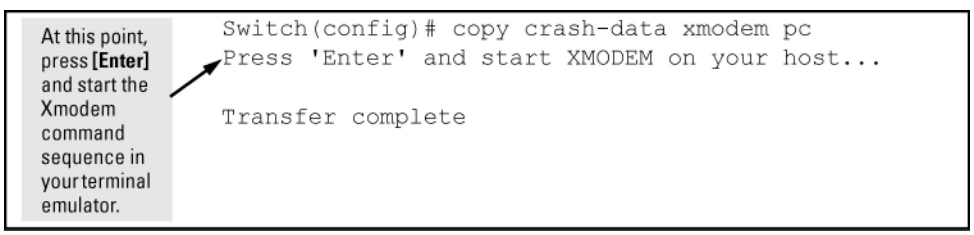

copy crash-data (redundant management)

Syntax

copy crash-data [<*SLOT-ID*>|mm] tftp <*IP-ADDRESS*> <*FILENAME*> [oobm]

Description

Copies the crash data of both the active and standby management modules to a user-specified file. With no parameter specified, concatenates files from all modules (management and interface).

Parameters and options

<SLOT-ID>

Retrieves the crash log or crash data from the module in the specified slot

mm

Retrieves the crash data from both management modules and concatenates them.

oobm

For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

copy crash-log

Syntax

```
copy crash-log [<SLOT-ID>|mm] tftp <IP-ADDRESS> <FILEPATH\FILENAME> [oobm]
```

```
copy crash-log [<SLOT-ID>|mm] usb <FILENAME>
```

```
copy crash-log [<SLOT-ID>|mm] xmodem
```

Description

Copies the crash log content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation. You can copy individual slot information or the management module (`mm`) switch information. If you do not specify either `mm` or `oobm`, the command defaults to `mm` data.

Parameters and options

<SLOT-ID>

a - h—Retrieves the crash log from the processor on the module in the specified slot.

mm

Retrieves the crash log from the switch's chassis processor. With mm specified, copies crash files from both management modules.

oobm

For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

Copy the crash log for slot C to a file in a PC connected to the switch

Figure 112: *Sending a crash log for slot C to a file on an attached PC*

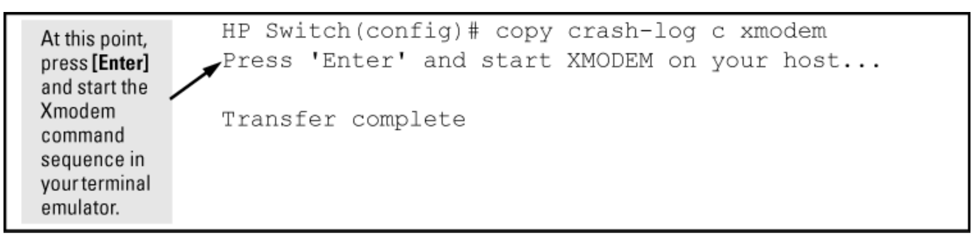

copy crash-log (redundant management)

Syntax

```
copy crash-log [<SLOT-ID>|mm] tftp <IP-ADDRESS> <FILENAME> [oobm]
```

Description

Copies the crash logs of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

Parameters and options

<SLOT-ID>

Retrieves the crash log or crash log from the module in the specified slot.

mm

Retrieves the crash data from both management modules and concatenates them.

oobm

For switches with a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

copy core-dump (standby module)

Syntax

`copy core-dump [mm usb <FILENAME>|standby flash|usb <FILENAME>]`

Description

Copies the management module coredump, or the standby management module coredump, to the active management module flash or to a USB flash drive, (see **Figure 113: Copying the standby coredump to flash** on page 436.) During the copy, the system displays the number of bytes transferred and the percentage of the total. Management module core files can be quite large. Use **Cntl-C** to cancel the transfer.

Make sure that the coredump files on the standby management module are accessible for diagnostic purposes.

Parameters and options

flash

Copies the core file of the standby management module to the flash of the active management module. The destination file is fixed as `dumpM1.cor` or `dumpM2.cor`, depending on which module is the standby management module.

usb *<FILENAME>*

Copies the management module's core file or the standby management module's core file to a USB flash drive. The optional filename defaults to `dumpM1.cor` or `dumpM2.cor`, depending on which module is the standby management module.

Copy the standby coredump to flash

Figure 113: *Copying the standby coredump to flash*

```
HP Switch(config)# copy core-dump standby flash
02816K of 26899K  (10%)
```

If there is no coredump on the standby management module, the following error message displays:

```
Standby MM coredump does not exist.
```

If there is not enough destination space before or during the transfer to flash or USB, the following error message displays:

```
Insufficient FLASH space to complete the file copy.
```

copy fdr-log

Syntax

```
copy fdr-log [slot <SLOT-LIST>|mm-active [current|previous]|mm-standby|all] tftp
[<HOSTNAME>|<IP-ADDR>] <FILENAME>
```

Description

Copies `fdr-log` files to a user-specified file. The FDR log collects information when the switch is not performing correctly, but has not crashed. Writes runtime logs to FDR memory while the switch is running. Crashtime logs are collected and stored in the FDR buffer during a switch crash.

Parameters and options

all

Copies all the log files from both management modules and all slots.

mn-active

Copies the active management module's log.

mn-standby

Copies the standby management module's log.

slot

Retrieves the crash log from the module in the identified slots.

Copy diagnostic data to a remote host, USB device, PC or UNIX workstation

You can use the CLI to copy the following types of switch data to a text file on a destination device:

Command output

Sends the output of a switch CLI command as a file on the destination device.

Event log

Copies the switch's Event Log into a file on the destination device.

Crash data

Software-specific data useful for determining the reason for a system crash.

Crash log

Processor-specific operating data useful for determining the reason for a system crash.

Flight data recorder (FDR) logs

Information that is “interesting” at the time of the crash, as well as when the switch is not performing correctly but has not crashed.

The destination device and copy method options include:

- Remote Host using `TFTP`.
- Physically connected USB flash drive using the `USB` port on the switch.
- Serially connected PC or UNIX workstation using Xmodem.

Transferring

Switch configuration transfer

Using CLI commands you can copy switch configurations to and from a switch, or copy a software image to configure or replace an ACL in the switch configuration.

For greater security, you can perform all TFTP operations using SFTP.

You can also use the `include-credentials` command to save passwords, secret keys, and other security credentials in the running config file.

TFTP

copy [startup-config|running-config]

Syntax

```
copy [startup-config|running-config] tftp <IP-ADDR> <REMOTE-FILE> [pc|unix][oobm]

copy config  <FILENAME>  tftp  IP-ADDR <REMOTE-FILE> [pc|unix][oobm]
```

Description

Copy a designated config file in the switch to a TFTP server. For more information, see the basic operation guide.

Parameters and options

oobm

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Upload current startup configuration

To upload the current startup configuration to a file named **`sw8200`** in the configs directory on drive **"d"** in a TFTP server having an IP address of 10.28.227.105:

```
ProCurve# copy startup-config tftp 10.28.227.105
   d:\configs\sw8200
```

copy tftp

Syntax

```
copy tftp [startup-config|running-config] tftp <IP-ADDR> <REMOTE-FILE> [pc|unix]
[oobm]

copy tftp config <FILENAME> <IP-ADDR> <REMOTE-FILE> [pc|unix][oobm]
```

Description

Copies a configuration from a remote host to a designated config file in the switch.

Parameters and options

`oobm`

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Download a configuration file

To download a configuration file named **`sw8200`** in the **`configs`** directory on drive **"d"** in a remote host having an IP address of 10.28.227.105:

```
switch# copy tftp startup-config 10.28.227.105
   d:\configs\sw8200
```

copy tftp show-tech

Exit the global config mode (if needed) before executing `show tech` commands.

Syntax

`copy tftp show-tech ipv4 or ipv6 address <filename> [oobm]`

Copies a customized command file to the switch. Using the `copy tftp` command with the `show-tech` option provides the ability to copy a customized command file to the switch.

Parameters and options

`show-tech`

Allows you to copy a customized command file to the switch.

`oobm`

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the out-of-band management interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port.

Upload a customized command file

```
switch(config)# copy tftp show-tech 10.10.10.3 commandfile1
```

show tech custom

Syntax

`show tech custom`

Description

Executes the commands found in a custom file instead of the hard-coded list. If no custom file is found, executes the current hard-coded list. This list contains commands to display data, such as the image stamp, running configuration, boot history, port settings, and so on. You can include `show tech` commands in the custom file, with the exception of `show tech custom`. For example, you can include the command `show tech all`.

No show-tech file found

If no custom file is found, a message displays stating "No SHOW-TECH file found." (No custom file was uploaded with the `copy tftp show-tech` command.)

```
switch# show tech custom
No SHOW-TECH file found.
```

copy tftp config

Syntax

```
copy tftp config <SOURCE CONFIG FILE NAME> <DESTINATION_IP-ADDRESS>
          <DESTINATION CONFIG FILE> [detail|oobm|pc|unix]
```

Description

Displays the progress, in lines and percentages, of the configuration file copied to or from the switch. A large configuration file takes several minutes to transfer. This feature allows the customer to watch the progress.

Parameters and options

detail

Display copy progress.

oobm

Use the OOBM interface to reach TFTP server.

pc

Change CR/LF to PC style.

unix

Change CR/LF to unix style

copy tftp config

```
HP-Switch-5406Rzl2# copy tftp config myConfig 10.100.0.12 myConfig.cfg oobm detail
Processing line 4968 of 20740 (23%)
```

Xmodem

Prerequisites

- Connect the switch to a PC or UNIX workstation using the serial port
- Determine the filename.
- Know the directory path you will use to store the configuration file.

copy config xmodem

Syntax

```
copy [startup-config|running-config] xmodem [pc|unix][oobm]
```

```
copy config <FILENAME> xmodem [pc|unix]
```

Description

Uses xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation.

Copy a configuration file to a PC

```
switch# copy startup-config xmodem pc
Press 'Enter' and start XMODEM on your host...
```

Execute the terminal emulator commands to begin the file transfer.

copy xmodem startup-config

Syntax

`copy xmodem startup-config [pc|unix]`

`copy xmodem config <FILENAME> [pc|unix]`

Description

Copies a configuration file from a serially connected PC or UNIX workstation to a designated configuration file on the switch.

> When the download finishes, you must reboot the switch to implement the newly downloaded software (see **boot system flash** on page 441 and **reload** on page 441).

Copy a configuration file from a PC

```
switch# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

Execute the terminal emulator commands to begin the file transfer.

boot system flash

Syntax

`boot system flash [primary|secondary]`

`boot system flash [config <FILENAME>]`

Description

Used to boot switches from the designated configuration file.

reload

Syntax

`reload`

Description

Reboots from the flash image currently in use.

USB

Be sure to connect a USB flash memory device to the USB port on the switch.

copy startup-config

Syntax

`copy startup-config usb <FILENAME>`

`copy running-config usb <FILENAME>`

Description

Copies the startup configuration or the running configuration to a USB flash drive.

copy startup-config

```
switch# copy startup-config usb HP Switch-config
```

`HP Switch-config` is the name given to the configuration file that you copy from the switch to the USB device.

copy usb startup-config

Syntax

`copy usb startup-config <FILENAME>`

Description

Copies a configuration file from a USB device to the startup configuration file on the switch. To execute the command, you must know the name of the file to copy.

copy usb startup-config

```
switch# copy usb startup-config HP Switch-config
```

ACL command file transfer

This section describes how to upload and execute a command file to the switch for configuring or replacing an ACL in the switch configuration. Such files should contain only access control entry (ACE) commands.

tftp

copy tftp command-file

Syntax

`copy tftp command-file <IP-ADDR> <FILENAME.TXT> [unix|pc][oobm]`

Description

Copies and executes the named text file from the specified TFTP server address and executes the ACL commands in the file. Depending on the ACL commands used, this action does one of the following in the `running-config` file:

- Creates a new ACL.
- Replaces an existing ACL.
- Adds to an existing ACL.

Parameters and options

<IP-ADDR>

The IP address of a TFTP server available to the switch.

<FILENAME.TXT>

A text file containing ACL commands and stored in the TFTP directory of the server identified by *<IP-ADDR>* .

unix|pc

The type of workstation used for serial, Telnet, or SSH access to the switch CLI.

oobm

For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)

Upload an ACL command file from a PC

```
switch(config)# copy tftp command-file 18.38.124.16
vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue
[y/n]?
```

If the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice (as shown in **Figure 114: Using the copy command to download and configure an ACL** on page 443), and continues to implement the remaining ACL commands in the file.

Figure 114: *Using the* `copy` *command to download and configure an ACL*

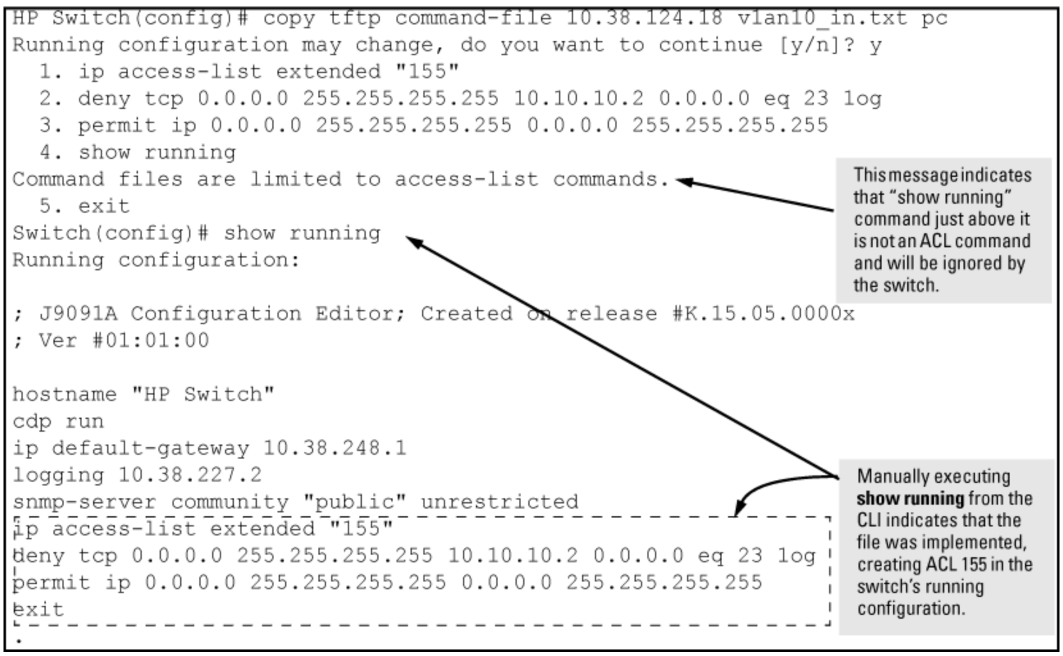

Xmodem

copy xmodem command-file

Syntax

`copy xmodem command-file [unix|pc]`

Description

Uses Xmodem to copy and execute an ACL command from a PC or UNIX workstation. Depending on the ACL commands used, this action does one of the following in the running-config file:

- Creates a new ACL.
- Replaces an existing ACL.
- Adds to an existing ACL.

USB

copy usb command-file

Syntax

```
copy usb command-file <FILENAME.TXT> [unix|pc]
```

Description

Copies and executes the named text file from a USB flash drive and executes the ACL commands in the file. Depending on the ACL commands used, this action does one of the following in the running-config file:

- Creates a new ACL.
- Replaces an existing ACL.
- Adds to an existing ACL.

Parameters

<FILENAME.TXT>

A text file containing ACL commands and stored in the USB flash drive.

unix|pc

The type of workstation used to create the text file.

Upload an ACL command file from USB

Using a PC workstation, execute the following from the CLI to upload the file to the switch and implement the ACL commands it contains:

```
switch(config)# copy usb command-file vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue
[y/n]?
```

If the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice, and continues to implement the remaining ACL commands in the file.

Switch software download

The terms **switch software and software image** refer to the downloadable software files the switch uses to operate its networking features. Other terms sometimes include **Operating System, or OS**.

Switch periodically provides switch software updates through the Switch Networking website. For more information, see the support and warranty booklet shipped with the switch, or visit **http://www.hpe.Com/Networking/Support**.

Switch software download rules

Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See **copy [startup-config|running-config]** on page 438.

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash.

- Switch software that you download using the menu interface always goes to primary flash.
- After a software download, reboot the switch to implement the new software. Until a reboot occurs, the switch continues to run on the software it was using before the download.

TFTP download failures

Symptom

When using the menu interface, if a TFTP download fails, the Download OS (Operating System, or software) screen indicates the failure.

Figure 115: *Example of message for download failure*

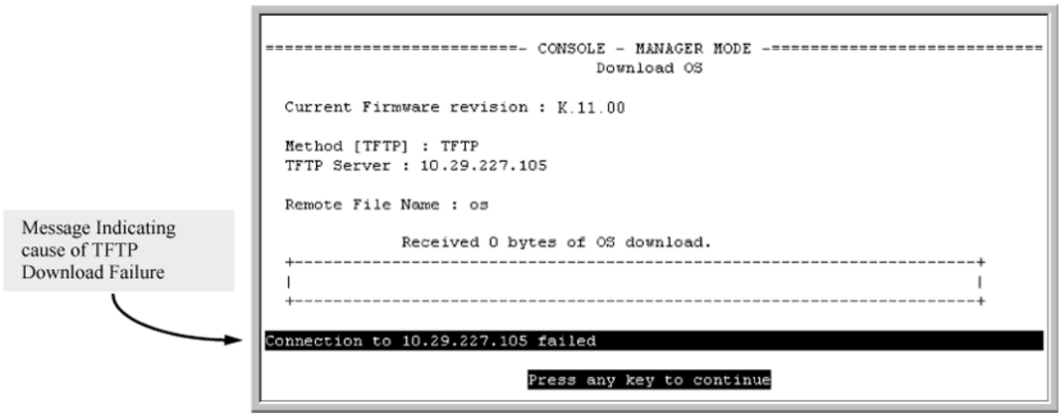

Cause

Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a UNIX machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the **Download OS** (Operating System, or software) screen.
- One or more of the switch's IP configuration parameters are incorrect.
- For a UNIX TFTP server, the file permissions for the software file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

Action

1.

Single copy command

When a switch crashes, five files relating to the crash; core-dump, crash-data, crash-log, fdr-log, and event-log are created and should be copied for review. All five files (core-dump, crash-data, crash-log, fdr-log, and event-log) should be copied to a destination specified under a directory by file name.

TFTP

A destination directory and files can be created for all crash files (core-dump, crash-data, crash-log, fdr-log, and event-log) on an TFTP server (with write permissions).

SFTP

Files are auto created on the SFTP server as a secured transfer. The destination directories however can be manually created on the server.

You can use specified directories for the TFTP/SFTP transfers in the `copy` command. If you specify a directory, the command copies all files under one directory. With no directory specified, the command copies all files to the TFTP/SFTP server home directory. You must specify a directory name.

copy source

Syntax

`copy <SOURCE> <DATA_FILE> <DESTINATION> <DATA_FILE> <OPTIONS>`

Description

Copies data files to and from the switch.

Parameters and options

<SOURCE>

Specify the source of data using any of the following destinations.

Flash

N/A

SFTP

For transfer of crash-files via SFTP, the destination directory must exist on the SFTP server with write permissions.

File creation is not mandatory as files are automatically created with the chassis serial number suffix to the filename when using SFTP.**See: example**

TFTP

For transfer of crash-files via TFTP, the destination directory along with the file names (core-dump, crash-data, crash-log, fdr-log, and event-log) must exist on the TFTP server with write permissions.

USB

For transfer of crash-files via USB, the destination directory along with the file names (core-dump, crash-data, crash-log, fdr-log, and event-log) must exist on the device with write permissions.

Xmodem

N/A

<DATA_FILES> **SHAWN/ERIN: Data Files was not part of the syntax, so I added it.**

Specify the data file to be copied from the source.

`command-output <COMMAND>`

Specify a command to copy output. When using `command-output`, place the desired CLI command in double-quotes. For example: "show system".

`config <FILE-NAME>`

Copy named configuration file. The `file-name` option is the source configuration file being copied.

`core-dump`

Copy core-dump file from flash.

crash-data

Copy the switch crash-data file.

crash-log

Copy the switch crash-log file.

`crash-files` <*A|B|C|D|E|F|G|H|MASTER*>

Copy core-dump, crash-data, crash-log, fdr-log, and event-log files to an SFTP/TFTP server, USB, or xmodem terminal.

When using the `crash-files` option, the destination directory alone must be specified as the destination path. Specifying the file names is not mandatory.

default-config

Copy custom default-config file.

event-log

Copy event-log file.

fdr-log

Copy FDR-og file from the switch to an SFTP/TFTP server, USB or xmodem terminal.

flash

Copy the switch system image file.

SFTP server

Copy data from a SFTP server.

startup-config

Copy in-flash configuration file.

ssh-client-known-hosts

Copy the known hosts file.

ssh-server-pub-key

Copy the switch's SSH server public key.

running-config

Copy running configuration file.

TFTP

Copy data from a TFTP server.

USB

Copy data from a USB flash drive.

xmodem

Use xmodem on the terminal as the data source.

<*DESTINATION* >

Specify the copy target.

SFTP
TFTP

USB
xmodem

<DATA_FILES>

Specify the data file name at the target.

autorun-cert-file
autorun-key-file
command file
config
default-config
flash
pub-key-file
show-tech
startup-config
ssh-client-key
ssh-client-known-hosts

<OPTIONS>

append

Add the keys for operator access.

directory

Directory name to upload. Required for TFTP, SFTP and USB transfers.

filename

File-name to upload/download. Required for TFTP, SFTP and USB transfers.

hostname

Hostname of the TFTP, SFTP server. Required for TFTP, SFTP transfers.

IPv4 address

TFTP, SFTP server IPv4 address. Required for TFTP, SFTP transfers.

IPv6 address

TFTP, SFTP server IPv6 address. Required for TFTP, SFTP transfers.

manager

Replace the keys for manager access; follow with the `append` option to add the keys.

operator

Replace the keys for operator access (default); follow with the `append` option to add the keys.

pc

N/A

unix

N/A

The listed crash-files captured for 3500 switch for both MM and slot using SFTP are as follows:

- MM crash-files:

  M-SG238TF00K.cor
  M-SG238TF00K.cdata

M-SG238TF00K.clog
M-SG238TF00K.evt
M-SG238TF00K.fdr

- Slot crash-files:

  I-SG238TF00K.cor
  I-SG238TF00K.cdata
  I-SG238TF00K.clog
  I-SG238TF00K.evt
  I-SG238TF00K.fdr

copy crash-files

Syntax

`copy crash-files [slot-id|mm-active|mm-standby|member]`

Description

Copies multiple management switches.

Parameters and options

`slot-id`

Copy interface management crash files to SFTP, TFTP, USB, and Xmodem.

`mm-active`

Copy active management module crash files to SFTP, TFTP, USB, and Xmodem.

`mm-standby`

Copy standby management module crash files to SFTP, TFTP, USB, and Xmodem.

copy crash-files member

Syntax

`copy crash-files member [management|interfaces}`

Description

Copies stacking or standalone switches.

Parameters and options

`management`

Copy stack member crash files to SFTP, TFTP, USB, and Xmodem.

`interfaces`

Copy stack member crash files to SFTP, TFTP, USB, and Xmodem.

copy crash-files crash-file-options

Syntax

`copy crash-files crash-file-options <HOST-NAME-STR> <IP-ADDR> <IPv6-ADDR> <SFTP> <DIRNAME-STRX> [oobm] <DESTINATION>`

Description

Copies crash files using various options.

Parameters and options

<HOST-NAME-STR>

Specify hostname of the SFTP server.

<IP-ADDR>

Specify SFTP server IPv4 address.

<IPv6-ADDR>

Specify SFTP server IPv6 address.

<USER>

Specify the username on the remote system.

<USERNAME@IP-STR>

Specify the username along with remote system. Information (hostname, IPv4 or IPv6 address).

<DIRNAME-STR>

Specify the destination directory name.

oobm

Use the OOBM interface to reach SFTP server.

<DESTINATION>

slot-id

Copy interface core-dump file.

mm-active

Copy active management module crash files.

mm-standby

Copy standby management module crash files.

member

Copy member crash files.

interfaces

Copy interfaces crash files.

management

Copy management crash files.

Switch and network operations

The switches have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- Status Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data.
- Counters Display details of traffic volume on individual ports (**Accessing port and trunk statistics (Menu)** on page 464.)
- Event Log Lists switch operating events. See the HPE ProVision switch software troubleshooting guide for troubleshooting information.
- Configurable trap receivers Uses SNMP to enable management stations on your network to receive SNMP traps from the switch.
- Port monitoring (mirroring) Copy all traffic from the specified ports to a designated monitoring port .

NOTE

Link test and ping test—analysis tools in troubleshooting situations—are described in the *ProVision Switch Software Troubleshooting Guide*.

Status and counters data

This section describes the status and counters screens available through the switch console interface and/or the WebAgent.

NOTE

You can access all console screens from the WebAgent via Telnet to the console. Telnet access to the switch is available in the **Device View** window under the **Configuration** tab.

Accessing status and counters (Menu)

Procedure

1. Beginning at the Main Menu, select **1. Status and Counters**.

   Figure 116: *The Status and Counters menu*

```
==========================- CONSOLE - MANAGER MODE -============================
                            Status and Counters Menu

   1. General System Information
   2. Switch Management Address Information
   3. Module Information
   4. Port Status
   5. Port Counters
   6. Vlan Address Table
   7. Port Address Table
   8. Spanning Tree Information
   0. Return to Main Menu...

Displays switch management information including software versions.
To select menu item, press item number, or highlight item and press <Enter>.
```

   Each of the above menu items accesses the read-only screens described on the following pages. See the online help for a description of the entries displayed in these screens.

show system

Syntax

`show system [chassislocate | information | temperature]`

Description

Shows global system information and operational parameters for the switch.

Command context

`manager and operator`

Parameters

`chassislocate`

Shows the chassis locator LED status. Possible values are ON, Off, and Blink. When the status is On or Blink, the number of minutes that the Locator LED will continue to be on or to blink is displayed.

`information`

Displays global system information and operational parameters for the switch.

`temperature`

Shows system temperature and settings.

Usage

- To show system fans, see **show system fan**.
- To show chassis power supply and settings, see **show system power-supply** .
- To show system fans for VSF members, see **show system fans**.

Examples

Locating the system chassis by LED blink using the `show system chassislocate` command.

```
HP Switch(config)# show system chassislocate

Chassis Locator LED: ON 5 minutes 5 seconds

HP Switch(config)# show system chassislocate

Chassis Locator LED: BLINK 10 minutes 6 seconds

HP Switch(config)# show system chassislocate

Chassis Locator LED: OFF
```

Showing the general switch system information by using the `show system` command.

```
HP Switch(config)# show system

 Status and Counters - General System Information

  System Name         : HP Switch Switch
  System Contact      :
  System Location     :

  MAC Age Time (sec) : 300

  Time Zone           : 0
  Daylight Time Rule : None


  Software revision   : T.13.XX           Base MAC Addr      : 001635-b57cc0
  ROM Version         : K.12.12           Serial Number      : LP621KI005

  Up Time             : 51 secs           Memory    - Total  : 152,455,616
  CPU Util (%)        : 3                             Free   : 110,527,264

  IP Mgmt   - Pkts Rx : 0                 Packet    - Total  : 6750
              Pkts Tx : 0                 Buffers     Free   : 5086
                                                      Lowest : 5086
                                                      Missed : 0
```

chassislocate

Syntax

Description

Identifies the location of a specific switch by activating the blue locator LED on the front panel of the switch.

`chassislocate [blink|on|off]`

Parameters and options

`blink` ***`<1-1440>`***

Blinks the chassis locate LED for a specified number of minutes (Default: 30 min.)

`on` ***`<1-1440>`***

Turns the chassis locate LED on for a specified number of minutes (Default: 20 min.)

`off`

Turns the chassis locate LED off.

Chassislocate at startup

The chassislocate command has an optional parameter that configures it to run in the future instead of immediately.

Syntax

`chassislocate [on|blink] <MINUTES> at [now|startup]`

```
chassislocate off
```

Parameters and options

<MINUTES>

Specify the number of minutes for the chassis locate LED to remain on or blink.

at

Specify when the command is applied (default immediately.)

now

Turn on the chassis locate LED immediately.

startup

Turn on the chassis locate LED at switch startup.

off

Turn off the chassis locate LED switch

---

chassislocate at startup

```
chassislocate blink 10 at startup
```

show system chassislocate

Syntax

```
show system chassislocate
```

Description

Displays the current status of the `chassislocate` settings.

Display locator LED status

```
Locator LED Status
            Current    Time
Member      State      Remaining  Configuration
------      -------    ---------  -------------
1           blink      00:27:05   blink 30 at startup
2           on         01:05:27
3           off
```

Collecting processor data with the task monitor

The task monitor feature allows you to enable or disable the collection of processor utilization data. The `task-monitor cpu` command is equivalent to the existing debug mode command `taskusage -d`. (The `taskUsageShow` command is also available.)

When the `task-monitor` command is enabled, the `show cpu` command summarizes the processor usage by protocol and system functions.

task-monitor cpu

Syntax

[no] `task-monitor cpu`

Description

Enables or disables the collection of processor utilization data, and requires a manager log in. Settings are not persistent; there are no changes to the configuration. Defaults to disabled.

task-monitor cpu command

Figure 117: *The `task-monitor cpu` command and `show cpu` output*

```
HP Switch(config)# task-monitor cpu
HP Switch(config)# show cpu

2 percent busy, from 2865 sec ago
1 sec ave: 9 percent busy
5 sec ave: 9 percent busy
1 min ave: 1 percent busy


 % CPU | Description
-------+-------------------------
    99 | Idle
```

Accessing system information (Menu)

From the console Main Menu, select **1. Status and Counters**, and then select **1. General System Information**.

Figure 118: *Example of general switch information*

```
==========================- CONSOLE - MANAGER MODE -============================
                Status and Counters - General System Information

  System Contact     :
  System Location    :

  Firmware revision  : K.11.00          Base MAC Addr      : 0001e7-a09900
  ROM Version        : K.11.Z4          Serial Number      : S2600017409

  Up Time            : 2 hours          Memory    - Total  : 24,588,136
  CPU Util (%)       : 1                            Free   : 19,613,568

  IP Mgmt   - Pkts Rx : 0               Packet    - Total  : 832
              Pkts Tx : 0               Buffers     Free   : 793
                                                    Lowest : 769
                          24,588,1 6                Missed : 0


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

This screen dynamically indicates how individual switch resources are being used. See the online help for details.

Switch management address information access

show management

Syntax

`show management`

Description

Displays switch management address information.

Accessing switch management address information (Menu)

From the Main Menu, select **1. Status and Counters ...** , and then select **2. Switch Management Address Information**.

Figure 119: *Example of management address information with VLANs configured*

```
==========================- CONSOLE - MANAGER MODE -============================
              Status and Counters - Management Address Information

  Time Server Address : Disabled

   VLAN Name     MAC Address           IP Address
  ------------  --------------------  --------------------
  DEFAULT_VLAN  0001e7-a09900         10.28.227.101
  VLAN-22       0001e7-a09900         Disabled
  VLAN-33       0001e7-a09900         Disabled


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

This screen displays addresses that are important for management of the switch. If multiple VLANs are **not configured**, this screen displays a single IP address for the entire switch. See the online help for details.

As shown in **Figure 119: Example of management address information with VLANs configured** on page 456, all VLANs on the switches use the same **MAC address**. (This includes both the statically configured VLANs and any dynamic VLANs existing on the switch as a result of GVRP operation.)

Also, the switches use a multiple forwarding database. When using multiple VLANs and connecting a switch to a device that uses a single forwarding database, such as a Switch 4000M, there are cabling and tagged port VLAN requirements.

Component information views

The CLI `show modules` command displays additional component information for the following:

- SSM—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, "J" number, and serial number (when available)

show modules

Syntax

`show modules`

Description

Displays information about the installed modules (**Figure 120: The show modules command output** on page 457), including:

- The slot in which the module is installed
- The module description
- The serial number

Additionally, this command displays the part number (J number) and serial number of the chassis. (See **Figure 121: The show modules details command for the 8212zl, showing SSM and mini-GBIC information** on page 457.)

show modules command

Figure 120: *The show modules command output*

```
HP Switch(config)# show modules

 Status and Counters - Module Information

  Chassis: 5406zl J8697A         Serial Number:   SG560TN124
  Slot  Module Description                       Serial Number
  ----- ---------------------------------------- --------------
  A     HP Switch J8706A 24p SFP zl Module        AD722BX88F
  B     HP Switch J8702A 24p Gig-T zl Module      FE999CV77F
  C     HP Switch J8707A 4p 10-Gbe zl Module      FB345DC99D
```

show modules details command

Figure 121: *The show modules details command for the 8212zl, showing SSM and mini-GBIC information*

```
HP Switch(config)# show modules details

 Status and Counters - Module Information

  Chassis: 8212zl J8715A         Serial Number:   SG560TN124
  Slot  Module Description                         Serial Number   Status
  ----- ----------------------------------------   --------------  -------
  MM1   HP Switch J9092A Management Module 8200zl   AD722BX88F      Active
  SSM   HP Switch J8784A System Support Module      AF988DC78G      Active
  C     HP Switch J8750A 20p +4 Mini-GBIC Module    446S2BX007      Active
         GBIC 1: J4859B  1GB LX-LC                 4720347DFED734
         GBIC 2: J4859B  1GB LX-LC                 4720347DFED735
```

Viewing port status (Menu)

From the Main Menu, select **1. Status and Counters ...** , and then select **3. Module Information**.

Compatibility mode for v2 zl and zl modules

In the following context, v2 zl modules are the second version of the current zl modules. Both v2 zl and zl modules are supported in the 5400zl series chassis switches.

Compatibility Mode allows the inter-operation of v2 zl modules with zl modules in a chassis switch. When in Compatibility Mode, the switch accepts either v2 zl or zl modules. The default is Compatibility Mode enabled. If Compatibility Mode is disabled by executing the `no allow-v1-modules` command, the switch will only power up v2 zl modules.

allow-v1-modules

Syntax

`[no] allow-v1-modules`

Enables Compatibility Mode for interoperation of v2 zl and zl modules in the same chassis. (See **Figure 122: Enabling compatibility mode** on page 458.) The `no` form of the command disables Compatibility Mode. Only the v2 zl modules are powered up. (See **Figure 123: Disabling compatibility mode** on page 458.) Defaults to enabled.

allow-v1–modules

Figure 122: *Enabling compatibility mode*

```
HP Switch(config)# allow-v1-modules
This will erase the configuration and reboot the switch.
Continue [y/n]?
```

no allow-v1–modules

Figure 123: *Disabling compatibility mode*

```
HP Switch(config)# no allow-v1-modules
All V1 modules will be disabled. Continue [y/n]?
```

Port status

You can view port status using either the CLI or the menu.

show interfaces brief

Syntax

`show interfaces brief`

Description

View the port status.

Viewing port status (menu)

From the Main Menu, select **1. Status and Counters ...** , and then select **4. Port Status**.

Figure 124: *Example of port status on the menu interface*

```
================================================================================
                      Status and Counters - Port Status

                         Intrusion                                 Flow
  Port     Type          Alert      Enabled  Status    Mode        Ctrl
  -----  ---------     ---------   -------  ------  ----------   -----
  A1                   No          Yes      Down                  off
  A2                   No          Yes      Down                  off
  A3                   No          Yes      Down                  off
  A4                   No          Yes      Down                  off
  B1     10/100TX      No          Yes      Up      100FDx        off
  B2     10/100TX      No          Yes      Down    10FDx         off
  B3     10/100TX      No          Yes      Down    10FDx         off
  B4     10/100TX      No          Yes      Down    10FDx         off
  B5     10/100TX      No          Yes      Down    10FDx         off
  B6     10/100TX      No          Yes      Down    10FDx         off
  B7     10/100TX      No          Yes      Down    10FDx         off

 Actions->    Back     Intrusion log     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Accessing port and trunk group statistics

Use the CLI to view port counter summary reports, and to view detailed traffic summary for specific ports.

Trunk bandwidth utilization

- Trunk interface counters display the accumulated statistics over the trunk members' ports since the time they are added into trunk.
- Bandwidth utilization for trunks is calculated by averaging the value of the sum of bandwidth utilization for each trunk member in the last 5 minute interval .

`show interfaces`

Syntax

`show interfaces [brief | config | <PORT-LIST>]`

Description

Shows interface information for ports or trunk groups in brief or configuration detail.

Command context

`operator or manager`

Parameters

`brief`

Shows the current operating status for all ports or trunk groups on the switch in brief detail.

`config`

Shows the configuration data for all ports or trunk groups on the switch.

`<PORT-LIST>`

Specifies the list of ports for which status information will be shown.

`<TRUNK-GROUP>`

Specifies the trunk group for which status information will be shown. The status information shown consists of total transmit and receive counters and weighted average rate for the trunk group specified. The weighted average rate is calculated in 5 minute intervals.

Usage

Both external and internal ports are supported on the same module. Internal ports have an “I” suffix in the output of the show command to indicate that they are internal ports.

- “10GbE-INT” – Internal 10G data-plane ports (1i-2i, 4i-5i)
- “1GbE-INT” – Internal 1G control-plane port (3i)
- Port 3i always shows as link-down.

Examples

Show brief information and status of all interfaces.

```
switch# show interfaces brief

Status and Counters - Port Status
               | Intrusion                              MDI   Flow  Bcast
Port  Type     | Alert     Enabled Status Mode          Mode  Ctrl  Limit
----- --------- + --------- ------- ------ ----------   ----- ----- ------
B1    100/1000T | No        Yes     Down   Auto-10-100  Auto  off   0
B2    100/1000T | No        Yes     Down   1000FDx      Auto  off   0
B3    100/1000T | No        Yes     Down   1000FDx      Auto  off   0
B4    100/1000T | No        Yes     Down   1000FDx      Auto  off   0
```

```
B5    100/1000T | No         Yes      Down   1000FDx     Auto  off   0
B6    100/1000T | No         Yes      Down   1000FDx     Auto  off   0
```

Show the configuration of the interfaces currently available.

```
switch# show interfaces config

Port Settings
  Port  Type       | Enabled Mode         Flow Ctrl MDI
  ----- ---------- + ------- ------------ --------- ----
  B1    100/1000T | Yes     Auto-10-100  Disable   Auto
  B2    100/1000T | Yes     Auto         Disable   Auto
  B3    100/1000T | Yes     Auto         Disable   Auto
  B4    100/1000T | Yes     Auto         Disable   Auto
  B5    100/1000T | Yes     Auto         Disable   Auto
  B6    100/1000T | Yes     Auto         Disable   Auto
```

Show brief information and status of interfaces for ports D1i, D2i, and D3i.

```
switch# show interfaces brief d1i-d3i

Status and Counters - Port Status
                   | Intrusion                           MDI  Flow Bcast
Port   Type        | Alert     Enabled Status Mode       Mode Ctrl Limit
------ ---------- + --------- ------- ------ ---------- ---- ---- -----
D1i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
D2i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
D3i    1GbE-INT   | No        Yes     Down   1000FDx    NA   off  0
```

Show the brief information and status of interfaces for ports B1 through internal port B3i.

```
switch# show interfaces brief b1-b3i

Status and Counters - Port Status
                   | Intrusion                           MDI  Flow Bcast
 Port   Type       | Alert     Enabled Status Mode       Mode Ctrl Limit
 ------ ---------- + --------- ------- ------ ---------- ---- ---- -----
 B1     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B2     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B3     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B4     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B5     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B6     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B7     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B8     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B9     100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B10    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B11    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B12    100/1000T  | No        Yes     Down   1000FDx    Auto off  0
 B1i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
 B2i    10GbE-INT  | No        Yes     Up     10GigFD    NA   off  0
 B3i    1GbE-INT   | No        Yes     Up     1000FDx    NA   off  0
```

Show detailed interface information for port trunk 1.

```
switch# show interface trk1

Status and Counters - Port Counters for port Trk1
  Name   :  Trk1
  MAC Address      : 3464a9-b1b8bf
  Link Status      : Up
  Totals (Since boot or last clear) :
   Bytes Rx        : 777,713,956          Bytes Tx        : 596,853,141
   Unicast Rx      : 1,154,693            Unicast Tx      : 0
```

```
   Bcast/Mcast Rx   : 48,563                Bcast/Mcast Tx   : 607,474,910
  Errors (Since boot or last clear) :
   FCS Rx           : 0                     Drops Tx         : 0
   Alignment Rx     : 0                     Collisions Tx    : 0
   Runts Rx         : 0                     Late Colln       : 0
   Giants Rx        : 0                     Excessive Colln : 0
   Total Rx Errors : 0                      Deferred Tx      : 0
  Others (Since boot or last clear) :
   Discard Rx       : 0                     Out Queue Len    : 0
   Unknown Protos   : 0
  Rates (5 minute weighted average) :
   Total Rx(Kbps)   : 76,800                  Total Tx(Kbps): 76,800
   Unicast Rx (Pkts/sec) : 21               Unicast Tx (Pkts/sec) : 0
   B/Mcast Rx (Pkts/sec) : 114,878          B/Mcast Tx (Pkts/sec) : 114,900
   Utilization Rx   : 00.76 %
Utilization Tx   : 00.76 %
```

show interfaces trunk-utilization

Syntax

`show interfaces trunk-utilization`

Description

Shows the bandwidth utilization calculations for all trunk members.

Command context

operator or manager

Example

Show bandwidth utilization for trunks.

```
Switch(config)# show interfaces trunk-utilization

Status and Counters - Port Utilization

                         Rx                            Tx
 Port     | ---------------------------- | ----------------------------
          | Kbits/sec   Pkts/sec  Util   | Kbits/sec  Pkts/sec   Util
 --------  ---------- ----- + ---------- ---------- ---------   -----
 Trk1     | 0          0          0      | 0          0          0
 Trk2     | 0          0          0      | 0          0          0
 Trk10    | 0          0          0      | 0          0          0
```

Statistic interactions of interface counters

Table 22: *Statistic interactions*

Interface counters are cleared using the command `clear statistics`. When certain actions are taken to ports and trunks, the outcome of the `clear` command differs.

Action taken	Trigger	Interaction with interface counter
Adding Port into trunk	CLI/SNMP	• Interface counters for this port will be cleared across all sessions. • Average rate counters are not cleared.
Removing Port from trunk	CLI/SNMP	• Interface counters for this port will be cleared across all sessions. • Average rate counters are not cleared.
Trunk port coming Up	CLI enable	• No change in counters. • Interface counters for this port are not cleared. • Average rate counters are not cleared. Counters will start from 0 when the trunk port comes up.
Trunk port coming Up	Cable connect	• No change in counters. • Interface counters for this port are not cleared. • Average rate counters are not cleared. Counters will start from 0 when the trunk port comes up.
Trunk port going Down	CLI disable	• Interface counters for this port are not cleared. • Average rate counters are cleared.
Trunk port going Down	Cable disconnect	• Interface counters for this port are not cleared. • Average rate counters are cleared.
Trunk port going Down	Module crash/ module reload	• Interface counters for this port are not cleared. • Average rate counters are cleared.
Trunk port going Down	Save power - off	• Interface counters for this port are not cleared. • Average rate counters are cleared.
Trunk port going Down	Stacking member reboot/ crash/shutdown.	• Interface counters for this port are not cleared. • Average rate counters of this port is not cleared. • Utilization for the port is cleared. • Utilization for the trunk group is updated accordingly
Trunk port going Down	Module remove/ member remove.	• Statistics for removed trunk port can not be accessed as the port is removed. • Interface counters for the trunk group is updated. • Utilization for the trunk group is updated.
Clear statistics on physical port which is part of trunk	CLI	• Not allowed. • The error message `Module not present for port or invalid port: <PORT-NUM>` displays when the command `clear statistics` is executed on a port which part of a trunk.
Clear statistics on trunk.	CLI	• Interface counters for physical ports which are part of trunk will be cleared. • Average rate counters are not cleared.

Reset port counters

When troubleshooting network issues, you can clear all counters and statistics without rebooting the switch using the `clear statistics global` command or using the menu.

SNMP displays the counter and statistics totals accumulated since the last reboot, and it is not affected by the `clear statistics global` command or the `clear statistics <PORT-LIST>` command. Clearing statistics initiates an SNMP trap.

Once cleared, statistics cannot be reintroduced.

clear statics

Syntax

`clear statistics [<PORT-LIST> | global | <TRUNK-LIST>]`

Description

Clears all interface counters and statistics for specified ports, specified trunks or clears statistics globally for all.

Command context

`manager`

Parameters

<PORT-LIST>

Specifies the list of ports.

global

Specifies all counters and statistics for all interfaces.

***<TRUNK-LIST>*]**

Specifies the list of trunks.

Usage

- The `clear statistics` command does not clear SNMP.

Accessing port and trunk statistics (Menu)

Procedure

1. From the Main Menu, select **1. Status and Counters ...** , and then select **4. Port Counters**.

   Figure 125: *Example of port counters on the menu interface*

```
==========================- CONSOLE - MANAGER MODE -============================
                      Status and Counters - Port Counters

                                                                        Flow
  Port      Total Bytes    Total Frames     Errors Rx      Drops Tx     Ctrl
  -------  -------------  -------------  -------------  -------------  ------
 A1              195,072            323              0              0   off
 A2              651,816            871              0              0   off
 A3-Trk1         290,163            500              0              0   off
 A4-Trk1         260,134            501              0              0   off
 C1              859,363           5147              0              0   off
 C2              674,574           1693              0              0   off
 C3               26,554            246              0              0   off
 C4              113,184            276              0              0   off
 C5                    0              0              0              0   off

 Actions->   Back     Show details     Reset     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

2. To view details about the traffic on a particular port, use the ↓ key to highlight that port number, and then select **Show Details**. For example, selecting port A2 displays a screen similar to this.

   Figure 126: *Example of the display for* `Show Details` *on a selected port*

```
==========================- CONSOLE - MANAGER MODE -============================
                 Status and Counters - Port Counters - Port A2

  Link Status      : Up

  Bytes Rx         : 630,746              Bytes Tx         : 21,070
  Unicast Rx       : 568                  Unicast Tx       : 285
  Bcast/Mcast Rx   : 18                   Bcast/Mcast Tx   : 0

  FCS Rx           : 0                    Drops Tx         : 0
  Alignment Rx     : 0                    Collisions Tx    : 0
  Runts Rx         : 0                    Late Colln Tx    : 0
  Giants Rx        : 0                    Excessive Colln  : 0
  Total Rx Errors  : 0                    Deferred Tx      : 0

 Actions->   Back     Reset     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

3. This screen also includes the **Reset** action for the current session.

MAC address tables

MAC address views and searches

You can view and search MAC addresses using the CLI or the menu.

show mac-address

Syntax

```
show mac-address [vlan <VLAN-ID>] [<PORT-LIST>] [<MAC-ADDR>]
```

Description

Lists all MAC addresses on the switch and their corresponding port numbers. You can also choose to list specific addresses and ports, or addresses and ports on a VLAN. The switches operate with a multiple forwarding database architecture.

List all learned MAC addresses on the switch and corresponding port numbers

```
HP Switch# show mac-address
```

List all learned MAC addresses on one or more ports and corresponding port numbers

```
HP Switch# show mac-address a1-a4,a6
```

List all learned MAC addresses on a VLAN and corresponding port numbers

```
HP Switch# show mac-address vlan 100
```

List the port on which the switch learned a specific MAC address

To find the port on which the switch learns a MAC address of 080009-21ae84:

Using the menu to view and search MAC addresses

To determine which switch port on a selected VLAN the switch uses to communicate with a specific device on the network:

Procedure

1. From the Main Menu, select **1. Status and Counters ...** , and then select **5. VLAN Address Table**.
2. Use the arrow keys to scroll to the VLAN you want, and then press **Enter** on the keyboard to select it.

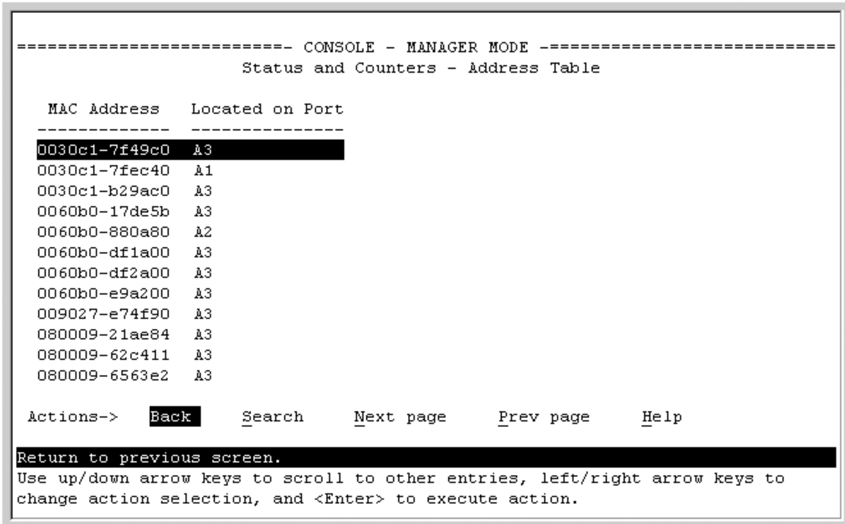

```
==========================- CONSOLE - MANAGER MODE -============================
                    Status and Counters - Address Table

  MAC Address     Located on Port
 -------------   ---------------
 0030c1-7f49c0   A3
 0030c1-7fec40   A1
 0030c1-b29ac0   A3
 0060b0-17de5b   A3
 0060b0-880a80   A2
 0060b0-df1a00   A3
 0060b0-df2a00   A3
 0060b0-e9a200   A3
 009027-e74f90   A3
 080009-21ae84   A3
 080009-62c411   A3
 080009-6563e2   A3

 Actions->   Back     Search     Next page     Prev page     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

The switch then displays the MAC address table for that VLAN (**Figure 127: Example of the address table** on page 466.)

Figure 127: *Example of the address table*

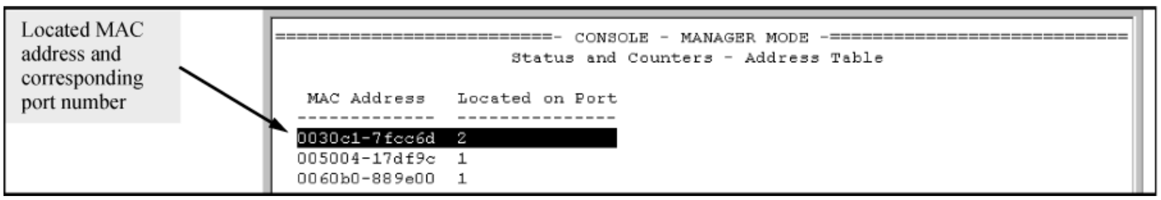


```
==========================- CONSOLE - MANAGER MODE -============================
                    Status and Counters - Address Table

  MAC Address     Located on Port
 -------------   ---------------
 0030c1-7fcc6d   2
 005004-17df9c   1
 0060b0-889e00   1
```

3. To page through the listing, use **Next page** and **Prev page** .

Finding the port connection for a specific device on a VLAN

This feature uses a device's MAC address that you enter to identify the port used by that device.

Procedure

1. Proceeding from **Figure 127: Example of the address table** on page 466, press **[S]** (for **Search**), to display the following prompt:

```
Enter MAC address: _
```

2. Enter the MAC address you want to locate and press **[Enter]**.

3. The address and port number are highlighted if found (**Figure 128: Example of menu indicating located MAC address** on page 467.) If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Figure 128: *Example of menu indicating located MAC address*

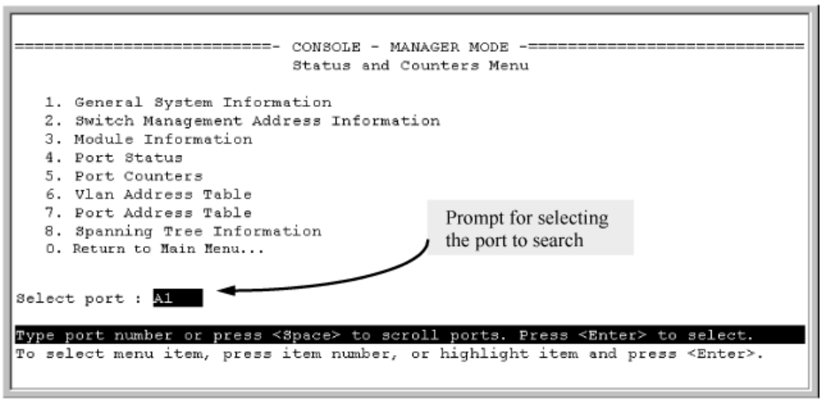
```
==========================- CONSOLE - MANAGER MODE -============================
                          Status and Counters Menu

   1. General System Information
   2. Switch Management Address Information
   3. Module Information
   4. Port Status
   5. Port Counters
   6. Vlan Address Table
   7. Port Address Table
   8. Spanning Tree Information
   0. Return to Main Menu...

Select port : A1

Type port number or press <Space> to scroll ports. Press <Enter> to select.
To select menu item, press item number, or highlight item and press <Enter>.
```


4. Press **[P]** (for **Prev page**) to return to the full address table listing.

Viewing and searching port-level MAC addresses

This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

Procedure

1. From the Main Menu, select:
2. **1. Status and Counters ...**
3. **7. Port Address Table**

Figure 129: *Listing MAC addresses for a specific port*

```
Switch-1(config)# show spanning-tree
 Multiple Spanning Tree (MST) Information

  STP Enabled    : Yes
  Force Version  : MSTP-operation
  IST Mapped VLANs : 1,66

  Switch MAC Address  : 0004ea-5e2000
  Switch Priority     : 32768
  Max Age   : 20
  Max Hops  : 20
  Forward Delay : 15

  Topology Change Count  : 0
  Time Since Last Change : 2 hours

  CST Root MAC Address : 00022d-47367f
  CST Root Priority    : 0
  CST Root Path Cost   : 4000000
  CST Root Port        : A1

  IST Regional Root MAC Address : 000883-028300
  IST Regional Root Priority    : 32768
  IST Regional Root Path Cost   : 200000
  IST Remaining Hops            : 19

                  |            Prio             | Designated    Hello
  Port  Type      | Cost       rity  State      | Bridge        Time  PtP Edge
  ----  --------- + ---------  ----- ---------- + ------------- ----- --- ----
  A1    10/100TX  | Auto       128   Forwarding | 000883-028300 9     Yes No
  A2    10/100TX  | Auto       128   Blocking   | 0001e7-948300 9     Yes No
  A3    10/100TX  | Auto       128   Forwarding | 000883-02a700 2     Yes No
  A4    10/100TX  | Auto       128   Disabled   |
  A5    10/100TX  | Auto       128   Disabled   |
   .        .          .         .        .
   .        .          .         .        .
   .        .          .         .        .
```

4. Use the Space bar to select the port you want to list or search for MAC addresses, then press **[Enter]** to list the MAC addresses detected on that port.

Determining whether a specific device is connected to the selected port

Proceeding from **4** on page 467, above:

Procedure

1. Press **[S]** (for **Search**), to display the following prompt:

```
Enter MAC address: _
```

2. Enter the MAC address you want to locate and press **[Enter]**.
3. The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.
4. Press **[P]** (for **Prev page**) to return to the previous per-port listing.

MSTP data

show spanning-tree

Syntax

`show spanning-tree`

Description

Displays the global and regional spanning-tree status for the switch, and displays the per-port spanning-tree operation at the regional level.

Values for the following parameters appear only for ports connected to active devices: `Designated Bridge`, `Hello Time`, `PtP`, and `Edge`.

show spanning-tree command output

Figure 130: *show spanning-tree command output*

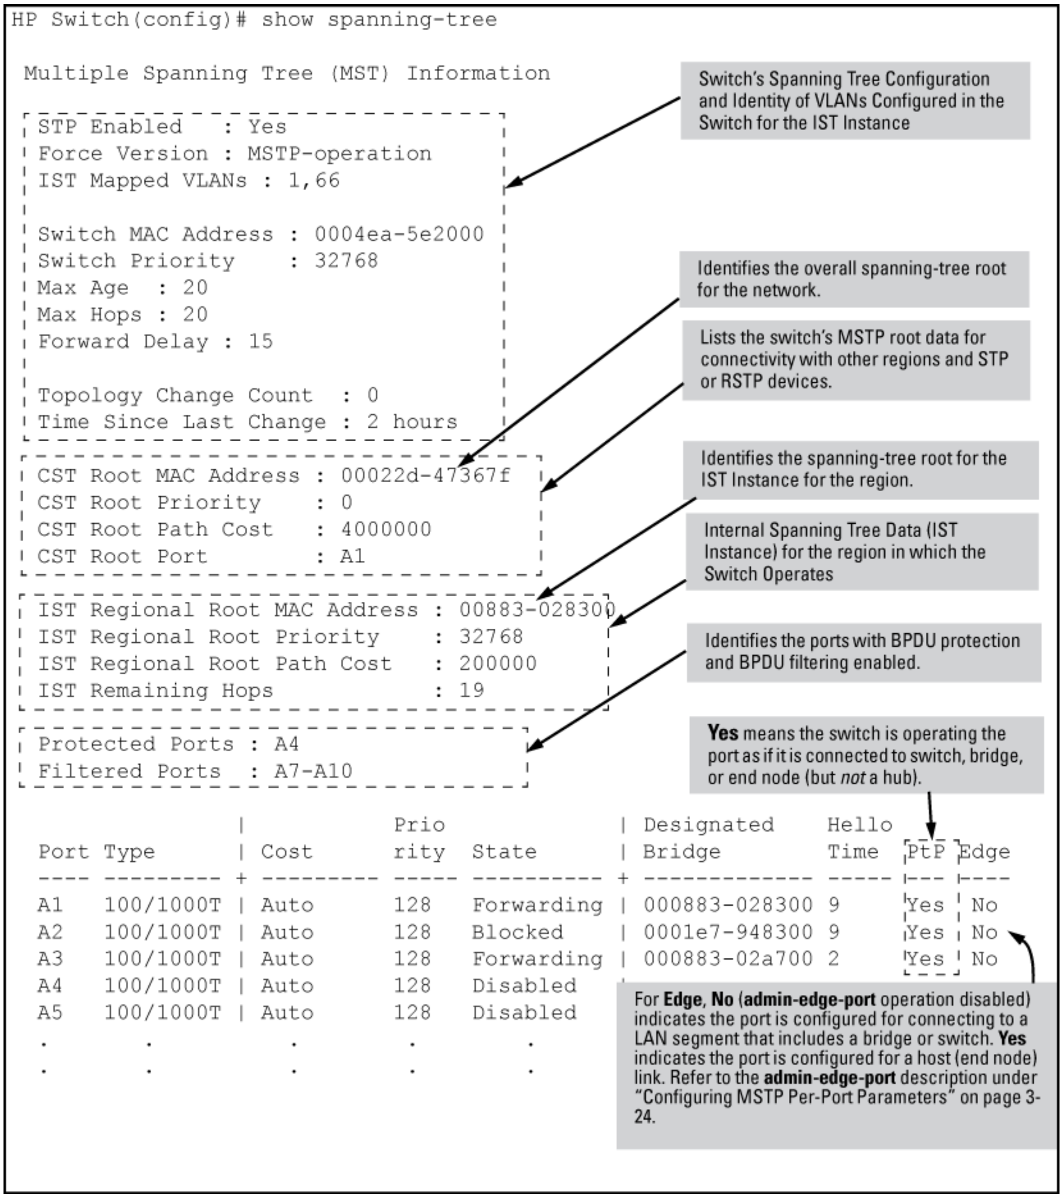

```
HP Switch(config)# show spanning-tree

 Multiple Spanning Tree (MST) Information

  STP Enabled   : Yes
  Force Version : MSTP-operation
  IST Mapped VLANs : 1,66

  Switch MAC Address : 0004ea-5e2000
  Switch Priority    : 32768
  Max Age  : 20
  Max Hops : 20
  Forward Delay : 15

  Topology Change Count  : 0
  Time Since Last Change : 2 hours

  CST Root MAC Address : 00022d-47367f
  CST Root Priority    : 0
  CST Root Path Cost   : 4000000
  CST Root Port        : A1

  IST Regional Root MAC Address : 00883-028300
  IST Regional Root Priority    : 32768
  IST Regional Root Path Cost   : 200000
  IST Remaining Hops            : 19

  Protected Ports : A4
  Filtered Ports  : A7-A10

                 |           Prio               | Designated    Hello
  Port Type      | Cost      rity  State        | Bridge        Time  PtP Edge
  ---- --------- + --------- ----- ---------- + ------------- ----- --- ----
  A1   100/1000T | Auto      128   Forwarding | 000883-028300 9     Yes  No
  A2   100/1000T | Auto      128   Blocked    | 0001e7-948300 9     Yes  No
  A3   100/1000T | Auto      128   Forwarding | 000883-02a700 2     Yes  No
  A4   100/1000T | Auto      128   Disabled   |
  A5   100/1000T | Auto      128   Disabled
  .       .          .        .        .
  .       .          .        .        .
```


IP IGMP status

show ip igmp

Syntax

`show ip igmp <VLAN-ID> [config] [group <IP-ADDR>|groups] [statistics]`

Description

Global command that lists IGMP status for all VLANs configured in the switch, including:

- VLAN ID (VID) and name
- Querier address
- Active group addresses per VLAN
- Number of report and query packets per group
- Querier access port per VLAN

Parameters and options

`config`

Displays the IGMP configuration information, including VLAN ID, VLAN name, status, forwarding, and Querier information.

`vlan-id`

Per-VLAN command listing above, IGMP status for specified VLAN (VID).

`group <`*`IP-ADDR`*`>`

Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.

`groups`

Displays VLAN-ID, group address, uptime, expiration time, multicast filter type, and the last reporter for IGMP groups.

`statistics`

Displays IGMP operational information, such as VLAN IDs and names, and filtered and flooding statistics.

Output from `show ip igmp config` command

```
HP Switch(config)# show ip igmp config

 IGMP Service

                         IGMP    Forward with  Querier Querier
  VLAN ID VLAN Name      Enabled High Priority Allowed Interval
  ------- ------------   ------- ------------- ------- --------
  1       DEFAULT_VLAN No        No            Yes     125
  2       VLAN2          Yes     No            Yes     125
  12      New_Vlan       No      No            Yes     125
```

IGMP statistical information

```
switch(vlan-2)# show ip igmp statistics

 IGMP Service Statistics

  Total VLANs with IGMP enabled               : 1
  Current count of multicast groups joined    : 1


 IGMP Joined Groups Statistics

  VLAN ID VLAN Name                        Filtered     Flood
  ------- -------------------------------- ------------ ------------
  2       VLAN2                            2            1
```

VLAN information

show vlan

Syntax

```
show vlan <VLAN-ID>
```

Description

Lists the maximum number of VLANs to support, existing VLANS, VLAN status (static or dynamic), and primary VLAN.

Parameters and options

<VLAN-ID>

Lists the following for the specified VLAN:

- Name, VID, and status (static/dynamic)
- Per-port mode (tagged, untagged, forbid, no/auto)
- "Unknown VLAN" setting (Learn, Block, Disable)
- Port status (up/down)

List data on specific VLANs

The next three figures show how you can list data for the following VLANs:

Ports	VLAN	VID
A1-A12	DEFAULT_VLAN	1

Table Continued

A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

Figure 131: *Listing the VLAN ID (vid) and status for specific ports*

```
HP Switch# show vlan ports A1-A2

 Status and Counters = VLAN Information - for ports A1,A2

  802.1Q VLAN ID Name          Status
  -------------- ------------- -------
  1              DEFAULT_VLAN  Static
  33             VLAN-33       Static
```

Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

Figure 132: *Example of VLAN listing for the entire switch*

```
HP Switch# show vlan
 Status and Counters - VLAN Information

  VLAN support : Yes
  Maximum VLANs to support : 9
  Primary VLAN: DEFAULT_VLAN

  802.1Q VLAN ID Name          Status
  -------------- ------------- -------
  1              DEFAULT_VLAN  Static
  33             VLAN-33       Static
  44             VLAN-44       Static
```

Figure 133: *Port listing for an individual VLAN*

```
HP Switch(config)# show vlan 1

 Status and Counters - VLAN Information - VLAN 1

  VLAN ID : 1
  Name : DEFAULT_VLAN
  Status : Static
  Voice : Yes
  Jumbo : No

  Port Information Mode     Unknown VLAN Status
  ---------------- -------- ------------ ----------
  A1               Untagged Learn        Up
  A2               Tagged   Learn        Up
  A3               Untagged Learn        Up
  A4               Untagged Learn        Down
  A5               Untagged Learn        Up
  A6               Untagged Learn        Up
  A7               Untagged Learn        Up
```

WebAgent status information

The WebAgent Status screen provides an overview of the status of the switch. Scroll down to view more details. For information about this screen, click on **?** in the upper right corner of the WebAgent screen.

Figure 134: *Example of a WebAgent status screen*

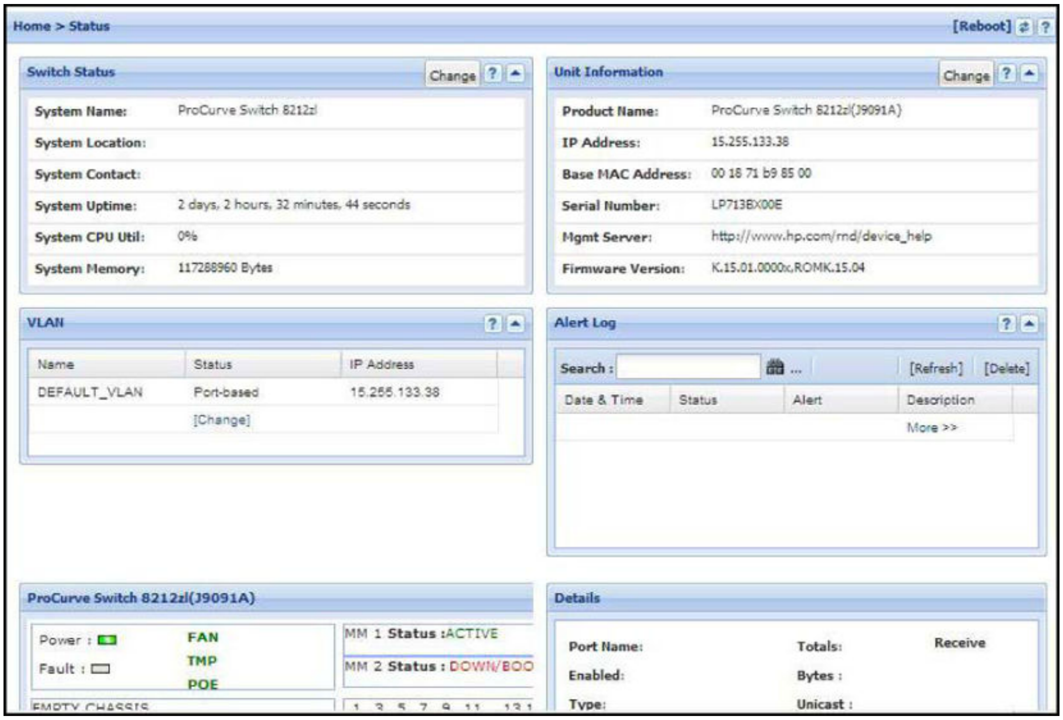

Configuring local mirroring

To configure a local mirroring session in which the mirroring source and destination are on the same switch, follow these general steps:

Procedure

1. Determine the session and local destination port:
   a. Session number (1-4) and (optional) alphanumeric name
   b. Exit port (any port on the switch except a monitored interface used to mirror traffic)

   Hewlett Packard Enterprise strongly discourages connecting a mirroring exit port to a network because doing so can result in serious network performance problems. Only connect an exit port to a network analyzer, IDS, or other network edge device that has no connection to other network resources.

2. Enter the `mirror` *`session-#`* `[name` *`session-name`* `] port` *`port-#`* command to configure the session.
3. Determine the traffic to be selected for mirroring by any of the following methods and the appropriate configuration level (VLAN, port, mesh, trunk, switch):
   a. Direction: inbound, outbound, or both
   b. Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic
   c. MAC source and/or destination address: inbound, outbound, or both
4. Enter the `monitor` command to assign one or more source interfaces to the session.

After you complete step 4, the switch begins mirroring traffic to the configured exit port.

The following commands configure mirroring for a local session in which the mirroring source and destination are on the same switch.

- The `mirror` command identifies the destination in a mirroring session.
- The `interface` and `vlan` commands identify the mirroring source, including source interface, traffic direction, and traffic-selection criteria for a specified session.

NOTE

With no

allow-v2-modules

specified in the configuration of a switch with V3 modules on KB firmware, Egress VLAN ACLs do not filter mirrored traffic. You must use a port ACL to filter mirrored traffic.

Local mirroring sessions

Syntax

```
[no] mirror  1 - 4  port  <EXIT-PORT-#> [name <NAME-STR>]
```

Description

Configure local mirroring sessions.

Parameters and options

`no`

When used with `no mirror session-# port` command, removes the mirroring session and any mirroring source previously assigned to that session by the following commands.

Traffic-direction criteria

interface monitor all

Syntax

Description

```
[no] [interface <PORT> |<TRUNK> |<MESH>]|vlan <VID-#>] monitor all in|out|both
mirror <SESSION> [session ...] [no-tag-added]
```

ACL criteria for inbound traffic — deprecated

NOTE

interface monitor ip

Syntax

```
[no] [interface <PORT> |<TRUNK> |<MESH>]|vlan <VID-#>] monitor ip access-group
<ACL-NAME> in mirror session [session ...]
```

Mirror policy for inbound traffic

class [ipv4|ipv6]

Syntax

`class [ipv4|ipv6] <CLASSNAME> [no] [seq-number] [match|ignore] <IP-PROTOCOL> <SOURCE-ADDRESS> <DESTINATION-ADDRESS>][precedence <PRECEDENCE-VALUE>][tos <TOS-VALUE>][ip-dscp <CODEPOINTS>][vlan <VLAN-ID>]`

Description

Configures the mirroring policy for inbound traffic on the switch.

Parameters and options

policy mirror

Syntax

`policy mirror <POLICY-NAME> [no] <SEQ-NUMBER> [class [ipv4|ipv6] <CLASSNAME> action mirror <SESSION>] [action mirror <SESSION>] [no] default-class action mirror <SESSION> [no] [interface <PORT/TRUNK>| vlan <VID-#>] service-policy <MIRROR-POLICY-NAME> in`

Description

The `[no] [interface <PORT/TRUNK>| vlan <VID-#>] service-policy <MIRROR-POLICY-NAME> in` command removes the mirroring policy from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Parameters and options

`mirror <SESSION>`

Accepts either a number (1 to 4) or a name. To use a name, you must first configure the `name <NAME-STR>` parameter option for the specified mirroring session using the `policy mirror` command.

MAC-based criteria to select traffic

monitor mac

Syntax

`[no] monitor mac  <MAC-ADDR> [src|dst|both] mirror session`

Description

Configures traffic using MAC-based criteria.

Parameters and options

`no`

Use the `no` form of the complete Command syntax (for example, `no monitor mac 112233-445566 src mirror 3`) to remove a MAC address as mirroring criteria from an active session on the switch without removing the session itself.

`mirror`

Enter the `monitor mac mirror` command at the global configuration level.

Remote mirroring destination on a remote switch

Syntax

`mirror endpoint ip <SRC-IP> <SRC-UDP-PORT > <DST-IP> <EXIT-PORT> [truncation]`

Description

Configures a remote mirroring destination on a remote switch.

Parameters and options

Remote mirroring destination on a local switch

mirror remote ip

Syntax

`mirror <SESSION> remote ip <SRC-IP> <SRC-UDP-PORT> <DST-IP>`

Description

Configures a remote mirroring destination on a local switch.

Parameters and options

Local mirroring destination on the local switch

mirror port

Syntax

`mirror <SESSION> port <EXIT-PORT>`

Description

Configures a local mirroring destination on a local switch.

Parameters and options

Monitored traffic

In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session `interface | vlan monitor ip access-group in mirror` command has been deprecated and is replaced with classifier-based mirroring policies.

interface

Syntax

`interface <PORT/TRUNK/MESH>`

Description

Parameters and options

monitor all

Syntax

`monitor all [in|out|both] mirror <SESSION> [no-tag-added]`

`monitor ip access-group ACL-NAME in mirror <SESSION>`

`monitor mac <MAC-ADDR> [src|dest|both] mirror`

`show monitor [endpoint|<SESSION-NUMBER>|name <SESSION-NAME>`

service-policy

Syntax

`service-policy <mirror-policy-name> in`

Configuring local mirroring (Menu)

If mirroring has already been enabled on the switch, the Menu screens appear different from the one shown in this section.

Procedure

1. From the Main Menu, select **1. Switch Configuration ...** , and then select **3. Network Monitoring Port**.

   Figure 135: *The default network mirroring configuration screen*

   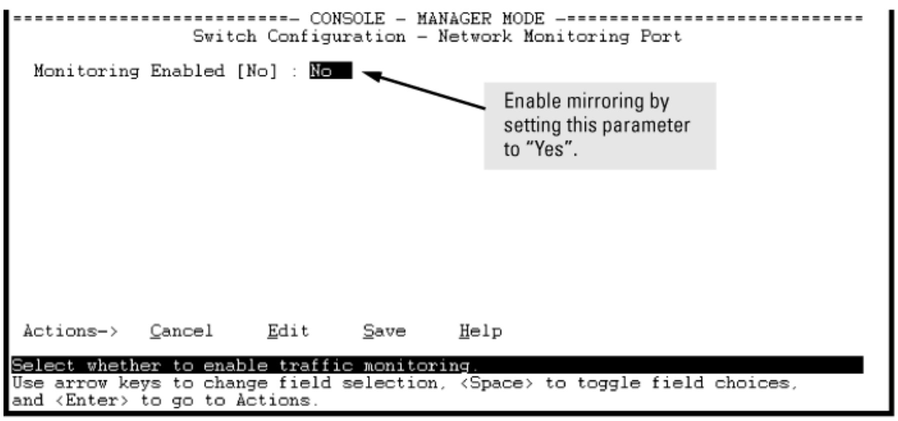

2. In the Actions menu, press **[E]** (for **Edit**.)
3. If mirroring is currently disabled for session 1 (the default), enable it by pressing the Space bar (or **[Y]**) to select **Yes**.

4. Press the down arrow key to display a screen similar to **Figure 136: How to select a local exit port** on page 478, and move the cursor to the **Monitoring Port** parameter.

Figure 136: *How to select a local exit port*

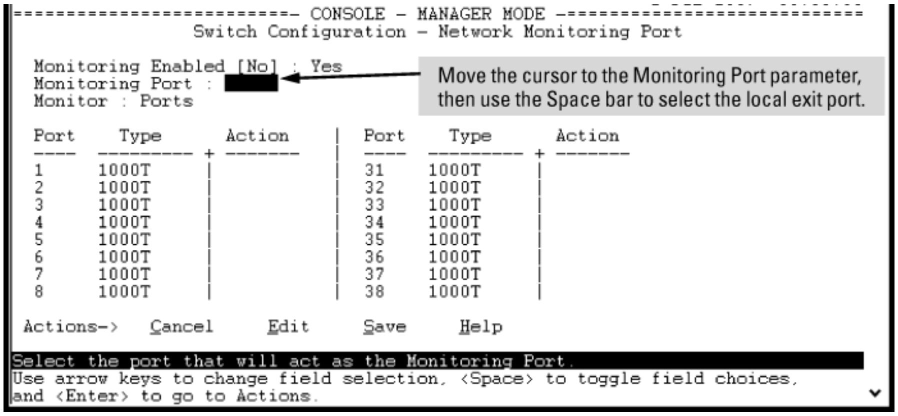

5. Use the Space bar to select the port to use for sending mirrored traffic to a locally connected traffic analyzer or IDS.
6. (The selected interface must be a single port. It cannot be a trunk or mesh.) In this example, port 5 is selected as the local exit port.
7. Highlight the Monitor field and use the Space bar to select the interfaces to mirror:
8. **Ports:** Use for mirroring ports, static trunks, or the mesh.
9. **VLAN:** Use for mirroring a VLAN.
10. Do one of the following:
    a. If you are mirroring ports, static trunks, or the mesh, go to
       11 on page 478

       .
    b. If you are mirroring a VLAN:
       I. Press **[Tab]** or the down arrow key to move to the **VLAN** field.

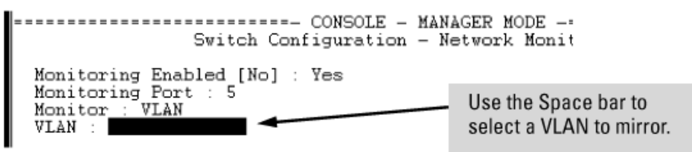

       II. Use the Space bar to select the **VLAN** you want to mirror.
       III. Go to **14** on page 479.
11. Use the down arrow key to move the cursor to the **Action** column for the individual port interfaces and position the cursor at a port, trunk, or mesh you want to mirror.

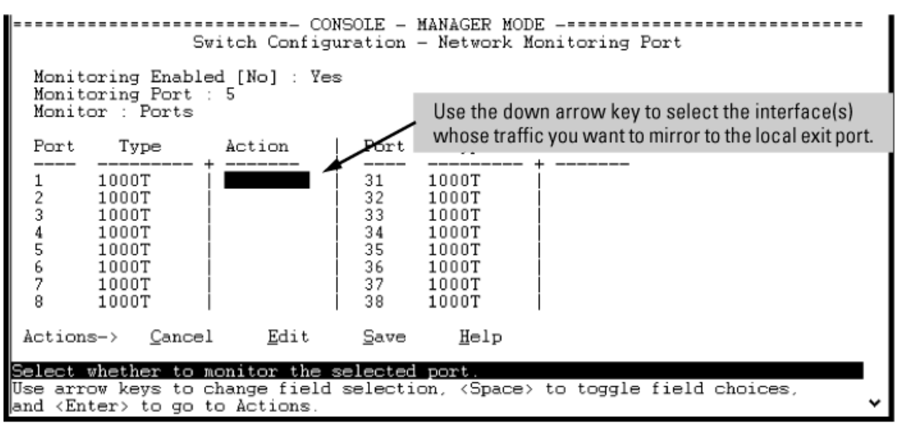

12. Press the Space bar to select **Monitor** for the ports, trunks, mesh, or any combination of these that you want mirrored.
13. Use the down arrow key to move from one interface to the next in the **Action** column. (If the mesh or any trunks are configured, they appear at the end of the port listing.)
14. When you finish selecting interfaces to mirror, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
15. Return to the Main Menu.

You can also use the CLI to configure a mirroring session for a destination device connected to an exit port on either:

- The same switch as the source interface (local mirroring.)
- A different switch (remote mirroring.) The remote switch must be an switch offering the full mirroring capabilities described in this chapter.

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port.

In a remote mirroring session that uses IPv4 encapsulation, if the exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic.

For this reason, Switch strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Destination mirror on a remote switch

mirror endpoint

Syntax

```
mirror endpoint ip <SRC-IP-ADDR> <SRC-UDP-PORT> <DST-IP-ADDR> port <EXIT-PORT>
```

Description

Enter this command on a remote switch to configure the exit port to use in a remote mirroring session. Use **Source mirror on the local switch** to configure the mirroring source on the local switch.

The `mirror endpoint ip` command configures:

- The unique UDP port number to be used for the mirroring session on the source switch. The recommended port range is from 7933 to 65535.
- The IP address of the source switch to use in the session.
- The IP address and exit-port number on the remote (endpoint) switch.

In a remote mirroring endpoint, the IP address of exit port and the remote destination switch can belong to different VLANs.

Source mirror on the local switch

mirror remote ip

Syntax

```
[no] mirror 1 - 4 [name <NAME-STR>] remote ip <SRC-IP> <SRC-UDP-PORT> <DST-IP> [truncation]
```

Description

Configures the mirroring source on the local switch.

Parameters and options

no mirror 1–4

Removes both the mirroring session and any mirroring sources previously assigned to the session by the following commands.

Traffic-direction criteria

Syntax

```
[no] [interface <PORT> <TRUNK> <MESH>|vlan <VID-#>] monitor all in|out|both mirror
1-4|<NAME-STR> [1 — 4|<NAME-STR . . .>]
```

Description

Configures traffic direction criteria for specific traffic

Configure ACL criteria to select inbound

NOTE

interface monitor ip access-group

Syntax

```
[no] [interface <PORT> <TRUNK> <MESH>|vlan <VID-#>] monitor ip access—group <ACL—
NAME> inmirror [1-4|<NAME-STR>] [1 — 4|<NAME-STR . . .>]
```

Mirror policy for inbound traffic

class [ipv4|ipv6]

Syntax

```
class [ipv4|ipv6] <CLASSNAME> [no] [seq-number] [match|ignore] <IP-PROTOCOL>
<SOURCE-ADDRESS> <DESTINATION-ADDRESS>][precedence <PRECEDENCE-VALUE>][tos <TOS-
VALUE>][ip-dscp <CODEPOINTS>][vlan <VLAN-ID>]
```

Description

Configures the mirroring policy for inbound traffic on the switch.

Parameters and options

policy mirror

Syntax

```
policy mirror <POLICY-NAME> [no] <SEQ-NUMBER> [class [ipv4|ipv6] <CLASSNAME> action
mirror <SESSION>] [action mirror <SESSION>] [no] default-class action mirror
<SESSION> [no] [interface <PORT/TRUNK>| vlan <VID-#>] service-policy <MIRROR-
POLICY-NAME> in
```

Description

The `[no] [interface <PORT/TRUNK>| vlan <VID-#>] service-policy <MIRROR-POLICY-NAME> in` command removes the mirroring policy from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Parameters and options

`mirror <SESSION>`

Accepts either a number (1 to 4) or a name. To use a name, you must first configure the `name <NAME-STR>` parameter option for the specified mirroring session using the `policy mirror` command.

Configuring a destination switch in a remote mirroring session

CAUTION

When configuring a remote mirroring session, **always** configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.

Syntax

`mirror endpoint ip src-ip src-udp-port dst-ip exit-port-# no mirror endpoint ip src-ip src-udp-port dst-ip`

Used on a destination switch to configure the remote endpoint of a mirroring session. The command uniquely associates the mirrored traffic from the desired session on a monitored source with a remote exit port on the destination switch. You must use the same set of source and destination parameters used when you configure the same session on both the source and destination switches.

For a given mirroring session, the same `src-ip` , `src-udp-port` and `dst-ip` values must be entered with the `mirror endpoint ip` command on the destination switch, and later with the `mirror remote ip` command on the source switch.

CAUTION

Do not remove the configuration of a remote mirroring endpoint support for a given session if there are source switches currently configured to mirror traffic to the endpoint.

src-ip	Must exactly match the *src-ip* address you configure on the source switch for the remote session.
src-udp-port	Must exactly match the `src-udp-port` value you configure on the source switch for the remote session. The recommended port range is 7933 to 65535.This setting associates the monitored source with the desired remote endpoint in the remote session by using the same, unique UDP port number to identify the session on the source and remote switches.
dst-ip	Must exactly match the `dst-ip` setting you configure on the source switch for the remote session.
exit-port-#	Exit port for mirrored traffic in the remote session, to which a traffic analyzer or IDS is connected.

The `no` form of the command deletes the mirroring endpoint for the configured session on the remote destination switch.

Configuring a source switch in a local mirroring session

Enter the `mirror port` command on the source switch to configure an exit port on the same switch. To create the mirroring session, use the information gathered in **High-level overview of the mirror configuration process** on page 507.

Syntax

`mirror 1 - 4 port` *`exit-port-#`* `[name` *`name-str`*`] no mirror 1- 4`

Assigns the exit port to use for the specified mirroring session and must be executed from the global configuration level.

`1 - 4`	Identifies the mirroring session created by this command. (Multiple sessions on the switch can use the same exit port.)
`name` *`name-str`*	Optional alphanumeric name string used to identify the session (up to 15 characters)
`port` *`exit-port-#`*	Exit port for mirrored traffic in the remote session. This is the port to which a traffic analyzer or IDS is connected.

The `no` form of the command removes the mirroring session and any mirroring source previously assigned to that session.

Configuring a source switch in a remote mirroring session

Syntax

`[no] mirror 1 - 4 [name` *`name-str`*`] remote ip` *`src-ip src-udp-port dst-ip`* `[truncation]`

Used on the source switch to uniquely associate the mirrored traffic in the specified session with a remote destination switch. You must configure the same source and destination parameters when you configure the same session on both the source and destination switches. (If multiple remote sessions use the same source and destination IP addresses, each session must use a unique UDP port value.)

When you execute this command, the following message is displayed:

```
Caution: Please configure destination switch first.
          Do you want to continue [y/n]?
```

- If you have not yet configured the session on the remote destination switch, follow the configuration procedure in **Configure a mirroring destination on a remote switch** on page 507 before using this command.
- If you have already configured the session on the remote destination switch, enter **y** (for "yes") to complete this command.

`1 - 4`	Identifies the mirroring session created by this command.
`name` *`name-str`*	Optional alphanumeric name string used as an additional session identifier (up to 15 characters.)

Table Continued

src-ip	The IP address of the VLAN or subnet on which the traffic to be mirrored enters or leaves the switch.
src-udp-port	Associates the remote session with a UDP port number. When multiple sessions have the same source IP address *src-ip* and destination IP address *dst-ip* , the UDP port number must be unique in each session. The UDP port number used for a given session should be in the range of 7933 to 65535. **CAUTION** UDP port numbers below 7933 are reserved for various IP applications. Using them for mirroring can result in the interruption of other IP functions and in non-mirrored traffic being received on the destination switch and sent to a device connected to the remote exit port. The configured UDP port number is included in the frames mirrored from the source switch to the remote destination switch (`mirror endpoint`), and enables the remote switch to match the frames to the exit port configured for the combined UDP port number, source IP address, and destination IP address..
dst-ip	For the remote session specified in the command, this is the IP address of the VLAN or subnet on which the remote exit port exists. (The exit port to which a traffic analyzer or IDS is connected is configured on the remote switch in section.) .)
`[truncation]`	Enables truncation of oversize frames, causing the part of the frame in excess of the MTU size to be truncated. Unless truncation is enabled, oversize frames are dropped. The frame size is truncated to a multiple of 18 bytes—for example, if the MTU is 1000 bytes, the frame is truncated to 990 bytes (55 * 18 bytes.)

The `no` form of the command removes the mirroring session and any mirroring source previously assigned to the session. To preserve the session while deleting a monitored source assigned to it.

Selecting all traffic on a port interface for mirroring according to traffic direction

Syntax

```
[no] interface port/trunk/mesh monitor all [in | out | both] [mirror 1 - 4 | name-str] [{1 - 4 | name-str} | {1 - 4 | name-str} | {1 - 4 | name-str}] [no-tag-added]
```

Assigns a mirroring source to a previously configured mirroring session on a source switch by specifying the port, trunk, and/or mesh sources to use, the direction of traffic to mirror, and the session.

`interface` *`port/trunk/mesh`*	Identifies the source ports, static trunks, and/or mesh on which to mirror traffic.Use a hyphen for a range of consecutive ports or trunks (`a5-a8`, `Trk2-Trk4`.)Use a comma to separate non-contiguous interfaces (`b11`, `b14`, `Trk4`, `Trk7`.)
`monitor all [in | out | both]`	For the interface specified by *port/trunk/mesh* , selects traffic to mirror based on whether the traffic is entering or leaving the switch on the interface: • `in:` Mirrors entering traffic. • `out:` Mirrors exiting traffic. • `both:` Mirrors traffic entering and exiting. If you enter the `monitor all` command without selection criteria or a session identifier, the command applies by default to session 1
`mirror [1 - 4 |` *`name-str`*`]`	Assigns the traffic specified by the interface and direction to a session by number or—if configured—by name. The session must have been previously configured.Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified source to up to four sessions, for example, `interface a1 monitor all in mirror 1 2 4.` • `1 - 4` : Configures the selected port traffic to be mirrored in the specified session number. • `[ name` *`name-str`* `]` Optional: configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session.
`[no-tag-added]`	Prevents a VLAN tag from being added to the mirrored copy of an outbound packet sent to a local or remote mirroring destination.

The `no` form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring source and adding another in its place.

Selecting all traffic on a VLAN interface for mirroring according to traffic direction

Syntax

```
vlan vid-# monitor all [in | out | both] [mirror 1 - 4 | name-str] [{1 - 4 | name-str} | {1 - 4 | name-str} | {1 - 4 | name-str}]
```

This command assigns a monitored VLAN source to a previously configured mirroring session on a source switch by specifying the VLAN ID, the direction of traffic to mirror, and the session.

`vlan` *`vid-#`*	Identifies the VLAN on which to mirror traffic.
`monitor all [in | out | both]`	Uses the direction of traffic on the specified *`vid-#`* to select traffic to mirror.If you enter the `monitor all` command without selection criteria or a session identifier, the command applies by default to session 1.
`mirror [1 - 4 |` *`name-str`*`]`	Assigns the VLAN traffic defined by the VLAN ID and traffic direction to a session number or name.Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified VLAN source to up to four sessions, for example, `interface a1 monitor all in mirror 1 2 4`. • `1 - 4` : Configures the selected VLAN traffic to be mirrored in the specified session number. • `[name` *`name-str`* `]:` Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session. To configure an alphanumeric name for a mirroring session, see the command description under **Configuring a source switch in a remote mirroring session** on page 482.

Assigning a VLAN to a mirroring session precludes assigning any other mirroring sources to the same session. If a VLAN is already assigned to a given mirroring session, using this command to assign another VLAN to the same mirroring session results in the second assignment replacing the first. Also, if there are other (port, trunk, or mesh) mirroring sources already assigned to a session, the switch displays a message similar to:

```
Mirror source port exists on session N. Can not add mirror
source VLAN.
```

The `no` form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This allows you to repurpose a session by removing an unwanted mirroring source and adding another in its place.

Configuring a MAC address to filter mirrored traffic on an interface

Enter the `monitor mac mirror` command at the global configuration level.

Syntax

[no] monitor mac *mac-addr* [src | dest | both] {mirror 1 - 4 | *name-str*} [1 - 4 | *name-str*] [1 - 4 | *name-str*] [1 - 4 | *name-str*]

Use this command to configure a source and/or destination MAC address as criteria for selecting traffic in one or more mirroring sessions on the switch. The MAC address you enter is configured to mirror inbound (`src`), outbound (`dest`), or both inbound and outbound (`both`) traffic on any port or learned VLAN on the switch.

`monitor mac` *`mac-addr`*	
Configures the MAC address as selection criteria for mirroring traffic on any port or learned VLAN on the switch.	
`{src | dest | both}`	Specifies how the MAC address is used to filter and mirror packets in inbound and/or outbound traffic on the interfaces on which the mirroring session is applied: • `src:` Mirrors all packets in inbound traffic that contain the specified MAC address as source address. • `dest:` Mirrors all packets in outbound traffic that contain the specified MAC address as destination address. NOTE: The MAC address of the switch is not supported as either the source or destination MAC address used to select mirrored traffic. • `both:` Mirrors all packets in both inbound and outbound traffic that contain the specified MAC address as either source or destination address.
`mirror [1 - 4 |` *`name-str`*`]`	Assigns the inbound and/or outbound traffic filtered by the specified MAC address to a previously configured mirroring session. The session is identified by a number or (if configured) a name.Depending on how many sessions are configured on the switch, you can use the same command to configure a MAC address as mirroring criteria in up to four sessions. To identify a session, you can enter either its name or number; for example: `mirror 1 2 3 traffsrc4` `1 - 4` **:** Specifies a mirroring session by number, for which the configured MAC address is used to select and mirror inbound and/or outbound traffic.

Packets that are sent or received on an interface configured with a mirroring session and that contain the MAC address as source and/or destination address are mirrored to a previously configured destination device.

To remove a MAC address as selection criteria in a mirroring session, you must enter the complete Command syntax, for example, `no monitor mac 998877-665544 dest mirror 4`.

The `no` form of the command removes the MAC address as a mirroring criteria from an active session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring criteria and adding another in its place.

Configuring classifier-based mirroring

For more information and a list of general steps for the process beginning with this command, see the information about restrictions on classifier-based mirroring.

Context: Global configuration

Syntax

`[no] class [ipv4 | ipv6` *`classname`*`]`

Defines the name of a traffic class and specifies whether a policy is to be applied to IPv4 or IPv6 packets, where *classname* is a text string (64 characters maximum.)

After you enter the `class` command, you enter the class configuration context to specify match criteria. A traffic class contains a series of `match` and `ignore` commands, which specify the criteria used to classify packets.

To configure a default traffic class, use the `default-class` command as described below. A default class manages the packets that do not match the match/ignore criteria in any other classes in a policy.

Context: Class configuration

Syntax

```
[no] seq-number
 [match | ignore ip-protocol source-address destination-address] [ip-dscp
codepoint] [precedence precedence-value] [tos tos-value] [vlan vlan-id]
```

For detailed information about how to enter `match` and `ignore` commands to configure a traffic class, the *Advanced Traffic Management Guide*.

Context: Global configuration

Syntax

```
[no] policy mirror policy-name
```

Defines the name of a mirroring policy and enters the policy configuration context.

A traffic policy consists of one or more classes and one or more mirroring actions configured for each class of traffic. The configured actions are executed on packets that match a `match` statement in a class. No policy action is performed on packets that match an `ignore` statement.

Context: Policy configuration

Syntax

```
[no] seq-number
 class [ipv4 | ipv6 classname] action mirror session
```

Defines the mirroring action to be applied on a pre-configured IPv4 or IPv6 traffic class when a packet matches the `match` criteria in the traffic class. You can enter multiple `class action mirror` statements in a policy.

`[`*`seq-number`*`]`	The (optional) `seq-number` parameter sequentially orders the mirroring actions that you enter in a policy configuration. Actions are executed on matching packets in numerical order.Default: Mirroring action statements are numbered in increments of 10, starting at 10.
`class [ipv4	ipv6` *`classname`*`]`	Defines the preconfigured traffic class on which the mirroring actions in the policy are executed and specifies whether the mirroring policy is applied to IPv4 or IPv6 traffic in the class. The classname is a text string (64 characters maximum.)
`action mirror` *`session`*	Configures mirroring for the destination and session specified by the `session` parameter.

Context: Policy configuration

Syntax

```
[no] default-class action mirror session [action mirror session ...]
```

Configures a default class that allows packets that are not matched nor ignored by any of the class configurations in a mirroring policy to be mirrored to the destination configured for the specified session.

Applying a mirroring policy on a port or VLAN interface

Enter one of the following `service-policy` commands from the global configuration context.

Context: Global configuration

Syntax

```
interface <PORT-LIST> service-policy policy-name in
```

Configures the specified ports with a mirroring policy that is applied to inbound traffic on each interface.

Separate individual port numbers in a series with a comma, for example, `a1,b4,d3`. Enter a range of ports by using a dash, for example, `a1-a5`.

The mirroring policy name you enter must be the same as the policy name you configured with the `policy mirror` command.

Syntax

```
vlan vlan-id service-policy policy-name in
```

Configures a mirroring policy on the specified VLAN that is applied to inbound traffic on the VLAN interface.

Valid VLAN ID numbers range from 1 to 4094.

The mirroring policy name you enter must be the same as the policy name you configured with the `policy mirror` command in the syntax **policy mirror**.

Viewing a classifier-based mirroring configuration

To display information about a classifier-based mirroring configuration or statistics on one or more mirroring policies, enter one of the following commands:

Syntax

```
show class [ipv4 class-name | ipv6 class-name | config]
```

Syntax

```
show policy [policy-name | config]
```

Syntax

```
show policy resources
```

Syntax

```
show statistics policy [policy-name] [interface port-num | vlan vid in]
```

Viewing all mirroring sessions configured on the switch

Syntax

```
show monitor
```

If a monitored source for a remote session is configured on the switch, the following information is displayed. Otherwise, the output displays: Mirroring is currently disabled.

Sessions	Lists the four configurable sessions on the switch.
Status	Displays the current status of each session: • active: The session is configured. • inactive: Only the destination has been configured; the mirroring source is not configured. • not defined: Mirroring is not configured for this session.
Type	Indicates whether the mirroring session is local (`port`), remote (`IPv4`), or MAC-based (`mac`) for local or remote sessions.
Sources	Indicates how many monitored source interfaces are configured for each mirroring session.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.

If a remote mirroring endpoint is configured on the switch, the following information is displayed. Otherwise, the output displays: There are no Remote Mirroring endpoints currently assigned.

Type	Indicates whether the mirroring session is local (`port`), remote (`IPv4`), or MAC-based (`mac`) for local or remote sessions.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.

Table Continued

UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Dest Port	Identifies the exit port for a remote session on a remote destination switch.

Figure 137: *Displaying the currently configured mirroring sessions on the switch*

```
HP Switch# show monitor

Network Monitoring

   Sessions  Status       Type    Sources  Policy
   --------  -----------  -----   -------  -----
   1         active        port   1        yes
   2         active        mac    2        no
   3         not defined
   4         inactive      IPv4   0        no

Remote Mirroring - Remote Endpoints

 Type  UDP Source Addr  UDP port  UDP Dest Addr
 ----  ---------------  --------  ---------------
 IPv4  10.10.30.1       7950      10.10.20.1
```

Local and Remote Mirroring Sources:
- **Session 1** is performing local mirroring using a classifier-based policy as traffic-selection criteria.
- **Session 2** is performing remote mirroring using MAC-based traffic-selection criteria.
- **Session 3** is not configured.
- **Session 4** is configured for remote mirroring from a non-policy source (for example, traffic direction), but is currently not mirroring any traffic.

Remote Mirroring Destination:

The switch is configured as a remote mirroring destination (endpoint) for a source at 10.10.30.1, using port B10 as the exit port.

Viewing the remote endpoints configured on the switch

Syntax

`show monitor endpoint`

Displays the remote mirroring endpoints configured on the switch. Information on local sessions configured on the switch is not displayed. (To view the configuration of a local session, use the

`show monitor [1-4 | name` *`name-str`* `]]`

command, as described on page 74 and page 77.)

Type	Indicates whether the session is a `port` (local) or `IPv4` (remote) mirroring session.
show monitor endpoint	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Dest Port	fies the exit port for a remote session on a remote destination switch.

Example

In **Figure 138: Displaying the configuration of remote mirroring endpoints on the switch** on page 491, the `show monitor endpoint` output shows that the switch is configured as the remote endpoint (destination) for two remote sessions from the same monitored source interface.

Figure 138: *Displaying the configuration of remote mirroring endpoints on the switch*

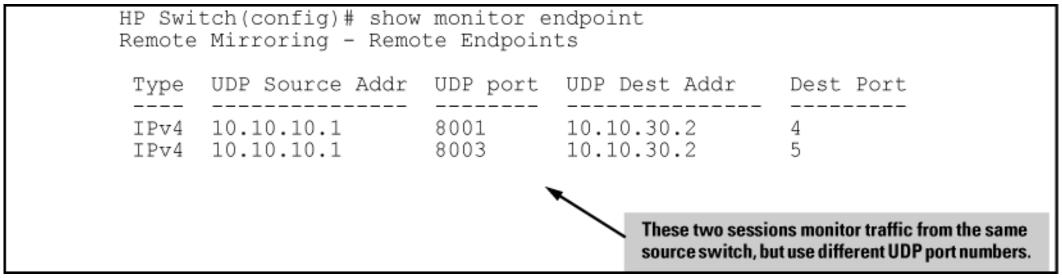

```
HP Switch(config)# show monitor endpoint
Remote Mirroring - Remote Endpoints

 Type  UDP Source Addr  UDP port  UDP Dest Addr    Dest Port
 ----  ---------------  --------  ---------------  ---------
 IPv4  10.10.10.1       8001      10.10.30.2       4
 IPv4  10.10.10.1       8003      10.10.30.2       5
```


Viewing the mirroring configuration for a specific session

Syntax

`show monitor [1 - 4 | name` *`name-str`*`]`

Displays detailed configuration information for a specified local or remote mirroring session on a source switch.

Session	Displays the number of the specified session.
Session Name	Displays the name of the session, if configured.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.
Mirroring Destination	For a local mirroring session, displays the port configured as the exit port on the source switch. For a remote mirroring session, displays `IPv4`, which indicates mirroring to a remote (endpoint) switch.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Status	For a remote session, displays current session activity: • active: The session is configured and is mirroring traffic. A remote path has been discovered to the destination. • inactive: The session is configured, but is not currently mirroring traffic. A remote path has **not** been discovered to the destination. • not defined: Mirroring is not configured for this session.

Table Continued

Monitoring Sources	For the specified local or remote session, displays the source (port, trunk, or VLAN) interface and the MAC address (if configured) used to select mirrored traffic.
Direction	For the selected interface, indicates whether mirrored traffic is entering the switch (`in`), leaving the switch (`out`), or `both`.

Viewing a remote mirroring session

After you configure session 2 for remote mirroring (**Figure 139: Configuring a remote mirroring session and monitored source** on page 492), you can enter the `show monitor 2` command to verify the configuration (**Figure 140: Displaying the Configuration of a Remote Mirroring Session** on page 492.)

Figure 139: *Configuring a remote mirroring session and monitored source*

```
HP Switch(config)# mirror 2 name test-10 remote ip 10.10.10.1 8010 10.10.30.2
Caution: Please configure destination switch first.
         Do you want to continue [y/n]? y
HP Switch(config)# interface b1 monitor all both mirror 2
```

Figure 140: *Displaying the Configuration of a Remote Mirroring Session*

```
HP Switch(config)# show monitor 2
Network Monitoring

   Session: 2    Session Name: test-10
   Policy: no policy relationship exists

      Mirror Destination:  IPv4
         UDP Source Addr  UDP port  UDP Dest Addr    Status
         ---------------  --------  ---------------  --------
         10.10.10.1       8010      10.10.30.2       active

      Monitoring Sources  Direction
      ------------------  ---------
      Port: B1            Both
```

If no monitored (source) interface is configured for a mirroring session, no information is displayed in the Monitoring Sources and Direction columns.

Viewing a MAC-based mirroring session

After you configure a MAC-based mirroring session (**Figure 141: Configuring a MAC-based mirroring session** on page 493), you can enter the `show monitor 3` command to verify the configuration (**Figure 142: Displaying a MAC-based mirroring session** on page 493.)

Figure 141: *Configuring a MAC-based mirroring session*

```
HP Switch(config)# mirror 3 port a1
HP Switch# monitor mac 112233-445566 src mirror 3
```

Figure 142: *Displaying a MAC-based mirroring session*

```
HP Switch(config)# show monitor 3
Network Monitoring

   Session: 3    Session Name:
   Policy: no policy relationship exists

      Mirror Destination:  A1    (Port)

      Monitoring Sources  Direction
      ------------------  ---------
      MAC:  112233-445566 Source
```

The MAC address used to select packets in a local mirroring session is displayed in these columns.

Viewing a local mirroring session

When used to display the configuration of a local session, the `show monitor` command displays a subset of the information displayed for a remote mirroring session.

Example

Figure 143: Displaying the configuration of a local mirroring session on page 493 displays a local mirroring configuration for a session configured as follows:

- Session number: 1
- Session name: Detail
- Classifier-based mirroring policy, "MirrorAdminTraffic", is used to select inbound traffic on port B1.
- Mirrored traffic is sent to exit port B3.

Figure 143: *Displaying the configuration of a local mirroring session*

```
HP Switch(config)# show monitor 1
Network Monitoring

   Session: 1    Session Name: Detail
   Policy: MirrorAdminTraffic

      Mirror Destination:  B3    (Port)

      Monitoring Sources  Direction
      ------------------  ---------
      Port: B1            In
```

Viewing information on a classifier-based mirroring session

In the following example, a classifier-based mirroring policy (`mirrorAdminTraffic`) mirrors selected inbound IPv4 packets on VLAN 5 to the destination device configured for mirroring session 3.

Figure 144: *Configuring a classifier-based mirroring policy in a local mirroring session*

```
HP Switch(config)# mirror 3 port c1
Caution: Please configure destination switch first.
         Do you want to continue [y/n]? y
HP Switch(config)# class ipv4 AdminTraffic
HP Switch(config-class)# match ip 15.29.61.1 0.63.255.255 0.0.0.0
255.255.255.255
HP Switch(config-class)# match ip 0.0.0.0 255.255.255.255 15.29.61.1
0.63.255.255
HP Switch(config-class)# exit
HP Switch(config)# policy mirror MirrorAdminTraffic
HP Switch(config-policy)# class ipv4 AdminTraffic action mirror 3
HP Switch(config-policy)# exit
HP Switch(config)# vlan 5 service-policy MirrorAdminTraffic in
```

Displaying a classifier-based policy in a local mirroring session

```
switch(config)# show monitor 3

Network Monitoring

   Session: 3   Session Name:
   Policy: MirrorAdminTraffic

     Mirror Destination:  C1   (Port)

     Monitoring Sources  Direction
     ------------------  ---------
     VLAN: 5             Source
```

Viewing information about a classifier-based mirroring configuration

Syntax

`show class ipv4` *`classname`* `show class ipv6` *`classname`* `show class config`

`ipv4` *`classname`*	Lists the statements that make up the IPv4 class identified by *`classname`* .
`ipv6` *`classname`*	Lists the statements that make up the IPv6 class identified by *`classname`* .
`config`	Displays all classes, both IPv4 and IPv6, and lists the statements that make up each class.

Additional variants of the `show class …` command provide information on classes that are members of policies that have been applied to ports or VLANs.

Figure 145: *`show class` output for a mirroring policy*

```
HP Switch(config)# show class ipv4 AdminTraffic

Statements for Class ipv4 "AdminTraffic"

 10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
 20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Viewing information about a classifier-based mirroring configuration

Syntax

`show policy` *`policy-name`* `show policy config`

`policy-name`	Lists the statements that make up the specified policy.
`config`	Displays the names of all policies defined for the switch and lists the statements that make up each policy.

Additional variants of the `show policy` command provide information on policies that have been applied to ports or VLANs.

Figure 146: *`show policy` output for a mirroring policy*

```
HP Switch(config)# show policy MirrorAdminTraffic

Statements for Policy "MirrorAdminTraffic"

     10 class ipv4 "AdminTraffic" action mirror 3
```

Viewing information about statistics on one or more mirroring policies

Syntax

`[show | clear] statistics policy` *`policy-name`* `port` *`port-num`*

`[show | clear] statistics policy` *`policy-name`* `vlan` *`vid`* `in`

`show`	Displays the statistics for a specified policy applied to a specified port or VLAN.
`clear`	Clears statistics for the specified policy and port or VLAN.
`policy-name`	The name of the policy.

Table Continued

`port-num`	The number of the port on which the policy is applied (single port only, not a range.)
`vid`	The number or name of the vlan on which the policy is applied. VLAN ID numbers range from 1 to 4094.
`in`	Indicates that statistics are shown for inbound traffic only.

Figure 147: show statistics policy output for a mirroring policy on page 496 shows the number of packets (in parentheses) that have been mirrored for each match/ignore statement in the mirroring policy.

Figure 147: *`show statistics policy` output for a mirroring policy*

```
HP Switch# show statistics policy MirrorAdminTraffic vlan 30 in

HitCounts for Policy MirrorAdminTraffic

10 class ipv4 "AdminTraffic" action mirror 3

(5244)  10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
(9466)  20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Viewing resource usage for mirroring policies

Syntax

```
show policy resources
```

Displays the number of hardware resources (rules, meters, and application port ranges) used by classifier-based mirroring policies (local and remote) that are currently applied to interfaces on the switch, as well as QoS policies and other software features.

NOTE

The information displayed is the same as the output of the `show qos resources` and `show access-list resources` commands.

Figure 148: *Displaying the hardware resources used by currently configured mirroring policies*

Includes the hardware resources used by classifier-based local and remote mirroring policies that are currently applied to interfaces on the switch.

```
HP Switch# show policy resources

 Resource usage in Policy Enforcement Engine
        |    Rules     |  Rules Used
  Ports |  Available   |  ACL  |  QoS  |  IDM  |   VT  | Mirror | Other |
  ------+--------------+-------+-------+-------+-------+--------+-------|
  1-24  |         3014 |    15 |    11 |     0 |     1 |      0 |     3 |
  25-48 |         3005 |    15 |    10 |    10 |     1 |      0 |     3 |
  A     |         3017 |    15 |     8 |     0 |     1 |      0 |     3 |

        |    Meters    |  Meters Used
  Ports |  Available   |  ACL  |  QoS  |  IDM  |   VT  | Mirror | Other |
  ------+--------------+-------+-------+-------+-------+--------+-------|
  1-24  |          250 |       |     5 |     0 |       |        |     0 |
  25-48 |          251 |       |     4 |     0 |       |        |     0 |
  A     |          253 |       |     2 |     0 |       |        |     0 |

        | Application  |
        | Port Ranges  |  Application Port Ranges Used
  Ports |  Available   |  ACL  |  QoS  |  IDM  |   VT  | Mirror | Other |
  ------+--------------+-------+-------+-------+-------+--------+-------|
  1-24  |         3014 |     2 |     0 |     0 |       |      0 |     0 |
  25-48 |         3005 |     2 |     0 |     0 |       |      0 |     0 |
  A     |         3017 |     2 |     0 |     0 |       |      0 |     0 |

 0 of 8 Policy Engine management resources used.
 Key:
 ACL = Access Control Lists
 QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
 IDM = Identity Driven Management
 VT  = Virus Throttling blocks
 Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
 Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

 Resource usage includes resources actually in use, or reserved for future
 use by the listed feature.  Internal dedicated-purpose resources, such as
 port bandwidth limits or VLAN QoS priority, are not included.
```

Viewing the mirroring configurations in the running configuration file

Use the `show run` command to view the current mirroring configurations on the switch. In the `show run` command output, information about mirroring sources in configured sessions begins with the `mirror` keyword; monitored source interfaces are listed per-interface.

Example

Figure 149: *Displaying mirroring sources and sessions in the running configurations*

```
HP Switch(config)# show run
Running configuration:
; J8697A Configuration Editor; Created on release #K.12.XX
max-vlans 300
ip access-list extended "100"
   10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
   exit
   no ip address
   exit
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface B1
   monitor ip access-group "100" In mirror 1
   monitor all Both mirror 2
   exit
. . .
```

Mirroring sessions with exit ports configured on the switch: B3 is an exit port for a local session; session 2 has a remote destination and exit port.

Selection criteria used to monitor traffic on port B1 for mirroring sessions 1 (ACL-based) and 2 (direction -based)

Information about remote endpoints configured for remote sessions on the switch begin with the `mirror endpoint` keywords. In the following example, two remote sessions use the same exit port:

Figure 150: *Displaying remote mirroring endpoints in the running configuration*

```
HP Switch(config)# show run
Running configuration:
; J8693A Configuration Editor; Created on release #K.12.XX
module 3 type J8694A
. . .

mirror endpoint ip 10.10.20.1 8010 10.10.30.2 port 4
mirror endpoint ip 10.10.51.10 7955 10.10.30.2 port 4
. . .
```

Remote endpoints configured on the switch, including source IP address, UDP port number, destination IP address, and remote exit port. Each remote session is identified by a unique UDP port number.

Compatibility mode

Table 23: *Modules compatibility mode (enabled versus disabled)*

Modules	Compatibility mode enabled	Compatibility mode disabled
v2 zl modules only	Can insert zl module and the module will come up. Any v2 zl modules are limited to the zl configuration capacities.	v2 zl modules are at full capacity.ZL modules are not allowed to power up.
Mixed v2 zl and zl modules	Can insert zl module and the module will come up. Any v2 zl modules are limited to the zl configuration capacities.If compatibility mode is disabled, the zl modules go down.	ZL modules are not allowed to power up.

Table Continued

Modules	Compatibility mode enabled	Compatibility mode disabled
ZL modules only	Same as exists already.If a v2 zl module is inserted, it operates in the same mode as the zl module, but with performance increases. In Compatibility Mode, no v2 zl features are allowed, whether the modules are all v2 zl or not.	The Management Module is the only module that powers up. If Compatibility Mode is disabled and then enabled, the startup config is erased and the chassis reboots.

Port and trunk group statistics and flow control status

The features described in this section enable you to determine the traffic patterns for each port since the last reboot or reset of the switch. You can display:

- A general report of traffic on all LAN ports and trunk groups in the switch, along with the per-port flow control status (On or Off.)
- A detailed summary of traffic on a selected port or trunk group.

You can also reset the counters for a specific port.

The menu interface provides a dynamic display of counters summarizing the traffic on each port. The CLI lets you see a static "snapshot" of port or trunk group statistics at a particular moment.

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting.

NOTE

The Reset action resets the counter display to zero for the current session, but does not affect the cumulative values in the actual hardware counters. (In compliance with the SNMP standard, the values in the hardware counters are not reset to zero unless you reboot the switch.) Exiting from the console session and starting a new session restores the counter displays to the accumulated values in the hardware counters.

Traffic mirroring overview

Starting in software release K.12.xx, traffic mirroring (Intelligent Mirroring) allows you to mirror (send a copy of) network traffic received or transmitted on a switch interface to a local or remote destination, such as a traffic analyzer or IDS.)

Traffic mirroring provides the following benefits:

- Allows you to monitor the traffic flow on specific source interfaces.
- Helps in analyzing and debugging problems in network operation resulting from a misbehaving network or an individual client. The mirroring of selected traffic to an external device makes it easier to diagnose a network problem from a centralized location in a topology spread across a campus.
- Supports remote mirroring to simultaneously mirror switch traffic on one or more interfaces to multiple remote destinations. (In remote mirroring, you must first configure the remote mirroring endpoint—remote switch and exit port—before you specify a mirroring source for a session.)

Mirroring overview

Figure 151: Local and remote sessions showing mirroring terms on page 500 shows an example of the terms used to describe the configuration of a sample local and remote mirroring session:

- In the local session, inbound traffic entering Switch A is monitored on port A2 and mirrored to a destination (host), traffic analyzer 1, through exit port A15 on the switch.A local mirroring session means that the monitored interface (A2) and exit port (A15) are on the same switch.
- In the remote session, inbound traffic entering Switch A is monitored on port A1. A mirrored copy of monitored traffic is routed through the network to a remote mirroring endpoint: exit port B7 on Switch B. A destination device, traffic analyzer 2, is connected to the remote exit port.A remote mirroring session means that:
  - The monitored interface (A1) and exit port (B7) are on different switches.
  - Mirrored traffic can be bridged or routed from a source switch to a remote switch.

Figure 151: *Local and remote sessions showing mirroring terms*

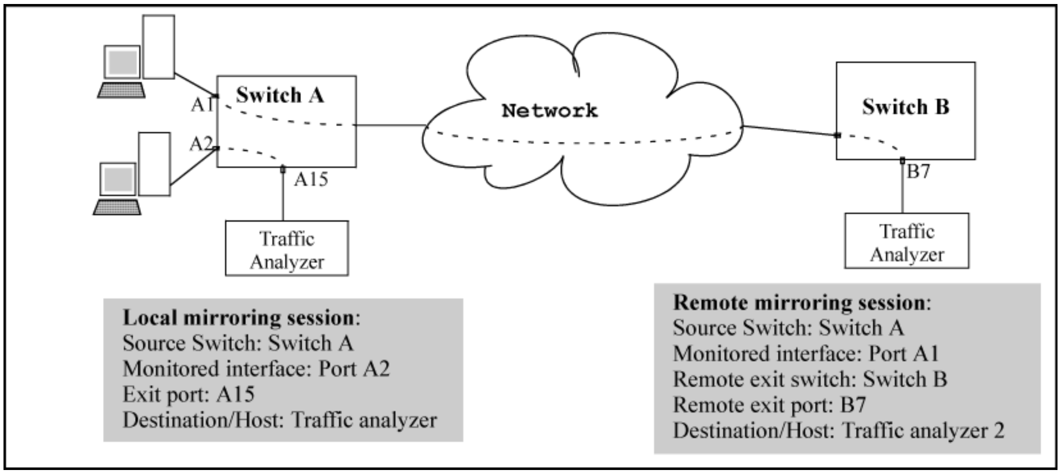

Mirroring destinations

Traffic mirroring supports destination devices that are connected to the local switch or to a remote switch:

- Traffic can be copied to a destination (host) device connected to the same switch as the mirroring source in a local mirroring session. You can configure up to four exit ports to which destination devices are connected.
- Traffic can be bridged or routed to a destination device connected to a different switch in a remote mirroring session. You can configure up to 32 remote mirroring endpoints (IP address and exit port) to which destination devices are connected.

Mirroring sources and sessions

Traffic mirroring supports the configuration of port and VLAN interfaces as mirroring sources in up to **four** mirroring sessions on a switch. Each session can have one or more sources (ports and/or static trunks, a mesh, or a VLAN interface) that monitor traffic entering and/or leaving the switch.

NOTE

Using the CLI, you can make full use of the switch's local and remote mirroring capabilities. Using the Menu interface, you can configure only local mirroring for either a single VLAN or a group of ports, static trunks, or both.

In remote mirroring, a 54-byte remote mirroring tunnel header is added to the front of each mirrored frame for transport from the source switch to the destination switch. This may cause some frames that were close to the MTU size to exceed the MTU size. Mirrored frames exceeding the allowed MTU size are dropped, unless the optional `[truncation]` parameter is set in the `mirror` command.

Mirroring sessions

A mirroring session consists of a mirroring source and destination (endpoint.) Although a mirroring source can be one of several interfaces, as mentioned above, for any session, the destination must be a single (exit) port. The exit port cannot be a trunk, VLAN, or mesh interface.

You can map multiple mirroring sessions to the same exit port, which provides flexibility in distributing hosts, such as traffic analyzers or an IDS. In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.

Mirroring sessions can have the same or a different destination. You can configure an exit port on the local (source) switch and/or on a remote switch as the destination in a mirroring session. When configuring a mirroring destination, consider the following options:

- Mirrored traffic belonging to different sessions can be directed to the same destination or to different destinations.
- You can reduce the risk of oversubscribing a single exit port by:
    - Directing traffic from different session sources to multiple exit ports.
    - Configuring an exit port with a higher bandwidth than the monitored source port.
- You can segregate traffic by type, direction, or source.

Mirroring session limits

A switch running software release K.12.**xx** or greater supports the following:

- A maximum of four mirroring (local and remote) sessions.
- A maximum of 32 remote mirroring endpoints (exit ports connected to a destination device that receive mirrored traffic originating from monitored interfaces on a different switch.)

Selecting mirrored traffic

You can use any of the following options to select the traffic to be mirrored on a port, trunk, mesh, or VLAN interface in a local or remote session:

- All trafficMonitors all traffic entering or leaving the switch on one or more interfaces (inbound and outbound.)
- Direction-based traffic selectionMonitors traffic that is either entering or leaving the switch (inbound or outbound.) Monitoring traffic in only one direction improves operation by reducing the amount of traffic sent to a mirroring destination.
- MAC-based traffic selectionMonitors only traffic with a matching source and/or destination MAC address in packet headers entering and/or leaving the switch on one or more interfaces (inbound and/or outbound.)
- Classifier-based service policyProvides a finer granularity of match criteria to zoom in on a subset of a monitored port or VLAN traffic (IPv4 or IPv6) and select it for local or remote mirroring (inbound only.)

Deprecation of ACL-based traffic selection

In software release K.14.01 or greater, the use of ACLs for selecting traffic in a mirroring session has been deprecated and is replaced by the use of advanced classifier-based service policies.

As with ACL criteria, classifier-based match/ignore criteria allow you to limit a mirroring session to selected inbound packets on a given port or VLAN interface (instead of mirroring all inbound traffic on the interface.)

The following commands have been deprecated:

- `interface` *`port/trunk/mesh`* `monitor ip access-group` *`acl-name`* `in mirror [1 - 4 |` *`name-str`*`]`
- `vlan` *`vid-#`* `monitor ip access-group` *`acl-name`* `in mirror [1 - 4 |` *`name-str`*`]`

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies.

If you are running software release K.13.XX or earlier, ACL permit/deny criteria are supported to select IP traffic entering a switch to mirror in a local or remote session, using specified source and/or destination criteria.

Mirrored traffic destinations

Local destinations

A local mirroring traffic destination is a port on the same switch as the source of the traffic being mirrored.

Remote destinations

A remote mirroring traffic destination is an switch configured to operate as the exit switch for mirrored traffic sessions originating on other switches.

CAUTION

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port. In a remote mirroring session that uses IPv4 encapsulation, if the intended exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, Switch strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Monitored traffic sources

You can configure mirroring for traffic entering or leaving the switch on:

- Ports and static trunks

  Provides the flexibility for mirroring on individual ports, groups of ports, static port trunks, or any combination of these..
- Meshed ports

  Enables traffic mirroring on all ports configured for meshing on the switch.
- Static VLANs

  Supports traffic mirroring on static VLANs configured on the switch. This option enables easy mirroring of traffic from all ports on a VLAN. It automatically adjusts mirroring to include traffic from newly added ports and to exclude traffic from ports removed from the VLAN.

Criteria for selecting mirrored traffic

On the monitored sources listed above, you can configure the following criteria to select the traffic you want to mirror:

- Direction of traffic movement (entering or leaving the switch, or both.)
- Type of IPv4 or IPv6 traffic entering the switch, as defined by a classifier-based service policy.
- Source and/or destination MAC addresses in packet headers.

Mirroring configuration

The table below shows the different types of mirroring that you can configure using the CLI, Menu, and SNMP interfaces.

Table 24: *Mirroring configuration options*

Monitoring interface and configuration level	Traffic selection criteria	Traffic direction		
		CLI config	Menu and web i/f config [1]	Snmp config
VLAN	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic)	See **About selecting inbound traffic using advanced classifier-based mirroring**.		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Port(s) Trunk(s) Mesh	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic)	See **About selecting inbound traffic using advanced classifier-based mirroring**.		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Switch (global)	MAC source/ destination address	Inbound only Outbound only Both directions	Not available	Inbound only Outbound only Both directions

[1] *Configures only session 1, and only for local mirroring.*

Configuration notes

Using the CLI, you can configure all mirroring options on a switch.

Using the Menu, you can configure only session 1 and only local mirroring in session 1 for traffic in both directions on specified interfaces. (If session 1 has been already configured in the CLI for local mirroring for inbound-only or outbound-only traffic, and you use the Menu to modify the session 1 configuration, session 1 is automatically reconfigured to monitor both inbound and outbound traffic on the assigned interfaces. If session 1 has been configured in the CLI with a classifier-based mirroring policy or as a remote mirroring session, an error message is displayed if you try to use the Menu to configure the session.)

You can use the CLI can configure sessions 1 to 4 for local or remote mirroring in any combination, and override a Menu configuration of session 1.

You can also use SNMP configure sessions 1 to 4 for local or remote mirroring in any combination and override a Menu configuration of session 1, **except** that SNMP cannot be used to configure a classifier-based mirroring policy.

Remote mirroring endpoint and intermediate devices

The remote mirroring endpoint that is used in a remote mirroring session must be an switch that supports the mirroring functions described in this chapter. (A remote mirroring endpoint consists of the remote switch and exit port connected to a destination device.) Because remote mirroring on an switch uses IPv4 to encapsulate mirrored traffic sent to a remote endpoint switch, the intermediate switches and routers in a layer 2/3 domain can be from any vendor if they support IPv4.

The following restrictions apply to remote endpoint switches and intermediate devices in a network configured for traffic mirroring:

- The exit port for a mirroring destination must be an individual port and **not** a trunk, mesh, or VLAN interface.
- A switch mirrors traffic on static trunks, but not on dynamic LACP trunks.
- A switch mirrors traffic at line rate. When mirroring multiple interfaces in networks with high-traffic levels, it is possible to copy more traffic to a mirroring destination than the link supports. However, some mirrored traffic may not reach the destination. If you are mirroring a high-traffic volume, you can reduce the risk of oversubscribing a single exit port by:
  - Directing traffic from different session sources to multiple exit ports.
  - Configuring an exit port with a higher bandwidth than the monitored source port.

Migration to release K.12.xx

On a switch that is running a software release earlier than K.12.xx with one or more mirroring sessions configured, when you download and boot release K.12.xx, the existing mirroring configurations are managed as follows:

- A legacy mirroring configuration on a port or VLAN interface maps to session 1.
- Traffic-selection criteria for session 1 is set to `both`; both inbound and outbound traffic (traffic entering **and** leaving the switch) on the configured interface is selected for mirroring.
- In a legacy mirroring configuration, a local exit port is applied to session 1.

Booting from software versions earlier than K.12.xx

If it is necessary to boot the switch from a legacy (pre-K.12.xx) software version after using version K.12.xx or greater to configure mirroring, remove mirroring from the configuration before booting with the earlier software.

Maximum supported frame size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU allowed in the path from the mirroring source to the mirroring destination, the frame is dropped, unless the optional `[truncation]` parameter is set in the `mirror` command.

Frame truncation

Mirroring does not truncate frames unless the `truncation` parameter in the `mirror` command is set. If that parameter is not set, oversized mirroring frames are dropped. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

Migration to release K.14.01 or greater

NOTE

If a switch is running software release K.12.xx, you must first upgrade to release K.13.xx before migrating the switch to release K.14.01 or greater.

When you download and boot software release K.14.01 or greater on a switch that is running release K.13.xx and has one or more mirroring sessions configured, an ACL-based mirroring configuration on a port or VLAN interface is mapped to a class and policy configuration based on the ACL.

The new mirroring policy is automatically configured on the same port or VLAN interface on which the mirroring ACL was assigned. The behavior of the new class and mirroring-policy configuration exactly matches the traffic-selection criteria and mirroring destination used in the ACL-based session.)

Figure 152: Mirroring configuration in show run output in release K.13.xx on page 505 and **Figure 153: Mirroring configuration in show run output in release K.14.01 or greater** on page 505 show how ACL-based selection criteria in a mirroring session are converted to a classifier-based policy and class configuration when you install release K.14.01 or greater on a switch.

Figure 152: *Mirroring configuration in* `show run` *output in release K.13.xx*

```
HP Switch(config)# show run
Running configuration:
. . .
ip access-list extended "100"
   10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
. . .
mirror 1 port C1
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface C1
   monitor ip access-group "100" In mirror 1
  exit
. . .
```

Configuration of ACL 100 that is used to select mirrored traffic in session 1

Existing mirror sessions configured on the switch for a local (port C1 in session 1) and remote (session 2) monitored interface

ACL-based traffic selection on monitored interface C1 in session 1

Figure 153: *Mirroring configuration in* `show run` *output in release K.14.01 or greater*

```
HP Switch(config)# show run
Running configuration:
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
class ipv4 "100MirrorClass"
   10 match icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
policy mirror "100MirrorPolicy"
   10 class ipv4 "100" action mirror 1
  exit
. . .
interface C1
   service-policy "100MirrorPolicy" In
  exit
. . .
```

After migration to release K.14.01 or greater, the existing mirroring configurations for sessions 1 (local) and 2 (remote) on the switch remain the same.

The traffic-selection criteria in ACL 100 (Figure B-B-27) applied to inbound traffic on port C1 in session 1 are converted to a class and policy configuration with the names, "100MirrorClass" and "100MirrorPolicy", which are applied to inbound traffic on port C1 in session 1 with the **service-policy** command.

Using the Menu to configure local mirroring

Menu and WebAgent limits

You can use the Menu and WebAgent to quickly configure or reconfigure local mirroring on session 1 and allow one of the following two mirroring source options:

- Any combination of source ports, trunks, and a mesh.
- One static, source VLAN interface.

The Menu and WebAgent also has these limits:

- Configure and display only session 1 and only as a local mirroring session for traffic in **both** directions on the specified interface. (Selecting inbound-only or outbound-only is not an option.)
- If session 1 has been configured in the CLI for local mirroring for inbound-only or outbound-only traffic on one or more interfaces, using the Menu to change the session 1 configuration **automatically reconfigures the session** to monitor both inbound and outbound traffic on the designated interface(s.)

- If session 1 has been configured in the CLI with an ACL/classifier-based mirroring policy or as a remote mirroring session, the Menu is not available for changing the session 1 configuration.
- The CLI (and SNMP) can be used to override any Menu configuration of session 1.

Remote mirroring overview

To configure a remote mirroring session in which the mirroring source and destination are on different switches, follow these general steps:

After you complete **6.b** on page 506, the switch begins mirroring traffic to the remote destination (endpoint) configured for the session.

1. Determine the IP addressing, UDP port number, and destination (exit) port number for the remote session:
   - **a.** Source VLAN or subnet IP address on the source switch.
   - **b.** Destination VLAN or subnet IP address on the destination switch.
   - **c.** Random UDP port number for the session (7933-65535.)
   - **d.** Remote mirroring endpoint: Exit port and IP address of the remote destination switch (In a remote mirroring endpoint, the IP address of the exit port and remote switch can belong to different VLANs. Any loopback IP address can be used except the default loopback address 127.0.0.1.)
2. Requirement: For remote mirroring, the same IP addressing and UDP port number must be configured on both the source and destination switches.
3. On the remote destination (endpoint) switch, enter the `mirror endpoint` command with the information from **1** on page 506 to configure a mirroring session for a specific exit port.
4. Determine the session (1 to 4) and (optional) alphanumeric name to use on the **source** switch.
5. Determine the traffic to be filtered by any of the following selection methods and the appropriate configuration level (VLAN, port, mesh, trunk, global):
   - **a.** Direction: inbound, outbound, or both.
   - **b.** Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic.
   - **c.** MAC source and/or destination address: inbound, outbound, or both.
6. On the **source** switch:
   - **a.** Enter the `mirror` command with the session number (1 to 4) and the IP addresses and UDP port number from **1** on page 506 to configure a mirroring session. If desired, enter the `[truncation]` parameter to allow oversize packets to be truncated rather than dropped.
   - **b.** Enter one of the following commands to configure one or more of the traffic-selection methods in **5** on page 506 for the configured session:

```
interface port/trunk/mesh [monitor | service-policy policy-name in] vlan vid
[monitor | service-policy policy-name in] monitor mac mac-addr
```

Quick reference to remote mirroring setup

The commands beginning with **Destination mirror on a remote switch** on page 479, configure mirroring for a remote session in which the mirroring source and destination are on different switches:

- The `mirror` command identifies the destination in a mirroring session.
- The `interface` and `vlan` commands identify the monitored interface, traffic direction, and traffic-selection criteria for a specified session.

CAUTION

When configuring a remote mirroring session, always configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.

High-level overview of the mirror configuration process

Determine the mirroring session and destination

For a local mirroring session

Determine the port number for the exit port (such as A5, B10, and so forth), then go to **Configure the monitored traffic in a mirror session**.

For a remote mirroring session

Determine the following information and then go to "Configure a mirroring destination on a remote switch" on page 498.

- The IP address of the VLAN or subnet on which the exit port exists on the destination switch.
- The port number of the remote exit port on the remote destination switch. (In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.)
- The IP address of the VLAN or subnet on which the mirrored traffic enters or leaves the source switch.

CAUTION

Although the switch supports the use of UDP port numbers from 1 to 65535, UDP port numbers below 7933 are reserved for various IP applications. Using these port numbers for mirroring can result in an interruption of other IP functions, and in non-mirrored traffic being received on the destination (endpoint) switch and sent to the device connected to the remote exit port.

- The unique UDP port number to use for the session on the source switch. (The recommended port range is from 7933 to 65535.)

Configure a mirroring destination on a remote switch

This step is required only if you are configuring a remote mirroring session in which the exit port is on a different switch than the monitored (source) interface. If you are configuring local mirroring, go to **Configure a mirroring session on the source switch** on page 507.

For remote mirroring, you must configure the **destination** switch to recognize each mirroring session and forward mirrored traffic to an exit port before you configure the **source** switch. Configure the destination switch with the values you determined for remote mirroring in **High-level overview of the mirror configuration process** on page 507.

NOTE

A remote destination switch can support up to 32 remote mirroring endpoints (exit ports connected to a destination device in a remote mirroring session.)

Configure a destination switch in a remote mirroring session

Enter the `mirror endpoint ip` command on the remote switch to configure the switch as a remote endpoint for a mirroring session with a different source switch.

Configure a mirroring session on the source switch

To configure local mirroring, only a session number and exit port number are required.

If the exit port for a mirroring destination is on a remote switch instead of the local (source) switch, you must enter the source IP address, destination IP address, and UDP port number for the remote mirroring session. You may also wish to enable frame truncation to allow oversize frames to be truncated rather than dropped.

Frames that exceed the maximum size (MTU) are either dropped or truncated, according to the setting of the `[truncation]` parameter in the `mirror` command. Frames that are near the MTU size may become oversize when the 54-byte remote mirroring tunnel header is added for transport between source switch and destination switch. (The addition of the header is a frequent cause for frames becoming oversize, but note that all oversize frames, whatever the cause of their excess size, are dropped or truncated.) If a frame is truncated, bytes are removed from the end of the frame. This may cause the checksum in the original frame header to fail. Some protocol analyzers may flag such a checksum mismatch as an alert.

NOTE

Note that if you enable jumbo frames to allow large frames to be transmitted, you must enable jumbo frames on all switches in the path between source and destination switches.

Configure a source switch in a remote mirroring session

Enter the `mirror remote ip` command on the source switch to configure a remote destination switch for a mirroring session on the source switch. The source IP address, UDP port number, and destination IP address that you enter must be the same values that you entered with the `mirror endpoint ip` command.

CAUTION

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to the destination device connected to each exit port. In a remote mirroring session that uses IPv4 encapsulation, if the remote (endpoint) switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, Hewlett Packard Enterprise strongly recommends that you configure the endpoint switch in a remote mirroring session, as described on the previous page in the section titled "For a remote mirroring session", before using the `mirror remote ip` command in this section to configure the mirroring source for the same session.

Configure the monitored traffic in a mirror session

This step configures one or more interfaces on a source switch with traffic-selection criteria to select the traffic to be mirrored in a local or remote session configured in section.

Traffic selection options

To configure traffic mirroring, specify the source interface, traffic direction, and criteria to be used to select the traffic to be mirrored by using the following options:

- Interface type
  - Port, trunk, and/or mesh
  - VLAN
  - Switch (global configuration level)
- Traffic direction and selection criteria
  - All inbound and/or outbound traffic on a port or VLAN interface
  - Only inbound IP traffic selected with an ACL (deprecated in software release K.14.01 and greater)
  - Only inbound IPv4 or IPv6 traffic selected with a classifier-based mirroring policy
  - All inbound and/or outbound traffic selected by MAC source and/or destination address

The different ways to configure traffic-selection criteria on a monitored interface are described in the following sections.

Mirroring-source restrictions

In a mirroring session, you can configure any of the following sources of mirrored traffic:

- Multiple port and trunk, and/or mesh interfaces
- One VLAN

  If you configure a VLAN as the source interface in a mirroring session and assign a second VLAN to the session, the second VLAN overwrites the first VLAN as the source of mirrored traffic.

- One classifier-based policy

  If you configure a mirroring policy on a port or VLAN interface to mirror inbound traffic in a session, you cannot configure a port, trunk, mesh, ACL, or VLAN as an additional source of mirrored traffic in the session.
- Up to 320 MAC addresses (used to select traffic according to source, destination MAC address, or both) in all mirroring sessions configured on a switch

About selecting all inbound/outbound traffic to mirror

If you have already configured session 1 with a local or remote destination, you can enter the `vlan` *`vid`* `monitor` or `interface` *`port`* `monitor` command without additional parameters for traffic-selection criteria and session number to configure mirroring for all inbound and outbound traffic on the specified VLAN or port interfaces in session 1 with the preconfigured destination.

Untagged mirrored packets

Although a VLAN tag is added (by default) to the mirrored copy of untagged outbound packets to indicate the source VLAN of the packet, it is sometimes desirable to have mirrored packets look exactly like the original packet. The `no-tag-added` parameter gives you the option of not tagging mirrored copies of outbound packets, as shown in **Figure 154: Mirroring commands with the no-tag-added option** on page 509 and **Figure 155: Displaying a mirror session configuration with the no-tag-added option** on page 509.

Figure 154: *Mirroring commands with the `no-tag-added` option*

```
HP Switch(config)#interface 3 monitor all in mirror 1 no-tag-added

HP Switch(config)#interface mesh monitor all both mirror 1 no-tag-added
```

Figure 155: *Displaying a mirror session configuration with the `no-tag-added` option*

```
HP Switch# show monitor 1

Network Monitoring

   Session: 1    Session Name:
   ACL: no ACL relationship exists

      Mirror Destination: 48
      Untagged traffic  : untagged    <--- Indicates the no-tag-added option is configured.
      Monitoring Sources  Direction
      ------------------  ---------
      Port: 3             Both
```

About using SNMP to configure `no-tag-added`

The MIB object hpicfBridgeDontTagWithVlan is used to implement the `no-tag-added` option, as shown below:

```
hpicfBridgeDontTagWithVlan OBJECT-TYPE
   SYNTAX INTEGER
     {
     enabled(1),
     disabled(2)
    }
   MAX-ACCESS  read-write
   STATUS   current
   DESCRIPTION
    "This oid mentions whether VLAN tag is part of the
   mirror'ed copy of the packet. The value 'enabled'
```

```
 denotes that the VLAN tag shouldn't be part
 of the mirror'ed copy; 'disabled' does put
 the VLAN tag in the mirror'ed copy. Only one
 logical port is allowed.
 This object is persistent and when written
 the entity
   SHOULD save the change to non-volatile storage."
 DEFVAL { 2 }
::= { hpicfBridgeMirrorSessionEntry 2 }
```

Operating notes

The following conditions apply for the `no-tag-added` option:

- The specified port can be a physical port, trunk port, or mesh port.
- Only a single logical port (physical port or trunk) can be associated with a mirror session when the `no-tag-added` option is specified. No other combination of ACL mirroring, VLAN mirroring, or port mirroring can be associated with the mirror session. If more than one logical port is specified, the following error message is displayed:

  Cannot monitor more than one logical port with no-tag-added option
- If a port changes its VLAN membership and/or untagged status within the VLAN, the "untagged port mirroring" associated with that port is updated when the configuration change is processed.
- Only four ports or trunks can be monitored at one time when all four mirror sessions are in use (one logical port per mirror session) without VLAN tags being added to a mirrored copy.
- The `no-tag-added` option can also be used when mirroring is configured with SNMP.
- A VLAN tag is still added to the copies of untagged packets obtained via VLAN-based mirroring.

About selecting inbound traffic using an ACL (deprecated)

Deprecation of ACL-based traffic selection

In release K.14.01 or greater, the use of ACLs to select inbound traffic in a mirroring session has been replaced with classifier-based mirroring policies.

The following commands have been deprecated:

- `interface` *port/trunk/mesh* `monitor ip access-group` *acl-name* `in mirror {1 - 4 |` *name-str*`}`
- `vlan` *vid-#* `monitor ip access-group` *<ACL-NAME>* `in mirror {1 - 4 |` *<NAME-STR>*`}`

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies.

About selecting inbound/outbound traffic using a MAC address

Use the `monitor mac mirror` command at the global configuration level to apply a source and/or destination MAC address as the selection criteria used in a local or remote mirroring session.

While classifier-based mirroring allows you to mirror traffic using a policy to specify IP addresses as selection criteria, MAC-based mirroring allows you monitor switch traffic using a source and/or destination MAC address. You can apply MAC-based mirroring in one or more mirroring sessions on the switch to monitor:

- Inbound traffic
- Outbound traffic
- Both inbound and outbound traffic

MAC-based mirroring is useful in Switch Network Immunity security solutions that provide detection and response to malicious traffic at the network edge. After isolating a malicious MAC address, a security administrator can mirror all traffic sent to and received from the suspicious address for troubleshooting and traffic analysis.

The MAC address that you enter with the `monitor mac mirror` command is configured to select traffic for mirroring from all ports and learned VLANs on the switch. Therefore, a suspicions MAC address used in wireless applications can be continuously monitored as it re-appears in switch traffic on different ports or VLAN interfaces.

You can configure MAC-based mirroring from the CLI or an SNMP management station and use it to mirror:

- All inbound and outbound traffic from a group of hosts to one destination device.
- Inbound and/or outbound traffic from each host to a different destination device.
- Inbound and outbound traffic from all monitored hosts separately on two destination devices: mirroring all inbound traffic to one device and all outbound traffic to another device.

Restrictions

The following restrictions apply to MAC-based mirroring:

- Up to 320 different MAC addresses are supported for traffic selection in all mirroring sessions configured on the switch.
- A destination MAC address is not supported as mirroring criteria for routed traffic, because in routed packets, the destination MAC address is changed to the next-hop address when the packet is forwarded. Therefore, the destination MAC address that you want to mirror will not appear in routed packet headers.

  This restriction also applies to the destination MAC address of a host that is directly connected to a routing switch. (Normally, a host is connected to an edge switch, which is directly connected to the router.)

  To mirror routed traffic, we recommend that you use classifier-based policies to select IPv4 or IPv6 traffic for mirroring, as described in.
- On a switch, you can use a MAC address only once as a source MAC address and only once as a destination MAC address to filter mirrored traffic.

  For example, after you enter the following commands:

  ```
  monitor mac 111111-222222 src mirror 1
  monitor mac 111111-222222 dest mirror 2
  ```

  The following commands are not supported:

  ```
  monitor mac 111111-222222 src mirror 3
  monitor mac 111111-222222 dest mirror 4
  ```

  In addition, if you enter the `monitor mac 111111-222222 both mirror 1` command, you cannot use the MAC address `111111-222222` in any other `monitor mac mirror` configuration commands on the switch.
- To re-use a MAC address that has already been configured as a source and/or destination address for traffic selection in a mirror session, you must first remove the configuration by entering the `no` form of the command and then re-enter the MAC address in a new `monitor mac mirror` command.

  For example, if you have already configured MAC address `111111-222222` to filter inbound and outbound mirrored traffic, and you decide to use it to filter only inbound traffic in a mirror session, you could enter the following commands:

  ```
  monitor mac 111111-222222 both mirror 1
  no monitor mac 111111-222222 both mirror 1
  monitor mac 111111-222222 src mirror 1
  ```
- A mirroring session in which you configure MAC-based mirroring is not supported on a port, trunk, mesh, or VLAN interface on which a mirroring session with a classifier-based mirroring policy is configured.

About selecting inbound traffic using advanced classifier-based mirroring

In software release K.14.01 or greater, in addition to the traffic selection options described in **Configure the monitored traffic in a mirror session** on page 508, traffic mirroring supports the use of advanced classifier-based functions that provide:

- A finer granularity for selecting the inbound IP traffic that you want to mirror on an individual port or VLAN interface (instead of mirroring all inbound traffic on the interface)
- Support for mirroring both IPv4 and IPv6 traffic
- The ability to re-use the same traffic classes in different software-feature configurations; for example, you can apply both a QoS rate-limiting and mirroring policy on the same class of traffic.

Deprecation of ACL-based traffic selection

In software release K.14.01 or greater, advanced classifier-based policies replace ACL-based traffic selection in mirroring configurations.

Like ACL-based traffic-selection criteria, classifier-based service policies apply only to inbound traffic flows and are configured on a per-port or per-VLAN basis. In a mirroring session, classifier-based service policies do not support:

- The mirroring of outbound traffic exiting the switch
- The use of meshed ports as monitored (source) interfaces

Classifier-based mirroring is **not** designed to work with other traffic-selection methods in a mirroring session applied to a port or VLAN interface:

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.
- If a mirroring session is configured to use a classifier-based mirroring policy, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session on the same or a different interface.

Classifier-based mirroring policies provide greater precision when analyzing and debugging a network traffic problem. Using multiple match criteria, you can finely select and define the classes of traffic that you want to mirror on a traffic analyzer or IDS device.

Classifier-based mirroring configuration

1. Evaluate the types of traffic in your network and identify the traffic types that you want to mirror.
2. Create an IPv4 or IPv6 traffic class using the `class` command to select the packets that you want to mirror in a session on a preconfigured local or remote destination device.

   A traffic class consists of match criteria, which consist of match and ignore commands.

   - `match` commands define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.
   - `ignore` commands define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class.

   NOTE

   Be sure to enter match/ignore statements in the **precise order** in which you want their criteria to be used to check packets.

   The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

   a. IP source address (IPv4 and IPv6)
   b. IP destination address (IPv4 and IPv6)
   c. IP protocol (such as ICMP or SNMP)
   d. Layer 3 IP precedence bits
   e. Layer 3 DSCP codepoint

f. Layer 4 TCP/UDP application port (including TCP flags)
g. VLAN ID

3. Enter one or more match or ignore commands from the class configuration context to filter traffic and determine the packets on which policy actions will be performed.
4. Create a mirroring policy to configure the session and destination device to which specified classes of inbound traffic are sent by entering the `policy mirror` command from the global configuration context.

NOTE: Be sure to enter each class and its associated mirroring actions in the **precise order** in which you want packets to be checked and processed.

5. To configure the mirroring actions that you want to execute on packets that match the criteria in a specified class, enter one or more class action mirror commands from the policy configuration context.

   You can configure only one mirroring session (destination) for each class. However, you can configure the same mirroring session for different classes.

   A packet that matches the match criteria in a class is mirrored to the exit (local or remote) port that has been previously configured for the session, where session is a value from 1 to 4 or a text string (if you configured the session with a name when you entered the `mirror` command.)

   Prerequisite: The local or remote exit port for a session must be already configured before you enter the `mirror` *`session`* parameter in a class action statement:

   - In a local mirroring session, the exit port is configured with the `mirror` *`<SESSION-NUMBER>`* `port` command.
   - In a remote mirroring session, the remote exit port is configured with the `mirror endpoint ip` and `mirror` *`<SESSION-NUMBER>`* `remote ip` commands.

   Restriction: In a policy, you can configure only one mirroring session per class. However, you can configure the same session for different classes.

   Mirroring is not executed on packets that match ignore criteria in a class.

   The execution of mirroring actions is performed in the order in which the classes are numerically listed in the policy.

   The complete no form of the `class action mirror` command or the `no` *`<SEQ-NUMBER>`* command removes a class and mirroring action from the policy configuration.
6. To manage packets that do not match the match or ignore criteria in any class in the policy, and therefore have no mirroring actions performed on them, you can enter an optional default class. The default class is placed at the end of a policy configuration and specifies the mirroring actions to perform on packets that are neither matched nor ignored.
7. (Optional) To configure a default-class in a policy, enter the `default-class` command at the end of a policy configuration and specify one or more actions to be executed on packets that are not matched and not ignored.

   Prerequisite: The local or remote exit port for a session must be already configured with a destination device before you enter the `mirror` *`<SESSION>`* parameter in a default-class action statement.
8. Apply the mirroring policy to inbound traffic on a port (`interface service-policy in` command) or VLAN (`vlan service-policy in` command) interface.

CAUTION

After you apply a mirroring policy for one or more preconfigured sessions on a port or VLAN interface, the switch immediately starts to use the traffic-selection criteria and exit port to mirror traffic to the destination device connected to each exit port.

In a remote mirroring session that uses IPv4 encapsulation, if the remote switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic.

For this reason, Switch strongly recommends that you first configure the exit switch in a remote mirroring session, as described in **Configure a mirroring destination on a remote switch** on page 507 and **Configure a mirroring session on the source switch** on page 507, before you apply a mirroring service policy on a port or VLAN interface.

Restrictions: The following restrictions apply to a mirroring service policy:

- Only one mirroring policy is supported on a port or VLAN interface.
- If you apply a mirroring policy to a port or VLAN interface on which a mirroring policy is already configured, the new policy replaces the existing one.
- A mirroring policy is supported only on inbound traffic.

Because only one mirroring policy is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

Classifier-based mirroring restrictions

The following restrictions apply to mirroring policies configured with the classifier-based model:

- A mirroring policy is supported only on **inbound** IPv4 or IPv6 traffic.
- A mirroring policy is not supported on a meshed port interface. (Classifier-based policies are supported only on a port, VLAN, or trunk interface.)
- Only one classifier-based mirroring policy is supported on a port or VLAN interface. You can, however, apply a classifier-based policy of a different type, such as QoS.
- You can enter multiple `class action mirror` statements in a policy.
    - You can configure only one mirroring session (destination) for each class.
    - You can configure the same mirroring session for different classes.

- If a mirroring session is configured with a classifier-based mirroring policy on a port or VLAN interface, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session.

Figure 156: *Mirroring configuration in which only a mirroring policy is supported*

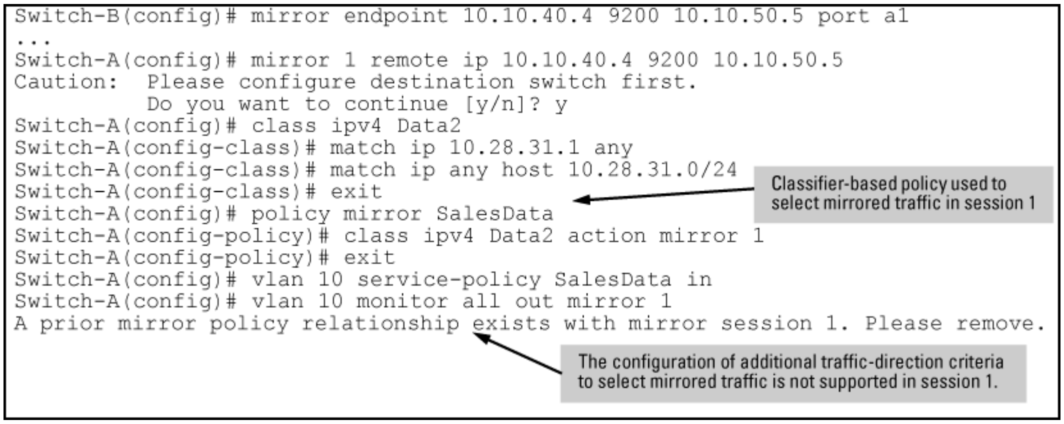

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution:  Please configure destination switch first.
          Do you want to continue [y/n]? y
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Switch-A(config)# vlan 10 monitor all out mirror 1
A prior mirror policy relationship exists with mirror session 1. Please remove.
```


- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.

Figure 157: *Mirroring configuration in which only traffic-selection criteria are supported*

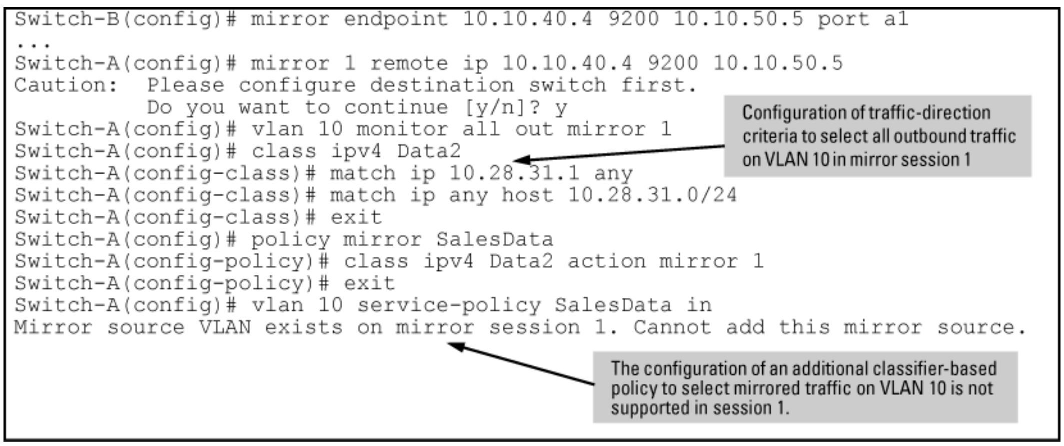

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution:  Please configure destination switch first.
          Do you want to continue [y/n]? y
Switch-A(config)# vlan 10 monitor all out mirror 1
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Mirror source VLAN exists on mirror session 1. Cannot add this mirror source.
```


About applying multiple mirroring sessions to an interface

You can apply a mirroring policy to an interface that is already configured with another traffic-selection method (MAC-based or all inbound and/or outbound traffic) for a different mirroring session.

The classifier-based policy provides a finer level of granularity that allows you to zoom in on a subset of port or VLAN traffic and select it for local or remote mirroring.

In the following example, traffic on Port b1 is used as the mirroring source for two different, local mirroring sessions:

- All inbound and outbound traffic on Ports b1, b2, and b3 is mirrored in session 4.
- Only selected voice traffic on Port b1 is mirrored in session 2.

Figure 158: *Example of applying multiple sessions to the same interface*

```
HP Switch(config)# mirror 4 port a2
HP Switch(config)# interface b1-b3 monitor all both mirror 4
HP Switch(config)# mirror 2 port b4
HP Switch(config)# class ipv4 voice
HP Switch(config-class)# match ip any any ip-dscp ef
HP Switch(config-class)# exit
HP Switch(config)# policy mirror IPphones
HP Switch(config-policy)# class ipv4 voice action mirror 2
HP Switch(config-policy)# exit
HP Switch(config)# interface b1 service-policy IPphones in
```

Mirroring configuration examples

Local mirroring using traffic-direction criteria

An administrator wants to mirror the inbound traffic from workstation "X" on port A5 and workstation "Y" on port B17 to a traffic analyzer connected to port C24 (see **Figure 159: Local mirroring topology** on page 516.) In this case, the administrator chooses "1" as the session number. (Any unused session number from 1 to 4 is valid.) Because the switch provides both the source and destination for the traffic to monitor, local mirroring can be used. In this case, the command sequence is:

- Configure the local mirroring session, including the exit port.
- Configure the monitored source interfaces for the session.

Figure 159: *Local mirroring topology*

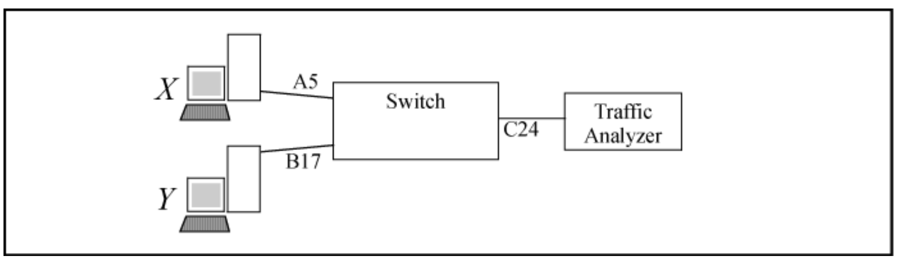

Figure 160: *Configuring a local mirroring session for all inbound and outbound port traffic*

Configures port C24 as the mirroring destination (exit port) for session 1.

```
HP Switch(config)# mirror 1 port c24
Caution: Please configure destination switch first.
         Do you want to continue [y/n]? y
HP Switch(config)# interface a5,b17 monitor all in mirror
1
```

Reminder to configure mirroring destination before configuring source.

Mirrors all inbound and outbound traffic on ports A5 and B17 to the mirroring destination configured for session 1.

Remote mirroring using a classifier-based policy

In the network shown in **Figure 161: Sample topology in a remote mirroring session** on page 517, an administrator has connected a traffic analyzer to port A15 (in VLAN 30) on switch C to monitor the TCP traffic to the server at 10.10.30.153 from workstations connected to switches A and B. Remote mirroring sessions are configured on switches A and B, and a remote mirroring endpoint on switch C. TCP traffic is routed through the network to the server from VLANs 10 and 20 on VLAN 30.

Figure 161: *Sample topology in a remote mirroring session*

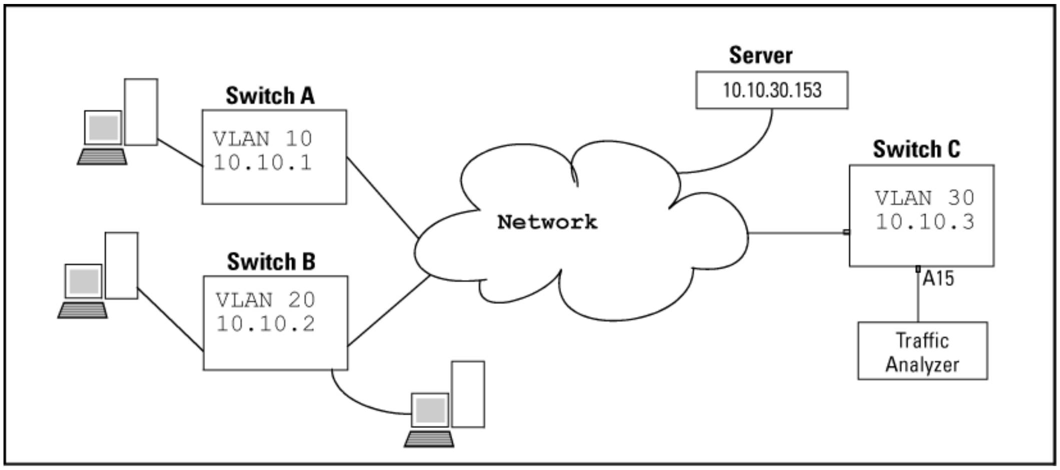

To configure this remote mirroring session using a classifier-based policy to select inbound TCP traffic on two VLAN interfaces, take the following steps:

1. On remote switch C, configure a remote mirroring endpoint using port A15 as the exit port (as described in **Configure a mirroring destination on a remote switch** on page 507.)

   Figure 162: *Configuring a remote mirroring endpoint: remote switch and exit port*

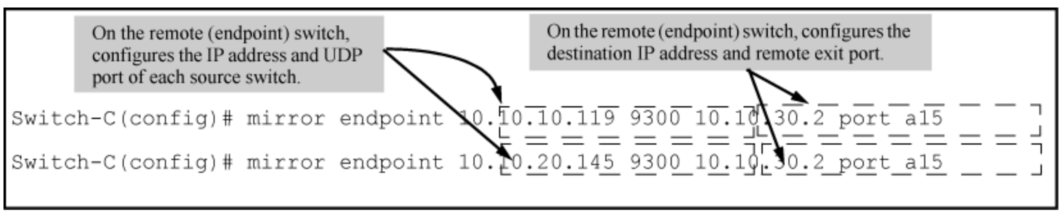

2. On source switch A, configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch A (as described in **Configure a mirroring session on the source switch** on page 507.)

3. On switch A, configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to the interfaces of VLAN 10 (as described in **About selecting inbound traffic using advanced classifier-based mirroring** on page 512.)

Figure 163: *Configuring a classifier-based policy on source switch A*

On a source switch, associates session number 1 with a source IP address and UDP port, and a remote destination IP address.

```
1 Switch-A(config)# mirror 1 remote ip 10.10.10.119 9300 10.10.30.2
  Caution: Please configure destination switch first.
           Do you want to continue [y/n]? y

2 Switch-A(config)# class ipv4 tcp7
  Switch-A(class-config)# match tcp any 10.10.30.153
  Switch-A(class-config)# match tcp any host 10.10.20.153/24
  Switch-A(class-config)# match tcp any any eq 80
  Switch-A(class-config)# exit

  Switch-A(config)# policy mirror mirrorTCP
  Switch-A(policy-config)# class ipv4 tcp7 action mirror 1
  Switch-A(policy-config)# exit

3 Switch-A(config)# vlan 10 service-policy mirrorTCP in
```

Class configuration that defines the matching TCP packets to be mirrored

Policy configuration that defines the preconfigured class and session/destination device to which matching packets are mirrored

Policy application to inbound traffic on a VLAN interface

1 The source IP address and UDP port number identify the mirroring source in session 1; the destination IP address identifies the remote switch to which traffic is mirrored. (The exit port for mirrored traffic, configured in Figure B-B-54, and the remote switch can belong to different VLANs.)

2 Configures a class that selects IPv4 TCP traffic destined to: the server at 10.10.30.153, a device in subnet 10.10.20.0, and any TCP traffic on port 80. (A packet that does not match these criteria is transmitted without being mirrored.)

3 Configures VLAN 10 as the source interface, and the mirroring policy as the selection criteria for inbound traffic on VLAN 10 in session 1.

4. On source switch B, repeat steps 2 and 3:
   a. Configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch B.
   b. Configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to a VLAN interface for VLAN 20.

   Because the remote session has mirroring sources on different switches, you can use the same session number (1) for both sessions.

Figure 164: *Configuring a classifier-based policy on source switch B*

The configuration of remote-mirroring session 1 on Switch B is the same as on Switch A (figure B-55), except for the difference in source VLAN and source IP address. Note that on different switches, the UDP port number (9300) can be the same.

```
Switch-B(config)# mirror 1 remote ip 10.10.20.145 9300 10.10.30.2
Caution: Please configure destination switch first.
        Do you want to continue [y/n]? y
Switch-B(config)# class ipv4 tcp7
Switch-B(class-config)# match tcp any 10.10.30.153
Switch-B(class-config)# match tcp any host 10.10.20.153/24
Switch-B(class-config)# match tcp any any eq 80
Switch-B(class-config)# exit
Switch-B(config)# policy mirror mirrorTCP
Switch-B(policy-config)# class ipv4 tcp7 mirror 1
Switch-B(policy-config)# exit
Switch-B(config)# vlan 20 service-policy mirrorTCP in
```

Remote mirroring using traffic-direction criteria

In the network shown in **Figure 165: Sample topology for remote mirroring from a port interface** on page 519, the administrator connects another traffic analyzer to port B10 (in VLAN 40) on switch C to monitor all traffic entering switch A on port C12. For this mirroring configuration, the administrator configures a mirroring destination (with a remote exit port of B10) on switch C, and a remote mirroring session on switch A.

If the mirroring configuration in the proceeding example is enabled, it is necessary to use a different session number (2) and UDP port number (9400.) (The IP address of the remote exit port [10.10.40.7] connected to traffic analyzer 2 [exit port B10] can belong to a different VLAN than the destination IP address of the VLAN used to reach remote switch C [10.20.40.1]).

Figure 165: *Sample topology for remote mirroring from a port interface*

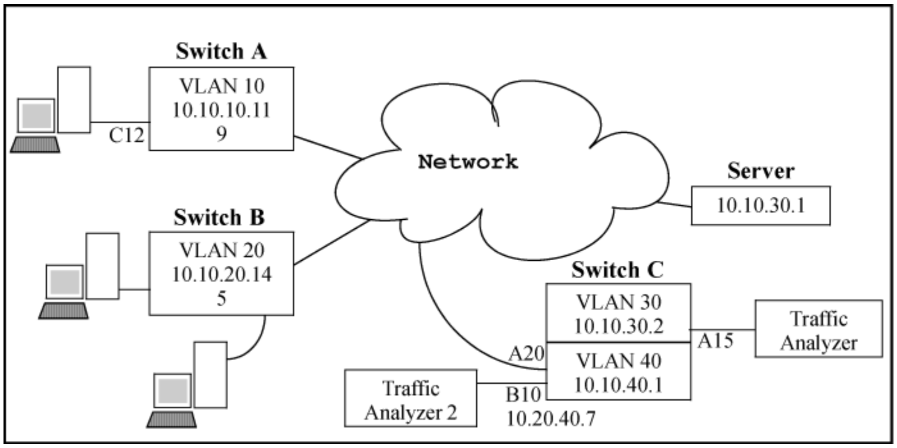

To configure this remote mirroring session using a directional-based traffic selection on a port interface, the operator must take the following steps:

1. On remote switch C, configure the remote mirroring endpoint using port B10 as the exit port for a traffic analyzer (as described in **Configure a mirroring destination on a remote switch** on page 507):

Figure 166: *Configuring a remote mirroring endpoint*

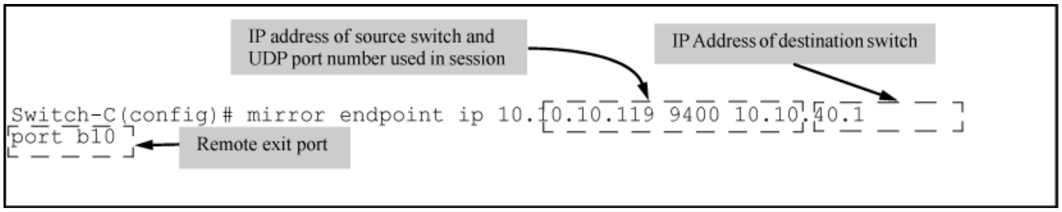

2. On source switch A, configure session 2 to use UDP port 9400 to reach the remote mirroring endpoint on switch C (10.10.40.1):`mirror 2 remote ip 10.10.10.119 9400 10.10.40.1`
3. On source switch A, configure the local port C12 to select all inbound traffic to send to the preconfigured mirroring destination for session 2:`interface c12 monitor all in mirror 2`

Figure 167: *Configuring a remote mirroring session for inbound port traffic*

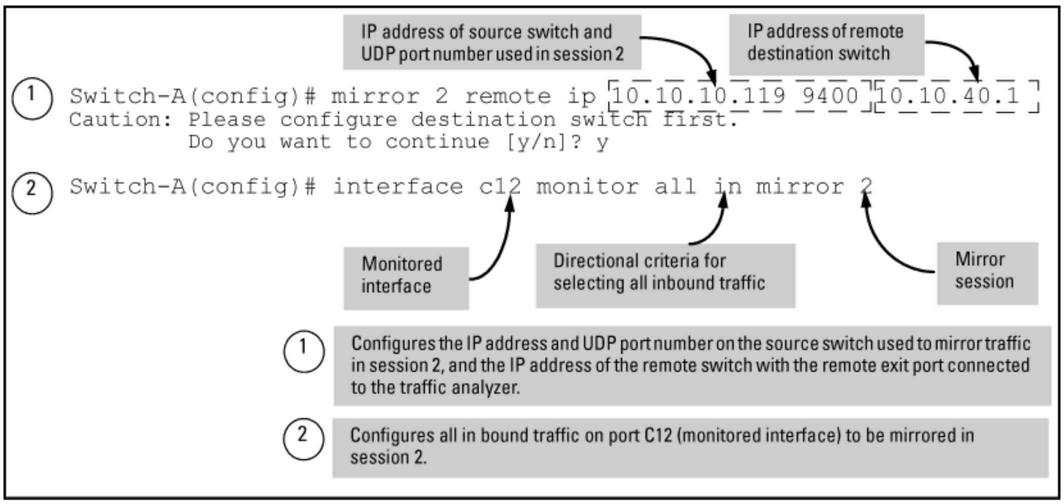

Maximum supported frame size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU allowed in the network, the frame is dropped or truncated.

NOTE

Oversized mirroring frames are dropped or truncated, according to the setting of the `[truncation]` parameter in the `mirror` command. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

If jumbo frames are enabled on the mirroring source switch, the mirroring destination switch and all downstream devices connecting the source switch to the mirroring destination must be configured to support jumbo frames.

Enabling jumbo frames to increase the mirroring path MTU

On 1-Gbps and 10-Gbps ports in the mirroring path, you can reduce the number of dropped frames by enabling jumbo frames on all intermediate switches and routers. (The MTU on the switches covered by this manual is 9220 bytes for frames having an 802.1Q VLAN tag, and 9216 bytes for untagged frames.)

Table 25: *Maximum frame sizes for mirroring*

	Frame type configuration	Maximum frame size	VLAN tag	Frame mirrored to local port	Frame mirrored to remote port	
				Data	Data	IPv4 header
Untagged	Non-jumbo (default config.)	1518	0	1518	1464	54
	Jumbo on all VLANs	9216	0	9216	9162	54
	Jumbo[11] On all but source VLAN	1518	0	n/a	1464	54
Tagged	Non-jumbo	1522	4	1522	1468	54
	Jumbo[11] on all VLANs	9220	4	9218	9164	54
	Jumbo[11] On all but source VLAN	1522	4	n/a[22]	1468	54

[1] Jumbo frames are allowed on ports operating at or above 1 Gbps

[2] For local mirroring, a non-jumbo configuration on the source VLAN dictates an MTU of 1518 bytes for untagged frames, and an MTU of 1522 for tagged frames, regardless of the jumbo configuration on any other VLANs on the switch.

Effect of downstream VLAN tagging on untagged, mirrored traffic

In a remote mirroring application, if mirrored traffic leaves the switch without 802.1Q VLAN tagging, but is forwarded through a downstream device that adds 802.1Q VLAN tags, the MTU for untagged mirrored frames leaving the source switch is reduced below the values shown in **Maximum frame sizes for mirroring**.

For example, if the MTU on the path to the destination is 1522 bytes, untagged mirrored frames leaving the source switch cannot exceed 1518 bytes. Likewise, if the MTU on the path to the destination is 9220 bytes, untagged mirrored frames leaving the source switch cannot exceed 9216 bytes.

Figure 168: *Effect of downstream VLAN tagging on the MTU for mirrored traffic*

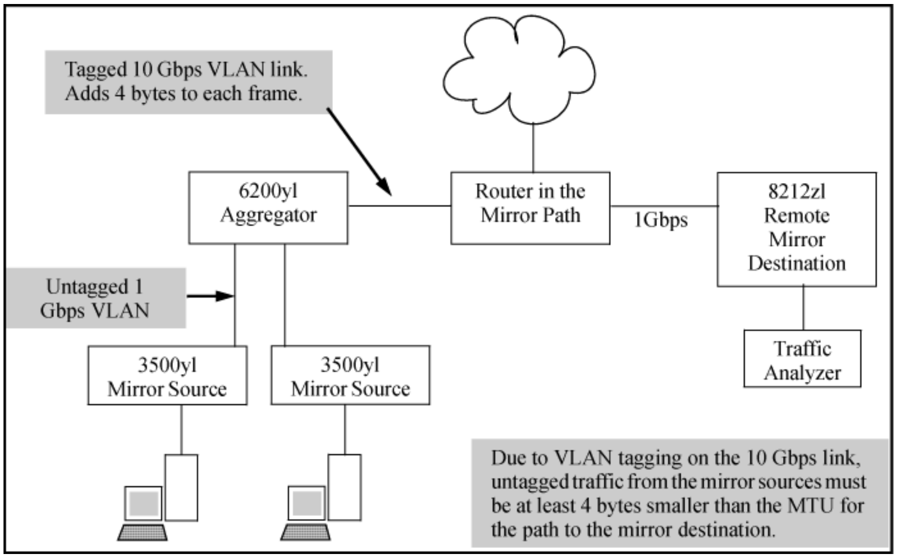

Operating notes for traffic mirroring

- Mirroring dropped traffic

  When an interface is configured to mirror traffic to a local or remote destination, packets are mirrored regardless of whether the traffic is dropped while on the interface. For example, if an ACL is configured on a VLAN with a `deny` ACE that eliminates packets from a Telnet application, the switch still mirrors the Telnet packets that are received on the interface and subsequently dropped.
- Mirroring and spanning tree

  Mirroring is performed regardless of the STP state of a port or trunk. This means, for example, that inbound traffic on a port blocked by STP can still be monitored for STP packets during the STP setup phase.
- Tagged and untagged frames

  For a frame entering or leaving the switch on a mirrored port, the mirrored copy retains the tagged or untagged state the original frame carried when it entered into or exited from the switch. (The tagged or untagged VLAN membership of ports in the path leading to the mirroring destination does not affect the tagged or untagged status of the mirrored copy itself.)

  Thus, if a tagged frame arrives on a mirrored port, the mirrored copy is also tagged, regardless of the status of ports in the destination path. If a frame exits from the switch on a mirrored port that is a tagged member of a VLAN, the mirrored copy is also tagged for the same reason.

  To prevent a VLAN tag from being added to the mirrored copy of an outbound packet sent to a mirroring destination, you must enter the `no-tag-added` parameter when you configure a port, trunk, or mesh interface to select mirrored traffic.
- Effect of IGMP on mirroring

  If both inbound and outbound mirroring is operating when IGMP is enabled on a VLAN, two copies of mirrored IGMP frames may appear at the mirroring destination.
- Mirrored traffic not encrypted

  Mirrored traffic undergoes IPv4 encapsulation, but mirrored encapsulated traffic is not encrypted.
- IPv4 header added

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the maximum MTU allowed in the network, it is dropped or truncated (according to the setting of the `[truncation]` parameter in the `mirror` command.)

To reduce the number of dropped frames, enable jumbo frames in the mirroring path, including all intermediate switches and/or routers. (The MTU on the switch is 9220 bytes, which includes 4 bytes for the 802.1Q VLAN tag.)

- Intercepted or injected traffic

  The mirroring feature does not protect against either mirrored traffic being intercepted or traffic being injected into a mirrored stream by an intermediate host.

- Inbound mirrored IPv4-encapsulated frames are not mirrored

  The switch does not mirror IPv4-encapsulated mirrored frames that it receives on an interface. This prevents duplicate mirrored frames in configurations where the port connecting the switch to the network path for a mirroring destination is also a port whose inbound or outbound traffic is being mirrored.

  For example, if traffic leaving the switch through ports B5, B6, and B7 is being mirrored through port B7 to a network analyzer, the mirrored frames from traffic on ports B5 and B6 will not be mirrored a second time as they pass through port B7.

- Switch operation as both destination and source

  A switch configured as a remote destination switch can also be configured to mirror traffic to one of its own ports (local mirroring) or to a destination on another switch (remote mirroring.)

- Monitor command note

  If session 1 is already configured with a destination, you can enter the `[no] vlan <VID>monitor` or `[no] interface <PORT> monitor` command without mirroring criteria and a mirror session number. In this case, the switch automatically configures or removes mirroring for inbound and outbound traffic from the specified VLAN or ports to the destination configured for session 1.

- Loss of connectivity suspends remote mirroring

  When a remote mirroring session is configured on a source switch, the switch sends an ARP request to the configured destination approximately every 60 seconds. If the source switch fails to receive the expected ARP response from the destination for the session, transmission of mirrored traffic in the session halts. However, because the source switch continues to send ARP requests for each configured remote session, link restoration or discovery of another path to the destination enables the source switch to resume transmitting the session's mirrored traffic after a successful ARP response cycle occurs.

  Note that if a link's connectivity is repeatedly interrupted ("link toggling"), little or no mirrored traffic may be allowed for sessions using that link. To verify the status of any mirroring session configured on the source switch, use the `show monitor` command.

Troubleshooting traffic mirroring

Cause

If mirrored traffic does not reach the configured remote destination (endpoint) switch or remote exit port, check the following configurations:

- In a remote mirroring session, the `mirror remote ip` command parameters configured on the source switch for source IP address, source UDP port, and destination IP address must be identical to the same parameters configured with the `mirror endpoint ip` command on the remote destination switch.
- The configured remote exit port must not be a member of a trunk or mesh.
- If the destination for mirrored traffic is on a different VLAN than the source, routing must be correctly configured along the path from the source to the destination.
- On the remote destination (endpoint) switch, the IP addresses of the remote exit port and the switch can belong to different VLANs.
- All links on the path from the source switch to the destination switch must be active.

CAUTION

A mirroring exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Configuring a mirroring exit port connection to a network can result in serious network performance problems, and is strongly discouraged by Switch Networking.

The Hewlett Packard Enterprise Virtual Technician is a set of tools aimed at aiding network switch administrators in diagnosing and caring for their networks. VT provides tools for switch diagnoses when faced with unforeseen issues.

To improve the Virtual Technician features of our devices, Hewlett Packard Enterprise has added the following tools:

- Cisco Discovery Protocol
- Enabling Debug tracing for MOCANA code
- User diagnostic crash via front panel security button
- User diagnostic crash via the serial console

Cisco Discovery Protocol (CDP)

show cdp traffic

Syntax

`show cdp traffic`

Description

Displays the number of Cisco Discovery Protocol (CDP) packets transmitted, received and dropped.

CDP frame Statistics

```
Port No|Transmitted Frames|Received Frames|Discarded Frames|Error Frames
-------  -----------------  --------------  ---------------  ---------
A1                46                  26              6            7
A2                30                  35              7            9
A3               120                 420            670          670
```

clear cdp counters

Syntax

`clear cdp counters`

Description

Allows a user to clear CDP statistics.

Clear cdp counters

```
Port No|Transmitted Frames|Received Frames|Discarded Frames|Error Frames
-------  -----------------  --------------  ---------------  ---------
A1           46                   26                 6               7
A2           30                   35                 7               9
A3          120                  420               670             670
```

Enable/Disable debug tracing for MOCANA code

debug security

Syntax

`debug security ssl`

Description

Enables the debug tracing for MOCANA code.

Use the [no] parameter to disable debug tracing.

ssl

Display all SSL messages.

User diagnostic crash via Front Panel Security (FPS) button

Allows the switch's front panel **Clear** button to manually initiate a diagnostic reset. In the case of an application hang, this feature allows you to perform reliable diagnostics by debugging via the front panel **Clear** button. Diagnostic reset is controlled via Front Panel Security (FPS) options.

front-panel-security password-clear

Syntax

`[no] front-panel-security [password-clear <RESET-ON-CLEAR>| factory-reset | password-recovery | diagnostic-reset <CLEAR-BUTTON | SERIAL-CONSOLE>]`

Description

Enables the ability to clear the passwords and/or configuration via the front panel buttons.

Parameters and options

no

Disables the password clear option.

password-clear

When disabled, you cannot reset the passwords using the clear button on the front panel of the device.

factory-reset

When disabled, you cannot reset the configuration/password using the clear and reset button combination at boot time.

password-recovery

When enabled (and the front panel buttons disabled), contact Hewlett Packard Enterprise customer support to recover a lost password. When disabled, there is no way to access a device after losing a password with the front panel buttons disabled.

diagnostic-reset

When disabled, the user cannot perform a diagnostic switch reset on those rare events where the switch becomes unresponsive to user input because of unknown reasons. When enabled, the user can perform a diagnostic hard reset which will capture valuable diagnostic data and reset the switch.

factory-reset

Enable/Disable factory-reset ability.

password-clear

Enable/Disable password clear.

password-recovery

Enable/Disable password recovery.

diagnostic-reset

Enable/Disable diagnostic reset.

front-panel-security diagnostic-reset

Syntax

`[no] front-panel-security diagnostic-reset <CLEAR-BUTTON | SERIAL-CONSOLE>`

Description

Enables the diagnostic reset so that the switch can capture diagnostic data.

- To initiate diagnostic reset via the clear button, press the clear button for at least 30 seconds but not more than 40 seconds.
- To initiate diagnostic switch reset via the serial console, enter the diagnostic reset sequence on the serial console.

Parameters and options

no

Disables the diagnostic reset feature so that the user is prevented from capturing diagnostic data and performing a diagnostic reset on the switch. Disables both serial console and the clear button. This is necessary if the switch becomes unresponsive (hangs) for unknown reasons.

<CLEAR BUTTON>

Enables diagnostic-reset using the clear button, allowing the user to perform diagnostic reset by pressing the clear button for 30 seconds and not more than 40 seconds.

<SERIAL CONSOLE>

Enables the diagnostics by choosing the serial console option.

Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

Front-panel-security diagnostic-rest clear-button

```
front-panel-security diagnostic-rest clear-button

Diagnostic Reset      - Enabled
clear-button         - Enabled
serial-console       -Disabled
```

No front-panel-security diagnostic-reset

```
no front-panel-security diagnostic-reset

Clear Password       - Enabled
Reset-on-clear     - Disabled
Factory Reset        - Enabled
Password Recovery    - Enabled
Diagnostic Reset     - Disabled
```

show front-panel-security

Syntax

`show front-panel-security`

Description

User initiated diagnostic reset defaults to enabled.

Show front-panel-security

```
Clear Password    - Enabled
Reset -on-clear   - Disabled
Factory Reset     - Enabled
Password Recovery - Enabled
Diagnostic Reset  - Enabled
```

Diagnostic table

To accomplish this	Do this	Result
Soft Reset (Standalone switch)	Press and release the Reset button	The switch operating system is cleared gracefully (such as data transfer completion, temporary error conditions are cleared), then reboots and runs self tests.
Hard Reset (Standalone switch)	Press and hold the Reset button for more than 5 seconds (until all LEDs turn on), then release.	The switch reboots, similar to a power cycle. A hard reset is used, for example, when the switch CPU is in an unknown state or not responding.
Soft Reset (Stacked switch)	Press and release the Reset button	Same as a standalone switch, except: • If the Commander, the Standby switch will become Commander. • If the Standby, a new Standby will be elected.

Table Continued

To accomplish this	Do this	Result
Hard Reset (Stacked switch)	Press and hold the Reset button for more than 5 seconds (until all LEDs turn on), then release.	Same as a standalone switch, except: • If the Commander, the Standby switch will become Commander. • If the Standby, a new Standby will be elected.
Delete console and management access passwords	Press Clear for at least one second, but not longer than 5 seconds.	The switch deletes all access password.
Restore the factory default configuration	**1.** Press Clear and Reset simultaneously. **2.** While continuing to press Clear, release Reset. **3.** When the Test LED begins blinking (after approximately 25 seconds), release Clear.	The switch removes all configuration changes, restores the factory default configuration, and runs self test.
Diagnostic reset	**1.** Press Clear to 30–40 seconds. **2.** When the test LED begins blinking (approximately after 30 seconds), release Clear. NOTE: Releasing the Clear button when TEST LED is not blinking (approximately after 40 seconds) will not honor the diagnostic reset request.	This initiates diagnostic reset, collects diagnostic information, and reboots the switch.

NOTE: These buttons are provided for the user’s convenience. If switch security is a concern, ensure that the switch is installed in a secure location, such as a locked writing closet. To disable the buttons, use the `front-panel-security` command.

Validation rules

Cause

Validation	Error
Extra ‘token’ passed after diagnostic-reset.	Invalid input: <token>.

FPS Error Log

Cause

Event	Message
RMON_BOOT_CRASH_RECORD1	Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. NOTE: On detection on local serial
RMON_BOOT_CRASH_RECORD1	SMM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. NOTE: On detection on SMM serial console and signaled to AMM
RMON_BOOT_CRASH_RECORD1	STKM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. NOTE: On detection on non-commander serial console and signaled to commander
RMON_BOOT_CRASH_RECORD1	User has initiated diagnostic reset via the serial console. NOTE: Sw_panic() message
RMON_BOOT_CRASH_RECORD1	SMM: User has initiated diagnostic reset via the serial console. NOTE: Sw_panic() message when triggered via SMM
RMON_BOOT_CRASH_RECORD1	STKM: User has initiated diagnostic reset via the serial console. NOTE: Sw_panic() message when triggered via non-commander

Table Continued

Event	Message
Console print	STKM: HA Sync in progress; user initiated diagnostic request via the serial console rejected. Retry after sometime. NOTE: Printed on the device console. When standby is in sync state,we don't want to crash the commander. So we report to the user to retry later
Console print	STKM: Member is booting; user initiated diagnostic request via the serial console rejected. Retry after sometime. NOTE: Printed on the device console. When the member is till booting, it doesn't have the commander member number, thus we can't issue UIDC on the commander. So we report to the user to retry later.

User initiated diagnostic crash via the serial console

Remotely triggers a diagnostic reset of the switch via a serial console. This reset reboots the switch and collects diagnostic data for debugging an application hang, a system hang or any other rare occurrence. Diagnostic reset is controlled via FPS options.

The serial sequence to initiate the User Initiated Diagnostic Reset via Serial console is Ctrl+S, Ctrl+T, Ctrl+Q, Ctrl +T, Ctrl+S.

front-panel-security diagnostic-reset serial-console

Syntax

```
[no] front-panel-security diagnostic-reset serial-console
```

Description

Enables the diagnostic-reset via serial console. Allows the user to perform diagnostic reset by keying-in diagnostic reset sequence.

Parameters and options

no

Disables the diagnostic-reset via serial console.

Front-panel-security diagnostic-reset serial-console

```
front-panel-security diagnostic-reset serial-console

Diagnostic Reset      - Enabled
clear-button        - Disabled
serial-console      - Enabled
```

No front-panel-security diagnostic-reset serial-console

```
no front-panel-security diagnostic-reset serial-console

Diagnostic Reset      - Disabled
```

Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

Serial console error messages

Error	Message
RMON_BOOT_CRASH_RECORD1	Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	SMM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	STKM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	User has initiated diagnostic reset via the serial console.
RMON_BOOT_CRASH_RECORD1	SMM: User has initiated diagnostic reset via the serial console.
RMON_BOOT_CRASH_RECORD1	STKM: User has initiated diagnostic reset via the serial console.
Console print	STKM: HA Sync in progress; user initiated diagnostic request via the serial console rejected. Retry after sometime.
Console print	STKM: Member is booting; user initiated diagnostic request via the serial console rejected. Retry after sometime.

The following table lists the switch scalability values for the areas of VLANs, ACLs, hardware, ARP, and routing.

Subject	Maximum
IPv4 ACLs	
total named (extended or standard)	Up to 2048 (minus any IPv4 numeric standard or extended ACL assignments and any RADIUS-assigned ACLs)[1]
total numbered standard	Up to 99
total numbered extended	Up to 100[1]
total ACEs in all IPv4 ACLs	Up to 3072[1]
IPv6 ACLs	
total IPv6 ACLs	Up to 2048[1]
total ACEs in all IPv6 ACLs	Up to 3072[1]
Layer-3	
VLANs with at least one IP Address	512
IP addresses per system	2048 IPv42048 IPv6
IP addresses per VLAN	32
Static routes (IPv4 and IPv6 combined)	256
IPv4 host hardware table	72 K (8K internal, 64K external)
IPv4 BMP hardware table	2 K
ARP	
ARP entries	25,000
Packets held for ARP resolution	25
Dynamic Routing	
Total routes supported	IPv4 only: 10,000 (including ARP)IPv4 and IPv6: 10 K (IPv4) and 3 K(IPv6)IPv6 only: 5 K
IPv4 Routing Protocol	

Table Continued

Subject	Maximum
RIP interfaces	128
OSPFv2	
Interfaces/subnets	512 (128 active)
Max. areas supported	16
ECMP next hops	4
IPv6 Routing Protocol	
DHCPv6 Helper Addresses	32 unique addresses; multiple instances of same address counts as 1 towards maximum
OSPFv3	
Interfaces/subnets	512 (128 active)
Max. areas supported	16
ECMP next hops	4

[1] Actual availability depends on combined resource usage on the switch.

Supported Platforms

Aruba 3810M Switch Series (JL071A, JL072A, JL073A, JL074A, JL075A, JL076A)
Aruba 5400Rzl2 Switch Series (J8698A, J8700A, J9823A-J9824A, J9825A, J9826A, J9868A, J9447A, J9448A)
Aruba 5406R Switch Series (JL002A, JL003A, JL095A,J9850A)
Aruba 5406zl Switch Series (J9821A, J9822A))
Aruba 5412R Switch Series (J9851A, JL001A)
HPE 3800 Switch Series (J9573A—J9576A, J9584A—J9588A)

Job Scheduler

The Job Scheduler feature enables the user to schedule commands or jobs on the switch for one time or multiple times. This is similar in concept to the UNIX ‘cron’ utility. The user can schedule any CLI command that the user would otherwise enter interactively. This includes commands to enable or disable ports, LEDs, and Power-Over-Ethernet. Jobs can also be scheduled to be triggered by certain pre-defined events such as switch reboot. The only major restriction on commands scheduled is that, it should not prompt/ask for any user inputs.

Commands

Job at | delay | enable | disable

Set schedule jobs using the options and set the count for the number of times the job is repeated.

Syntax

```
job <JOB NAME> at | delay | enable | disable
```

Description

Schedule a command to run automatically. Jobs can be scheduled to run once, multiple times on a recurring basis, or after certain events such as reboots. All commands run with manager privilege in configuration context.

The [no] form of the command deletes a scheduled job.

By default, jobs will be repeated an infinite number of times.

Restrictions

Jobs scheduled at any event will not be counted.

Jobs that are scheduled at the event “reboot” will not work in some multi management switches.

Range

- <1-1000>: is the value range for the `count` option.
- ([[DD:]HH:]MM): is the format used for the specific delay.

Options

count

Specify the number of times the job should run.

delay

Specify the delay before running the job.

enable

Enable a job that is disabled or expired.

disable

Disable a job. By default, a job is enabled.

Usage

```
job <JOB NAME> at <([DD:]HH:]MM on <WEEKDAY-LIST>)> config-save <COMMAND> count
<1-1000>
```

```
job <JOB NAME> at <[HH:]MM on [MM/]DD> config-save <COMMAND> count <1-1000>
```

```
job <JOB NAME> at <EVENT> config-save <COMMAND>
```

```
job <JOB NAME> delay <([DD:]HH:]MM> config-save <COMMAND> count <1-1000>
```

```
job <JOB NAME> enable | disable
```

```
[no] job <JOB NAME>
```

Show job

Syntax

```
show job
```

Description

Show the jobs scheduled.

Show job

```
switch# show job

 Job Scheduler Status and Configuration

  Scheduler Status : Waiting for the system time to be set

                     Event or            Repeat  Save
  Name               Time                Count   Cfg  Command
  ------------------ ---------------- -------- ---- ------------
  Burrrrrrrrrrr...   reboot              --    Yes chassislocate blink
  baz                reboot              --     No  show time
  foo                17:00 SxTWTxS       --     No  savepower led
  a1                 12:00               2      Yes sh time
  a2                 Every 2:14:30 days  75     Yes vlan 3
  a3                 Every 00:00:25 days 1      No  vlan 4
```

Caution

The scheduler does not run until the system time is set.

Show job <Name>

Syntax

`show job <JOB NAME>`

Description

Show the job by name.

Show job <JOB NAME>

```
Aruba-3810M-16SFPP-2s # show job a1

Job Information

Job Name     : a1
Runs At      : 01:24
Config Save : No
Repeat Count: --
Job Status   : Enabled
Run Count    : 1
Error Count : 0
Command      : show time
Job Status   : Enabled

 Output from Last Run
 --------------------
 Tue Dec 15 01:24:00 2015


HP-2530-24 # show job a2

Job Information

Job Name     : a2
Runs At      : Every 2:14:30 days
Config Save : Yes
Repeat Count: 75
Run Count    : 0
Error Count : 0
Command      : vlan 3
Job Status   : Disabled

HP-2530-24 # show job foo


Job Information

Job Name     : foo
Runs At      : 17:00 SxTWTxS
Config Save : Yes
Repeat Count: --
Run Count    : 0
Error Count : 0
```

```
Command      : savepower led
Job Status   : Enabled
```

Show job <JOB NAME>

```
Aruba-3810M-16SFPP-2s # show job a1

Job Information

Job Name       : a1
Runs At        : 01:24
Config Save    : No
Repeat Count   : --
Job Status     : Enabled
Running Status: Active
Run Count      : 1
Error Count    : 0
Command        : show time
Last Run Time : Tue Mar  4 12:00:00 2003
Job Status     : Enabled

 Output from Last Run
 --------------------
 Tue Dec 15 01:24:00 2015


HP-2530-24 # show job a2

Job Information

Job Name       : a2
Runs At        : Every 2:14:30 days
Config Save    : Yes
Repeat Count   : 75
Job Status     : Enabled
Running Status: Expired
Run Count      : 0
Error Count    : 0
Command        : vlan 3
Last Run Time : Thu Nov 26 12:00:00 2015
Job Status     : Disabled

HP-2530-24 # show job foo


Job Information

Job Name     : foo
Runs At      : 17:00 SxTWTxS
Config Save : Yes
Repeat Count: --
Run Count    : 0
Error Count : 0
Command      : savepower led
Job Status   : Enabled
```

To create a profile and associate a device-type to that profile, the specified device-name must first be defined in the device-identity. Device-identity uses discovery protocols like LLDP to identify the device.

Procedure for creating a device identity and associating a device type

Procedure

1. Create a device identity using the command:

```
switch# device-identity name <DEVICE-NAME>
```

2. Specify the OUI used in LLDP's organization using specific TLV, (type =127). OUI should be in XXXXXX format. The default OUI "000000" indicates that device-identity will not use LLDP to identify device:

```
switch(config)# device-identity name <DEVICE-NAME> lldp oui <MAC_OUI>
 sub-type <SUBTYPE>
```

To add new device on switch:

```
switch(config)# device-identity name abc lldp oui a1b2c3 sub 2
```

To remove device from switch:

```
switch(config)# no device-identity name abc
```

3. Show device identity configuration:

```
switch(config)# show device-identity lldp

Device Identity Configuration

 Index  Device name             Oui        Subtype
 ------ ----------------------- ---------- -------
 1      abc                     a1b2c3     2
```

LLDP device profile detection dynamically uses organization-specific TLV to detect and apply profiles to devices. Organizational Unique Identifiers (OUI) and subtypes are detected based on the configuration of the switch. A maximum of 16 devices can be detected and defined using LLDP.

Requirements

- The device-identity must be configured with a name.

device-profile

Syntax

`device-profile <PROFILE-NAME>`

`no device-profile <PROFILE-NAME>`

Description

Create and name a new device profile.

Command context

`config`

Parameters

<PROFILE-NAME>

Specifies the name assigned to the profile.

Example

When adding a new device-identify name to a device, the new identity will show in the command context.

```
switch-Stack# config
switch-Stack(config)# device-identity name test
switch-stack(identity-test)#
```

device-profile type-device

Syntax

`device-profile type [aruba-ap | aruba-switch | scs-wan-cpe | device-name <DEVICE-NAME> associate <PROFILE-NAME> | enable | disable]`

`no device-profile type [aruba-ap | aruba-switch | scs-wan-cpe | device-name <DEVICE-NAME> associate <PROFILE-NAME> | enable | disable]`

Description

Associates the device profile with the type of device by identity.

The `no` form of this command removes the device profile from the device type.

Command context

`config`

Parameters

associate *<PROFILE-NAME>*

Selects the profile name associated with the device-type.

enable

Selects the profile of the device being enabled.

disable

Selects the profile of the device being disabled.

Usage

- The command `device-profile type aruba-ap enable` enables an Aruba-AP within the device profile.
- Device Name is defined the same as Device Identity.

device-profile device-type enable

Syntax

```
device-profile device-type <DEVICE-NAME> enable

device-profile device-type <DEVICE-NAME> disable
```

Description

Enables the newly created device profile and associates it with the device-type and device-name.

The `disable` form of this command removes the association.

Command context

`config`

Parameters

<DEVICE-NAME>

Specifies the device name.

NOTE

The device identification may be used in other subsystems (like device profile).

Examples

Enable the device d1 using the command `device-profile device-type <d1> enable`.

```
switch(config)# device-profile device-type d1 enable
```

Disable the device d1 using the command `device-profile device-type d1 disable`.

```
switch (config)# device-profile device-type d1 disable
```

Associating a profile with a device

Associate a profile with a device by using the command `device-profile device-type` *<DEVICE-NAME>* `associate` *<PROFILE-NAME>*. Associated devices can be Aruba Access Points, ArubaOS-Switch Switches, scs-wan-cpe, or association can be by the device profile.

The feature is disabled by default.

device-profile device-type associate

Syntax

`device-profile device-type` *<DEVICE-NAME>* `associate` *<PROFILE-NAME>*

`no device-profile device-type` *<DEVICE-NAME>* `associate` *<PROFILE-NAME>*

Description

Associate the profile to newly added device .

The `no` form of this command removes the profile association.

Usage

- The device identification may be use in other subsystems like device profile.

Command context

`config`

Parameters

<DEVICE-NAME>

Specifies the name of the device.

<PROFILE-NAME>

Specifies the name given to the profile.

Examples

Use the command to associate the device `d1` with the profile `abc`.

```
switch(config)# device-profile device-type d1 associate abc
```

show device-profile status

Syntax

`show device-profile status`

Description

Shows the profile status of the device.

Command context

`config`

Example

Use the `show decive-profile status` command to view status.

```
switch(config)# show device-profile status
Port     Device-type          Applied device profile
 -------- -------------------- ----------------------
 A1         <device-name>           abc
```

show device-profile config

Syntax

```
show device-profile config
```

Description

Shows the device profile configuration.

Command context

`config`

Examples

Use the command `show device-profile config` to display the device profile configuration.

```
switch(config)# show device-profile config

Device Profile Configuration
 Configuration for device-profile : default-ap-profile
  untagged-vlan     : 1
  tagged-vlan       : None
  ingress-bandwidth : 100%
  egress-bandwidth  : 100%
  cos               : 0
  speed-duplex      : auto
  poe-max-power     : Class/LLDP
  poe-priority      : critical
  allow-jumbo-frames: Disabled
 Configuration for device-profile : abc
  untagged-vlan     : 1
  tagged-vlan       : None
  ingress-bandwidth : 100%
  egress-bandwidth  : 100%
  cos               : None
  speed-duplex      : auto
  poe-max-power     : Class/LLDP
  poe-priority      : critical
  allow-jumbo-frames: Disabled
 Configuration for device-profile : default-aos-profile
  untagged-vlan     : 1
  tagged-vlan       : None
  ingress-bandwidth : 100%
  egress-bandwidth  : 100%
  cos               : 0
  speed-duplex      : auto
  poe-max-power     : Class/LLDP
  poe-priority      : critical
  allow-jumbo-frames: Disabled
 Configuration for device-profile : default-scs-profile
  untagged-vlan     : 1
  tagged-vlan       : None
```

```
  ingress-bandwidth : 100%
  egress-bandwidth  : 100%
  cos               : 0
  speed-duplex      : auto
  poe-max-power     : Class/LLDP
  poe-priority      : critical
  allow-jumbo-frames: Disabled
  allow-jumbo-frames: Disabled
  Configuration for device-profile : default-device-profile
  untagged-vlan     : 1
  tagged-vlan       : None
  ingress-bandwidth : 100%
  egress-bandwidth  : 100%
  cos               : 0
  speed-duplex      : auto
  poe-max-power     : Class/LLDP
  poe-priority      : critical
  allow-jumbo-frames: Disabled

  Device Profile Association

  Device Type   : aruba-ap
  Profile Name  : default-ap-profile

  Device Status : Disabled
  Device Type   : aruba-switch
  Profile Name  : default-aos-profile
  Device Status : Disabled

  Device Type   : scs-wan-cpe
  Profile Name  : abc
  Device Status : Disabled

  Device Type   : <lldp-device-name>
  Profile Name  : default-device-profile
  Device Status : Disabled
```

show device-identity

Syntax

`show device-identity <LLDP>`

Description

Show the device-identity configuration.

Command context

`config`

Parameters

<LLDP>

Shows the configuration of the LLDP device.

Examples

Shows the device identity configuration.

```
switch(config)# show device-identity lldp
```

```
Device Identity Configuration

 Index   Device name          Oui        Subtype
 ------ --------------------- ---------- -------
 1      cpe                   33 bb cc   0
 2      phone                 11 22 33   2


switch(config)# show running configuration

Running configuration:

; J9625A Configuration Editor; Created on release #RA.16.04.0000x
; Ver #0f:02.43.18.82.34.61.1c.28.f3.84.9c.63.ff.37.2f:da
hostname "HP-2620-24-PoEP"
snmp-server community "public" unrestricted
device-identity name "cpe" lldp oui 33bbcc
device-identity name "cpe" lldp sub-type 1
device-identity name "phone" lldp oui 112233

vlan 1
   name "DEFAULT_VLAN"
   untagged 1-28
   ip address dhcp-bootp
   exit
device-profile name "ram"
   exit
device-profile type “scs-wan-cpe”
   associate "ram"
   enable
   exit
device-profile type-device "cpe"
   associate "ram"
   enable
   exit
device-profile type-device "phone"
   associate "default-device-profile"
   exit
```

Supported devices

Code	Switch
KB	Aruba 5400R Switch Series

NOTE

VSF is supported only on V3 blades. Once VSF is enabled, the switch will reboot in V3-only mode.

Overview of VSF

HPE-Aruba Virtual Switching Framework (VSF) technology virtualizes two physical devices in the same layer into one Virtual Fabric which provides high availability and scalability. VSF allows supported switches connected to each other through normal ethernet connections (copper or fiber) to behave like a single switch.

Figure 169: *Two devices using VSF technology appearing as a single node to the upper-layer and lower-layer devices*

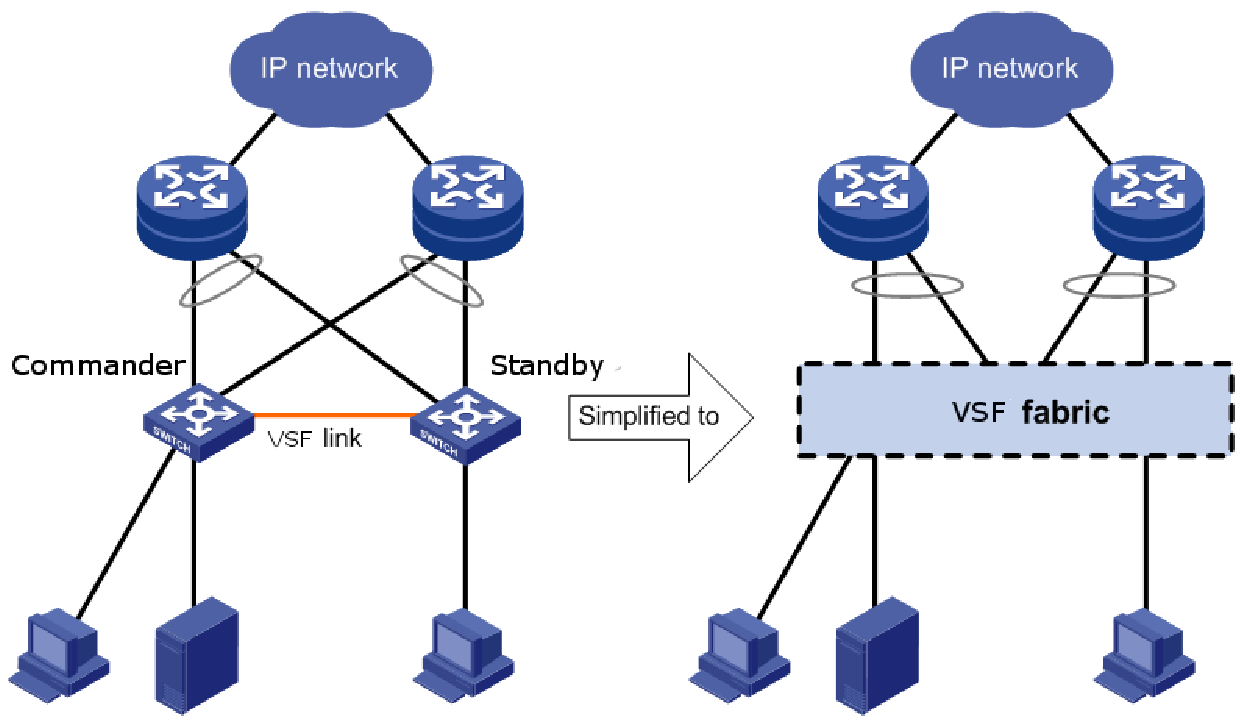

Benefits of VSF

Simplified topology and ease of management

A VSF fabric appears and behaves as one logical switch and is accessible by the network through a single IP address.

NOTE

Spanning tree features are not necessary among VSF members.

1:1 redundancy

One member acts as the Commander to manage and control the entire VSF fabric. The other switch acts as a Standby and backs up the commander, and takes over in case of commander failure.

VSF port aggregation

A VSF link can aggregate upto eight VSF ports with immediate neighboring member. This provides redundancy till failure of seven VSF ports.

Distributed trunking

The Ethernet link aggregation feature can be used to aggregate physical links between the VSF and its upstream or downstream devices across the VSF members. This eliminates the need for spanning tree and also provides load balancing across across all ports of the link aggregate.

Network scalability

The processing power is equal to the Commander. The forwarding capacity is equal to both the Commander and the Standby combined.

Member roles

VSF uses two member roles: Commander and Standby.

NOTE

All switches in the VSF stack will have the same software version. During stack formation, switches which do not have the same software version as the commander, will be updated to the commander's software. This will cause a reboot of the updated switch.

Commander

Control and management plane protocols run on the Commander, which is responsible for managing the forwarding databases, synchronizing them with the Standby and controlling all line cards including that of the Standby.

Standby

Standby is a stateful backup device for the Commander and is ready to take control of the VSF virtual chassis if the Commander device fails. This enables the VSF virtual chassis to continue its operations seamlessly in the event of a failure.

Management module

The switch has two management module (MM) slots available. It is recommended that you have only one MM for each Aruba 5400R switch when VSF is enabled. A second MM, if present, will be shutdown, and should be removed from the chassis to prevent it from becoming active inadvertently.

VSF member ID

VSF uses member IDs to uniquely identify and manage its members. The first part of the interface module number is the Member ID information, which identifies interfaces in a VSF fabric.

The device that wins election and becomes Commander will keep its member ID while the other will automatically be assigned a different unassigned member ID from the pool and reboot.

NOTE

If the VSF member ID changes when joining a VSF virtual chassis it will cause a reboot of that member not the whole VSF virtual chassis.

VSF link

A VSF link is a logical interface that connects VSF member devices. The VSF link is named `I-Link<Member ID>_1`. To configure a VSF link, bind a minimum of one physical interface to it. The bound physical interfaces automatically aggregate to form a VSF link. A VSF link goes down only if all its VSF physical interfaces are down.

`I-Link<Member ID>_1` is the default name.

vsf member link

Syntax

```
[no] vsf member <MEMBER-ID> link <LINK-ID> [[ethernet] <PORT-LIST> | name <LINK-NAME>]
```

Description

Create the VSF links. A set of physical ports between any 2 members, carrying VSF traffic, is collectively referred to as a VSF link.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<LINK-ID>

VSF link ID value. For HPE-Aruba 5400R devices, allowed Link ID value is 1.

<PORT-LIST>

The port number or a list of ports. Upto 8 ports can be assigned into a VSF link.

<LINK-NAME>

The VSF link name. Upto 10 characters are allowed for link name.

Operating Notes

- A VSF link is a logical port dedicated to the internal connection of an VSF virtual device.
- A VSF link is effective only after it is bound to a physical port.
- When an Ethernet port is bound to a VSF link, it carries VSF data traffic and VSF protocol packets.

Validation rules for VSF member

Validation	Error/Warning/Prompt
When trunk static/manual and mesh is getting configured as VSF port	• Cannot configure VSF on port “A1” because that port is an LACP trunk. • Cannot configure VSF on port “A1” because that port is a Mesh. • Cannot configure VSF on port “A1” because that port is a Distributed LACP trunk. • Cannot configure VSF on port “A1” because that port is a Distributed trunk. • Cannot configure VSF on port “A1” because that port is a Dynamic trunk. • Cannot configure VSF on port “A1” because that port is an InterSwitch Connect (ISC) portError configuring VSF on port “A1”: An unsupported trunking mode is already configured on this port.
Adding a 1G port to a VSF link	Cannot enable VSF on a port operating at other than 10G or 40G.
Adding both 10G and 40G ports to a VSF link	Port 1/A1 does not support the current VSF port speed of 10G and may not be activated. Use compatible transceiver or reconfigure port speed to utilize this port. All configuration on this port has been removed and port is placed in VSF mode.
Max 8 ports per link.	Cannot configure more than 8 physical ports as an VSF link.
For other than physical ports.	VSF capabilities are not supported on port “A1”.
Cannot set a link name which is having more than 10 characters.	Cannot configure the VSF link name. The name is not a valid UI display string, or is blank, or exceeds 10 characters.
Direct VSF port removal case	Cannot remove VSF link when it has physical ports associated with it. First remove the associated physical ports and then remove the VSF link.
Removing the last VSF port in a VSF link that is "Up" is forbidden.	Removing of binding between physical ports and VSF link is not allowed since it would result in a stack split.
Using a port reserved for internal use as a VSF port.	Cannot use stolen/reserved ports as VSF ports.

Physical VSF ports

VSF ports connect VSF member devices and must be assembled using a VSF link. These VSF ports forward VSF protocol packets and data traffic.

VSF domain ID

VSF uses VSF domain IDs to uniquely identify VSF fabrics and prevent VSF fabrics from interfering with one another. One VSF fabric forms one VSF domain.

Figure 170: *Two VSF domains*

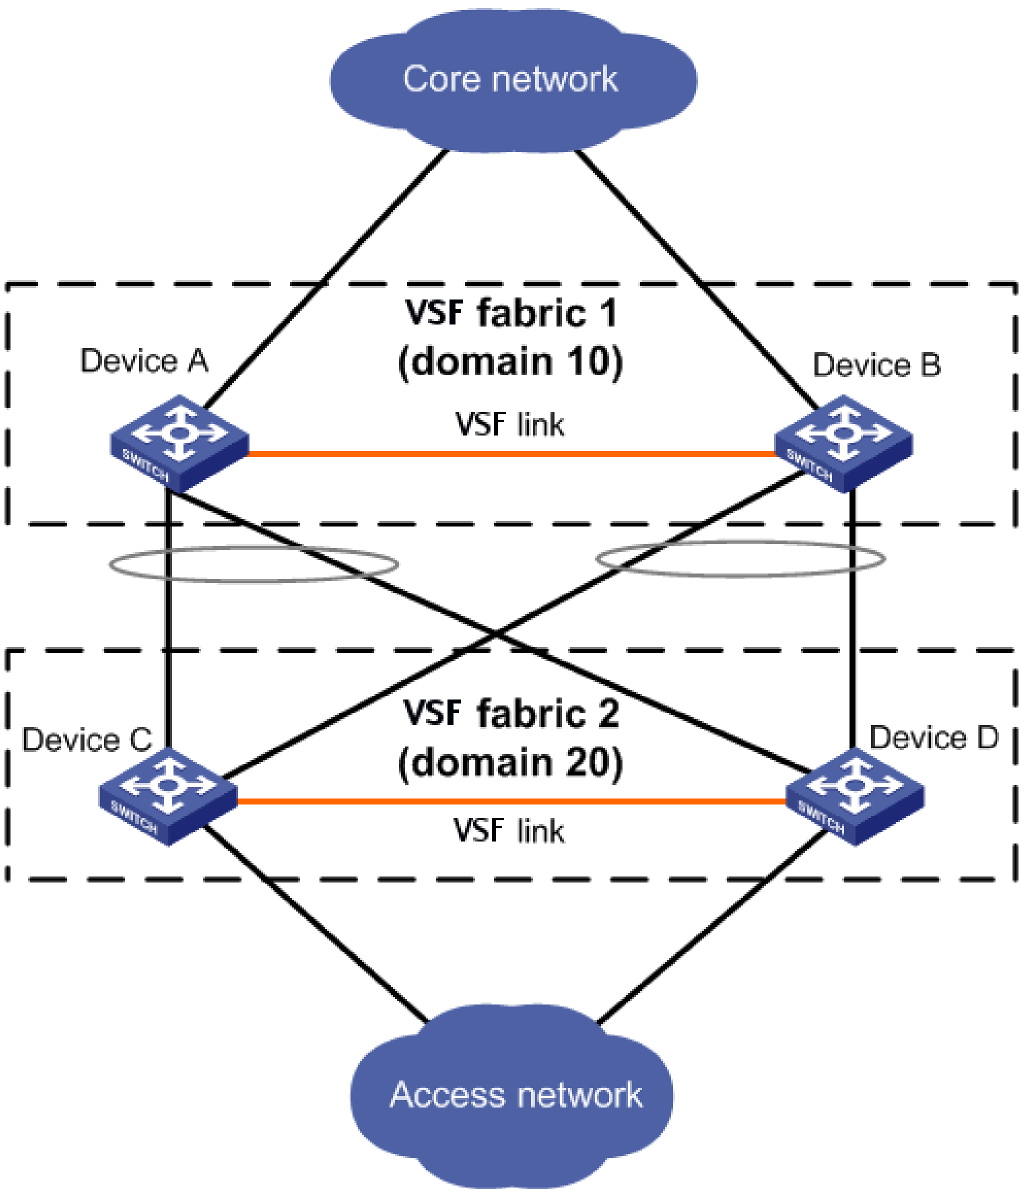

VSF split

A VSF split can occur due to a VSF link failure where all ports in the VSF link go down. This failure results in independent VSF fabric fragments each having its own Commander role. HPE-Aruba recommends configuring a Multiple Active Detection (MAD) mechanism to avoid duplicate IP addresses, routing issues, and traffic forwarding problems when a VSF split occurs.

Figure 171: *VSF split*

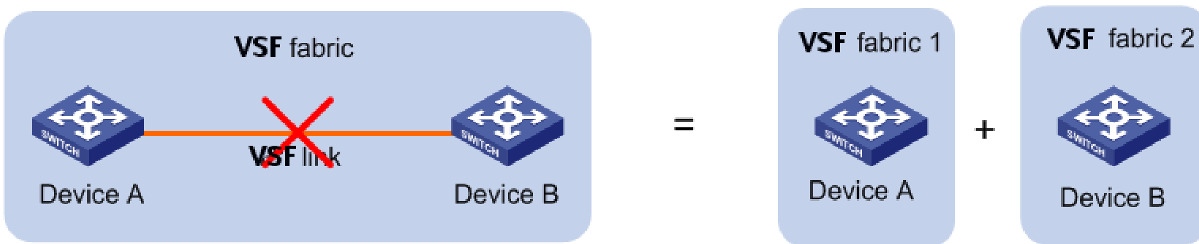

VSF merge

VSF merge happens on connecting two different VSF fragments having same domain-id via VSF links. One of the VSF fragment will go for reboot to merge with the other fragment and forms one VSF fabric. VSF fabric can only be formed among same type of devices. For example Aruba 5406R cannot form a VSF with Aruba 5412R.

Figure 172: *VSF Merge*

Member priority

Member priority determines the possibility of a member device to get elected as the Commander. A member with higher priority is more likely to be elected as the Commander. The default priority is 128, and ranges from 1 to 255.

NOTE

Changing the priority does not affect the commander immediately. It is only used at the next election, which will be at the next stack reboot

Interface naming conventions

An interface is named in the following format:

Interface name

<MEMBER-ID>/<interface-module><port-index>

Example

1/A1, 2/L24

Parameters

<MEMBER-ID>

VSF member ID of the switch. The VSF member ID always takes effect, whether or not the device has formed a VSF fabric with other devices. If the device is alone, the device is considered to be a standalone VSF fabric.

This argument defaults to 1.

<interface-module >

Slot letter of the front panel. Letter can be A-F for Aruba 5406R switch and A-L for Aruba 5412R switch.

<port-index >

Index of the port on the device. Port index depends on the number of ports available on the linecard (or Interface Module).

Interface name

On VSF, an interface name would take this form:

```
<member ID>/<interface-module><port-index>
```

```
2/B4
```

where 2 is the member ID and B4 is the port-index.

Running-configuration synchronization

VSF uses a strict running-configuration synchronization mechanism. In a VSF fabric, a device obtains and runs the configuration of the Commander. Commander manages and retains the configuration of all the devices.

VSF deployment methods

There are two ways to configure a VSF: Discovered Configuration Mode and Provisioned Configuration Mode.

Discovered configuration mode procedure

The following procedure configures devices into VSF members:

1. Configure VSF memberID and ports on one switch and enable VSF on that switch
2. After reboot, the device will come up as standalone VSF member. Connect a new device (with the factory default configuration) to this VSF device. The commander should be connected to the VSF ports of the first device.
3. The new device will reboot and join as standby.

Provisioned configuration mode procedure

The following procedure configures devices into VSF members:

Procedure

1. Configure VSF memberID and VSF links on one switch and enable VSF on that switch.
2. After reboot, the device will come up as standalone VSF commander. Provision the second device with its memberID, J-Number, and, optionally, the MAC address.

3. Provision the module with its J-Number. Then, provision VSF link and ports on that module.
4. Connect the new device (with the factory default configuration) to this VSF device. The new device should be connected to the configured VSF ports.
5. The new device will reboot and join as standby.

Configuration commands

copy core-dump

Copy core-dump from the specified VSF member. User can copy available core-dump file from interface module or management module.

Syntax

```
copy core-dump vsf member <MEMBER-ID> <SLOT-ID> | mm-active sftp | tftp | usb |
xmodem <HOST-NAME-STR> | <IP-ADDR> | <IPV6-ADDR> <FILENAME-STR>
```

Description

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

Parameters

<MEMBER-ID>

Specify the VSF member.

<SLOT-ID>

Copy interface module core-dump file.

copy crash-data

Copy the crash data file of the switch.

Syntax

```
copy crash-data vsf member <MEMBER-ID> <SLOT-ID-RANGE> | mm | sftp | tftp | usb |
xmodem sftp | tftp | usb | xmodem <HOST-NAME-STR> | <IP-ADDR> | <IPV6-ADDR>
<FILENAME-STR>
```

Description

Copy the switch crash data file.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<SLOT-ID-RANGE>

Enter the single slot identifier.

copy crash-files

Syntax

```
copy crash-files vsf member <MEMBER-ID> [<SLOT-ID-RANGE> | mm-active sftp | tftp |
usb | xmodem] <HOST-NAME-STR> | <IP-ADDR> | <IPV6-ADDR> <FILENAME-STR>
```

Description

Copy the switch crash files from the specific VSF member

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<SLOT-ID-RANGE>

Enter single slot identifier.

mm-active

Copy active management module crash files.

Usage

```
Switch(config)# copy crash-files vsf member
Switch(config)# copy crash-files vsf member 1
```

copy crash-log vsf-member

Syntax

```
copy crash-log vsf member <MEMBER-ID> | <SLOT-ID-RANGE> | mm | sftp | tftp | usb |
xmodem sftp | tftp | usb | xmodem <HOST-NAME-STR> | <IP-ADDR> | <IPV6-ADDR>
<FILENAME-STR>
```

Description

Copy the switch log file.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<SLOT-ID-RANGE>

Enter the single slot identifier.

copy fdr-log

Copy FDR (Flight data recorder) logs. User can either copy from management module or interface module or both.

Syntax

```
copy fdr-log vsf member <MEMBER-ID> all | mm-active sftp | tftp | usb | xmodem
<HOST-NAME-STR> | <IP-ADDR> | <IPV6-ADDR> <FILENAME-STR>
```

Description

Copy FDR logs from the switch to an SFTP/TFTP server, USB or xmodem terminal.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

erase fdr-log vsf

Erase FDR log from the specified member.

Syntax

`erase fdr-log [vsf member <MEMBER-ID>] [slot | mm-active]`

Description

Erase the FDR log files.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

power-over-ethernet vsf-member

Syntax

`[no] power-over-ethernet vsf member <MEMBER-ID> pre-std-detect [slot <SLOT-LIST>] [ports <PORT-LIST>]`

Description

Set Power Over Ethernet (PoE) configuration parameters. Pre-standard detection and redundancy can be configured only at a per-member level when VSF is enabled.

Parameters

<MEMBER-ID>

The VSF member-ID for the 'member' command/parameter. Member ID value can be in the range from 1 to 2.

SLOT-ID-RANGE

Alphabetic device slot identifier or slot range.

<PORT-LIST>

Ethernet port number, a list of ethernet ports or 'all' for all ports.

<SLOT-ID-RANGE>

Enter an alphabetic device slot identifier or slot range preceded with the VSF member-ID [VSF-MEMBER/SLOT].

n+1

One of the highest power supplies will be held inreserve. In the event of a simple power supply failure,no powered devices will be shut down.

full

Half of the available power supplies will be held on reserve.

<THRESHOLD-LEVEL>

Threshold level in the range from 1 to 99.

Usage

`power-over-ethernet vsf member <MEMBER-ID> redundancy [n+1 | full]`

`[no] power-over-ethernet vsf member <MEMBER-ID> redundancy [n+1 | full]`

`power-over-ethernet vsf member <MEMBER-ID> slot <SLOT-ID-RANGE> threshold <THRESHOLD-LEVEL>`

redundancy switchover

Redundancy configuration for management modules.

Syntax

```
redundancy switchover
```

Description

The command causes the VSF Commander switch to immediately switch over to the standby switch.

snmp-server enable traps vsf

Syntax

```
[no] snmp-server enable traps vsf
```

Description

Enable traps for the VSF functionality.

Validation rules

Validation	Error/Warning/Prompt
This command cannot be executed if VSF is not enabled.	VSF is not enabled.

vsf domain

Syntax

```
vsf domain <DOMAIN-ID>
```

Description

Change domain ID for the VSF virtual chassis. Once VSF is enabled and virtual chassis is formed, VSF domain ID can be changed using this command.

Parameters

<DOMAIN-ID>

The VSF virtual chassis domain ID. Domain ID value can be in the range from 1 to 4294967296.

Validation rules

Validation	Error/Warning/Prompt
Domain ID must be a 32-bit unsigned integer.	The domain ID cannot be zero.

vsf enable

From the config context:

Syntax

```
vsf enable domain <DOMAIN-ID>
```

Description

Enable VSF on the switch. Allows for switches to be stacked using Ethernet ports.

Parameters

<DOMAIN-ID>

The VSF virtual chassis domain ID. Domain ID value can be in the range from 1 to 4294967296.

The command `vsf enable domain` reboots only one switch to form the VSF fabric. Upon reboot, the switches come up in the “VSF enabled” mode. Port numbers are prefixed with member numbers, such as “1/A1,”. The configuration on the switch becoming Commander will be retained, but any pre-existing configuration on other switches **will** be over-written.

The switches will inherit the same switch software as the member becoming Commander. If the software image of a switch needs to be updated, the switch will reboot twice.

vsf disable

Syntax

```
vsf disable
```

Description

Disable VSF on the virtual chassis.

Restriction

This command will not be available until VSF is enabled.

Validation rules

Validation	Error/Warning/Prompt
When `vsf enable` is executed on an VSF disabled switch following warning message will be displayed.	This will save the current configuration and reboot the switch. Continue [y/n]?
	Run `vsf enable` on VSF virtual chassis.
Run VSF disable when VSF links is UP.	VSF cannot be disabled when the VSF virtual chassis is active.
	Run `vsf disable` command on VSF disabled switch.

vsf member reboot

Syntax

```
boot vsf member <MEMBER-ID>
```

Description

Reboot the VSF member and rejoin the virtual chassis with the current configuration.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

Restriction

This command will not be available until VSF is enabled.

vsf member remove

This command removes the entire configuration for a specified member. After issuing the command, the specified member-ID is available for re-use and may be provisioned or assigned to another device.

If the member physically exists, its configuration will be erased. It will then be powered down by default.

Syntax

```
vsf member <MEMBER-ID> remove
```

Description

This command erases the VSF member configuration and VSF member will be powered off.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

Restriction

This command will not be available until VSF is enabled.

Validation rules

Validation	Error/Warning/Prompt
If VSF not enabled, this command is not allowed.	VSF is not enabled.
VSF member neither exists nor provisioned	The specified VSF virtual chassis member either does not exist or is not provisioned.
VSF standby syncing add remove member blocked	VSF virtual chassis members cannot be added or removed while the standby is booting.
VSF member remove causes VSF virtual chassis split	Removing this member is not allowed since it would result in a VSF virtual chassis split.
VSF missing member remove	The specified VSF virtual chassis member will be removed and its configuration will be erased. The resulting configuration will be saved. Continue [y/n]?
VSF VC member does not exist	The specified VSF virtual chassis member does not exist.
VSF provision member remove	The specified VSF virtual chassis member configuration will be erased. The resulting configuration will be saved. Continue [y/n]?
VSF remove commander	The VSF virtual chassis commander cannot be removed. Please fail over to standby before trying to remove the commander.

Table Continued

Validation	Error/Warning/Prompt
VSF member remove	The specified VSF virtual chassis member will be removed and its configuration will be erased. The resulting configuration will be saved. The VSF member will be shut down. Continue [y/n]?
VSF standby remove	The specified VSF virtual chassis member will be removed and its configuration will be erased. The resulting configuration will be saved. The VSF member will be shut down. Continue [y/n]?

vsf member shutdown

For a switch that physically exists, this command will cause the switch to shut down. `shutdown` is used in preparation to remove the switch from the virtual chassis. The switch will not become a deciding member of the virtual chassis again until it is rejoined.

The `shutdown` command can not be used on the Commander. The `shutdown` command will succeed only if the switch physically exists and is an active member of the virtual chassis.

Syntax

`vsf member <MEMBER-ID> shutdown`

Description

Shut down the VSF virtual chassis member.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

Restriction

Shutdown will not be available until VSF is enabled.

Validation rules

Validation	Error/Warning/Prompt
If member switch physically exists	The specified VSF virtual chassis member will be shut down. Continue [y/n]?
If member switch physically exists and is the commander	The VSF virtual chassis commander cannot be shut down. Please fail over to the standby first.
If member switch does not physically exist	The specified VSF virtual chassis member does not exist.
If shutting down a member will cause a VC –split	Shutting down this VSF virtual chassis member is not allowed since it would result in a VSF virtual chassis split.
If VSF not enabled, this command is not allowed.	VSF is not enabled.

vsf member priority

Syntax

`[no] vsf member <MEMBER-ID> priority <PRIORITY>`

Description

Assign a priority to the specified VSF virtual chassis member. The higher the priority, the more likely that the virtual chassis member will become the commander at the next virtual chassis reboot. The default priority value is 128.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<PRIORITY>

The priority value for VSF member. Priority value can be in the range from 1 to 255.

vsf member type

Syntax

`vsf member <MEMBER-ID> type <TYPE> [mac <MAC-ADDR>]`

Description

Configure the family of the VSF member-switch being provisioned. After provisioning, the VSF member-switch can be configured as if it were physically present. When an VSF member-switch matching the provisioned details joins the VSF, it is provided this configuration. A new or missing VSF member can be configured as a provisioned device by using this command.

Parameters

<MEMBER-ID >

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<TYPE>

Configure the family of the VSF member-switch being provisioned.

<MAC-ADDR >

Configure the MAC address of the VSF member switch being provisioned.

Restrictions

- If switch “N” physically exists, the command will fail.
- If switch “N” is provisioned, the command can be used to change the MAC or type.
- If the J-Number is known to not support stacking, or the J-Number is unknown, the command will fail.
- If the same MAC address is already provisioned or exists on another member ID, the command will fail.

Usage

- `vsf member <2> type <J9850A> mac <001122-334455>`

  Updates the strict provisioning for VSF member 2, and changes the MAC address to 001122-334455.

- `vsf member <2> type <J9850A>`

Changes the “strict” provisioning for VSF member 2 to “loose” provisioning. The configured MAC address is then removed.

- `vsf member <2> type <J9850A> mac <00aabb-cceedd>`

  Changes “loose” provisioning for VSF member 2 to “strict” provisioning with MAC address 00aabb-cceedd.

Validation rules

Validation	Error/Warning/Prompt
If the member-ID is not between 1 to 2 for 5400R, then this command will return an error.	The VSF member-ID value is not in range.
The member-ID must physically exist or already be provisioned.	The specified VSF virtual chassis member either does not exist or is not provisioned.
When each time new member is configured, write mem is called.	This will save the current configuration. Continue [y/n]?

Show commands

show vsf

Shows the current status and all current configurations of the provisioned VSF configuration on a switch.

Syntax

`show vsf [detail]`

Description

Shows detailed information of VSF virtual chassis members that are provisioned.

```
Switch# show vsf
VSF Domain ID : 44444
MAC Address : 3464a9-b2533f
VSF Topology : Chain
VSF Status : Active
Uptime : 32d 4h 28m
VSF Oobm-MAD : Enabled
Software Version : KB.16.01.0004
Mbr
ID      Mac Address   Model                     Pri Status
----------------------------------------------------------
1       3464a9-b24300 HP J9850A Switch 5406Rzl2 255 Commander
2       288023-98ae00 HP J9850A Switch 5406Rzl2 100 Standby
```

```
Switch(config)# show vsf detail

 VSF Domain ID    : 44444
 MAC Address      : 3464a9-b2533f
 VSF Topology     : Chain
 VSF Status       : Active
 Uptime           : 32d 4h 28m
 VSF Oobm-MAD     : Enabled
 VSF Lldp-MAD     : Disabled
 Software Version : KB.16.01.0004

Name              : HP-VSF-Switch
```

```
Contact          :
Location         :


Member ID        : 1
MAC Address      : 3464a9-b2533f
Type             : J9850A
Model            : HP J9850A Switch 5406Rzl2
Priority         : 255
Status           : Commander
ROM Version      : KB.16.01.0006
Serial Number    : ff ff ff ff ff ff ff ff ff ff
Uptime           : 32d 4h 28m
CPU Utilization  : 0%
Memory - Total   : 704,471,040 bytes
Free             : 522,821,984 bytes
VSF Links -
#1 : Active, Peer member 2



Member ID        : 2
MAC Address      : 288023-98ae00
Type             : J9850A
Model            : HP J9850A Switch 5406Rzl2
Priority         : 100
Status           : Standby
ROM Version      : KB.16.01.0006
Serial Number    : SG46G4904K
Uptime           : 32d 4h 20m
CPU Utilization  : 2%
Memory - Total   : 704,471,040 bytes
Free             : 536,014,360 bytes
VSF Links -
#1 : Active, Peer member 1
```

Validation rules

Validation	Error/Warning/Prompt
If VSF is not enabled, this command is not applicable.	VSF is not enabled.

show vsf link

Syntax

`show vsf link [detail]`

Description

Shows the VSF port state of the VSF links for each VSF member, in detail.

show vsf link

```
Switch# show vsf link

 VSF Member 1
                Link     Peer   Peer
 Link Link-Name  State    Member Link
 ---- ---------- -------- ------ ----
 1    I-Link1_1  Up       2      1
```

```
VSF Member 2
                 Link     Peer   Peer
 Link Link-Name  State    Member Link
 ---- ---------- -------- ------ ----
 1    I-Link2_1  Up       1      1
```

show vsf link detail

```
Switch# show vsf link detail
 vsf Member: 1     Link: 1
 Vsf-Port   Port-State
 --------   -------------------------
 1/E1       Up: Connected to port 2/E1

vsf Member: 2     Link: 1
 Vsf-Port   Port-State
 --------   -------------------------
 2/E1       Up: Connected to port 1/E1
```

show vsf member

Syntax

`show vsf member <MEMBER-ID>`

Description

Shows the virtual chassis information for member.

Parameters

<MEMBER ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show vsf member 1

```
Switch# show vsf member 1
Member ID         : 1
MAC Address       : a01d48-8f6700
Type              : J9850A
Model             : HP J9850A Switch 5406Rzl2
Priority          : 128
Status            : Standby
ROM Version       : KB.16.01.0005
Serial Number     : SG4ZG95321
Uptime            : 21d 19h 5m
CPU Utilization   : 2%
Memory - Total    : 698,957,824 bytes
Free              : 528,240,524 bytes
VSF Links -
#1 : Active, Peer member 2
```

show vsf member 2

```
Switch# show vsf member 2
Member ID          : 2
Mac Address        : 288023-98ae00
Type               : J9850A
Model              : HP J9850A Switch 5406Rzl2
Priority           : 100
Status             : Standby
ROM Version        : KB.16.01.0005
Serial Number      : SG46G4906P
Uptime             : 32d 4h 11m
CPU Utilization    : 0%
Memory - Total     : 709,357,568 bytes
Free               : 546,939,520 bytes
VSF Links -
#1 : Active, Peer member 1
```

show system information vsf member

Syntax

`show system information [vsf member <MEMBER-ID>]`

Description

Show global configured and operational system parameters of the specified VSF members.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show system

```
Switch(config)# show system
 Status and Counters - General System Information
  System Name          : bolt-vsf-sws
  System Contact       :
  System Location      :
  Allow V2 Modules     : No
  MAC Age Time (sec)   : 300
  Time Zone            : -480
  Daylight Time Rule   : Continental-US-and-Canada
  Software revision    : KB.16.01.0004
  Base MAC Addr        : 3464a9-b2533f

 VSF-Member :1
  ROM Version          : KB.16.01.0005
 Up Time              : 32 days
  CPU Util (%)         : 2
  MAC Addr             : 3464a9-b24300
  Serial Number        : SG4BG491BL
  Memory   - Total     : 709,357,568    Free    : 529,020,080

 VSF-Member :2
  ROM Version          : KB.16.01.0005
  Up Time              : 32 days
  CPU Util (%)         : 0
```

```
  MAC Addr            : 288023-98ae00
  Serial Number       : SG46G4906P
  Memory   - Total    : 709,357,568    Free    : 546,939,520
```

show system information VSF member 2

```
Switch# show system information vsf member 1
 Status and Counters - General System Information
  System Name         : bolt-vsf-sws
  System Contact      :
  System Location     :
  Allow V2 Modules    : No
  MAC Age Time (sec) : 300
  Time Zone           : -480
  Daylight Time Rule : Continental-US-and-Canada
  Software revision   : KB.16.01.0004
  Base MAC Addr       : 3464a9-b2533f

 VSF-Member :1
  ROM Version         : KB.16.01.0005
  Up Time             : 32 days
  CPU Util (%)        : 0
  MAC Addr            : 3464a9-b24300
  Serial Number       : SG4BG491BL
  Memory   - Total    : 709,357,568      Free    : 529,413,568

Switch# show system information vsf member 2
 Status and Counters - General System Information
  System Name         : bolt-vsf-sws
  System Contact      :
  System Location     :
  Allow V2 Modules    : No
  MAC Age Time (sec) : 300
  Time Zone           : -480
  Daylight Time Rule : Continental-US-and-Canada
  Software revision   : KB.16.01.0004
  Base MAC Addr       : 3464a9-b2533f

 VSF-Member :2
  ROM Version         : KB.16.01.0005
  Up Time             : 32 days
  CPU Util (%)        : 0
  MAC Addr            : 288023-98ae00
  Serial Number       : SG46G4906P
  Memory   - Total    : 709,357,568    Free    : 546,939,520
```

show system chassislocate

Syntax

`show system chassislocate [vsf member <MEMBER-ID>]`

Description

Show locator LED information. If VSF is enabled, this shows locator LED information for all the VSF members.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show system chassislocate

```
Switch# show system chassislocate
 Locator LED Status
  VSF     Current  Time
  Member  State    Remaining  Configuration
  ------  -------  ---------  -------------
  1       off
  2       blink    00:29:10
```

show system chassislocate vsf member 2

```
Switch# show system chassislocate vsf member 2
Locator LED Status
Switch# show system chassislocate vsf member 2
 Locator LED Status
  VSF     Current  Time
  Member  State    Remaining  Configuration
  ------  -------  ---------  -------------
  2       blink    00:29:45
```

show boot-history

Syntax

`show boot-history [vsf member <MEMBER-ID>]`

Description

Display the system boot log for VSF.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show system temperature

Syntax

`show system temperature [vsf member <MEMBER-ID>]`

Description

Show current temperature sensor information. If VSF is enabled, this shows the temperature sensor information for all VSF members.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show system temperature

```
Switch# show system temperature
```

```
System Air Temperatures
  VSF-Member  1
  Temp    Current  Max   Min                        Avg
  Sensor  Temp     Temp  Temp  Threshold  OverTemp  Temp
  ------- -------- ----- ----- ---------- --------- -------
  Chassis 31C      33C   27C   55C        NO         29.46C

  VSF-Member  2
  Temp    Current  Max   Min                        Avg
  Sensor  Temp     Temp  Temp  Threshold  OverTemp  Temp
  ------- -------- ----- ----- ---------- --------- -------
  Chassis 30C      32C   28C   55C        NO         29.08C
```

show system temperature vsf member 2

```
Switch# show system temperature vsf member 2

 System Air Temperatures
  VSF-Member  2
  Temp    Current  Max   Min                        Avg
  Sensor  Temp     Temp  Temp  Threshold  OverTemp  Temp
  ------- -------- ----- ----- ---------- --------- -------
  Chassis 30C      32C   28C   55C        NO         29.08C
```

show system fans

Syntax

`show system fans [vsf member <MEMBER-ID>]`

Description

Show system fan status. If VSF is enabled, this shows the system fan status for all VSF members.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show system fans

```
show system fans

 Fan Information
 VSF-Member  1
    Num   | State        | Failures
   -------+--------------+----------
   Sys-1  | Fan OK       |   0
   Sys-2  | Fan OK       |   0
   Sys-3  | Fan OK       |   0
   Sys-4  | Fan OK       |   0
 0 / 4 Fans in Failure state
 0 / 4 Fans have been in Failure state

 VSF-Member  2
    Num   | State        | Failures
   -------+--------------+----------
   Sys-1  | Fan OK       |   0
   Sys-2  | Fan OK       |   0
```

```
  Sys-3  | Fan OK       |    0
  Sys-4  | Fan OK       |    0
0 / 4 Fans in Failure state
0 / 4 Fans have been in Failure state
```

show system fans vsf member 1

```
show system fans VSF member 1

Fan Information
 VSF-Member  1
     Num  | State        | Failures
   -------+--------------+----------
   Sys-1  | Fan OK       |    0
   Sys-2  | Fan OK       |    0
   Sys-3  | Fan OK       |    0
   Sys-4  | Fan OK       |    0
 0 / 4 Fans in Failure state
 0 / 4 Fans have been in Failure state
```

show CPU

Syntax

`show cpu <SECONDS>`

Description

Show average CPU utilization.

Parameters

<INTEGER>

Refresh time.

<SLOT-ID-RANGE>

Alphabetic device slot identifier or slot range.

Usage

`show cpu slot process refresh <INTEGER>`

`show cpu process slot <SLOT-ID-RANGE> refresh <INTEGER>`

show cpu slot all

```
Switch# show cpu slot all
  VSF slot 1/a:
  -------------
  12 percent busy, from 18 sec ago

  1 sec ave: 14 percent busy
  5 sec ave: 12 percent busy
  1 min ave: 12 percent busy
  VSF slot 1/f:
  -------------
  16 percent busy, from 17 sec ago
  1 sec ave: 27 percent busy
  5 sec ave: 16 percent busy
  1 min ave: 15 percent busy

  VSF slot 2/a:
```

```
  -------------
  12 percent busy, from 18 sec ago
  1 sec ave: 14 percent busy
  5 sec ave: 12 percent busy
  1 min ave: 12 percent busy

  VSF slot 2/f:
  -------------
  16 percent busy, from 17 sec ago
  1 sec ave: 27 percent busy
  5 sec ave: 16 percent busy
  1 min ave: 15 percent busy
```

show cpu slot 1/A

```
Switch# show cpu slot 1/A

  VSF slot 1/a:
  -------------
  12 percent busy, from 18 sec ago
  1 sec ave: 14 percent busy
  5 sec ave: 12 percent busy
  1 min ave: 12 percent busy
```

show CPU process slot

Syntax

`show cpu <SECONDS>`

Description

Show average CPU utilization.

Parameters

<SECONDS >

The time in seconds over which to average CPU utilization.

Usage

`show cpu slot <SLOT-LIST> <SECONDS>`

`show cpu process slot <SLOT-LIST> refresh <COUNT>`

show cpu process slot all

```
Switch# show cpu process slot all

 VSF slot 1/A:
 -------------
 Process tracker state: ACTIVE
 Process tracking time: 30 seconds

                  Total      %            Time Since   Times    Max
Process Name      Priority   Time    CPU  Last Ran     Ran      Time
----------------+----------+-------+-----+------------+-------+-------
Hardware Mgmt-3     192        3 s    6      234 ms    214      35 ms
System Services-2   156        3 s    5       55 ms    110      50 ms
Idle-3               1        12 s    24     731 us    245918   193 us
Idle-1              226       25 s    51     770 us    123627   319 us
```

```
Idle-0                226       5 s    10     459 us    122921  170 us

 VSF slot 2/F:
 -------------
 Process tracker state: ACTIVE
 Process tracking time: 30 seconds

                      Total       %           Time Since    Times   Max
Process Name          Priority   Time    CPU   Last Ran      Ran     Time
------------------+----------+------+-----+------------+-------+-----
Hardware Mgmt-3        192        3 s      8      54 ms       189     41 ms
System Services-2      156        3 s      8      2 s         131     50 ms
Idle-3                  1         9 s      23     870 us      160197 199 us
Idle-0                 226        4 s      10     926 us      80053   162 us
Idle-1                 226       19 s      48     1 ms        80545   395 us
```

show cpu process slot 1/A

```
Switch# show cpu process slot 1/A

 VSF slot 1/A:
 -------------
 Process tracker state: ACTIVE
 Process tracking time: 30 seconds

                              Total     %      Time Since    Times    Max
Process Name      Priority   Time      CPU    Last Ran      Ran      Time
----------------+----------+--------+-----+-----------+-------+-----
Hardware Mgmt-3    192        3 s       6      234 ms       214     35 ms
System Services-2  156        3 s       5      55 ms        110     50 ms
Idle-3              1         12 s      24     731 us       245918 193 us
Idle-1             226        25 s      51     770 us       123627 319 us
Idle-0             226        5 s       10     459 us       122921 170 us
```

show modules

Syntax

`show modules details vsf member <MEMBER-ID> MM1 | MM2 | slot <SLOT-LIST>`

Description

Show module details for VSF members.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

MM1

Show MM1 module information of the specified VSF member.

MM2

Show MM2 module information of the specified VSF member.

slot

Show SLOT module information of the specified VSF member.

<SLOT-LIST>

Enter an alphabetic device slot identifier or a slot range.

show modules

```
Switch# show modules

 Status and Counters - Module Information
  Chassis: 5406Rzl2 J9850A         Serial Number:   SG4BG491BL
  Allow V2 Modules:   No
                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  1/MM1 HP J9827A Management Module 5400Rzl2   SG4BG4C0C0     Active   YES   1
  1/MM2 HP J9827A Management Module 5400Rzl2   A123456789     Offline  YES   1
  1/A   HP J9992A 20p PoE+ / 1p 40GbE QSFP+... B123456789     Up       YES   3
  1/F   HP J9991A 20p PoE+ / 4p 1/2.5/5/XGT... SG5ZGPH190     Up       YES   3
  2/MM1 HP J9827A Management Module 5400Rzl2   SG45G4C0VZ     Active   YES   1
  2/A   HP J9992A 20p PoE+ / 1p 40GbE QSFP+... c123456789     Up       YES   3
  2/F   HP J9991A 20p PoE+ / 4p 1/2.5/5/XGT... SG5ZGPH183     Up       YES   3

Switch# show modules details vsf member 1
 MM1                   Show MM1 module information of the specified VSF member.
 MM2                   Show MM2 module information of the specified VSF member.
 slot                  Show SLOT module information of the specified VSF member.

Switch# show modules details vsf member 1
 Status and Counters - Module Information
  Chassis: 5406Rzl2 J9850A         Serial Number:   SG4BG491BL
  Allow V2 Modules:   No
                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  1/MM1 HP J9827A Management Module 5400Rzl2   SG4BG4C0C0     Active   YES   1

                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  1/MM2 HP J9827A Management Module 5400Rzl2   D123456789     Offline  YES   1

                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  1/A   HP J9992A 20p PoE+ / 1p 40GbE QSFP+... E123456789     Up       YES   3

                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  1/F   HP J9991A 20p PoE+ / 4p 1/2.5/5/XGT... SG5ZGPH190     Up       YES   3

Switch# show modules details vsf member 2
 Status and Counters - Module Information
  Chassis: 5406Rzl2 J9850A         Serial Number:   SG4BG491BL
  Allow V2 Modules:   No

                                                                 Core  Mod
  Slot  Module Description                     Serial Number  Status   Dump  Ver
  ----- -------------------------------------- -------------- -------- ----- ---
  2/MM1 HP J9827A Management Module 5400Rzl2   SG45G4C0VZ     Active   YES   1

                                                                 Core  Mod
```

```
  Slot  Module Description                      Serial Number  Status   Dump  Ver
  ----- --------------------------------------- -------------- -------- ----- ---
  2/A   HP J9992A 20p PoE+ / 1p 40GbE QSFP+... H123456789      Up       YES   3

                                                                        Core  Mod
  Slot  Module Description                      Serial Number  Status   Dump  Ver
  ----- --------------------------------------- -------------- -------- ----- ---
  2/F   HP J9991A 20p PoE+ / 4p 1/2.5/5/XGT... SG5ZGPH183     Up       YES   3
```

show modules details vsf member 1 slot 1/a

```
Switch# show modules details vsf member 1 slot 1/a
 Status and Counters - Module Information
  Chassis: 5406Rzl2 J9850A         Serial Number:   SG4BG491BL
  Allow V2 Modules:    No

                                                                        Core  Mod
  Slot  Module Description                      Serial Number  Status   Dump  Ver
  ----- --------------------------------------- -------------- -------- ----- ---
  1/A   HP J9992A 20p PoE+ / 1p 40GbE QSFP+... A123456789      Up       YES   3
```

show power-over-ethernet

Syntax

`show power-over-ethernet vsf member <MEMBER-ID>`

`show power-over-ethernet slot all`

Description

Show power-over-ethernet for named slots or specified VSF member switches.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show power-over-ethernet slot all

```
Switch# show power-over-ethernet slot all

Status and Counters - System Power Status for slot 1/A
Maximum Power        :   0 W           Operational Status   : On
Power In Use         :   0 W +/- 6 W   Usage Threshold (%) : 80

Status and Counters - System Power Status for slot 2/A
Maximum Power        :   0 W           Operational Status   : On
Power In Use         :   0 W +/- 6 W   Usage Threshold (%) : 80
```

show power-over-ethernet slot 1/A

```
Switch# show power-over-ethernet slot 1/A

Maximum Power        :   0 W           Operational Status   : On
Power In Use         :   0 W +/- 6 W   Usage Threshold (%) : 80
```

show power-over-ethernet vsf member 1

```
Switch(config)# show power-over-ethernet vsf member 1
Status and Counters - System Power Status for member 1
      Maximum Operational                      Usage
Slot  Power   Status      Power In Use      Threshold (%)
----- ------- ----------- ---------------- -------------
1/A   266 W   On            0 W +/- 6 W       80
1/L     0 W   Faulty        0 W +/- 6 W       80
```

show power-over-ethernet vsf member 2

```
Switch# show power-over-ethernet vsf member 2
 Status and Counters - System Power Status for member 2
       Maximum Operational                      Usage
 Slot  Power   Status      Power In Use      Threshold (%)
 ----- ------- ----------- ---------------- -------------
 2/A   266 W   On            0 W +/- 6 W        80
 2/C     0 W   On            0 W +/- 6 W        80
```

show system power-supply

Syntax

`show system power-supply [detailed | fahrenheit]`

Description

Shows power supply information in either full detail or full detail in Fahrenheit only. Default temperature is displayed in degrees Celsius.

Command context

`manager and operator`

Parameters

`detailed`

Shows detailed switch power supply sensor information.

`fahrenheit`

Shows detailed switch power supply sensor information with temperatures in degrees Fahrenheit.

Usage

- The `show system power-supply detailed` command shows detailed information for the local power supplies only.
- The `show system power-supply detailed` command shows detailed information for power supplies in the `powered` state only.

Examples

Use of the command `show system power-supply` shows the power supply status for all active switches.

```
Switch# show system power-supply

Power Supply Status:
```

```
PS#    Model      Serial          State         AC/DC   + V    Wattage
---- --------- ------------ ----------------- -------------- ---------
  1   J9828A    IN30G4D009   Powered          AC 120V/240V      700
  2   J9828A    IN30G4D00C   Powered          AC 120V/240V      700
  3                          Not Present      -- ---------        0
  4   J9830A    IN43G4G05H   Powered          AC 120V/240V     2750

  3 /  4 supply bays delivering power.
  Total power: 4150 W
```

Use of the command `show system power-supply detailed` shows the power supply status in detail for all active switches.

```
Switch# show system power-supply detailed

 Status and Counters - Power Supply Detailed Information

  PS# Model   Serial      State        Status
  --- ------- ----------- ------------ -------------------------------------
  1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 31 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.0C/80.6F
                                       Internal Temp (C/F)  : 30.5C/86.0F
                                       Fan 1 Speed          : 1600 RPM (47%)
                                       Fan 2 Speed          : 1600 RPM (47%)

  2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 21 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.7C/80.6F
                                       Internal Temp (C/F)  : 32.5C/89.6F
                                       Fan 1 Speed          : 1600 RPM (47%)
                                       Fan 2 Speed          : 1600 RPM (47%)

  3                       Not Present

  4   J9830A  IN43G4G05H  Powered      AC Power Consumption : 90 Watts
                                       AC MAIN/AUX Voltage  : 210/118 Volts
                                       Power Supplied       : 16 Watts
                                       Power Capacity       : 2750 Watts
                                       Inlet Temp (C/F)     : 30.9C/86.0F
                                       Internal Temp (C/F)  : 65.6C/149.0F
                                       Fan 1 Speed          : 2000 RPM (37%)
                                       Fan 2 Speed          : 1950 RPM (36%)



   3 /  4 supply bays delivering power.
   Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply fahrenheit` shows the power supply status in Fahrenheit for all active switches.

```
Switch# show system power-supply detailed fahrenheit
Power Supply Status:
 Mem PS# Model   Serial      State        Status
 --- --- ------- ----------- ------------ -------------------------------------
 1   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                          AC MAIN/AUX Voltage : 118/208 Volts
```

```
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)

 1   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
                                               AC Input Voltage    : 208 Volts
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)

 1   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                               AC Input Voltage    : 119 Volts
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)

 1   4                       Not Present

 2   1   J9830A  IN5BGZ81KZ  Powered      Power Consumption   : 95 Watts
                                               AC MAIN/AUX Voltage : 118/208 Volts
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)

 2   2   J9829A  IN5BGZ81KX  Powered      Power Consumption   : 51 Watts
                                               AC Input Voltage    : 208 Volts
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)

 2   3   J9828A  IN5BGZ81KY  Powered      Power Consumption   : 43 Watts
                                               AC Input Voltage    : 119 Volts
                                               Inlet/Internal Temp : 85.6F/87.7F
                                               Fan 1 Speed (util)  : 1650RPM (20%)
                                               Fan 2 Speed (util)  : 1600RPM (19%)
 2   4                       Not Present
------------------------------------------------------------------------------
   6 / 8 supply bays delivering power.
   Total Input Power: 378 Watts
```

Use of the command `show system power-supply detailed` shows the power supply status all active switches including a nonpowered J9830A PSU.

```
switch# show system power-supply detailed

 Status and Counters - Power Supply Detailed Information

  PS# Model   Serial      State        Status
  --- ------- ----------- ------------ -------------------------------------
  1   J9828A  IN30G4D009  Powered      AC Power Consumption : 44 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 31 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.0C/80.6F
                                       Internal Temp (C/F)  : 30.5C/86.0F
                                       Fan 1 Speed          : 1600 RPM
                                       Fan 2 Speed          : 1600 RPM

  2   J9828A  IN30G4D00C  Powered      AC Power Consumption : 46 Watts
                                       AC MAIN Voltage      : 209 Volts
                                       Power Supplied       : 21 Watts
                                       Power Capacity       : 700 Watts
                                       Inlet Temp (C/F)     : 27.7C/80.6F
                                       Internal Temp (C/F)  : 32.5C/89.6F
```

```
                                                   Fan 1 Speed              : 1600 RPM
                                                   Fan 2 Speed              : 1600 RPM

  3                        Not Present

  4   J9830A   IN43G4G05H  Aux Not Powered


   2 /  4 supply bays delivering power.
   Currently supplying 68 W / 4150 W total power.
```

Use of the command `show system power-supply` shows the power supply status all active switches with power supply #2 showing permanent failure.

```
switch# show system power-supply

Power Supply Status:

 PS#    Model      Serial          State          AC/DC  + V    Wattage
 ---- --------- ------------ ---------------- -------------- ---------
   1                          Not Present       -- ---------       0
   2   J9829A    IN30G4D00C   Permanent Failure  AC 120V/240V     1100
   3   J9829A    IN30G4D00D   Powered           -- ---------       1100
   4   J9829A    IN43G4G05H   Powered           AC 120V/240V      1100

   3 /  4 supply bays delivering power.
   Total power: 3300 W
```

Table 26: *Field key for output of* `show system power-supply detailed`

Field	Description
AC Power Consumption	Actual power consumed from AC input
AC MAIN/AUX Voltage	Actual voltage measured on AC Input: • Two voltages are displayed for PS#4, as the J9830A includes two AC input IEC connectors. • Most power-supplies contain a single AC Input IEC connector and are labeled MAIN.
Power Supplied	Actual voltage being supplied from the power-supply to the switch for general power and PoE.
Power Capacity	The maximum power that the power-supply can provide to the switch.
Inlet Temp (C/F)	The thermal sensor at the inlet of the power-supply - shown in both Celsius and Fahrenheit
Internal Temp (C/F)	The thermal sensor internal to the power-supply (will vary depending upon the model) - shown in both Celsius and Fahrenheit. NOTE: There is no "Output Temperature Sensor" on either the 5400R or 3810M switches.

Table Continued

Field	Description
Fan Speed	Shows the current fan speed in RPM and the percent of total fan speed utilization. For PSUs that contain more than one fan, a separate line will be included for each.
Currently Supplying	A summary of the total power being supplied and the total capacity (same summary as seen on the command `show system power-supply`).

OOBM-MAD commands

vsf oobm-mad

Syntax

`[no] vsf oobm-mad`

Description

Enable OOBM-MAD (Multi-Active Detection) on the VSF device.

Validation rules

Validation	Error/Warning/Prompt
This command cannot be executed if VSF is not enabled.	`invalid input: oobm-mad`

oobm vsf member

Syntax

`oobm vsf member <MEMBER-ID> ip address { <IP-ADDR>/<MASK-LENGTH> | dhcp-bootp }`

Description

This command is used to configure the IPV4 address for the VSF OOBM.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

<IP-ADDR>/<MASK-LENGTH>

Interface IP address/mask.

`oobm vsf member <MEMBER-ID> ip default-gateway IP-ADDR`

Specify the default gateway using this form of the command. Configure the IPv4 default gateway address, which will be used when routing is not enabled on the switch. The *<IP-ADDR>* must be specified if the command is not preceded by `[no]`. Preceding the command with `[no]` deletes the default gateway address.The `[no]` form of this command does not take effect on default gateway address obtained via DHCP.

Usage

`oobm vsf member <VSF-MEMBER> ip`

`oobm vsf member <VSF-MEMBER> ip address`

oobm vsf member interface speed-duplex

Syntax

`oobm vsf member <MEMBER-ID> interface { disable | enable | <SPEED-DUPLEX> }`

Description

Configure various interface parameters for OOBM. The `interface` command must be followed by a feature-specific keyword. This is an OOBM context command. It can be called directly from the OOBM context.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

enable

Enable OOBM port.

disable

Disable OOBM.

<SPEED-DUPLEX>

Define mode of operation for the oobm port.

10-half

10 Mbps, half duplex.

100-half

100 Mbps, half duplex.

10-full

10 Mbps, full duplex.

100-full

100 Mbps, full duplex.

1000-full

1000 Mbps, full duplex.

auto

Use Auto Negotiation for speed and duplex mode.

Usage

```
oobm vsf member <VSF-MEMBER-ID> interface enable
oobm vsf member <VSF-MEMBER-ID> interface disable
```

show oobm vsf member

Syntax

`show oobm vsf member <MEMBER-ID>`

Description

Show OOBM VSF member.

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show oobm vsf member 1

```
Switch# show oobm vsf member 1

 VSF Member 1
  OOBM Port Type          : 100/1000T
  OOBM Interface Status   : Up
  OOBM Port               : Enabled
  OOBM Port Speed         : Auto

  MAC Address             : 3464a9-b24301
```

show oobm ip

Syntax

`show oobm ip [vsf member <MEMBER-ID>]`

Description

Show OOBM IP for specific members in VSF

Parameters

<MEMBER-ID>

The VSF member-ID for the member command or parameter. Member ID value can be in the range from 1 to 2.

show oobm ip

```
Switch# show oobm ip

IPv4 Status     : Enabled
IPv4 Default Gateway : 120.93.49.1

            |                                      Address  Interface
 VSF-member | IP Config IP Address/Prefix Length  Status   Status
 ---------- + --------- ------------------------- -------- ---------
 Global     | dhcp      120.93.49.9/24             Active   Up
 1          | dhcp      120.93.49.9/24             Active   Up
 2          | disabled                             Inactive Down
```

show oobm ip vsf member 1

```
Switch# show oobm ip vsf member 1
 IPv4 Status     : Enabled
  IPv4 Default Gateway : 15.212.178.1
            |                                      Address  Interface
 VSF-member | IP Config IP Address/Prefix Length  Status   Status
 ---------- + --------- ------------------------- -------- ---------
 1          | dhcp      15.212.178.244/24          Active   Up
```

show oobm ip vsf member 1,2

```
Switch(config)# show oobm ip vsf member 1,2
IPv4 Status : Enabled
IPv4 Default Gateway :
             |                                       Address   Interface
 VSF-member | IP Config IP Address/Prefix Length  Status    Status
 ---------- + --------- ------------------------- -------- ---------
 1          | dhcp                                 Active    Down

Switch(config)# sho oobm ip detail
Internet (IP) Service for OOBM Interface
Global Configuration
IPv4 Status : Enabled
IPv6 Status : Disabled
IPv4 Default Gateway :
IPv6 Default Gateway :




Origin | IP Address/Prefix Length Status
----------------------------------------
dhcp   |

VIPv4 SF Member 1
Status : Enabled
IPv6 Status : Disabled
IPv4 Default Gateway :
IPv6 Default Gateway :

Origin | IP Address/Prefix Length Status
----------------------------------------
dhcp   |
```

show oobm discovery

Syntax

`show oobm discovery`

Description

Show the discovered virtual chassis information.

show oobm discovery

```
Switch# show oobm discovery

Active Stack (This fragment)
 VSF-member Mac Address    Status
 ID
 ---------- -------------- ----------------
 2          10604b-b7a140  Global Commander
 1          10604b-b66980  Global Member
```

show running-config oobm

Syntax

`show running-config oobm`

Description

Show running-config OOBM.

show running-config oobm

```
Switch# show running-config oobm

Running configuration:
oobm
   ip address dhcp-bootp
   VSF-member 1
      ip address dhcp-bootp
      exit

   VSF-member 2
      ip address 192.168.10.1 255.255.255.0
      exit

   exit
```

show vsf trunk-designated-forwarder

Syntax

`show vsf trunk-designated-forwarder`

Description

Show the designated forwarders for each trunk.

For each trunk, only one member of the trunk will forward L2 flood traffic (unknown destination, Broadcast & Multicast). Use the `show vsf trunk-designated-forwarder` command to know which member will forward flood frames for a given trunk.

For known unicast traffic, trunks will always forward using local member links when possible and traffic will cross the VSF links to the other member only when local links of a trunk are down.

show vsf trunk designated forwarder

```
Switch(config)# show vsf trunk-designated-forwarder

 Trunk Designated Forwarders
  NAME  TYPE  Member
  ----- ----- ------
  Trk1  TRK   1
  Trk2  LACP  0
  Trk3  TRK   0
  Trk10 TRK   1
```

Validation rules

Validation	Error/Warning/Prompt
Downloading a non-VSF config on a VSF switch or downloading an invalid VSF-config for a current VSF must be blocked.	The configuration file for this VSF device is incorrect.
Upon enabling VSF, the hostname of the virtual chassis would change to a different string than it is when VSF is disabled.	HP-5406R-VSF

LLDP-MAD

LLDP-MAD is used to detect multiple-active VSF fragments.

VSF split and MAD operation

The following sequence explains a MAD scheme for a simple 2-member, VSF virtual chassis split scenario.

1. When the VSF link goes down and the VSF virtual chassis splits:
    - The Commander member (Fragment-A for this example) would continue to stay active.
    - The Standby member (Fragment B) would failover and become another commander.
2. Fragment-B sends an SNMP request to the downstream device seeking port status information of all non-local ports of the LACP Trunk. Non-local ports on Fragment-B refers to ports that are part of Fragment-A's member.
    - The downstream device responds to the SNMP request with the appropriate port status information.
    - If Fragment-A receives an unsolicited response to the SNMP request, it is ignored. This is because Fragment A has the pre-split Commander as part of its fragment and therefore will remain active.
3. Fragment-B sends 2 more SNMP queries downstream. If no response is received, the frontplane ports are shut down and turned inactive.
4. Alternatively, if Fragment-B receives an SNMP response:
    - If Fragment A links are UP, the frontplane ports will be shut down.
    - If Fragment-A links are DOWN, Fragment-B would stay UP.
5. Consider that Fragment-A is actually DOWN which has caused the split:
    - Request made to Fragment-B will be received by the downstream device and response will return to Fragment-B.
    - The downstream links to Fragment-A are DOWN therefore Fragment-B will remain UP.
    - Alternately, if Fragment-B is DOWN and caused the split then Fragment-A will neither send a request or act on an unsolicited response and will remain UP.

MAD readiness check

The MAD assist device must be connected over a LACP trunk interface to the VSF device. Once you configure the IP address of a MAD assist device, the VSF switch will perform a MAD readiness check to determine:

- If the MAD assist device is reachable.
- If a trunk interface is used to reach the device.
- If the trunk interface has at least one, linked —up, physical port on each member of the VSF switch.

If the above three conditions are not met, MAD will fail to detect dual active fragments in the event of a VSF split. This error will create a log message.

NOTE

The MAD readiness check is repeated periodically. If MAD-probe parameters have changed, an appropriate log message will be created.

vsf lldp-mad ipv4

Syntax

```
[no] vsf lldp-mad ipv4 <IPV4_ADDR> v2c <COMMUNITY-STR>
```

Description

Enable LLDP-MAD on the VSF device.

NOTE

The command `vsf lldp mad` requires a peer switch to be configured as the “assist” device.

Parameters

<IPV4_ADDR>

The IPv4 address of the MAD (Multi-Active Detection) device.

<COMMUNITY-STR>

The SNMP community string for the MAD (Multi-Active Detection) device.

Usage

```
Switch(config)# vsf lldp-mad ipv4
Switch(config)# vsf lldp-mad ipv4 <IPv4_ADDR>
Switch(config)# vsf lldp-mad ipv4 <MAD-IP-ADDRESS> v2c
Switch(config)# vsf lldp-mad ipv4 210.10.0.12 v2c <COMMUNITY-STR>
```

Validation rules

Validation	Error/Warning/Prompt
This command cannot be executed if VSF is not enabled.	VSF is not enabled. Cannot configure VSF LLDP MAD IP address because the specified IP address is a multicast IP address. Cannot configure VSF LLDP MAD IP address because the specified IP address is a link-local IP address. Cannot configure VSF LLDP MAD IP address because the specified IP addressis configured on the loopback interface. Cannot configure VSF LLDP MAD IP address because the specified IP address is configured on a local interface.
The MAD assist device and the VSF device must be on a common IP subnet for LLDP-MAD to work.	The MAD (Multi-Active Detection) device and the VSF device are not on the same network.

show vsf lldp-mad

Syntax

`show vsf lldp-mad [parameters | status]`

Description

Show the VSF LLDP-MAD information on the switch.

Parameters

parameters

Displays the MAD-assist configuration as well as the readiness state.

status

Displays the current state of the MAD probe.

Usage

`show vsf lldp-mad parameters`

`show vsf lldp-mad status`

show vsf lldp-mad parameters

```
Switch# show vsf lldp-mad parameters
    MAD device IP           : 210.10.0.12
    MAD readiness status    : Success
    MAD device MAC          : 5065f3-128cc5
    Reachable via Vlan      : 916
    Local LAG interface     : Trk10
    MAD-probe portset       : 1/A21,2/A21,
    LAG connectivity        : Full
```

show vsf lldp-mad status

```
Switch# show vsf lldp-mad status

    MAD device IP                           : 210.10.0.12
    MAD-probe portset                       : 1/A21,2/A21,
    VSF split                               : No
    MAD probe originator                    : No
    Number of probe requests sent           : 0
    Number of probe responses received      : 0
    MAD Active Fragment                     : Yes
```

VSF re-join after a split

If split fragment(s) re-join the VSF and become a single device, MAD readiness checks will be re-run and a fresh set of readiness parameters determined.

NOTE

One of the devices will reboot to join the VSF.

Mad assist device requirements

Limitations of MAD

VSF restrictions

- VSF is mutually exclusive with distributed trunking, mesh, and Q-in-Q.
- VSF port restrictions:
    - Must be 10Gbps/40Gbps. 1Gbps links are not supported.
    - A VSF link can only comprise ports with the same speed; either all 10G or all 40G
    - Maximum 8 ports in 1 VSF link.
    - VSF ports must be directly connected and there should be no transit devices between members.
- In a VSF virtual chassis, flow-control is not supported between ports on different chassis across VSF links.

Updates for a VSF virtual chassis

To update the firmware on a VSF virtual chassis, copy the new firmware to the VSF virtual chassis and reboot the VSF virtual chassis with the `boot system flash <primary | secondary>` command.

Chapter 25 VSF Fast Software Upgrade

Fast Software Upgrade (FSU) is a feature in VSF stacking which enables users to upgrade switch software with minimal down-time of links by rebooting the standby to boot with the upgraded image while the current commander and its IMs are still intact and forwarding traffic. The traditional software upgrade of a VSF stack consists of:

- Copying the switch software image through a CLI command on to the switch primary or secondary bank.
- Rebooting the entire stack for all members to boot with the newly upgraded SW image.

The down-time of links is longer because when the whole stack reboots, all IMs, on all members, are down during the time of reboot and stack formation.

NOTE

FSU is supported only on two member VSF stacks based on 5400R devices.

Upgrading the VSF stack software

Upgrading the software on a VSF stack consists of two basic steps:

Procedure

1. Copy the new software image into primary or secondary bank in flash.
2. Reboot the stack with new image. See **vsf sequenced-reboot** on page 586

FSU differs from the normal software upgrade in the second step where only the standby of the VSF stack boots, using the upgraded image. The current commander is intact, thus ensuring that the IMs and their ports on commander switch are still available for traffic.

Following is the sequence used by FSU:

1. The user downloads new firmware into the primary or secondary bank and issues the CLI command to perform FSU.
2. The standby switch reboots to the new firmware and boots up as a commander switch.
3. The old commander switch reboots to the new firmware when the VSF link is brought up and joins the other switch as a new standby.
4. Traffic starts passing through the new commander as it comes up. However, the switchover is not stateful. L2 traffic is guaranteed to be within the timeline of 3 secs but not L3.

vsf sequenced-reboot

Syntax

```
vsf sequenced-reboot {primary|secondary}
```

Description

Initiates a reboot of the VSF virtual chassis. First the standby reboots and becomes the commander. Then, the existing commander reboots and becomes the standby.

Example output

```
Switch# vsf

sequenced-reboot      Initiate a reboot of the VSF virtual chassis.
```

```
Switch# vsf sequenced-reboot
primary Reboot from the primary flash image.
secondary Reboot from the secondary flash image.
```

Validation rules for VSF sequenced reboot

Validation	Error/Warning/Prompt
Sequenced reboot is not allowed if the update version is KB.16.02 or lower.	Sequenced reboot is supported only if the updated software version is KB.16.03 or higher.
When software update is in progress, configuration commands will be blocked.	Software update is in progress; configuration commands are not allowed.
If the software image sync to standby is in progress and if `sequenced-reboot` command is executed, the command will fail with an error message.	Software image synchronization to the VSF standby is in progress; try to execute the command later.

Overview

IP Service Level Agreement (IP SLA) is a feature that helps administrators collect information about network performance in real time. With increasing pressure on maintaining agreed-upon Service Level Agreements on Enterprises and ISPs alike, IP SLA serves as a useful tool.

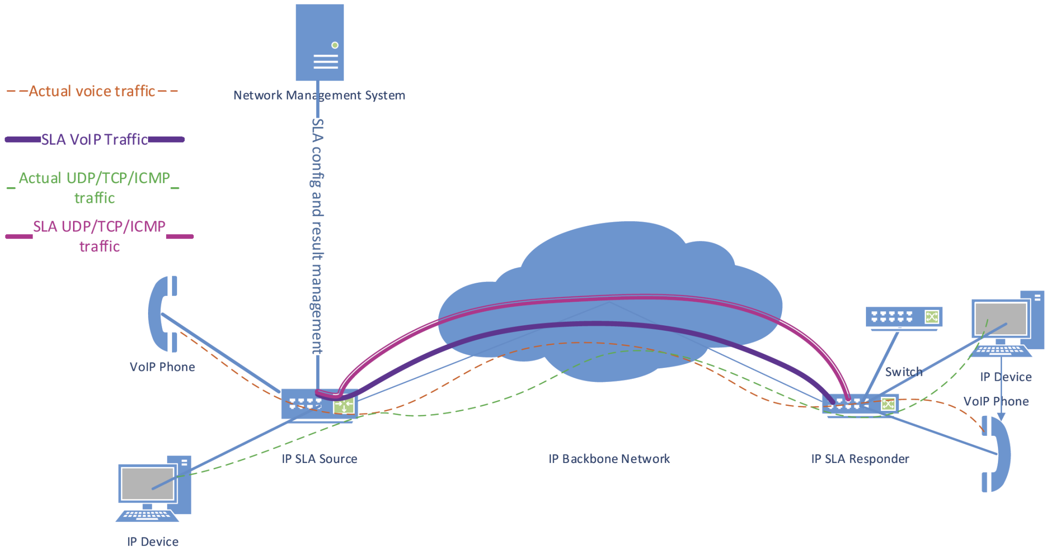

Any IP SLA test involves a source node and a destination node. For all discussions in this document, the source is always an ArubaOS-switch with IP SLA support. As shown in the diagram above, a destination can, in most cases, be any IP-enabled device. For some SLA types that expect a nonstandard response to a test packet, an “SLA responder” must be configured. An “SLA responder” is nothing but an ArubaOS-switch with IP SLA configurations on it that enable it to respond to the test packet.

The IP SLA feature provides:

- Application-aware monitoring that simulates actual protocol packets.
- Predictable measures that aid in ease of deployment and help with assessment of existing network performance.
- Accurate measures of delay and packet loss for time-sensitive applications.
- End-to-end measurements to represent actual user experience.

We support the following SLA types:

- UDP Echo, including connectivity testing of transport layer (UDP) services, Round-Trip-Time (RTT) measurement, one-way delay, and packet loss details.
- ICMP Echo, including connectivity testing, RTT measurement, and packet loss details.
- TCP Connect, including connectivity testing of transport layer (TCP) services, and handshake time measurement.
- DHCP, which measures the round-trip time taken to discover a DHCP Server and obtain a leased IP address from it.
- DNS, which measures the time taken for a DNS resolution. This measures the difference between the time taken to send a request to the DNS server and the time the IP SLA source receives a reply.

- User Datagram Protocol (UDP) Jitter, which measures RTT, one way jitter and one way delays.
- UDP Jitter for VoIP, which measures RTT, one way jitter, one way delays, ICPIF (Impairment Calculated Planning Impairment Factor) and MOS (Mean Opinion Score).

Limitations for IPSLA support on HPE Aruba switches:

- IP SLA is not enabled for IPv6.
- DHCP SLA supports DHCPv4 only.
- IP SLA tests cannot be initiated over OOBM interfaces.
- History results for the configured IP SLAs will not be available after a switchover or a reboot.
- Maximum number of IP SLAs that can be configured varies based on the type of SLA test.
- When there are multiple IP SLAs configured with destination as hostname, the DNS resolution happens serially. There can be a delay in sending the test probe (which will be sent only after successful DNS resolution).
- For TCP Connect SLA type, the four-tuple (source IP/port, destination IP/port) must be unique.
- System clocks between the source and the responder must be synchronized with NTP if One Way Delay parameters have to be calculated for UDP Echo tests.
- Timeout for probes is 3 seconds for all SLA types and is not configurable.
- Transient spikes in RTT occur during the tests (in the source and the responder) if processor usage is high. Consider average result values over a period of time rather than point-in-time results. This is not applicable for UDP Jitter nor Jitter for VoIP.

Entity	Limit
Maximum number of SLAs enabled.	50
Maximum history bucket size per SLA.[1]	50
Number of responders that can be configured.	10

[1] *Not applicable for UDP Jitter and Jitter for VoIP.*

The following are operational restrictions with respect to IP SLA jitter implementation:

- Feature is not supported on 2540.
- Feature is supported only on v3-based platforms (5400R, 3810, 2930F).
- No history results are stored.
- IPSLA Jitter and Jitter for VoIP initiator and responder is only supported on 5400R with v3 modules (noncompatibility mode), 3810, and 2930F switches.
- The maximum number of Jitter responder sessions (UDP Jitter + Jitter For VoIP) supported is 10. The maximum number of Jitter initiator sessions (UDP Jitter + Jitter For VoIP) supported is 5.
- IMC (Intelligent Management Center) supports IP SLA jitter configuration and monitoring on the initiator.
- Measurement of RTT and jitter values is in milliseconds.
- IPv6 SLA for UDP jitter and VoIP is not supported.
- UDP jitter and UDP jitter for VoIP tests are not supported over Tunnel, Trunk, and OOBM interfaces.
- UDP jitter and UDP jitter for VoIP results are not carried forward across failover or a device reboot.
- History bucket size cannot be configured for UDP jitter and VoIP tests. Results are aggregated for the last 25 probes.
- System clocks between the source and the responder must be synchronized with NTP if One Way Delay parameters have to be calculated for UDP Jitter & UDP Jitter for VoIP tests.

- The UDP jitter and UDP jitter for VoIP feature on ArubaOS-Switches has the following limited interoperability with Comware 7 SLA v2 version:
    - One Way packet drops (SD packet loss and DS packet loss) on the Comware Jitter initiator is not reported when interoperating with HPE Aruba Jitter Responder.
- IP SLA responder or initiator implementation is not interoperable with Cisco's IP SLA feature.

How IP SLA works

1. The source originates a test packet to the destination.
2. The destination responds to the test packet, at times embedding the needed information in the response packet.
3. Upon receiving the response, the source calculates the test results based on the timestamp, other packet parameters, and so on.
4. The source stores the results and updates the history records for the SLA.
5. The source reschedules the SLA for the next run.

NOTE

For one-way delay calculations, the IP SLA sender and IP SLA responder must be NTP Time Synchronized.

Configuration commands

[no] ip-sla <ID>

Syntax

```
[no] ip-sla <ID>
```

Description

Configure the IP Service Level Agreement (SLA) parameters. The value of ID can range from 1-255.

Options

clear

Clear history records, message statistics, and threshold counters of particular SLA entry.

dhcp

Configure DHCP as the IP SLA test mechanism.

disable

Disable the IP SLA.

dns

Configure DNS as the IP SLA test mechanism.

enable

Enable the IP SLA.

history-size

Configure the number of history records to be stored for the IP SLA.

icmp-echo

Configure ICMP echo as the IP SLA test mechanism.

monitor

Configure monitoring parameters and respective threshold-action values.

schedule

Configure the start time, stop time, lifetime, and frequency of run for the IP SLA.

tcp-connect

Configure TCP connect as the IP SLA test mechanism.

tos

Configure the Type of Service value to be set in the test packet for the IP SLA.

udp-echo

Configure UDP echo as the IP SLA test mechanism.

On platforms that support Jitter and VOIP, the following options are also provided:

udp-jitter

Configure UDP jitter as the IP SLA test mechanism.

udp-jitter-voip

Configure UDP jitter for VoIP as the IP SLA test mechanism.

ip-sla <ID> clear

Syntax

```
ip-sla <ID> clear
```

Description

Clear history records, message statistics, and threshold counters of a particular SLA entry.

Options

records

Clear history records, message statistics, and threshold counters of particular SLA entry.

[no] ip-sla <ID> history-size

Syntax

```
[no] ip-sla <ID> history-size
```

Description

Configure the number of history records to be stored for the IP SLA. The maximum supported size is 50 and the default value for history-size is 25.

[no] ip-sla <ID> icmp-echo

Syntax

```
[no] ip-sla <ID> icmp-echo [<IP-ADDR> | <HOST-NAME>] [source <IP-ADDR>
| source-interface vlan <VLAN-ID>] [payload-size <SIZE>]
```

Description

Configure ICMP echo as the IP SLA test mechanism. Requires destination address/hostname and source address/vlan id for the IP SLA of ICMP-Echo SLA type.

- **payload-size**

: Value can range from 1-1440. By default, payload-size is not set.

[no] ip-sla <ID> udp-echo

Syntax

```
[no] ip-sla <ID> udp-echo [destination [<IP-ADDR> | <HOST-NAME>]
<PORT-NUM>] [source <IP-ADDR> | <VLAN-ID>] [payload-size <SIZE>]
```

Description

Configure UDP echo as the IP SLA test mechanism. Requires destination address/hostname and source address/VLAN ID for the IP SLA of UDP-Echo SLA type.

- **PORT-NUM**: Value can range from 1024–65535.
- **payload-size**: Value can range from 1-1440. By default, payload-size is not set.

[no] ip-sla <ID> tcp-connect

Syntax

```
[no] ip-sla <ID> tcp-connect [destination [<IP-ADDR> | <HOST-NAME>]
<PORT-NUM>] [source [<IP-ADDR> | <VLAN-ID>] <PORT-NUM>]
```

Description

Configure TCP connect as the IP SLA test mechanism. Requires destination address/hostname and source address/VLAN ID for the IP SLA of TCP connect SLA type. The value of PORT-NUM can range from 1024-65535.

[no] ip-sla <ID> monitor threshold-config

Syntax

```
[no] ip-sla <ID> monitor threshold-config [rtt | srcTodstTime | dstToSrcTime]
threshold-type [immediate | consecutive <COUNT>] threshold-value <UPPER-LIMIT>
<LOWER-LIMIT> action-type [trap | log | trap-log | none]
```

Description

Set upper and lower threshold parameters.

- **threshold-type immediate**: Take action immediately when the monitored parameters cross the threshold upper limit (subsequent notifications for upper thresholds are not generated until the parameter values go lower than the configured lower threshold value).
- **threshold-type consecutive**: Take action after threshold is hit consecutively for number of times.
- **action-type**: Describes action to be taken when the upper threshold is crossed.
- **trap**: Send snmp-trap when configured threshold is hit.
- **log**: Only log the event when configured threshold is hit.
- **trap-log**: Send snmp-trap and log the event when configured threshold is hit.
- **none**: Take no action.

NOTE

The command option threshold-config can be individually set for rtt, srcTodstTime, and dstToSrcTime.

[no] ip-sla <ID> monitor packet-loss

Syntax

```
[no] ip-sla <ID> monitor packet-loss threshold-type [immediate | consecutive
<COUNT>] action-type [trap | log | trap-log | none]
```

Description

Configure threshold-action values when packet loss happens.

- **threshold-type immediate**: Take action immediately when the monitored parameters cross the threshold upper limit (subsequent notifications for upper thresholds are not generated until the parameter values go lower than the configured lower threshold value).
- **threshold-type consecutive**: Take action after threshold is hit consecutively for number of times.
- **action-type**: Describes action to be taken when the upper threshold is crossed.
- **trap**: Send snmp-trap when configured threshold is hit.
- **log**: Only log the event when configured threshold is hit.
- **trap-log**: Send snmp-trap and log the event when configured threshold is hit.
- **none**: Take no action.

[no] ip-sla <ID> monitor test-completion

Syntax

```
[no] ip-sla <ID> monitor test-completion action-type [trap | log | trap-log | none]
```

Description

Configure action to be taken when test gets completed.

- **trap**: Send snmp-trap when configured threshold is hit.
- **log**: Only log the event when configured threshold is hit.
- **trap-log**: Send snmp-trap and log the event when configured threshold is hit.
- **none**: Take no action.

[no] ip-sla <ID> schedule

Syntax

```
[no] ip-sla <ID> schedule [[now | startTime <START-TIME>] [forever | stopTime
<STOP-TIME>
| repetitions <NUM>] [frequency <FREQUENCY>
```

Description

Configure the start time, stop time, lifetime, and frequency of run for the IP SLA. The default value for the frequency of operation is 60 seconds.

[no] ip-sla <ID> tos

Syntax

```
[no] ip-sla <ID> tos <VALUE>
```

Description

Configure the Type of Service value to be set in the test packet for the IP SLA.

- **Valid values**: 0–255

[no] ip-sla responder

Syntax

```
[no] ip-sla responder
```

Description

Configure SLA responder to respond to probe packets.

- **IP address**: local interface IP address
- **port**: takes L4 port numbers.
- **SLA types supported**: udp-echo, tcp-connect, UDP Jitter & Jitter For VoIP.

Show commands

show ip-sla <ID>

Syntax

`show ip-sla <ID>`

Description

Show IP SLA configurations.

Options

history

Show the IP SLA results history.

message-statistics

Show the IP SLA message statistics.

results

Show the IP SLA results for UDP Jitter and UDP Jitter VoIP.

aggregated-results

Show the IP SLA aggregated results for UDP Jitter and UDP Jitter VOIP.

show ip-sla <ID>

```
SLA ID: 1
Status: [Enabled | Admin-disabled | Scheduled | Expired | Running]

SLA Type: [ICMP-echo | tcp-connect | UDP-echo | DHCP | DNS | udp-jitter | voip]

Destination Hostname: www.hp.com
Destination Address : 20.0.0.2
Source Address      : 20.0.0.1
History Bucket Size : 5
TOS: 32
Schedule:
      Frequency (seconds)        : 60
      Life                       : [Forever | 144 seconds]
      Start Time                 : Tue Oct 27 22:12:16 2015
      Next Scheduled Run Time    : Tue Oct 27 22:43:16 2015

Threshold-Monitor is    : Enabled
      Threshold Config: RTT
      Threshold Type   : immediate
      Upper Threshold : 500 ms
     Lower Threshold : 100 ms
     Action Type     : Trap and Log

     Threshold Config: packet-loss
     Threshold Type  : consecutive (5)
     Action Type     : Trap
```

```
    Threshold Config: test-completion
    Action Type: None
```

show ip-sla <ID> history

Syntax

`show ip-sla <ID> history`

Description

Show the IP SLA results history.

show ip-sla <ID> history

```
SLA ID : 1

SLA Type : UDP-Echo

Minimum RTT (ms)         : 1
Maximum RTT (ms)         : 4294967282
Average RTT (ms)         : 3
Total RTT (ms)           : 315
RTT2 (sum of RTT squared): 63681

Start Time                Status   RTT      Description
------------------------ -------- ----      ------------
Mon Jan 1 00:51:28 1990 Failed  -        DMA tail drop detected.
Mon Jan 1 00:51:30 1990 Failed  -        SLA disabled before probe response
arrived.
```

show ip-sla <ID> message-statistics

Syntax

`show ip-sla <ID> message-statistics`

Description

Show the IP SLA message statistics.

show ip-sla <ID> message-statistics

```
SLA ID : 1
Status : Running
SLA Type : UDP-Echo
Destination Address : 10.0.0.2
Source Address : 10.0.0.1
Destination Port : 2000
History Bucket Size : 25
Payload Size : 500
TOS : 0
Messages:
Destination Address Unreachable : 0
Probes Skipped Awaiting DNS Resolution : 0
DNS Resolution Failed : 0
No Route to Target : 0
Internal Error : 0
Local Interface is Down : 0
```

```
No Response from Target : 0
Successful Probes Sent : 3
Probe Response received : 3
Possibly Tail Dropped : 0
```

show ip-sla responder

Syntax

`show ip-sla responder`

Description

Show the IP SLA responder details.

show ip-sla responder

```
SLA type          : UDP-echo
Listening Address: 1.1.1.1
Listening Port    : 5555
```

show ip-sla responder statistics

Syntax

`show ip-sla responder statistics`

Description

Show the IP SLA responder statistics details.

Options

udp-jitter

Show the IP SLA responder statistics for UDP Jitter SLA type.

udp-jitter-voip

Show the IP SLA responder statistics for UDP Jitter VoIP SLA type.

show ip-sla responder statistics

```
IP SLA Responder  : Active
Number of packets received       : 31
Number of error packets received : 0
Number of packets sent           : 0

Recent Sources :
 10.12.80.100 [07:23:49.085 UTC Sun Oct 25 2015] UDP
 10.12.80.100 [07:22:49.003 UTC Sun Oct 25 2015] TCP
 10.12.80.100 [07:20:48.717 UTC Sun Oct 25 2015] TCP
 10.12.80.100 [07:18:48.787 UTC Sun Oct 25 2015] TCP
 10.12.80.100 [07:17:48.871 UTC Sun Oct 25 2015] TCP
```

show tech ip-sla

Syntax

`show tech ip-sla`

Description

Display output of a predefined command sequence used by technical support.

show tech ip-sla

```
switch# sh tech ip-sla

ipslaShowTech

 ======== IP SLA show tech BEGIN ============

 GLOBALS:
 Hash Handle:                        1e7bab20
 Struct Mem Handle for hash:         1e7ba2a8
 Struct Mem Handle for SLA ID LL:    1e7c9430
 Struct Mem Handle for FD List:      1e7bd690
 FastLog Handle:                     dfabf5c
 IPSLACtrl task ID:                  1068091456
 IPSLA Sender ID:                    1068092544
 IPSLA Listener ID:                  1068091840
 Number of enabled SLA's:            1
 SLA ID List Handle:                 1ec1ffd4
 FD ID List Handle:                  0
 Ring Full Counter:                  0

 Details for SLA ID: 1


SLA ID: 1
Status: Running

SLA mechanism: ICMP-Echo

Destination address: 192.168.1.2
Source address: 192.168.1.1
History bucket size: 25
Payload size: 0
TOS: 0
Schedule:
      Frequency (seconds)       : 60
      Life                      : Forever
      Start Time                : Mon Jun 13 10:42:52 2016
      Next Scheduled Run Time   : Mon Jun 13 10:46:52 2016


Threshold-Monitor is  : Enabled
        Threshold Config     : RTT
        Threshold Type       : Immediate
        Upper Threshold      : 10
        Lower Threshold      : 2
        Action Type          : Log



SLA ID: 1
Status: Running

SLA mechanism: ICMP-Echo

Destination address: 192.168.1.2
Source address: 192.168.1.1
```

```
History bucket size: 25
Payload size: 0
TOS: 0
Messages:
       Destination address unreachable        : 0
       Probes skipped awaiting DNS resolution : 0
       DNS resolution failed                  : 0
       No route to target                     : 0
       Internal error                         : 0
       Local interface is down                : 0
       No response from target                : 0
       Successful probes sent                 : 9
       Probe response received                : 9
       Possibly tail dropped                  : 0

Count of Threshold hits:
                RTT               : 0
                packetLoss        : 0


SLA ID: 1

Minimum RTT (ms)          : 1
Maximum RTT (ms)          : 1
Average RTT (ms)          : 1
Total RTT (ms)            : 9
RTT2 (sum of RTT squared): 9

     Start Time                    Status       RTT                 Description
     ----------                    ------       ---                 -----------
Tue Jun 14 10:43:12 2016          Passed          1
Mon Jun 13 10:39:05 2016          Passed          1
Mon Jun 13 10:40:05 2016          Passed          1
Mon Jun 13 10:41:05 2016          Passed          1
Mon Jun 13 10:42:05 2016          Passed          1
Mon Jun 13 10:42:52 2016          Passed          1
Mon Jun 13 10:43:52 2016          Passed          1
Mon Jun 13 10:44:52 2016          Passed          1
Mon Jun 13 10:45:52 2016          Passed          1

 ICMP ID hash walk:


 ========== IP SLA show tech END ==============

 ======== IP SLA Server show tech BEGIN ============
Responder not active
IP SLA Responder: Inactive

 ======== IP SLA Server show tech END ============


=== The command has completed successfully. ===
```

clear ip-sla responder statistics

Syntax

clear ip-sla responder statistics *<SLA-TYPE>* *<LOCAL-IP-ADDR>* *<LOCAL-PORT-NUM>*
source *<SOURCE-IP-ADDR>*

Description

Clear IP SLA responder statistics for either UDP jitter or VoIP UDP jitter.

Command context

`config`

Parameters

<SLA-TYPE>

Specifies the SLA type.

`udp-jitter`

Selects standard UDP jitter.

`udp-jitter-voip`

Selects UDP VOIP jitter.

<LOCAL-IP-ADDR>

Specifies the local interface IP address

<LOCAL-PORT-NUM>

Specifies the local interface port number. Range: *1024 to 65535*.

<SOURCE-IP-ADDR>

Specifies the Source IP address.

Examples

Clear IP SLA responder statistics for UDP jitter:

```
switch(config)# clear ip-sla responder statistics udp-jitter 1.1.1.1 1100 source
1.1.1.2
```

Validation rules

Validation	Error/Warning/Prompt
Enabling SLA without configuring SLA type.	Cannot enable IP SLA, no valid source/destination configured.
IP address given for source or destination is multicast or broadcast.	Invalid IP address.
Configure the SLA type with a source IP which is configured in the same switch.	Destination IP cannot be configured as the same as one of the local interface IP addresses.
Configure threshold with invalid value.	Invalid threshold count value. For threshold type 'Immediate', count must be 1 and for 'Consecutive', count must be greater than or equals to 2.
Configure threshold value for 'PacketLoss' or 'TestCompletion'	Configuration is not applicable when threshold is configured for 'PacketLoss' or 'TestCompletion'.

Table Continued

Validation	Error/Warning/Prompt
Configure threshold type for TestCompletion.	Configuration is not applicable when threshold is configured for 'TestCompletion'.
Configure schedule with proper end time with a frequency which is out of end time.	Invalid endtime. Endtime is not enough to run the tests for configured frequency and repetitions.
Configuring ‘srcTodstTime’ or ‘dstTosrcTime’ threshold configuration for ‘icmp-echo’ or ‘tcp-connect’.	Invalid threshold configuration for configured SLA type.
Enabling the IP SLA which is already in enabled state.	IP SLA is already enabled.
Disabling the IP SLA which is already in disabled state.	IP SLA is already disabled.
Show IP SLA history of un-configured SLA.	IP SLA is not configured for this ID.
Enable more number (currently decided 50 as limit) of IP SLA.	Maximum number of enabled IP SLAs at a time is limited to 50.
Removing IP SLA type/tos/history size/ schedule/ threshold configuration with un-configured value.	IP SLA configuration does not exist.
Configuring scheduler with a frequency value which is not satisfying the condition frequency > number of packets per probe * packet interval.	Frequency value is insufficient to configure the scheduler.
Scheduler already configured and try to configure SLA type with a value of ‘number of packets per probe’ and ‘packet interval’ which is not satisfying the condition frequency > number of packets per probe * packet interval.	Number of packets/packet interval is insufficient to configure IP SLA type.
Configuring IP SLA with invalid values.	Invalid configuration for IP SLA.
Change the IP SLA configuration when the SLA is enabled.	Configuration changes not allowed when IP SLA is enabled.
When IP address vs port number configured for an SLA is already in use	Error: Socket for configured address, port is already in use, choose different port number
When Source IP address given in SLA configuration is not configured in the switch	Error: Source IP address is not configured in switch
Invalid SLA ID given in show command	Error: Invalid IP SLA ID
Configure SLA more than allowed limit	Warning: The maximum number of IP SLAs allowed is 50.
Configure Responder more than allowed limit	Error: IP SLA Responder configurations reached max limit. No more configurations accepted.

Table Continued

Validation	Error/Warning/Prompt
Configure inter-packet interval when number of packets to be sent out is one.	Error Not applicable as Number of packets to be sent out is 1.
Upper threshold value is less than lower threshold value.	Error: Upper threshold value X is less than lower threshold value Y.
Configure schedule with start time greater than stop time.	Error: Stop time must be greater than start time.
Configure schedule with past stop time.	Error: Stop time must be greater than current time.
Configure schedule with invalid frequency value.	Error: Schedule frequency is out of range. Valid range is 5 to 604800.
Configuring history size with invalid value.	Error: IP SLA History size is out of range. Valid range is 1–50.
Configuring SLA type with invalid payload value.	IP SLA Payload value is out of range. Valid range is 1–1440.
Configuring SLA type with invalid port number.	Invalid port number. Valid range is 1024 to 65535.
Configure the IPSLA parameters without configuring SLA type.	No valid IP SLA type configuration found.
Configuring the responder with existing details.	IP SLA Responder with same configuration exist.
Configure management VLAN as source VLAN.	Error: Not allowed to configure management VLAN as source interface.
Enabling IP SLA without required configuration parameters.	Configuration is incomplete to enable the entry.
Configure IPSLA UDP Jitter/VoIP session when an OpenFlow custom match mode instance is already enabled.	Cannot enable IP SLA: Initiator cannot co-exist with OpenFlow custom Cannot enable IP SLA: Initiator cannot co-exist with OpenFlow custom
Configure UDP Jitter/VoIP IPSLA Initiator session with a source IP same as tunnel overlay VLAN	Cannot enable IP SLA: The source VLAN cannot be a tunnel overlay VLAN.
Configure UDP Jitter/VoIP IPSLA Initiator session with a source IP same as a service tunnel endpoint.	Cannot enable IP SLA: The source VLAN cannot be a service tunnel endpoint.

Event log messages

Cause

Event	Message
User adds IP SLA endpoint configuration.	I 10/28/15 02:47:12 05020 ipsla: The IP SLA 1 of SLA Type: UDP-Echo, Source IPv4 Address: 10.0.0.1, Destination IPv4 Address: 10.0.0.5, Destination Port: 54563 added.
User removes the endpoint configuration.	I 10/28/15 02:47:12 05021 ipsla: The IP SLA 1 of SLA Type: UDP-Echo, Source IPv4 Address: 10.0.0.1, Destination IPv4 Address: 10.0.0.5, Destination Port: 54563 removed.
When the SLA state changes (can be either system initiated or done by the user)	I 10/28/15 01:42:22 05027 ipsla: IP SLA 1 state changed to Expired.I 10/28/15 01:42:22 05021 ipsla: IP SLA 1 state changed to Enabled.I 10/28/15 01:42:22 05021 ipsla: IP SLA 1 state changed to Scheduled.I 10/28/15 01:42:22 05021 ipsla: IP SLA 1 state changed to Admin-disabled.
User configures a responder	I 10/28/15 01:42:22 05025 ipsla: IP SLA responder configured for SLA Type: TCP-Connect, Listen Address: 10.0.0.7, Listen Port: 38425
User removes a responder	I 10/28/15 01:42:22 05026 ipsla: IP SLA responder removed for SLA Type: TCP-Connect, Listen Address: 10.0.0.7, Listen Port: 38425
SLA test results cross configured threshold	I 10/28/15 01:42:22 05022 ipsla: IP SLA 1, threshold is crossed. Monitored Param: RTT, Threshold Type: immediate, Upper threshold: 500, Lower threshold: 100, Action Type: Trap and Log. Actual Threshold: 600
User adds DNS IP-SLA configuration	I 08/09/16 02:47:12 05029 ipsla: The IP SLA 1 of SLA Type: DNS, Name server IPv4 Address: 10.0.0.1, Target Hostname: a.hp.com added
User removes DNS IP-SLA configuration	I 08/09/16 02:47:12 05030 ipsla: The IP SLA 1 of SLA Type: DNS, Name server IPv4 Address: 10.0.0.1, Target Hostname: a.hp.com removed.
The packet loss threshold for the SLA has reached	I 08/09/16 02:47:12 05023 ipsla: The IP SLA 1 of SLA Type: DNS, Packet loss is observed. Threshold type: immediate, Action type: Trap and Log
The SLA test has completed and the threshold config for test completion is present	I 08/09/16 02:47:12 05024 ipsla: The IP SLA 1, Probe for SLA Type: UDP-Echo, Source IPv4 Address:1.1.1.2, Destination IPv4 Address:1.1.1.3, Destination Port:3000 Completed

Table Continued

Event	Message
TCP-Connect Duplicate configuration is present	W 08/09/16 02:47:12 05028 ipsla: Duplicate configuration detected; the IP SLA 10 configuration is the same as the IP SLA 1. For TCP, 4-tuple combination has to be unique.
User adds DHCP SLA configuration	I 08/09/16 02:47:12 05031 ipsla: The IP SLA 1 of SLA Type: DHCP, Vlan: 1 added.
User removes DHCP SLA configuration	I 08/09/16 02:47:12 05032 ipsla: The IP SLA 1 of SLA Type: DHCP, Vlan: 1 removed.
DHCP SLA test has completed and the threshold config for test completion is present	I 08/09/16 02:47:12 05033 ipsla: The IP SLA 1, Probe for SLA Type: DHCP, Source Vlan : 1 Completed

Interoperability

NOTE

Packet loss is expected when H3C TCP-CONNECT source, with a frequency less than or equal to 10 ms, interoperates with HPE TCP-CONNECT responder.

IP SLA UDP Jitter and Jitter for VoIP

Overview

The UDP Jitter and Jitter for VoIP SLA types enable the user to assess the suitability of the network for voice & video related traffic. These SLA's basically calculate parameters like RTT, one way delay, one way positive and negative Jitter etc.

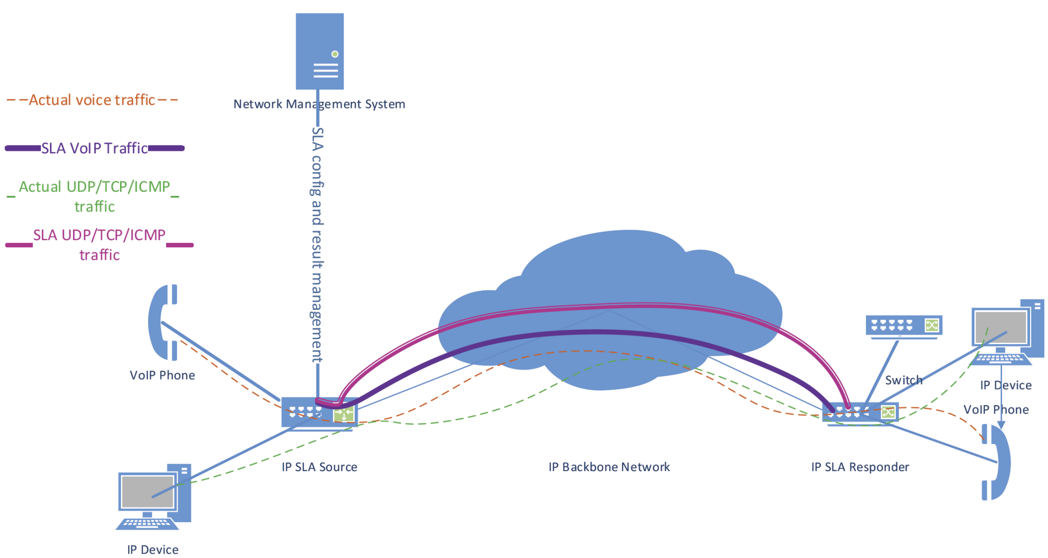

The above diagram shows a typical deployment, where voice & video traffic are exchanged between branch offices of an enterprise over the backbone network. Assessment of the network readiness is always helpful for hosting such services. Parameters like RTT, Jitter and one way delay are a good indicator of network health which assist a network administrator to diagnose latency related issues in the network. VoIP traffic is generally sensitive to delays in the network.

Jitter stands for inter-packet delay variance. If the inter-packet delay increases between successive probe packets, jitter is said to be positive. If the inter-packet delay decreases, jitter is said to be negative. Positive jitter values are undesirable for a network as they indicate increased latencies. A value of 0 jitter is desirable.

Significance of jitter

Consider a media player which plays video streams from a server. Assume that packets take 1 second in flight to reach the media player. This means the moment a user requests a video from the server, the very first packet will arrive after one second and successive packets will be sent immediately (ideally). In real world scenarios, intermediate node latencies, different return paths for different packets and network congestion can contribute to varying delays. To counter such effects, packets are buffered at the media player. The amount of packet buffering needed can be derived from the jitter values.

NOTE

The above analogy is applicable for other voice & video services and can be a good measure to assess the possibility for hosting a service on a network.

Solution components

IP SLA responder

This device receives IP SLA probe packets from a configured initiator, timestamps the frame at a pre-defined location in the packet upon receipt and sends the same frame back to the initiator.

IP SLA initiator

This device initiates IP SLA probe packets to multiple destinations each with a certain user configurable packet content and periodicity.

The initiator timestamps the frame at a pre-defined location before sending the frame out to the configured destinations and re-timestamps the frame at a different location once it receives the same back from the responder.

IP SLA measurement engine

This is an application running on the initiator. It processes response frames received from the IP SLA responder and computes one-way delay, jitter and RTT based on the timestamps present in the packet.

This application aggregates this computed information across multiple probe samples and stores this for consumption by an NMS via SNMP or via the device CLI.

It also supports asynchronous user configurable threshold breach notification to an NMS (via SNMP Traps).

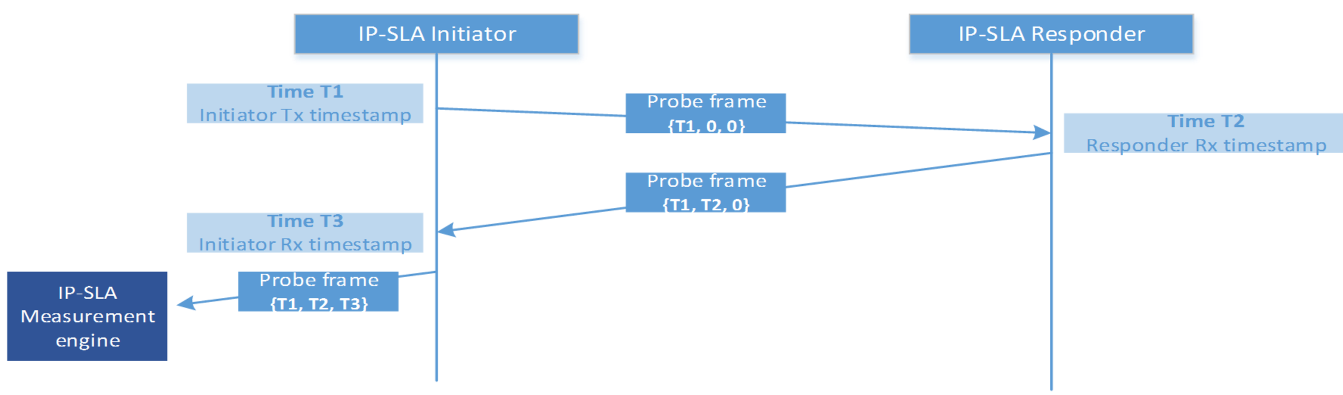

Aruba Central Network Management Solution Overview

The Aruba Central network management solution, a software-as-a-service subscription in the cloud, provides streamlined management of multiple network devices. Aruba switches are able to talk to Aruba Central and utilize cloud-based management functionality. Cloud-based management functionality allows for the deployment of network devices at sites with no IT personnel (branch offices, retail stores, and so forth). The communication channel used to connect the devices with the cloud portal is outside the control of end users. It adheres to corporate standards like the use of firewalls.

This feature provides:

- Zero-touch provisioning
- Network Management/Remote monitoring
- Events/alerts notification
- Configuration
- Firmware management

Table 27: *Features supported by Aruba Central*

Configuration	
VLANS	Create VLANs IP address assignment Tag/Untag
Ports	Admin status [Up/Down] PoE [Enable/Disable for an interface]
User Management	Create/delete Operator User Create Manager user IP address assignment to Uplink Vlan [Static to DHCP and vice versa] Name server configuration

RCD
Establish a remote console session to Switch

Monitoring
Events
Statistics

Firmware Management
Update the switch firmware

Troubleshooting
Execute troubleshooting commands from central

Restricted commands in switch when connected to Central

boot
recopy
erase
reload
startup-default
upgrade-software
setup
delete
reboot
restore
menu
write memory
amp-server

Limited Management Interface on switch

WebUI
REST
SNMP
TR-69
Menu

You can provision the switch in Aruba Activate. For more information about provisioning, See *Aruba Networks and AirWave Switch Configuration Guide*.

LED Blink feature

Central connectivity loss is indicated by LEDs. If connectivity is broken and Aruba-Central is enabled, the USR/FDX and Locator LEDs will blink. The LEDs will stop blinking once the switch is connected back to Central.

Configuration commands

aruba-central

Syntax

```
aruba-central {enable | disable | support-mode {enable | disable}}
```

Description

Configure Aruba Central server support. When enabled, and when a server web address has been obtained using Aruba Activate, the system will connect to an Aruba Central server. The system will obtain configuration updates and most local configuration commands will be disabled. This mode is enabled by default.

Enter support mode to enable all configuration commands. Normally, when the system is connected to an Aruba Central server, the configuration is updated from that server and most local configuration commands are disabled.

Support mode enables those commands for use in troubleshooting problems. Support mode is disabled by default. When the system is not connected to Aruba Central server, the full command set is enabled for local configuration.

Restrictions

- Switch communication to Aruba Central is not supported via OOBM.
- Aruba-central is not supported in FIPS switches and it will be disabled by default.
- Aruba-central is not supported in Stack switches and it will be disabled by default.

CAUTION

To avoid broadcast storm or loops in your network while configuring ZTP, do not have redundant links after you complete ZTP and Airwave registration. Authorize the new switch and then push the Golden Configuration template from Airwave.

Example

Enable Aruba Central server support

```
switch(config)# aruba-central enable
```

Disable Aruba Central server support

```
switch(config)# aruba-central disable
```

Enter support mode to enable all CLI configuration commands

```
switch(config)# aruba-central support-mode enable

This mode will enable all CLI configuration commands, including those normally
reserved by the Aruba Central
service.
Use of this mode may invalidate the configuration provisioned through Aruba
Central server.
Continue (y/n)?
```

Show commands

show aruba-central

Syntax

`show aruba-central`

Description

Show Aruba Central server information.

show aruba-central

```
switch# show aruba-central

 Configuration and Status - Aruba Central
  Server URL            : https://hpsw-jenkins-soa-qa-
build-1404-250.test.pdt1.arubathena.com/ws
  Connected             : Yes
  Mode                  : Managed
  Last Disconnect Time : Tue Jun 14 16:01:15 2016
```

Overview

Auto device detection

The command `device-profile` enables the user to define profiles and configure the associations of profiles to each device type. By creating a device profile, parameters will be defined for a connection interface by device type. To configure each parameter under a profile name, a context level is provided.

The command `device-profile name` *<PROFILE NAME>* configures for the default values. The default value is permissible when no user-defined profile is created.

To associate each device type with a device profile, a context level is created which authorizes the user to enable or disable the profile by device-type. Only the device type `aruba-ap` is supported.

Rogue AP isolation

The command `rogue-ap-isolation` configures each device and blocks, logs, or allows a rogue AP when detected. The command enables or disables rogue AP isolation.

The command `clear rogue-ap-isolation` is provided to clear the detected rogue AP device MAC address.

Show commands are provided to display the configuration and status of the profiles. Another show command will display the list of rogue APs detected.

Jumbo frames on a device port

Configure jumbo frame support for the device port. Jumbo frames are not permissible by default.

Enabling jumbo frame support in a profile might affect other ports with different profiles. When a profile has jumbo frame enabled and is applied to any port, all other ports that are members of any VLAN listed in the profile will also have jumbo frame support.

Applicable products

Aruba 2530 Switch (JL070A, J9772A, J9773A, J9774A, J9775A, J9776A, J9777A, J9778A, J9779A, J9780A, J9781A, J9782A, J9783A, J9853A, J9854A, J9855A, J9856A)
HPE 2620 Switch (J9624A, J9625A, J9623A, J9627A, J9626A)
Aruba 2920 (J9726A, J9727A, J9728A, J9729A, J9836A)
Aruba 2930F (JL253A, JL254A, JL255A, JL256A, JL259A, JL260A, JL261A, JL262A, JL263A, JL263A, JL264A)
HPE 3800 (J9573A, J9574A, J9575A, J9576A, J9584A)
Aruba 3810M (JL075A, JL071A, JL073A, JL076A, JL072A, JL074A)
HPE 5406v2zl Switch Series (J9866A, J8697AX, J9642A, J9533A, J9539A, J9447A, J8699A)
Aruba 5406R Switch Series (J9850A, JL002A, JL003A, JL095A, J9821A)
Aruba 5406zl Switch Series (J9821A, J9822A)
HPE E5406 zl Switch (J8697A)
Aruba 5412R Switch Series (JL001A, J9822A, J9851A)
HPE 5412zl Switch Series (J9643A, J9532A, J9540A, J9448A, J8700A, J9809A)
HPE E5412 zl Switch (J8698A)

Configuration commands

allow-jumbo-frames

Syntax

`allow-jumbo-frames`

Description

Configure jumbo frame support for the device port. Jumbo frames are not enabled by default.

Enabling jumbo frame support in a profile affects other ports with different profiles. When a profile has jumbo frames enabled and is applied to any port, all other ports that are members of any VLAN listed in the profile will also have jumbo frame support.

Validation rules

Validation	Error/Warning/Prompt
Invalid jumbo command.	Invalid input.
If jumbo frame support is configured on a VLAN for which the device profile had overridden the configuration, display the existing warning.	This configuration change will be delayed because a device profile that enables jumbo frame support is applied to a port in this VLAN.

Default AP Profile

Creates a user-defined profile.

The profile name is a valid character string with the maximum permissible length of 32. The default profile is named `default-ap-profile` and cannot be modified.

The default configuration parameters may be modified using the command `device-`*`<PROFILE NAME>`* `default-ap-profile` . Up to four different profiles may be configured.

The `[no]` command removes the user-defined profiles.

device-profile

From within the configure context:

Syntax

`device-profile` *`<PROFILE-NAME>`* *`<DEVICE-TYPE>`*

Description

Create port configuration profiles and associate them with devices. When a configured device type is connected on a port, the system will automatically apply the corresponding port profile. When the device is disconnected, the profile is removed after a 2 minute delay. Connected devices are identified using LLDP.

Options

<PROFILE-NAME>

Specify the name of the profile to be configured.

<DEVICE-TYPE>

Specify an approved device-type to configure and attach a profile to.

Parameters

allow-jumbo-frames

Configure jumbo frame support for the device port.

untagged-vlan <*VLAN-ID*>

Configure this port as an untagged member of specified VLAN.

tagged-vlan <*VLAN-LIST*>

Configure this port as a tagged member of the specified VLANs.

cos <*COS-VALUE*>

Configure the Class of Service (CoS) priority for traffic from the device.

ingress-bandwidth <*PERCENTAGE*>

Configure ingress maximum bandwidth for the device port.

egress-bandwidth <*PERCENTAGE*>

Configure egress maximum bandwidth for the device port.

poe-max-power <*WATTS*>

Configure the maximum PoE power for the device port (in watts).

poe-priority

Configure the PoE priority for the device port.

Usage

```
[no] device-profile name <PROFILE-NAME>

[no] device-profile type <DEVICE>
```

Associating a profile with a device

Associate a profile with a device by using the command `device-profile device-type <DEVICE-NAME> associate <PROFILE-NAME>`. Associated devices can be Aruba Access Points, ArubaOS-Switch Switches, scs-wan-cpe, or association can be by the device profile.

The feature is disabled by default.

device-profile type-device

Syntax

```
device-profile type [aruba-ap | aruba-switch | scs-wan-cpe |
device-name <DEVICE-NAME> associate <PROFILE-NAME> | enable | disable]

no device-profile type [aruba-ap | aruba-switch | scs-wan-cpe |
device-name <DEVICE-NAME> associate <PROFILE-NAME> | enable | disable]
```

Description

Associates the device profile with the type of device by identity.

The `no` form of this command removes the device profile from the device type.

Command context

`config`

Parameters

associate *<PROFILE-NAME>*

Selects the profile name associated with the device-type.

enable

Selects the profile of the device being enabled.

disable

Selects the profile of the device being disabled.

Usage

- The command `device-profile type aruba-ap enable` enables an Aruba-AP within the device profile.
- Device Name is defined the same as Device Identity.

Configuring the rogue-ap-isolation command

Used to configure the `rogue-ap-isolation` command. A block/log option may be configured for when a rogue AP is identified by the switch. The block/log option may be enabled or disabled. The default action is to block a rogue AP.

The `whitelist` command is used to configure any specific MAC addresses excluded from the rogue AP list. The whitelist configuration is saved in the configuration. The whitelist supports 128 MACs.

The `[no]` form the command is used to remove the MAC address individually by specifying the MAC.

rogue-ap-isolation

Within the configure context:

Syntax

`rogue-ap-isolation`

Description

Configure rogue AP isolation and rogue AP Whitelist MAC addresses for the switch. When enabled, the system detects the MAC address of rogue access points and takes the specified action for traffic or from that address. The whitelist is used to add MAC addresses of approved access points to the whitelist.

Options

action

Configure the action to take for rogue AP packets. Actions available are enable, disable, block, log, and whitelist.

block

Block and logs traffic to or from any rogue access points.

log

Log traffic to or from any rogue access points.

enable

Enable the rogue AP Isolation.

disable

Disable the rogue AP Isolation.

whitelist <*MAC-ADDRESS*>

Configures rogue AP Whitelist MAC addresses for the switch. This option is used to add MAC addresses of approved access points to the whitelist.

<*MAC-ADDR*>

Specify the MAC address of the device to be moved from the Rogue AP list to the whitelist.

Usage

```
rogue-ap-isolation [enable | disable]
```

```
rogue-ap-isolation action [log | block]
```

```
[no] rogue-ap-isolation whitelist <MAC-ADDRESS>
```

Show commands

show device-profile

Syntax

Within the configure context:

```
show device-profile
```

Description

Show device profile configuration and status.

config

Show the device profile configuration details for a single, or all, profiles.

status

Show currently applied device profiles.

Usage

```
show device-profile config <PROFILE-NAME>

show device-profile status
```

show device-profile config

```
Switch# Show device-profile config
 Device Profile Configuration
  Configuration for device profile : default-ap-profile
   untagged-vlan    : 1
   tagged-vlan      : None
   ingress-bandwith : 100%
   egress-bandwith  : 100%
   cos       : 0
   speed-duplex     : auto
   poe-max-power    : 33W
   poe-priority     : High
   allow-jumbo-frames: Enabled

  Configuration for device profile : profile1
   untagged-vlan    : 10
   tagged-vlan      : 40,50,60
```

```
   ingress-bandwith : 10%
   egress-bandwith  : 95%
   cos       : 4
   speed-duplex     : auto-10
   poe-max-power    : 20W
   poe-priority     : Low
```

show device-profile config profile1

```
Switch# Show device-profile config profile1
 Device Profile Configuration
  Configuration for device profile : profile1
   untagged-vlan    : 10
   tagged-vlan      : 40,50,60
   ingress-bandwith : 10%
   egress-bandwith  : 95%
   cos       : 4
   speed-duplex     : auto-10
   poe-max-power    : 20W
   poe-priority     : Low
```

show command device-profile status

Syntax

`show device-profile [config | status]`

Description

Displays the device-profile configuration or device-profile status.

Options

config

Show device profile configuration details for a single profile or all profiles.

status

Show currently applied device profiles status.

show device-profile status

```
Switch# show device-profile status

 Device Profile Status
  Port          Device Type      Applied Device Profile
  ----          -----------      ----------------------
  5             aruba-ap         profile1
  10            aruba-ap         profile1
```

Show rogue-ap-isolation

Syntax

`show rogue-ap-isolation`

Description

Show rogue access point information.

Options

whitelist

Show rogue access point whitelist information.

Usage

`show rogue-ap-isolation whitelist`

show rogue-ap-isolation

```
Switch# show rogue-ap-isolation

 Rogue AP Isolation
  Rogue AP Status : Enable
  Rogue AP Action : Block
   Rogue AP MAC                    Neighbor Device
   -----------------               -----------------
   11:22:33:44:55:66               00:12:34:56:67:89
   aa:bb:cc:dd:ee:ff               00:98:45:56:67:89
```

show rogue-ap-isolation whitelist

```
Switch# show rogue-ap-isolation whitelist

 Rogue AP Whitelist Configuration
  Rogue AP MAC
  -----------------
  11:22:33:44:55:66
  aa:bb:cc:dd:ee:ff
```

Overview

This feature supports secure communication between ArubaOS-Switches and the Aruba mobility controller (VPN concentrator) for Network Management Server (AirWave) traffic. The switch also provides the necessary support for Zero Touch Provisioning (ZTP) by establishing a secure tunnel between an ArubaOS-Switch and the Network Management Server (AirWave) which are provided for by a DHCP Server or Activate.

IPsec ensures that communication between ArubaOS-Switch-based switches and AirWave Server (management traffic) is protected by establishing a secure channel between the switches and the Aruba VPN Controller (connected to AirWave server).

Applicable products

Aruba Switch 3800 Series (J9573A, J9574A, J9575A, J9576A, J9584A, J9585A, J9586A, J9587A, J9588A)
Aruba 3810M Switch Series (JL071A, JL072A, JL073A, JL074A, JL075A,JL076A)
Aruba 5400R zl2 Switch Series (J8698A, J8700A, J9823A-J9824A, J9825A, J9826A, J9868A, J9447A, J9448A)
Aruba 5406R Switch Series (JL002A, JL003A, JL095A,J9850A)
Aruba 5412R Switch Series (J9851A, JL001A)

AirWave details

ZTP discovers switches in their respective management stations (AirWave) during initial boot up which enables the automatic configuration and management of the switches.

- ZTP checks if AirWave details are provided along with IP via DHCP.
  - If AirWave details are missing from DHCP, ZTP will try to connect to Activate to receive AirWave details.

IPsec for Management Traffic

IPsec provides support for Zero Touch Provisioning in deployment scenarios less restrictive than private LANs. This enables switches to be configured and managed automatically without administrator intervention. In a deployment scenario where a switch and AirWave are located in different branches connected via an untrusted public network (the Internet), communication between the switch and the AirWave server is now protected.

IPsec Tunnel Establishment

- IPsec tunnel for AirWave is auto-configured. The switch decides to create IPSec tunnel only when an Aruba controller IP is present in the device before establishing the connection to AirWave.
- If the controller IP is not provided, the switch will try to establish a direct connection to AirWave.
- If the controller IP is present, the ArubaOS-Switch auto configures and initiates an IPsec tunnel interface. Once the tunnel is established, the Aurba controller provides an inner IP which the switch will then use as source IP to send any AirWave bound traffic. The switch then creates a static route to AirWave with the IPsec tunnel interface as the gateway.

IPSec Tunnel Failures

The following behaviors can cause an IPSec tunnel creation failure:

- Time

The time in the switch has to be valid and correct. Time issues have been observed on the Aruba 2930F and Aruba 2920 24G Switches.

- Authentication

  The switch MAC addresses for both members must be added to the Aruba controller whitelist.

- Controller IP

  The controller IP must be reachable from the switch.

- Static Route

  There must not be any conflicting static route in the system for the AirWave IP configured.

AirWave IP after discovery

AirWave IP and Aruba Controller IP (either from the Activate Server or from a DHCP server) are established and auto configured in an IPSEC-IPv4 Tunnel. Once received, the IPsec tunnel is auto configured and established to send AirWave traffic securely. The Aruba Controller provides an inner-ip to the switch which then can communicate with AirWave.

Configuring the Aruba controller

On the Aruba Controller, configure via CLI:

Procedure

1. Disable control-plane-security (CPSEC).

   ```
   control-plane-security
   ```

   ```
   no cpsec-enable
   ```

2. Add the switch MAC address to whitelist and for authentication.

   ```
   whitelist-db rap add mac-address <Switch Mac add> ap-group default [remote-ip <ip
   address for Switch>]
   ```

   ```
   local-userdb add username <Switch Mac Add> password <switch mac add>
   ```

3. Add an IP address pool that can be assigned to switch after tunnel creation. IP range must be in the same subnet through which AirWave is reachable from Controller.

   ```
   ip local pool "ipsec" 2.0.0.100 2.0.0.255
   ```

4. Create access lists that permit AirWave traffic and assign them to ap-roles.

   ```
   ip access-list session hpe-acl
   ```

   ```
   any any tcp 22 permit
   ```

   ```
   any any tcp 443 permit !
   ```

   ```
   user-role ap-role
   ```

   ```
   access-list session hpe-acl !
   ```

5. View the whitelist.
   show whitelist-db cpsec

   ```
   (host) #show whitelist-db cpsec
             ap-group <ap_group>
             ap-name <ap_name>
             cert-type {factory-cert|switch-cert}
             mac-address <name>
             page <num>
             start &lt;offset>
             state {approved-ready-for-cert|certified-factory-cert|
               unapproved-factory-cert|unapproved-no-cert}
   ```

show whitelist-db cpsec-status

```
(host) #show whitelist-db cpsec-status
(host) #show whitelist-db rap
          apgroup <rap-group>
          apname <rap-name>
          fullname <rap-fullname>
          long
          mac-address <mac-address>
          page <page-number>
          start &lt;offset>
```

show whitelist-db rap-status

```
(host) #show whitelist-db rap-status
```

show ip interface brief

```
(Aruba7210) #show ip interface brief

Interface    IP Address / IP Netmask       Admin   Protocol   VRRP-IP   (VRRP-Id)
vlan 1     172.16.0.254 / 255.255.255.0      up      up         none      (none)
vlan 30      30.30.30.2 / 255.255.255.0      up      up         none      (none)
vlan 17        17.0.0.5 / 255.255.255.0      up      up         none      (none)
loopback     unassigned / unassigned         up      up
```

show vlan

```
(Aruba7210) #show vlan

VLAN CONFIGURATION
------------------
VLAN  Description  Ports              AAA Profile
----  -----------  -----              -----------
1     Default      GE0/0/2-0/5 Pc0-7  N/A
17    VLAN0017     GE0/0/1            N/A
30    VLAN0030     GE0/0/0            N/A
amp ip is : 30.30.30.1
```

show running-config | begin "0/0/0"

```
#show running-config | begin "0/0/0"
(Aruba7210) #show running-config | begin "0/0/0"
interface gigabitethernet 0/0/0
        description "GE0/0/0"
        trusted
        trusted vlan 1-4094
        switchport access vlan 30

interface gigabitethernet 0/0/1
        description "GE0/0/1"
        trusted
        trusted vlan 1-4094
        switchport access vlan 17

interface gigabitethernet 0/0/2
        description "GE0/0/2"
        trusted
        trusted vlan 1-4094
```

```
interface gigabitethernet 0/0/3
        description "GE0/0/3"
        trusted
        trusted vlan 1-4094

interface gigabitethernet 0/0/4
        description "GE0/0/4"
        trusted
        trusted vlan 1-4094

interface gigabitethernet 0/0/5
        description "GE0/0/5"
        trusted
        trusted vlan 1-4094

interface vlan 1
        ip address 172.16.0.254 255.255.255.0
        ipv6 address 2001::1/64

interface vlan 30
        ip address 30.30.30.2 255.255.255.0

interface vlan 17
        ip address 17.0.0.5 255.255.255.0

no uplink wired vlan 1
uplink disable
ip nexthop-list pan-gp-ipsec-map-list

crypto isakmp policy 20
  encryption aes256

crypto isakmp policy 10001

crypto isakmp policy 10002
  encryption aes256
  authentication rsa-sig


crypto isakmp policy 10003
  encryption aes256

crypto isakmp policy 10004
  version v2
  encryption aes256
  authentication rsa-sig

crypto isakmp policy 10005
  encryption aes256

crypto isakmp policy 10006
  version v2
  encryption aes128
  authentication rsa-sig

crypto isakmp policy 10007
  version v2
  encryption aes128

crypto isakmp policy 10008
  version v2
  encryption aes128
  hash sha2-256-128
```

```
  group 19
  authentication ecdsa-256
  prf prf-hmac-sha256

crypto isakmp policy 10009
  version v2
  encryption aes256
  hash sha2-384-192
  group 20
  authentication ecdsa-384
  prf prf-hmac-sha384

crypto isakmp policy 10012
  version v2
  encryption aes256
  authentication rsa-sig

crypto isakmp policy 10013
  encryption aes256

crypto ipsec transform-set default-ha-transform esp-3des esp-sha-hmac
crypto ipsec transform-set default-boc-bm-transform esp-aes256 esp-sha-hmac
crypto ipsec transform-set default-1st-ikev2-transform esp-aes256 esp-sha-hmac
crypto ipsec transform-set default-3rd-ikev2-transform esp-aes128 esp-sha-hmac
crypto ipsec transform-set default-rap-transform esp-aes256 esp-sha-hmac
crypto ipsec transform-set default-aes esp-aes256 esp-sha-hmac
crypto dynamic-map default-rap-ipsecmap 10001

  version v2
  set transform-set "default-gcm256" "default-gcm128" "default-rap-transform"

crypto dynamic-map default-dynamicmap 10000
  set transform-set "default-transform" "default-aes"

crypto map GLOBAL-IKEV2-MAP 10000 ipsec-isakmp dynamic default-rap-ipsecmap
crypto map GLOBAL-MAP 10000 ipsec-isakmp dynamic default-dynamicmap
crypto isakmp eap-passthrough eap-tls
crypto isakmp eap-passthrough eap-peap
crypto isakmp eap-passthrough eap-mschapv2

ip local pool "ipsec" 30.30.30.100
```

AirWave Controller IP configuration commands

aruba-vpn type

From within the configure context:

Syntax

`aruba-vpn type amp peer-ip <ip-addr>`

`no aruba-vpn type amp peer-ip <ip-addr>`

Description

Configure the Aruba VPN type, peer IP address.

The no form of the command removes the aruba-vpn type statement from the configuration.

Options

amp

Configure the AirWave Management Platform (AMP) server.

type

Configure the controller IP.

<ip-addr>

IP address of the VPN.

Usage

```
switch(config)# aruba-vpn type
switch(config)# aruba-vpn type amp
switch(config)# aruba-vpn type amp peer-ip
```

Show commands

show aruba-vpn

Syntax

`show aruba-vpn type <VPN-TYPE>`

Description

Show Aruba-VPN configuration information.

Switch(config)# show aruba-vpn

```
show aruba-vpn
 Aruba VPN details
  Aruba VPN Type              : amp
  Aruba VPN Peer IP           : 171.0.0.3
  Aruba VPN Config Status     : Configured
  Aruba VPN tos               : Value from IPv4 header
  Aruba VPN ttl               : 64
```

show aruba-vpn type amp

```
show aruba-vpn type amp

Aruba VPN details
  Aruba VPN Type              : amp
  Aruba VPN Peer IP           : 2.2.2.2
  Aruba VPN Config Status     : Configured
  Aruba VPN tos               :   32
  Aruba VPN ttl               :   54
```

show ip route

Syntax

`show ip route`

Description

Show the IP route.

show ip route

```
                              IP Route Entries
  Destination        Gateway          VLAN Type      Sub-Type   Metric     Dist.
  ------------------ ---------------- ---- --------- ---------- ---------- -----
  0.0.0.0/0          192.168.20.31    1    static               250        1
  2.0.0.25/32*        aruba-vpn             connected            1          0
  2.0.0.199/32**       aruba-vpn             static               1          1
  127.0.0.0/8        reject                static               0          0
  127.0.0.1/32       lo0                   connected            1          0
  192.168.20.0/24    DEFAULT_VLAN     1    connected            1          0
```

*The inner IP received from the Aruba Controller.

**Static Route for AirWave IP. Added automatically by the switch after tunnel establishment.

show interfaces tunnel aruba-vpn

Syntax

`show interfaces tunnel aruba-vpn`

Description

Auto-configured tunnel interface before creating IPSec. The tunnel ID is auto generated and to avoid conflict with user generated tunnel interface, the tunnel id is always the max tunnel supported by the switch + 1.

aruba-vpn

Display the configuration and status details of aruba-vpn tunnel.

brief

Display brief configuration and status for all tunnels.

Usage

`show interfaces tunnel aruba-vpn`

`show interfaces tunnel brief`

`show interfaces [tunnel] [<TUNNEL-LIST> | <TUNNEL-NAME> | brief | type]`

show interfaces tunnel aruba-vpn

```
Aruba-3810M-24G-PoEP-1-slot(config)# show interfaces tunnel aruba-vpn
 Tunnel Configuration :
  Tunnel                 : tunnel-129
  Tunnel Name            : aruba-vpn-tunnel
  Tunnel Status          : Enabled
  Source Address         : 17.0.0.30
  Destination Address    : 17.0.0.5
  Mode                   : IPsec IPv4
  TOS                    : Value from IPv4 header
  TTL                    : 64
  IPv6                   : Disabled
  MTU                    : 1280
```

```
Current Tunnel Status :
  Tunnel State                  : Up
  Destination Address Route     : 17.0.0.0/24
  Next Hop IP                   : 17.0.0.5
  Next Hop Interface            : vlan-1
  Next Hop IP Link Status       : Up
  Source Address                : Configured on vlan-1
  IP Datagrams Received         : 9732
  IP Datagrams Transmitted      : 13129
```

show interfaces tunnel brief

```
Aruba-3810M-24G-PoEP-1-slot(config)#  show interfaces tunnel brief
 Status - Tunnel Information Brief
  Tunnel                    : tunnel-129
  Mode                      : IPsec IPv4
  Source Address            : 17.0.0.30
  Destination Address       : 17.0.0.5
  Configured Tunnel Status : Enabled
  Current Tunnel State      : Up
```

show ip counters tunnel aruba-vpn

Syntax

`show ip counters tunnel aruba-vpn`

Description

Show IP counters for a tunnel.

Options

aruba-vpn

Show counters for aruba-vpn tunnel.

ipv4

Show IPv4 only.

ipv6

Show IPv6 only.

<TUNNEL-ID>

Show specified tunnel only.

Usage

`show ip counters tunnel ipv4`

`show ip counters tunnel ipv6`

`show ip counters tunnel <TUNNEL-ID>`

show ip counters tunnel aruba-vpn

```
sh ip counters tunnel
Address Family :    IPv4
Interface      :    Tunnel 129
```

```
 IP In Datagrams Received                                : 2439
 IP In Octets Received                                   : 362736
 IP In Datagrams Broadcast Received                      : 0
 IP In Octets Broadcast Received                         : 0
 IP In Datagrams Multicast Received                      : 0
 IP In Octets Multicast Received                         : 0
 IP In Datagrams Discarded Datagram Header Error         : 0
 IP In Datagrams Discarded No Route                      : 0
 IP In Datagrams Discarded Invalid Address               : 0
 IP In Datagrams Discarded Unknown Protocol              : 0
 IP In Datagrams Discarded Truncation                    : 0
 IP In Datagrams Discarded Processing Error              : 0
 IP In Datagrams Forwarding Required                     : 0
 IP In Datagrams Delivery to Protocols Successful        : 2439
 IP Datagrams Reassembly Required                        : 0
 IP Datagrams Reassembly Successful                      : 0
 IP Datagrams Reassembly Failed                          : 0
 IP Out Datagrams Transmitted                            : 2514
 IP Out Octets Transmitted                               : 1197348
 IP Out Datagrams Broadcast Transmitted                  : 0
 IP Out Octets Broadcast Transmitted                     : 0
 IP Out Datagrams Multicast Transmitted                  : 0
 IP Out Octets Multicast Transmitted                     : 0
 IP Out Datagrams Discarded Processing Error             : 0
 IP Out Datagrams Forwarded                              : 0
 IP Out Datagrams Transmit Requests from Protocols       : 2509
 IP Out Datagrams Fragmentation Required                 : 0
 IP Out Datagrams Fragmentation Successful               : 5
 IP Out Datagrams Fragmentation Failed                   : 0
 IP Out Datagrams Fragments Created                      : 0

Address Family :    IPv6
Interface      :    Tunnel 129
 IP In Datagrams Received                                : 0
 IP In Octets Received                                   : 0
 IP In Datagrams Broadcast Received                      : 0
 IP In Octets Broadcast Received                         : 0
 IP In Datagrams Multicast Received                      : 0
 IP In Octets Multicast Received                         : 0
 IP In Datagrams Discarded Datagram Header Error         : 0
 IP In Datagrams Discarded No Route                      : 0
 IP In Datagrams Discarded Invalid Address               : 0
 IP In Datagrams Discarded Unknown Protocol              : 0
 IP In Datagrams Discarded Truncation                    : 0
 IP In Datagrams Discarded Processing Error              : 0
 IP In Datagrams Forwarding Required                     : 0
 IP In Datagrams Delivery to Protocols Successful        : 0
 IP Datagrams Reassembly Required                        : 0
 IP Datagrams Reassembly Successful                      : 0
 IP Datagrams Reassembly Failed                          : 0
 IP Out Datagrams Transmitted                            : 0
 IP Out Octets Transmitted                               : 0
 IP Out Datagrams Broadcast Transmitted                  : 0
 IP Out Octets Broadcast Transmitted                     : 0
 IP Out Datagrams Multicast Transmitted                  : 0
 IP Out Octets Multicast Transmitted                     : 0
 IP Out Datagrams Discarded Processing Error             : 0
 IP Out Datagrams Forwarded                              : 0
 IP Out Datagrams Transmit Requests from Protocols       : 0
 IP Out Datagrams Fragmentation Required                 : 0
 IP Out Datagrams Fragmentation Successful               : 0
 IP Out Datagrams Fragmentation Failed                   : 0
 IP Out Datagrams Fragments Created                      : 0
```

show crypto-ipsec sa

Syntax

```
show crypto ipsec sa
```

Description

Show crypto-IPsec statistics.

Switch(config)# show crypto-ipsec sa

```
Aruba-2930F-48G-4SFPP# show crypto ipsec sa

  Crypto IPSec Status
  Interface               : 1
  Source Address          : 192.168.20.14
  Destination Address     : 171.0.0.3
  Source Port             : 0              Destination Port    : 0
  SPI                     : 3767553536
  Encapsulation Protocol  : ESP
  Encryption              : AES            Hash                : SHA1
  PFS                     : 0              PFS Group           :
  Mode                    : tunnel
  Key Life                : 3600           Remaining key Life : 3303
  Key Size                : 0              Remaining key Size : 0
  Interface               : 2
  Source Address          : 171.0.0.3
  Destination Address     : 192.168.20.14
  Source Port             : 0              Destination Port    : 0
  SPI                     : 4173307552
  Encapsulation Protocol  : ESP
  Encryption              : AES            Hash                : SHA1
  PFS                     : 0              PFS Group           :
  Mode                    : tunnel
  Key Life                : 3600           Remaining key Life : 3301
  Key Size                : 0              Remaining key Size : 0
```

Usage

```
show crypto ipsec statistics
```

show running-configuration

Syntax

```
show running-configuration
```

NOTE

IP route or tunnel interface will not be displayed in show run as they are auto created.

show running-configuration

```
show running-configuration

; JL254A Configuration Editor; Created on release #WC.16.02.0000x
; Ver #0e:01.b3.ef.7c.5f.fc.6b.fb.9f.fc.f3.ff.37.ef:ab

hostname "Aruba-2930F-48G-4SFPP"
```

```
module 1 type jl254a
snmp-server community "public" unrestricted

vlan 1
   name "DEFAULT_VLAN"
   untagged 1-52
   ip address dhcp-bootp
   exit

amp-server ip 2.0.0.199 group "aw_group" folder "fold" secret "secr"
aruba-vpn type amp peer-ip 171.0.0.3
```

Overview

Every client is associated with a user role. User roles associate a set of attributes for authenticated clients (clients with authentication configuration) and unauthenticated clients, applied to each user session. User roles must be enabled globally.

NOTE: Local user roles are only supported when running YA software.

Examples of user roles are:

- Employee = All access
- Contractor = Limited access to resources
- Guest = Browse Internet

Each user role determines the client network privileges, frequency of reauthentication, applicable bandwidth contracts, and other permissions. There are a maximum of 32 administratively configurable user roles available with one predefined and read-only user role called **denyall**.

A user role consists of optional parameters such as:

- Captive portal profileSpecifies the URL via:
  - **captive-portal profile**

    or
  - Vendor Specific Attribute (VSA). RADIUS: HP **HP-Captive-Portal-URL** = <http://...>
- Ingress user policy

  L3 (IPv4 and/or IPv6) ordered list of Classes with actions, with an implicit deny all for IPv4 and IPv6.
- Reauthentication period

  The time that the session is valid for. The default is 0 unless the user role is overridden. The default means that the reauthentication is disabled.

  NOTE: Reauthentication period is required to override the default of 0.
- Untagged VLAN (either VLAN ID or VLAN-name)

  VLAN precedence order behavior:
  - If configured, untagged VLAN specified in the user role (VSA Derived Role, UDR, or Initial Role).
  - Statically configured untagged and/or tagged VLANs of the port the user is on.

Operational notes

- When user roles are enabled, all users that are connecting on ports where authentication is configured will have a user role applied. User role application happens even if the user fails to authenticate. If the user cannot be authenticated, the “Initial Role” will be applied to that user.
- The user role may be applied in one of two ways:
  - Vendor Specific Attribute (VSA)

Type: RADIUS: Hewlett-Packard-Enterprise

Name: HPE-User-Role

ID: 25

Value: <myUserRole>

The RADIUS server (ClearPass Policy Manager) determines application of the VSA Derived Role. The role is sent to the switch via a RADIUS VSA. The VSA Derived Role will have the same precedence order as the authentication type (802.1x, WMA).

- User Derived Role (UDR

)The User Derived Role is part of Local MAC authentication (LMA) and is applied when user roles are enabled and LMA is configured.

  UDR will have the same precedence as LMA. Precedence behavior of the authentication types will be maintained, (802.1x -> LMA -> WMA (highest to lowest)).

Restrictions

- User roles cannot be enabled when BYOD redirect, MAC authentication failure redirect, or enhanced web-based authentication are enabled.
- Web-based authentication is not supported on the same port with other authentication methods when user roles are enabled.
- `show port-access <AUTH-TYPE>` commands are not supported when user-roles are enabled. The command `show port-access clients [detail]` is the only way to see authenticated clients with their associated roles.
- `aaa port-access auth <port> control` commands are not supported when user roles are enabled.
- `unauth-vid` commands are not supported when user roles are enabled.
- `auth-vid` commands are not supported when user roles are enabled.

Limitations for web-based authentication

Cannot be combined with other authentication types on same port.

Limitations for LMA

Reauthentication period and captive portal profile are not supported.

Error messages

Action	Error message
Attempting to enable BYOD Redirect when user roles are enabled.	BYOD redirect cannot be enabled when user roles are enabled.
Attempting to enable MAFR when user roles are enabled.	MAC authentication failure redirect cannot be enabled when user roles are enabled.
Attempting to enable enhanced web-based authentication when user roles are enabled.	Enhanced web-based authentication cannot be enabled when user roles are enabled.
Attempting to enable web-based authentication when other authentication types are enabled for the same port, and user roles are enabled.	Web-based authentication cannot be enabled with other authentication types on this port when user roles are enabled.
`switch (config)# show port-access mac-based clients`	User roles are enabled. Use `show port-access clients` to view client information.

Table Continued

Action	Error message
`switch (config)# aaa port-access authenticator e8 control autho`	802.1x control mode, **Force Authorized/ Unauthorized**, cannot be set when user roles are enabled.
`Attempting to enable local user role when MAFR, BYOD, or EWA are enabled.`	User roles cannot be enabled when BYOD redirect, MAC authentication failure redirect, or enhanced web-based authentication are enabled.

Captive-portal commands

Overview

The Captive Portal profile defines the web address that a user is redirected to for Captive Portal authentication. If the url is blank, a RADIUS VSA will be used.

NOTE

There is a predefined profile called **use-radius-vsa** that is already configured to use the RADIUS VSA.

Two captive portal profiles are supported:

- Predefined and read-only
  - Predefined and read-only profile name is `use-radius-vsa`.
- Customized

[no] aaa authentication captive-portal profile

Syntax

`[no] aaa authentication captive-portal profile` *`<PROFILE-STR>`* `[url` *`<URL-STR>`*`]`

Description

Create a captive-portal profile. Profiles are used in user roles to direct the user to a designated captive portal server. When the profile includes a web address, that web address is always used to contact the server. When no web address is specified, it is obtained from the RADIUS VSA.

NOTE

A profile does not have to be pre-existing in the switch for it to be configured to a user role.

Options

profile

Configure a captive portal profile.

<PROFILE-STR>

Configure a captive portal profile string 64 characters long.

url

Configure the captive portal server web address.

<URL-STR>

Configure the captive portal server web address string.

Usage

```
Switch# aaa authentication captive-portal profile <NAME>

Switch# aaa authentication captive-portal profile <NAME> url <URL>
```

Validation rules

Validation	Error/Message/Prompt
Attempts made to remove a nonexisting profile will return an error: `switch# no aaa authentication captive-portal profile NON_EXISTING_PROFILE`	`Captive portal profile NON_EXISTING_PROFILE not found.`
When including the configured web address after the web address parameter: `[no] aaa authentication captive-portal profile myCaptivePortalProfile url http://myCPPM.local/guest/captive_portal_login.php`	`Invalid input: http://blablabla.com`
A profile name with invalid syntax produces an error: `Switch# aaa authentication captive-portal-profile “this is an invalid name”`	`#aaa authentication captive-portal-profile “this is an invalid name” Invalid character ' ' in name.`
When trying to modify a profile that is predefined, `switch# aaa authentication captive-portal-profile name use-radius-vsa`	Captive portal profile `use-radius-vsa` is read only and cannot be modified
A profile name that is too long produces an error: `switch# aaa authentication captive-portal-profiletest342...;ldklsdjflkdsjflk`	The name must be fewer than 64 characters.
When attempting to configure more than the number of admin configured profiles,`switch# aaa authentication captive-portal-profile profileNumber2`	No more captive portal profiles may be created.

Policy commands

Overview

These commands create a context that may be used to classify the policy. From the existing `policy` command, a new policy type called **user** was added. The new actions are specific to **policy user**:

- redirect
- permit
- deny

NOTE

Only L3 classes (IPv4 and IPv6) are currently supported.

The user policy includes “implicit deny all rules” for both IPv4 and IPv6 traffic.

policy user

Syntax

```
policy user <POLICY-NAME>
```

Description

Create and enter newly created user policy context.

Usage

```
Switch (config)# policy user employee
```

[no] policy user

Syntax

```
[no] policy user <POLICYNAME>
```

Description

Delete and remove specified user policy from switch configuration.

Operating notes

- The user policy will include implicit “deny all” rules for both IPv4 and IPv6 traffic.
- ipv4 or ipv6 classes must specify source address as *any*. Specifying host addresses or subnets will result in the following error message:

  ```
  Switch (policy-user)# class ipv4 class25 action priority 0
  User policies cannot use classes that have a source IP address specified.
  ```
- *permit* and *deny* are mutually exclusive.
- *ip-precedence* and *dscp* are mutually exclusive.

Usage

```
switch (config)# no policy user employee
```

policy resequence

Syntax

```
policy resequence <POLICYNAME> <START><INCREMENT>
```

Description

Resequence classes and remarks configured within specified user policy. The usage shows resequencing classes and remarks within user policy “employee” starting at 200 and incrementing by 2.

Usage

```
Switch (config)# policy user employee 200 2
```

Commands in the policy-user context

Create classes inside of the **policy** context before you apply actions to them.

(policy-user)# class

Within the **policy-user** context:

Syntax

```
(policy-user)# [no] [<SEQUENCE-NUMBER>] class ipv4 | ipv6 <CLASS-NAME> [action
permit | deny | redirect captive portal] | [action dscp | ip—precedence <CODEPOINT
| PRECEDENCE>] [action priority <PRIORITY>] | [action rate-limit kbps <RATE>]
```

Description

Associate a class with ACL or QoS actions for this policy.

Options

Options

deny

Deny all traffic.

DSCP

Specify an IP DSCP.

IP-precedence

Specify the IP precedence.

permit

Permit all traffic.

priority

Specify the priority.

rate-limit

Configure rate limiting for all traffic.

redirect

Specify a redirect destination.

Usage

```
Switch(policy-user)# class ipv6 employeeIpv6Http action deny
Switch(policy-user)# class ipv4 http action redirect captive-portal
Switch(policy-user)# class ipv4 dnsDhcp action permit
```

User role configuration

aaa authorization user-role

Syntax

```
aaa authorization user-role [enable | disable| [initial-role <ROLE-STR>] |[name
<ROLE>]]
```

Description

Configure user roles. A user role determines the client network privileges, the frequency of reauthentication, applicable bandwidth contracts, along with other permissions. Every client is associated with a user role or the client is blocked from access to the network.

Options

enable

Enable authorization using user roles.

disable

Disable authorization using user roles.

initial-role

The default initial role “denyall” is used when no other role applies. If a client connects to the switch and does not have a user role associated, then the initial role is used. Any role can be configured as initial role using this option.

The initial role may be assigned if:

- captive-portal profile is configured with a web address, but the Captive Portal VSA is sent from RADIUS
- captive-portal profile is configured to use the RADIUS VSA but no Captive Portal VSA is sent.
- captive-portal feature is disabled when the captive-portal profile is referenced in the applied user role to the client.
- The user role feature is enabled with RADIUS authentication, but no user role VSA is returned.
- User role does not exist.
- Not enough TCAM resource available.
- Access-Reject from RADIUS.
- User role VSA is sent along with invalid attributes.
- RADIUS not reachable.
- VLAN configured on the user role does not exist.
- Captive Portal profile does not exist.
- User policy configured on the user role does not exist.
- Reauthentication period is enabled (nonzero) in the user role for LMA.
- Captive Portal profile is included in the user role for LMA.

name *<NAME-STR>*

Create or modify a user-role. Role name identifies a user-role. When adding a user-role, a new context will be created. The context prompt will be named “user-role” (user-role)#.

Usage

```
Switch# aaa authorization user-role enable
Switch# aaa authorization user-role disable
Switch# aaa authorization user-role name <ROLE1>

Switch# [no] aaa authorization user-role enable
Switch# [no] aaa authorization user-role name <ROLE1>

Switch# aaa authorization user-role initial-role <ROLE1>

Switch# aaa authorization user-role name <MYUSERROLE> policy <MYUSERPOLICY>

Switch# aaa authorization user-role name <MYUSERROLE> captive-portal-profile
<MYCAPTPORTPROFILE>

Switch# aaa authorization user-role name <MYUSERROLE> vlan-id <VID>

Switch# aaa authorization user-role name <MYUSERROLE> reauth-period <0-999999999>
```

Error log

Scenario	Error Message
If the user tries to delete a user-role configured as the initial role	User role *<INITIAL_ROLE_NAME>* is configured as the initial role and cannot be deleted.
If the user attempts to configure more than the number of administrator configured roles	`#aaa authorization user-role name roleNumber33`. No more user roles can be created.
If the user enters a role name that is too long	`switch# aaa authorization user-role test342....jflkdsjflk`. The name must be fewer than 64 characters long.
If the user enters a role name with invalid syntax	`switch# aaa authorization user-role name "this is an invalid name"`. Invalid character ' ' in name.
If the user tries to delete a nonexisting user-role	User role *<NON_EXISTING_ROLE_NAME>* not found.
`Switch# aaa authorization user-role name <DENYALL>`	User role *<DENYALL>* is read only and cannot be modified.

captive-portal-profile

From within the **user-role** context:

Syntax

```
captive-portal-profile <PROFILE_NAME>
```

Description

Assigns a captive portal profile to the user role. The predefined captive portal profile, `use-radius-vsa`, indicates that the redirect web address must be sent via RADIUS.

To clear a captive portal profile from the user role, use the [no] version of the command.

policy

From within the **user-role** context:

Syntax

```
policy <POLICY_NAME>
```

Description

Assigns a user policy to the user role. To clear a policy from the user role, use the [no] version of the command.

NOTE

Modification of the user policy, or class contained in a user policy, will force users consuming that user policy via a user role to be deauthenticated.

reauth-period

From within the user-role context:

Syntax

```
reauth-period <VALUE>
```

Description

Set the reauthentication period for the user role. Use [0] to disable reauthentication. For RADIUS-based authentication methods, it will override the RADIUS session timeout. It also overrides any port-based reauth-period configuration with the exception that LMA does not support a reauth-period.

Options

<VALUE>

Valid values are 0 – 999,999,999; a required configuration in user roles and it defaults to 0.

(user-role)# reauth-period 100

Set the reauthentication value for the current user role:

```
(user-role)# reauth-period 100
```

(user-role)# reauth-period 0

0 is used to disable reauthentication, and it is the default value.

```
(user-role)# reauth-period 0
```

Validation rules

Validation	Error/Warning/Prompt
`(user-role)# reauth-period 10000000`	`Invalid input: 1000000000000000000`

VLAN commands

NOTE

The VLAN must be configured on the switch at the time the user role is applied. Only one of VLAN-name or VLAN-ID is allowed for any user role.

Modification of the VLAN will force users assigned to that VLAN via a user role to be deauthenticated.

vlan-id

From within the user-role context:

Subcommand syntax

```
vlan-id <VLAN-ID>
```

Description

Create a VLAN with id VLAN-ID.

Use the `[no]` version of the command when clearing the VLAN-ID from the user role:

Usage

```
(user-role)# no vlan-id
```

vlan-name

From within the **user-role** context:

Subcommand syntax

```
vlan-name <VLAN-NAME>
```

Description

Create a VLAN with the name VLAN-NAME. Only one of VLAN-NAME or VLAN-ID is allowed for any user role.

Use the `[no]` version of the command when clearing the VLAN from the user role, by name:

Usage

```
(user-role)# no vlan-name
```

vlan-id 100

```
(user-role)# vlan-id 100
```

vlan-name vlan100

```
(user-role)#vlan-name VLAN100
```

VLAN range commands

This command is executed from a global configuration context.

VLANs specified by VLAN-ID-LIST

Syntax

```
[no] vlan <VLAN-ID-LIST>
```

Description

Creates VLANs specified by the VLAN-ID-LIST and returns to the global configuration context. Use the [no] version of the command to delete the VLANs specified by the VLAN-ID-LIST.

Examples

```
config# vlan 2-15
config# vlan 5,10,13-20,25
config# no vlan 2-10
config# no vlan 2,5,15-18,25
```

VLANs specified by VLAN-ID-LIST and tag specified ports specified by PORT-LIST

Syntax

```
[no] vlan <VLAN-ID-LIST> tagged <PORT-LIST>
```

Description

Creates VLANs specified by the VLAN-ID-LIST and tags the ports specified by the PORT-LIST to the VLAN-ID-LIST. If VLANs already exist, the tagging of ports specified by the PORT-LIST is performed.

Use the [no] version of the command to remove the tagged PORT-LIST from a range of VLANs specified by the VLAN-ID-LIST. After command execution, CLI returns to the global configuration context.

Examples

```
config# vlan 2-15 tagged A1-A20
config# vlan 5,10,13-20,25 tagged A1-A5,L2,L5-L10
config# vlan 2-20 tagged all
config# no vlan 2-15 tagged A1-A5
config# no vlan 5,10,13-20 tagged A1-A5,L6
```

Applying a UDR

UDR can be used to assign user roles locally (that is, without RADIUS). LMA has been extended to allow applying a user role to a MAC address, MAC group, MAC mask, or MAC OUI.

aaa port-access local-mac apply user-role

Syntax

```
[no] aaa port-access local-mac apply user-role <Role-Name> [ mac-oui <MAC-OUI> |
mac-mask <MAC-MASK> |mac-addr <MAC-ADDR> | mac-group <MAC-GROUP-NAME>]
```

Description

Apply user roles.

Options

mac-addr

To apply user role with MAC address.

mac-group

To apply user role with MAC group.

mac-mask

To apply user role with MAC Mask.

mac-oui

To apply user role with MAC OUI.

Usage

```
[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-oui <MAC-OUI>]

[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-mask <MAC-MASK>]

[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-addr <MAC-ADDR>]

[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-group <MAC-GROUP-NAME>]
```

Show commands

show captive-portal profile

Syntax

```
show captive-portal profile
```

Description

Show Captive Portal profile configuration.

show captive-portal profile

```
(config)# show captive-portal profile

 Captive Portal Profile Configuration
  Name : use-radius-vsa
  Type : predefined
  URL  :

  Name : myCaptivePortalProfile
  Type : custom
  URL  : http://mycppm.local/guest/captive_portal_login.php
```

show user-role

Syntax

```
show user-role [<ROLE-NAME>] [detailed]
```

Description

Show users role configuration.

Options

<ROLE-NAME>

Show user roles by role-name.

***<ROLE-NAME>* detailed**

Show user roles in detail by role-name.

show user-role

```
Switch# show user-role

 User Roles

  Enabled      : <Yes/No>
  Initial Role : denyall

  Type          Name
  ----------    ------------
  local         Employee
  local         Guest
  predefined    denyall
```

show user-role <ROLE-NAME>

```
Switch# show user-role captivePortalwithVSA

User Role Information

Name                                 : captivePortalwithVSA
  Type                                 : local
  Reauthentication Period (seconds) : 0
  Untagged VLAN                        : 610
  Captive Portal Profile               : use-radius-vsa
  Policy                               : cppolicy
```

show user-role detailed

The example shows how to configure user roles to use Clearpass as a Captive Portal. The Captive Portal URL is specified in a RADIUS VSA.

```
Switch# show user-role captivePortalwithVSA detailed

User Role Information
   Name                                 : captivePortalwithVSA
   Type                                 : local
   Reauthentication Period (seconds) : 0
   VLAN                                 : 610
   Captive Portal Profile               : use-radius-vsa
     URL                                : (use RADIUS VSA)
   Policy                               : cppolicy

Statements for policy "cppolicy"
policy user "cppolicy"
     10 class ipv4 "cppm" action permit
     20 class ipv4 "steal" action redirect captive-portal
     30 class ipv4 "other" action permit
```

```
   exit

Statements for class IPv4 "cppm"
class ipv4 "cppm"
     10 match tcp 0.0.0.0 255.255.255.255 1.0.9.15 0.0.0.0 eq 80
     20 match tcp 0.0.0.0 255.255.255.255 1.0.9.15 0.0.0.0 eq 443
   exit

Statements for class IPv4 "steal"
class ipv4 "steal"
     10 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80
     20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 443
   exit

Statements for class IPv4 "other"
class ipv4 "other"
     10 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 53
     20 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 67
     30 match icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
```

show port-access clients

Syntax

`show port-access clients [detailed]`

Description

Use this command to display the status of active authentication sessions.

show port-access clients

```
Port Access Client Status

Port  Client Name   MAC Address   IP Address      User Role         Type VLAN
----- ------------- ------------- --------------- ----------------- ------
1/A18 001517581ec4  001517-581ec4 10.108.1.201    ixia1             MAC   108
A7                  000c29-5121fc n/a             denyall           LOCAL
A8                  000c29-d12996 n/a             myrole            LOCAL 42
```

show port-access clients detailed

```
Switch (config)# show port-access clients detailed

 Port Access Client Status Detail
  Client Base Details :
   Port            : 1/A18              Authentication Type : mac-based
   Client Status   : authenticated      Session Time        : 11 seconds
   Client Name     : 001517581ec4       Session Timeout     : 60 seconds
   MAC Address     : 001517-581ec4
   IP              : 10.108.1.201

 User Role Information
   Name                               : ixia1
   Type                               : local
   Reauthentication Period (seconds)  : 60
   Untagged VLAN                      : 108
   Tagged VLANs                       :
   Captive Portal Profile             :
```

```
   Policy                              : policyIxia1

Statements for policy "policyIxia1"
policy user "policyIxia1"
     10 class ipv4 "classIxia1" action rate-limit kbps 11000
   exit

Statements for class IPv4 "classIxia1"
class ipv4 "classIxia1"
     10 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
```

Per-user tunneled node

Aruba per-user tunneled node (PUTN) is similar to the MAS equivalent per-port tunneled node where all ingress traffic is tunneled through a specified switch port to a controller. Aruba PUTN gives ArubaOS-Switch the ability to tunnel traffic per client through Generic Routing Encapsulation (GRE) from an interface on a switch (tunneled-node-port) to an Aruba controller (tunneled-node-server).

Once PUTN is enabled, the Aruba controller provides a centralized security policy, authentication, and access-control. The decision to tunnel client traffic is based on the user-role. User-role authorizations are given permission to redirect traffic to an Aruba controller when the PUTN status is UP. A secondary role, provided by the authentication subsystem (DCA), when present in the user-role authorizations, notifies the PUTN and provides a secondary role.

PUTN, combined with CPPM/LMA policies, can be used to indicate per client whether that clients traffic is to be tunneled to a controller or forwarded locally.

For example when a PC Phone is deployed to a network, the administrator can use PUTN to route the PC traffic through additional security and allow the VoIP traffic to move freely through the wired network.

Flow

- Authenticate user.
- Apply user role to authenticated user.
- Redirect user traffic to controller
- Apply secondary user-role to user traffic on controller.

Limitations

- 1,024 users with redirect policy per switch (or stack).
- 32 users supported per port.
- Tunneled and nontunneled users are not allowed to belong to same VLAN.
- Tunneled node is not supported on port-channel.

Tunneled node server prerequisites

Commands necessary to configure a tunneled node server for a tunneled node ArubaOS-switch:

- `switch(config)#tunneled-node-server`
- `switch(tunneled-node-server)# controller-ip <IP-ADDR>`
- Optional: `switch(tunneled-node-server)# backup-controller-ip <IP-ADDR>`
- Optional: `switch(tunneled-node-server)# keepalive interval <Integer>`
- `switch(tunneled-node-server)# mode role-based`
- `switch(tunneled-node-server)# enable`

These commands are available in detail in the ArubaOS-Switch Configuration Guide at **http://support.arubanetworks.com/**.

NOTE

Per-port tunneled-node encapsulates incoming packets from end-hosts in GRE packets and forwards them to an Aruba Mobility Controller to be processed further. The Aruba Mobility Controller, upon receiving the GRE packets, strips the GRE header and further processes the packet for additional purposes such as authentication, stateful firewall, and others.

Cluster mode

If the controller is in a cluster mode, see the Mobility Master Configuration managed controller information in the ArubaOS-Switch User Guide at **%3C?xm-replace_text%20http://support.arubanetworks.com/?%3E**

Note

- PUTN works for Standalone Controller or with the Aruba 8.x controller cluster.

tunneled-node-server-redirect

Syntax

```
tunneled-node-server-redirect [secondary-role <ROLE-NAME>]

no tunneled-node-server-redirect [secondary-role <ROLE-NAME>]
```

Description

This command is used to instruct the switch to redirect traffic for a particular user to the user-based tunnel.

The `no` form of this command configures traffic redirect to the user-based tunnel. Secondary role is the new user role that will be applied to the tunneled traffic by the controller.

Command context

`user-role`

Parameters

secondary-role *<ROLE-NAME>*

Specifies the secondary role applied on the user traffic by the controller.

Example

User role configuration example on TN switch. The tunneled-node-server-redirect attribute instructs the switch to redirect all traffic with user-role “testrole” to the controller. The secondary-role “authenticated” specified with the redirect attribute should be configured and present on the controller. The switch sent VLAN (client VLAN) has to be present on the controller.

```
class ipv4 "testclass"
     10 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
     20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
policy user "testpolicy"
     10 class ipv4 "testclass" action permit
   exit
aaa authorization user-role name "testrole"
   policy "testpolicy"
   vlan-id 100
   tunneled-node-server-redirect secondary-role "authenticated"
   exit
```

Show the tunneled-node-server status for all users.

```
switch-PoEP# show tunneled-node-users all

PORT     MAC-ADDRESS      TUNNEL-STATUS   SECONDARY-USERROLE      FAILURE-REASON
1        000ffe-c8ce92    UP              authenticated
5        082e5f-263518    UP              authenticated
```

show user-role

Syntax

`show user-role <role-name>`

Description

Displays the user-role information for the specified user-role name.

Command context

`manager or operator`

Parameters

<role-name>

Specifies the user-role name.

Examples

Shows the user-role by specific name.

```
switch# show user-role XYZ

 User Role Information

   Name                                                   : XYZ
   Type                                                   : local
   Reauthentication Period (seconds)                      : 0
   Untagged VLAN                                          :
   Tagged VLAN                                            :
   Captive Portal Profile                                 :
   Policy                                                 :
   Tunnelednode Server Redirect                           : Enabled
   Secondary Role Name                                    : ROLE NAME
```

show tunneled-node-server information

Syntax

`show tunneled-node-server [information | statistics <controller> | state <controller>]`

Description

Shows the tunneled-node-server information, statistics, or state.

NOTE

Information for SAC and SSAC Switch Anchor Controller (SAC), Standby Anchor Controller (SSAC) and User Anchor Controller (UAC) is available in the ArubaOS-Switch Controller Guide at **https://support.arubanetworks.com/**.

Command context

`manager or operator`

Parameters

`information`

Specifies tunneled-node-server information, node-list, and the bucket map information.

`statistics <controller>`

Specifies the data plane statistics with respect to a controller for each port.

`state <controller>`

Specifies the data plane state with respect to a controller.

Example

Show the cluster, node-list, and the bucket map information.

```
Switch# show tunneled-node-server information

 SAC Information

  SAC : 121.34.211.90
  Standby-SAC : 192.168.34.111

 Node List Information

  Cluster Name: abcd
  Cluster Status: Enabled
  Node List
  10.1.1.1
  10.1.1.2
  10.1.1.3

 Bucket Map Information

  Bucket name: xyz
  Bucket ID   A-UAC          S-UAC         Connectivity
  0           10.1.1.1       10.1.1.2      L2
  1           10.1.1.2       10.1.1.1      L2
  2           10.1.1.3       10.1.1.4      L2
```

Example

Show the tunneled-node-server state. If the PUTN is active, the output of this command will change.

```
Switch# show tunneled-node-server state

 Local Master Server (LMS) State

  LMS Type    IP Address         State          Capability    Role
  Primary   : 120.121.111.221   Complete       Per User      Operational
Primary
  Secondary : 120.121.111.222   In Progress    Per User      Operational
Secondary

 Switch Anchor Controller (SAC) State
                  IP Address       MAC Address    State
  SAC         :    120.121.111.222  68b599-1e6200  Registered
  Standby-SAC :    120.121.111.221  68b599-1e6201  Registered

  User Anchor Controller (UAC): 123.222.212.90
  User                Port         VLAN     State          Bucket ID
  68b599-1e6100       5            1000     Registering    2
```

```
 68b599-1e6200     7            2000     Registered     7

 User Anchor Controller (UAC): 123.222.212.70
 User              Port         VLAN     State          Bucket ID
 68b599-1e6300     15           800      Unregistering  3
 68b599-1e6400     17           3000     Registered     9
```

Example

Show the users information with respect to a UAC. The output is displayed when per user tunneled node feature is active.

```
Switch# show tunneled-node-server state <UAC-IP>

 User              Port         VLAN     State          Bucket ID
 68b599-1e6300     15           800      Unregistering  3
 68b599-1e6400     17           3000     Registered     9
```

Example

Show the tunneled-node-server statistics. The output of this command will not change if the PUTN is active.

```
Switch# show tunneled-node-server statistics

 Tunneled Node Statistics

  Control Plane Statistics

  SAC              : 123.234.122.222
  Bootstrap Tx     : 10                     Bootstrap Rx        : 10
  Nodelist Rx      : 25               Nodelist Ackd       : 6
  Bucketmap Rx     : 21               Bucketmap Ackd      : 10
  Failover Tx      : 4                Failover Ack Rx     : 3
  Unbootstrap Tx   : 7                Unbootstrap Ack Rx : 5
  Heartbeat Tx     : 5                Heartbeat Rx        : 3

  Standby-SAC      : 112.234.111.222
  Bootstrap Tx     : 10                     Bootstrap Rx        : 10
  Nodelist Rx      : 25               Nodelist Ackd       : 6
  Bucketmap Rx     : 21               Bucketmap Ackd      : 12
  Failover Tx      : 4                Failover Ack Rx     : 3
  Unbootstrap Tx   : 5                Unbootstrap Ack Rx : 3
  Heartbeat Tx     : 7                Heartbeat Rx        : 4

  UAC              : 101.111.231.222
  Bootstrap Tx     : 10               Bootstrap Ack       : 5
  Unbootstrap Tx   : 5                Unbootstrap Ack     : 5
  Keepalive Tx     : 2                Keepalive Ack       : 2

  UAC              : 101.222.111.142
  Bootstrap Tx     : 5                Bootstrap Ack       : 5
  Unbootstrap Tx   : 0                Unbootstrap Ack     : 0
  Keepalive Tx     : 0                Keepalive Ack       : 0

  Data Plane Statistics

  UAC             Packets Tx     Packets Rx
  10.1.1.1        45678          23456
  10.1.1.2        34567          23457

  User Statistics
```

```
  UAC            User Count
  10.1.1.1         10
  10.1.1.2         20
```

Example

Show the tunneled-node-server statistics of a data plane IP address. The output will be displayed when PUTN feature is active.

```
Switch# show tunneled-node-server statistics <UAC-IP>

  Port       Packets Tx     Packets Rx
  A1         45678          23456
  A2         34567          23457
```

show tunneled-node-users

Syntax

` show tunneled-node-users [all | count | down |  mac <MAC-ADDRESS> |  port <PORT-ADDR> | up ]`

Description

Shows the status of a client after configuring and enabling tunneled-node-server-redirect.

Command context

`manager or operator`

Parameters

<ALL>

Specifies all clients and their status.

<COUNT>

Specifies the total number of clients configured to tunnel their traffic to the controller.

<DOWN>

Specifies the clients which are not able to tunnel their traffic.

PORT <PORT-ADDR>

Specifies the port client status.

MAC <MAC-ADDRESS>

Specifies the client status based on the MAC address desired by the user.

<UP>

Specifies the status of clients with tunnels up.

Example

Show tunneled-node-users information for all clients.

```
switch(user-role)# show tunneled-node-users all

PORT   MAC-ADDRESS  TUNNEL-STATUS   SECONDARY-USERROLE   FAILURE-REASON
123      XYZ123       UP                 ROLE NAME1
234      XYZ1234      DOWN               ROLE NAME2        UAC_DOWN
```

Example

Show the total number of clients configured with the user-based tunneled-node.

```
switch(user-role)# show tunneled-node-users count

Total number of clients configured with user-based tunneled node: <NUMBER>
```

Example

Show the port access client status in detail.

```
switch(user-role)# show tunneled-node-users port 1/7

PORT    MAC-ADDRESS     TUNNEL-STATUS   SECONDARY-USERROLE      FAILURE-REASON
1/7     d48564-940c46   UP              authenticated
```

Example

Show the port access clients in detail.

```
switch# show port-access clients

Downloaded user roles are preceded by *
 Port Access Client Status

  Port  Client Name   MAC Address       IP Address      User Role         Type
VLAN
  ----- ------------- ----------------- --------------- ----------------- -----
----
  1/7   2c41387f35b9  2c4138-7f35b9     n/a             Voice_HP          MAC   171
  1/7   d48564940c46  d48564-940c46     n/a             *DUR_prof2_PUT... MAC   100
  1/16  dur1          001517-857121     n/a             *DUR_prof2_PUT... 8021X 100
  1/17  dur1          d48564-a8afa0     n/a             *DUR_prof2_PUT... 8021X 100
  3/7   2c41387fe7f8  2c4138-7fe7f8     n/a             Voice_HP          MAC   171
  3/7   e8393537b4a5  e83935-37b4a5     n/a             *DUR_prof2_PUT... MAC   100
```

Example

Show the detailed port-access clients for port 1/7.

```
switch# show port-access clients detailed 1/7

 Port Access Client Status Detail

  Client Base Details :
   Port             : 1/7                   Authentication Type : mac-based
   Client Status    : authenticated         Session Time        : 18972 seconds
   Client Name      : 2c41387f35b9          Session Timeout     : 0 seconds
   MAC Address      : 2c4138-7f35b9
   IP               : n/a

Downloaded user roles are preceded by *
 User Role Information

   Name                              : Voice_HP
   Type                              :
   Reauthentication Period (seconds) : 0
   Untagged VLAN                     : 171
   Tagged VLANs                      :
```

```
   Captive Portal Profile             :
   Policy                             :
   Tunnelednode Server Redirect       : Disabled
   Secondary Role Name                :

  Client Base Details :
   Port              : 1/7                     Authentication Type : mac-based
   Client Status     : authenticated           Session Time        : 18947 seconds
   Client Name       : d48564940c46            Session Timeout     : 0 seconds
   MAC Address       : d48564-940c46
   IP                : n/a

Downloaded user roles are preceded by *
 User Role Information

   Name                               : *DUR_prof2_PUTN-3037-12
   Type                               : downloaded
   Reauthentication Period (seconds)  : 0
   Untagged VLAN                      : 100
   Tagged VLANs                       :
   Captive Portal Profile             :
   Policy                             : upol2_DUR_prof2_PUTN-3037-12

Statements for policy "upol2_DUR_prof2_PUTN-3037-12"
policy user "upol2_DUR_prof2_PUTN-3037-12"
     10 class ipv4 "remark2_DUR_prof2_PUTN-3037-12" action rate-limit kbps 1000000
action priority 2 action permit
   exit
Statements for class IPv4 "remark2_DUR_prof2_PUTN-3037-12"
class ipv4 "remark2_DUR_prof2_PUTN-3037-12"
     10 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit

   Tunnelednode Server Redirect       : Enabled
   Secondary Role Name                : authenticated
```

Tagged VLAN for user role

Tagged VLAN for user-role supports the flow of tagged traffic. Once the client is authenticated and a user-role applied, tagged traffic is allowed to flow through a user-based tunnel.

vlan-id-tagged

Syntax

`vlan-id-tagged [<VLAN-ID> | < VLAN-STR>]`

`no vlan-id-tagged [<VLAN-ID> | < VLAN-STR>]`

Description

Supports tagged client traffic by directing the user-role. When the user-role is applied, tagged traffic is allowed to pass.

The `no` form of this command removes the [<*VLAN-ID*> | < *VLAN-STR*>] from the tagged VLAN ID.

Command context

`user-role`

Parameters

<VLAN-ID>

Specifies the tagged VLAN ID assigned to users.

<VLAN-STR>

Specifies the tagged VLAN name assigned to users.

Example

Show the user-role information for PUTN-emp.

```
switch-PoEP# show user-role PUTN-emp

 User Role Information

   Name                              : PUTN-emp
   Type                              : local
   Reauthentication Period (seconds) : 0
   Untagged VLAN                     :
   Tagged VLAN                       : 11
   Captive Portal Profile            :
   Policy                            :
   Tunnelednode Server Redirect      : Enabled
   Secondary Role Name               : authenticated
```

user-role vlan-id

Syntax

vlan-id *<VLAN-ID>*

no vlan-id *<VLAN-ID>*

Description

Creates the VLAN within the `user-role` context.

The `no` form of this command removes the VLAN from the `user-role` context.

Command context

`user-role`

Parameters

<VLAN-ID>

Specifies the ID assigned to the VLAN for when the client is authenticated and the `user-role` applied.

Example

Show the detailed status of the clients port-access.

```
switch-PoEP# show port-access clients detailed

 Port Access Client Status Detail

  Client Base Details :
   Port            : 1                     Authentication Type : 802.1x
   Client Status   : authenticated         Session Time        : 58081 seconds
   Client name     : emp4                  Session Timeout     : 0 seconds
   MAC Address     : 000ffe-c8ce92
```

```
  IP               : n/a

User Role Information

  Name                                 : PUTN-emp
  Type                                 : local
  Reauthentication Period (seconds)    : 0
  Untagged VLAN                        :
  Tagged VLANs                         : 11
  Captive Portal Profile               :
  Policy                               :
  Tunnelednode Server Redirect         : Enabled
  Secondary Role Name                  : authenticated
```

Downloadable user-roles

Downloadable user-role enables ArubaOS-Switches to download user-roles, policy, and class from the Clear Pass Policy Manager (CPPM) server. The download facilitates the setup of policies and attributes for a specific user-role which can then be stored on the switch. New users can be configured and assigned the same stored version of the user-role in ClearPass, saving the administrator time reconfiguring each user individually.

- The command `radius-server cppm identity <IDENTITY> key <KEY>` can only be used if the user role is going to be downloaded from the CPPM server.
- See commands `radius-server host` *`<server-ip>`* and `radius-server key` *`<key-string>`*.

aaa authorization user-role enable download

Syntax

`aaa authorization user-role enable download`

`no aaa authorization user-role enable download`

Description

Enables the downloadable user-role.

The `no` form of this command disables the downloadable user-role in ArubaOS-Switches.

Command context

config

Parameters

`download`

Enables the switch to download the user role from a CPPM server.

Warning

- The command `no aaa authorization user-role enable` disables both the user-role and the downloadable user-role.

Usage

- A CPPM user name and password must be configured for download user-role to work.

Example

Download the user role from a CPPM server.

```
switch(config)# aaa authorization user-role enable download

Some legacy secure client access functionality is not supported when user roles
are enabled.
Please refer to the end user documentation for details.

CPPM user name and password must be configured for downloading the user role.
CPPM HTTPS root certificate must be installed for downloading the user role.
```

radius-server cppm identity

Syntax

`radius-server cppm identity <IDENTITY> key <KEY>`

`no radius-server cppm identity <IDENTITY>`

Description

User name and logon password combination of the CPPM allows access to the download user-roles.

The `no` form of the command removes the user name and password combination from CPPM.

Command context

`user-role`

Parameters

<IDENTITY>

Specifies the user name of the CPPM.

<KEY>

Specifies the password of the CPPM user.

Example

Set the password key for the CPPM user.

```
switch(config)# radius-server cppm identity admin key xxxxxxxx
```

downloadable-role-delete

Syntax

`downloadable-role-delete <USER-ROLE-NAME>`

Description

Deletes the downloaded user-role, the associated class, and policy information from the user-role.

Command context

`manager`

Parameters

<USER-ROLE-NAME>

Specifies the user-role name to be deleted.

Example

Delete the user-role for the user DUR_prof2_PUTN-3037-12.

```
switch(config)# downloadable-role-delete DUR_prof2_PUTN-3037-12
```

show user-role <XYZ>

Syntax

```
show user-role <XYZ> [detailed | downloaded]
```

Description

Shows the named user-role in either detail or downloaded information.

Command context

`operator or manager`

Parameters

<detailed>

Specifies the downloaded user-role in detail.

<downloaded>

Specifies the downloaded user-role.

Example

Show the download information for user-role XYZ.

```
Switch(config)# show user-role XYZ downloaded

 User Role Information

   Name                                              : XYZ
   Type                                              : Downloaded
   Reauthentication Period (seconds) : 0
   Untagged VLAN                                     :
   Tagged VLAN                                       :
   Captive Portal Profile                            :
   Policy                                            :
   Tunnelednode Server Redirect                      : Enabled
   Secondary Role Name                               : ROLE NAME
```

show port-access clients

Syntax

```
show port-access clients [ethernet <PORT-LIST> |<PORT>| detailed]
 [ethernet] PORT-LIST  Show information for specified ports only.
 detailed              Show detailed client authentication status.
```

Description

Shows the general port-access status information for clients by detailed status information for a specified port, list of ports or ethernet.

Command context

`operator or config`

Parameters

ethernet ***<PORT-LIST>***

Specifies the detailed port-access information by a list of ethernet ports.

detailed

Specifies the detailed port-access information for all.

<PORT>

Specifies the detailed port-access information for a client by port name.

Example

Show the port access status of a client.

```
switch# show port-access clients

Downloaded user roles are preceded by *
 Port Access Client Status

  Port  Client Name   MAC Address       IP Address      User Role         Type
VLAN
  ----- ------------- ----------------- --------------- ----------------- -----
----
  1/7   2c41387f35b9  2c4138-7f35b9     n/a             Voice_HP          MAC   171
  1/7   d48564940c46  d48564-940c46     n/a             *DUR_prof2_PUT... MAC   100
  1/16  dur1          001517-857121     n/a             *DUR_prof2_PUT... 8021X 100
  1/17  dur1          d48564-a8afa0     n/a             *DUR_prof2_PUT... 8021X 100
  3/7   2c41387fe7f8  2c4138-7fe7f8     n/a             Voice_HP          MAC   171
  3/7   e8393537b4a5  e83935-37b4a5     n/a             *DUR_prof2_PUT... MAC   100
```

Example

Show the detailed port-access for clients.

```
switch# show port-access clients detailed 1/7

 Port Access Client Status Detail

  Client Base Details :
   Port            : 1/7                    Authentication Type : mac-based
   Client Status   : authenticated          Session Time        : 22203 seconds
   Client Name     : 2c41387f35b9           Session Timeout     : 0 seconds
   MAC Address     : 2c4138-7f35b9
   IP              : n/a

Downloaded user roles are preceded by *
 User Role Information

   Name                                 : Voice_HP
   Type                                 :
   Reauthentication Period (seconds)    : 0
   Untagged VLAN                        : 171
   Tagged VLANs                         :
   Captive Portal Profile               :
   Policy                               :
   Tunnelednode Server Redirect         : Disabled
   Secondary Role Name                  :
```

```
  Client Base Details :
   Port             : 1/7                    Authentication Type : mac-based
   Client Status    : authenticated          Session Time        : 1259 seconds
   Client Name      : d48564940c46           Session Timeout     : 0 seconds
   MAC Address      : d48564-940c46
   IP               : n/a

Downloaded user roles are preceded by *
 User Role Information

   Name                               : *DUR_prof2_PUTN-3037-12
   Type                               : downloaded
   Reauthentication Period (seconds) : 0
   Untagged VLAN                      : 100
   Tagged VLANs                       :
   Captive Portal Profile             :
   Policy                             : upol2_DUR_prof2_PUTN-3037-12

Statements for policy "upol2_DUR_prof2_PUTN-3037-12"
policy user "upol2_DUR_prof2_PUTN-3037-12"
     10 class ipv4 "remark2_DUR_prof2_PUTN-3037-12" action rate-limit kbps 1000000
action priority 2 action permit
   exit
Statements for class IPv4 "remark2_DUR_prof2_PUTN-3037-12"
class ipv4 "remark2_DUR_prof2_PUTN-3037-12"
     10 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit

   Tunnelednode Server Redirect      : Enabled
   Secondary Role Name               : authenticated
```

debug usertn

Syntax

`debug usertn`

Description

Enables the debug for the user tunneled node.

Command context

`manager`

Example

Debug session for the user tunneled node.

```
0007:23:11:47.44 TNT mtnodeUserCtrl:userTNodeProcAddUserReq User Add for num
clients: 1
0007:23:11:47.54 TNT mtnodeUserCtrl:AMM: Sync user add
0007:23:11:47.60 TNT mtnodeUserCtrl:Bucket 148 initialized
0007:23:11:47.68 TNT mtnodeUserCtrl:userTNodeDecapReserveResource: reservetgTcam
   0 reserve 1 port's uitgReserveCnt: 0 uac->uitgReserve:0

0007:23:11:47.83 TNT mtnodeUserCtrl:1664a8c0 uac->uitgReserve ZERO reserve 1

0007:23:11:47.92 TNT mtnodeUserCtrl:Entering userTNodeTCAMDecapReserveResources

0007:23:11:48.07 TNT mtnodeUserCtrl:userTNodeDecapReserveResource: reservetgTcam
   0 reserve 1 port's uitgReserveCnt: 0 uac->uitgReserve:0
```

```
0007:23:11:48.22 TNT mtnodeUserCtrl:1564a8c0 uac->uitgReserve ZERO reserve 1

0007:23:11:48.30 TNT mtnodeUserCtrl:Entering userTNodeTCAMDecapReserveResources

0007:23:11:48.45 TNT mtnodeUserCtrl:User 000ffe-c8ce92 added to tree
0007:23:11:48.52 TNT mtnodeUserCtrl:UAC Bootstrap to node: 192.168.100.22
0007:23:11:48.60 TNT mtnodeUserCtrl:userTNodeSendUACBStrapReq: Packet Sent
Successfully
0007:23:11:48.70 TNT mtnodeUserCtrl:User 000ffe-c8ce92 bootstrapping
0007:23:11:48.82 TNT mtnodeUserCtrl:create tunnel: 1 port 1664a8c0 uac 0 ifindex 0
count
0007:23:11:48.91 TNT mtnodeUserCtrl:AMM: Active Sync State uac ifindex:
   318767788 port: 1 cluster: 0 uac: 1
0007:23:11:49.03 TNT mtnodeUserCtrl:create tunnel: 1 port 1564a8c0 uac 0 ifindex 0
count
0007:23:11:49.17 TNT mtnodeUserCtrl:AMM: Active Sync State uac ifindex:
   318767789 port: 1 cluster: 0 uac: 0
0007:23:11:49.29 TNT mtnodeUserCtrl:User 000ffe-c8ce92 tunneling to 318767788
0007:23:11:49.37 TNT mtnodeUserCtrl:AMM: Sync papi payload from the controller
event
```

Overview

The Port QoS Trust feature restricts which packet QoS information may be used to determine inbound queue servicing and any priority information to be permitted into the local hop.

Port QoS Trust Mode configuration allows preservation or removal of the inbound QoS priorities carried in Layer 2 (the VLAN cos or Priority CodePoint (PCP) value, known as the 802.1p priority tag) and/or in Layer 3 (the IP-ToS byte, in IP-Precedence or IP-Diffserv mode). The different modes let the customer trust all, some, or no packet priority fields.

The per-port configuration enables the customer to trust some sources or devices and not others. This feature is mutually exclusive with any active port-priority configuration.

Configuration commands

qos trust

Syntax

```
qos trust [default|dot1p|dscp|ip-prec|none|device [none|<DEVICE-TYPE>]]
```

Description

Set the QoS Trust Mode configuration for the port.

Options

default

Trust 802.1p priority and preserve DSCP or IP-ToS.

device *<DEVICE-TYPE>*

On approved devices, trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated priority, the priority will be remarked to 0. On unapproved devices, trust 802.1p priority and preserve any IP- ToS values.

dot1p

Trust 802.1p priority and preserve DSCP or IP-ToS.

dscp

Trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated 802.1p priority, the priority will be remarked to 0.

ip-precedence

Trust IP-ToS IP-Precedence mode in IP packets and remark the 802.1p priority.

none

Do not trust either the 802.1p priority or the IP-ToS values.

QoS trust devices

aruba-ap

Aruba Access point device.

none

Clear all trusted devices from port.

NOTE

Both SNMP and the CLI will verify that the current QoS Port Priority and desired QoS Trust Mode configuration are not mutually exclusive (and conversely).

qos dscp-map

Syntax

`qos dscp-map <CODEPOINT> priority <PRIORITY> [name <NAME> | default | legacy]`

Description

Modifies DSCP mapping.

Options

default

Returns switch to the fully mapped factory-default configuration.

legacy

Restore the legacy default behavior (partial mapping) used in earlier code releases.

Show commands

show qos trust

Syntax

`show qos trust [device] <PORT>`

Description

Shows port-based QoS trust configuration

Options

device

Show list of trusted devices per-port.

<port>

Show trusted devices on a single port.

Usage

`show qos trust [device | [ethernet <PORT-LIST> ]`

show qos trust

```
switch# show qos trust
```

```
Port-based qos Trust Configuration

  Port    Trust Mode    | Device Trust State ---- --- ----
  A1      Default       |
  A2      Default       |
  A3      Device**      | Trusted
  A4      IP-Prec       |
  A5      Dot1p         |
  A5      None          |
  A5      DSCP          |
  A5      Device**      |
  A5      Dot1p         |

** For a list of trusted devices per-port, use the command show qos trust device.
To show trusted devices on a single port, use the command show qos trust device
<PORT>.
```

show qos trust device

```
switch# show qos trust device

Port-Based QoS Trust Configuration

  Port    Trusted Devices
  ------- ---------------
  A1      aruba-ap
  A2      aruba-ap
  A4      aruba-ap
```

show qos trust device <PORT>

```
switch# show qos trust device <PORT>

Port A4 QoS Trust Configuration
   Current state: Trusted

   Trusted Devices: aruba-ap
```

Validation rules

Validation	Error/Warning/Prompt
`qos trust` *<UNSUPPORTEDDEVICETYPE>*	Invalid input: %s
`no qos trust` *<ANYVALUE>*	Invalid command. To disable trust for a port, use `qos trust none`. To return to the default configuration and leave priority information unchanged, use `qos trust default`.

Table Continued

Validation	Error/Warning/Prompt
QoS priority when trust mode is anything other than <*NONE*> or <*DEFAULT*>.	The port QoS trust mode must be <*DEFAULT*> or <*NONE*> to configure the QoS port priority feature.
QoS DSCP when trust mode is anything other than <*NONE*> or <*DEFAULT*>.	The port QoS trust mode must be <*DEFAULT*> or <*NONE*> to configure the QoS port priority feature.
`QoS trust dot1.p` when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust ip-prec when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust DSCP when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust device when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.

Overview

The tunneled node feature encapsulates incoming packets from end-hosts in Generic Routing Encapsulation (GRE) and forwards them to a Mobility Controller for additional processing. The Mobility Controller strips the GRE header and processes the packet for authentication and stateful firewall, which enables centralized security policy, authentication, and access control.

The tunneled node feature is enabled on a per-port basis. Any traffic coming from nontunneled node interfaces is forwarded without being tunneled to a Mobility Controller.

Operating notes

- Tunneled node profile may be created using CLI and SNMP.
- The tunneled node profile supports configuring of:
  - Primary controller (IPv4 only).
  - Backup controller (IPv4 only).
  - Heartbeat keepalive timeout – range 1-8 seconds.
- Only one tunneled node profile may be created.
- The tunneled-node profile may be applied to a physical port only via CLI and SNMP.
- The maximum number of physical ports to which the profile may be applied is:
  - Aruba 5400R Switch Series (non-VSF): 256 physical ports.
  - Aruba 5400R Switch Series (VSF): 512 physical ports.
- The configuration related to the tunneled node profile will be stored in the flash and restored after a boot.
- High availability (HA) will be supported for the tunneled-node related configuration.
- A tunnel, associated with a port, is “up” when both conditions are met. A tunnel is “down” when either of the conditions are not met.
  - Either the primary or backup controller is reachable.
  - A boot strap message response is received from the controller.
- Heartbeat between the switch and controller has failed when the controller does not respond after five attempts. All tunnels are brought down with a heartbeat failure.
- A tunnel “up or down” status will be logged for each tunnel node port in the event log.
- The `show tech` command dumps all user-mode and test-mode command outputs.
- To reach the Aruba controller, the VLAN must have a manual IP configured.
- With the exception of the 802.1x BPDU, the switch consumes all other BPDUs.

Protocol Application Programming Interface (PAPI)

The PAPI Enhanced Security configuration provides protection to Aruba devices, AirWave, and ALE against malicious users sending fake messages that results in security challenges.

Starting from ArubaOS-Switch version 16.02, a minor security enhancement has been made to Protocol Application Programming Interface (PAPI) messages. Protocol Application Programming Interface endpoint authenticates the sender by performing a check of the incoming messages using MD5 (hash). All PAPI endpoints — APs, Controllers, Mobility Access Switches, AirWave, and ALE — must use the same secret key. The switch software currently uses a fixed key to calculate the MD5 digest and cooperate with the controller for PAPI enhanced security.

NOTE

To use this functionality, the PAPI security profile must be configured on the controller. For more information on the Aruba controller, see the **Aruba Networks Controller Configuration Manual**.

Configuration commands

tunneled-node-server

Syntax

```
tunneled-node-server [controller-ip <IP-ADDR> | backup-controller-ip
<IP-ADDR> | [keepalive <TIMEOUT>] | enable | fallback-local-switching]


no tunneled-node-server [controller-ip <IP-ADDR> | backup-controller-ip
<IP-ADDR> | [keepalive <TIMEOUT>] | enable | fallback-local-switching]
```

Description

Configure tunneled node server information.

The no form of the command removes the tunneled node server configuration.

Options

controller-IP

Configure the controller IP address for the tunneled node.

backup-controller-IP

Configure the backup controller IP address for the tunneled node.

keepalive

Configure the keepalive timeout for the tunneled node in seconds [1-40]. The default is 8 seconds.

enable

Enter the manager command context.

fallback-local-switching

Apply fallback option when communication with the controller fails. When the tunneled node is applied to a port and the tunnel cannot be established with the controller, the fallback-local-switching option allows port traffic to be switched locally. When the option fallback-local-switching is not specified, the port traffic is dropped when the tunnel reestablishment fails.

Usage

```
switch(config)# tunneled-node-server controller-ip 15.255.133.148

switch(config)# tunneled-node-server backup-controller-ip 15.255.133.148

switch(config)# tunneled-node-server keepalive 40

switch(config)# tunneled-node-server fallback-local-switching
```

interface tunneled-node-server

Syntax

```
interface <PORT> tunneled-node-server
```

```
no interface <PORT> tunneled-node-server
```

Description

Enable tunneled node on a port.

The `no` command disables the tunneled node on the port.

controller-ip

From within the **tunneled-node-profile** context:

Syntax

```
[no] controller-ip <IP-ADDR>
```

Description

Configure the Controller IP address for the tunneled node.

Usage

```
[no] controller-ip <IP-ADDR>
```

controller-ip

Configure the Controller IP address for the tunneled node.

keepalive

From within the **tunneled-node** context:

Syntax

```
[no] keepalive <TIMEOUT>
```

Description

Configure the keepalive timeout for the tunneled node in seconds.

Keepalive timeout seconds [1-40].

Default: 8 seconds.

Options

keepalive

Configure the keepalive timeout for the tunneled node in seconds.

backup-controller-ip

From within the **tunneled-node-profile** context:

Syntax

```
[no] backup-controller-ip <IP-ADDR>
```

Description

Configure the backup controller IP address for the tunneled node.

Options

backup-controller-ip

Configure the backup controller IP address for the tunneled node.

Usage

`[no] backup-controller-ip <IP-ADDR>`

fallback-local-switching

From within the **interface** context:

Syntax

`fallback-local-switching`

Description

To switch traffic locally upon losing connectivity to the controller, you must configure the fallback option before connectivity fails. When the tunneled node is applied to a port and the tunnel cannot be established with the controller, the fallback-local-switching option allows port traffic to be switched locally. When the option fallback-local-switching is not specified, the port traffic is dropped when the tunnel reestablishment fails.

Show commands

show tunneled-node-server

Syntax

`show tunneled-node-server [state | statistics]`

Description

Display switch operation information.

Options

state

Display the tunneled node port state.

statistics

Display the tunneled node statistics.

show tunneled-node-server state

```
Tunneled node Port State
Active Controller IP Address  :
Port   State
------ -------------------------
2      Port down
```

show tunneled-node-server statistics

```
Tunneled node Statistics

 Port : 2

  Control Plane Statistics
```

```
    Bootstrap packets sent      : 0
    Bootstrap packets received  : 0
    Bootstrap packets invalid   : 0

   Tunnel Statistics
    Rx Packets                                        : 0
    Tx Packets                                        : 0
    Rx 5 Minute Weighted Average Rate (Pkts/sec)      : 0
    Tx 5 Minute Weighted Average Rate (Pkts/sec)      : 0

 Aggregate Statistics
  Heartbeat packets sent           : 0
  Heartbeat packets received       : 0
  Heartbeat packets invalid        : 0
  Fragmented Packets Dropped (Rx)  : 0
  Packets to Non-Existent Tunnel   : 0
  MTU Violation Drop               : 0
```

clear statistics tunneled-node-server

Syntax

`clear statistics tunneled-node-server`

Description

Clear statistics from the tunneled node server.

Interaction table

Features enabled with tunneled node:

Feature
Mirrors (MAC, VLAN, port)
PVST/RPVST/STP
DLDP
UDLD
LLDP/CDP
GVRP/MVRP
LACP
UFD
Sflow
Loop protect
Smartlink
Global QoS (VLAN, port, rate limit)

Table Continued

Feature
Mac lockout/lockdown
ACL/Classifiers (ingress/egress)
IGMP/MLD
GMB
Broadcast-limit
energy-efficient-Ethernet
flow-control
power-over-ethernet • poe-allocate-by • poe-lldp-detect
Rogue Mac detection
LLDP auto-provisioning

Restrictions

- Once a tunneled-node profile is applied to a port, the controller IP (primary and backup) cannot be changed.
- IP address cannot be assigned to VLANs that the tunnel-node port belongs to.
- No support for fragmentation and reassembly for encapsulated frames that result in an MTU violation. Such frames will be dropped. HPE recommends configuring the switch-controller path for Jumbo MTU. No support for PMTU detection for tunnel traffic.
- The packets from nontunneled node ports (in the same VLAN as tunnel-node port) will not be bridged to the tunneled-node ports and conversely.

Features not allowed on a tunneled node port/VLAN with tunneled node ports/globally:

Feature	Blocked globally/per port/ VLAN with tunneled-node-ports
IP multicast routing	Global
Openflow	Global
Q-in-Q	Global
Distributed Trunking	Global
Mesh	Global
VXLAN	Global
IP address: manual and dhcp	VLAN

Table Continued

Feature	Blocked globally/per port/ VLAN with tunneled-node-ports
802.1x, mac auth, webauth, LMA, port security	port
DIPLD (IPv4/IPv6)	port
DSNOOP (IPv4/IPv6)	VLAN
ARP protect	VLAN
RA guard	port
Virus throttling	port
BYOD	VLAN
Trunk	Profile cannot be applied to a trunk
PBR policies	VLAN
Tunneled-node port	port
Src port/Mcast filters	port
DHCP client/Server/Relay	VLAN

PAPI security

Protocol Application Programming Interface (PAPI)

The PAPI Enhanced Security configuration provides protection to Aruba devices, AirWave, and ALE against malicious users sending fake messages that results in security challenges.

Starting from ArubaOS-Switch version 16.02, a minor security enhancement has been made to Protocol Application Programming Interface (PAPI) messages. Protocol Application Programming Interface endpoint authenticates the sender by performing a check of the incoming messages using MD5 (hash). All PAPI endpoints — APs, Controllers, Mobility Access Switches, AirWave, and ALE — must use the same secret key. The switch software currently uses a fixed key to calculate the MD5 digest and cooperate with the controller for PAPI enhanced security.

NOTE

To use this functionality, the PAPI security profile must be configured on the controller. For more information on the Aruba controller, see the **Aruba Networks Controller Configuration Manual**.

PAPI configurable secret key

To support enhanced PAPI security, a command is available to configure a MD5 secret key.

papi-security

Syntax

```
switch# (config) papi-security
```

Description

Configure MD5 key for enhanced PAPI security.

Parameters

enhanced-security

The enhanced-security CLI must be enabled in Aruba controller for the connection to be truly secured.

<KEY-STR>

Configure MD5 key for enhanced PAPI security using a key-string parameter.

<KEY-VALUE>

Configure MD5 key for enhanced papi security using a key-value parameter.

Restrictions

- To view the status of the PAPI security, using the `show run` command with the option `include credentials` enabled, the PAPI security key will show in the output as an encrypted form.
- Key length has to be between 10-64.
- By default the enhanced-security is disabled.
- When enhanced-security mode is disabled, any AP can obtain the current shared secret key.
- When enhanced-security mode is enabled, an AP is not updated with the new shared secret key unless the AP knows the previous key and the AP is updated with the new key within one hour of the key creation.
- Key length has to be between 10-64 or the following message will appear:

```
Minimum key-value length allowed is 10 characters and maximum allowed is 64
characters.
```

Usage

```
Switch(config)# papi-security key-value <KEY-VALUE>
Switch(config)# [no] papi-security <KEY-VALUE>
```

papi-security key-value

```
HP-2920-24G(config)# papi-security key-value TestKey12345678
HP-2920-24G(config)# no papi-security key-value

HP-2920-24G(config)# papi-security key-value Test
Minimum key-value length allowed is 10 characters and maximum allowed is 64
characters.
```

show run with encrypted key

```
 Switch(config)# sh run
Running configuration:
;J9576A Configuration Editor
;Created on release #KA.16.02.0000x
;Ver #0e:01.f0.92.34.5f.3c.6b.fb.ff.fd.ff.ff.3f.ef:78
;encrypt-cred +NXT3w7ky2IXNXadlJblS/1ZRi/o73Qq28XXcLkSCZq9PU30Kl+KMLMva8rQri5g
```

```
hostname "HP-3800-48G-4SFPP"
module 1 type j9576y
module 2 type j9576x
encrypt-credentials
papi-security encrypted-key <"encrypted-key">
snmp-server community "public" unrestricted
snmpv3 engineid "00:00:00:0b:00:00:50:65:f3:b4:a6:c0"
oobm
ip address dhcp-bootp
exit

vlan 1
name "DEFAULT_VLAN"
untagged 1-52
ip address dhcp-bootp
exit

activate provision disable
```

show run with include key

```
show run
Running configuration:
; J9576A Configuration Editor
; Created on release #KA.16.02.0000x
; Ver#0e:01.f0.92.34.5f.3c.6b.fb.ff.fd.ff.ff.3f.ef:78

 hostname "HP-3800-48G-4SFPP"
 module 1 type j9576y
 module 2 type j9576x
 include-credentials
 papi-security key-value <"key">
 snmp-server community "public" unrestricted
 snmpv3 engineid "00:00:00:0b:00:00:50:65:f3:b4:a6:c0"
 oobm
 ip address dhcp-bootp
 exit

 vlan 1
 name "DEFAULT_VLAN"
 untagged 1-52
 ip address dhcp-bootp
 exit

 activate provision disable
```

The Time Domain Reflectometry (TDR) is a new port feature supported on Aruba 3810M switches and Aruba 5400R v3 blades. TDR is introduced to detect cable faults on 100BASE-TX and 1000BASE-T ports.

Supported Platforms

- Aruba 2930F switches
- Aruba 3810M switches
- Aruba 5400R v3 blades (J9986A, J9987A, J9989A, J9990A, J9991A [applicable only for ports 1–20, rest of the four ports are Smart Rate ports], and J9992A)

Virtual cable testing

The Virtual Cable Test (VCT) uses the same command as TDR. It is applicable only for GigT transceivers like copper transceiver (J8177C–ProCurve Gigabit 1000Base-T Mini-GBIC). The VCT test results include distance to the fault, but not the cable length.

Test cable-diagnostics

Syntax

```
test cable-diagnostics <PORT-LIST>
```

Description

Use the command to test for cable faults.

Option

PORT-LIST

Specify copper port as a input port number.

Test cable-diagnostics C21

```
test cable-diagnostics C21

The 'test cable-diagnostics' command will cause a loss of link and will take a few
seconds per interface to complete.
Continue [Y/N]? y

MDI         Cable    Distance    Pair   Pair        MDI
Port Pair   Status   to Fault    Skew   Polarity    Mode
---------   ------   --------    -----   ------   ------
C21 1-2       Open       0 m       0 ns
    3-6       Open       0 m       0 ns
    4-5       Open       0 m       0 ns
    7-8       Open       1 m       0 ns
```

Test cable-diagnostics 1/1-1/10

```
switch# test cable-diagnostics 1/1-1/10

This command will cause a loss of link on all tested ports and will take
several seconds per port to complete.  Use the 'show cable-diagnostics' command to
view
 the results.

Continue (y/n)? Y

switch# show cable-diagnostics 1/1-1/10

Cable Diagnostic Status - Copper Ports

       MDI    Cable       Cable Length or
  Port Pair   Status      Distance to Fault
  ---- ------ ----------- ---------------------
  1/1  1-2    OK          5m
       3-6    OK          5m
       4-5    OK          7m
       7-8    OK          7m
  1/2  1-2    OK          7m
       3-6    OK          7m
       4-5    OK          7m
       7-8    OK          7m
  1/3  1-2    OK          5m
       3-6    OK          7m
       4-5    OK          5m
       7-8    OK          7m
  1/4  1-2    OK          7m
       3-6    OK          7m
       4-5    OK          7m
       7-8    OK          5m
  1/5  1-2    OK          4m
       3-6    OK          5m
       4-5    OK          5m
       7-8    OK          4m
  1/6  1-2    OK          4m
       3-6    OK          4m
       4-5    OK          4m
       7-8    OK          4m
  1/7  1-2    OK          5m
       3-6    OK          4m
       4-5    OK          5m
       7-8    OK          4m
  1/8  1-2    OK          4m
       3-6    OK          5m
       4-5    OK          4m
       7-8    OK          4m
  1/9  1-2    OK          5m
       3-6    OK          5m
       4-5    OK          5m
       7-8    OK          5m
  1/10 1-2    OK          7m
       3-6    OK          5m
       4-5    OK          5m
       7-8    OK          5m
```

Good cable tests

```
switch# test cable-diagnostics 51

This command will cause a loss of link on all tested ports and will take
several seconds per port to complete.  Use the 'show cable-diagnostics'
command to view the results.

Continue (y/n)? Y

switch# show cable-diagnostics 51

Cable Diagnostic Status - Transceiver Ports

         MDI    Cable        Distance   Pair    Pair       MDI
  Port   Pair   Status       to Fault   Skew    Polarity   Mode
  ---- ------   ---------    --------   -----   -------   -----
  51     1-2      OK           0 m       8 ns   Normal     MDI
         3-6      OK           0 m       8 ns   Normal
         4-5      OK           0 m       8 ns   Normal     MDIX
         7-8      OK           0 m       0 ns   Normal


switch# test cable-diagnostics 52

This command will cause a loss of link on all tested ports and will take
several seconds per port to complete.  Use the 'show cable-diagnostics'
command to view the results.

Continue (y/n)?  Y

switch# show cable-diagnostics 52

Cable Diagnostic Status - Transceiver Ports

         MDI    Cable        Distance   Pair    Pair       MDI
  Port   Pair   Status       to Fault   Skew    Polarity   Mode
  ---- ------   ---------    --------   -----   -------   -----
  52    1-2       OK          0 m        0 ns   Normal     MDI
        3-6       OK          0 m        0 ns   Normal
        4-5       OK          0 m        0 ns   Normal     MDIX
        7-8       OK          0 m        0 ns   Normal
```

Faulty cable test

```
switch# test cable-diagnostics 51
This command will cause a loss of link on all tested ports and will take
several seconds per port to complete.  Use the 'show cable-diagnostics'
command to view the results.

Continue (y/n)?  y
switch# show cable-diagnostics 51

Cable Diagnostic Status - Transceiver Ports

       MDI    Cable        Distance   Pair   Pair       MDI
  Port Pair   Status       to Fault   Skew   Polarity   Mode
  ---- ------ -----------  ---------- ------ ---------- ------
  51    1-2    OK            0 m        0 ns
```

```
        3-6    Short        1 m         0 ns
        4-5    Short        1 m         0 ns
        7-8    OK           0 m         0 ns


switch# test cable-diagnostics 52

This command will cause a loss of link on all tested ports and will take
several seconds per port to complete.  Use the 'show cable-diagnostics'
command to view the results.

Continue (y/n)?  Y

switch# show cable-diagnostics 52

Cable Diagnostic Status - Transceiver Ports

       MDI    Cable        Distance    Pair   Pair        MDI
  Port Pair   Status       to Fault    Skew   Polarity    Mode
  ---- ------ ----------- ---------- ------ ---------- ------
  52   1-2     Open         0 m         0 ns
       3-6     Open         0 m         0 ns
       4-5     Open         1 m         0 ns
       7-8     Open         0 m         0 ns
```

Error message

Error Message	Cause
The transceiver on port 1/A1 does not support cable diagnostics.	• usage of invalid(fiber-SFP+) port • The selected range includes an entry for an invalid port.

show cable-diagnostics

Syntax

`show cable-diagnostics <PORT-LIST>`

Description

Use the command to generate results of completed tests on single or multiple ports. For incomplete tests, a warning is displayed.

Option

`PORT`

Specify one copper port as an input port number.

clear cable-diagnostics

Syntax

`clear cable-diagnostics`

Description

Use the command to clear the result buffer.

Example

```
switch(config)# clear cable-diagnostics
```

Limitations

TDR has the following limitations:

- TDR length accuracy is ± 5 m
- Does not work on Smart Rate Interfaces with 10GBASE-T and NGBASE-T (2.5G, 5G copper) ports available on:
    - v3 blades
        - J9991A — Aruba 20-port 10/100/1000BASE-T PoE+ / 4-port 1/2.5/5/10GBASE-T PoE+ MACsec v3 zl2 Module
        - J9995A — Aruba 8-port 1/2.5/5/10GBASE-T PoE+ MACsec v3 zl2 Module
    - 3810M (JL076A — Aruba 3810M 40G 8 HPE Smart Rate PoE+ 1-slot Switch)
- Not supported on v2 zl modules
- Valid only on 100BASE-TX and 1000BASE-T ports

Overview

The Link Layer Discovery Protocol (LLDP) is a vendor-neutral link layer protocol in the Internet Protocol Suite used by Aruba network devices for advertising their identity, capabilities, and neighbors on an IEEE 802 local area network, principally wired ethernet. The LLDP-bypass authentication feature provides zero touch provisioning of Aruba 802.11ac wireless access points (APs).

In an LLDP module, the packet is parsed and inspected for the presence of an Aruba Organizational Unit Identifier (OUI) Type-Lenght-Value (TLV). The Aruba OUI TLV, once detected, will bypass the authentication and permit traffic to pass on the port. If the Aruba OUI TLV is absent, the packet will be dropped and processing of the packet or LLDP transmission for that device will not pass.

In ZTP environments, when an Aruba AP is plugged into the switch port, the device profiles will be applied on the AP without any user intervention. After discovery of an Aruba AP, the switch will dynamically provision the AP connected port without initiating any authentication needs. This feature is enabled at the port-level or on a range of ports.

Features not supported

- Authorization parameters configured in RADIUS and the switch are not supported by the LLDP-bypass authentication feature.

Configuration commands

aaa port-access lldp-bypass

From within the configure context:

Syntax

`[no] aaa port-access lldp-bypass`

Description

The command configures lldp-bypass authentication on the switch ports.

Configure lldp-bypass on the switch ports to bypass authentication for Aruba-APs which sends special LLDP TLVs. When LLDP-BYPASS is enabled on a port, it will behave as all other devices and bypass LLDP PDU for all device-profile enabled device.

By default, lldp-bypass is disabled on the switch ports.

Usage

`[no] aaa port-access lldp-bypass <PORT-LIST>`

Description

Validation rules

Validation	Error/Warning/Prompt

When the lldp-bypass is enabled on the port, different error messages are displayed.	If MAC lockdown is enabled on the port:
	`Error configuring port A1: lldp-bypass cannot be enabled on a port with MAC lock-enabled.`
	If learn-mode is configured on the port:
	`A1: lldp-bypass cannot be enabled on the port with learn-mode configured.`
	If MACsec is configured on the port:
	`Error configuring port A1: lldp-bypass cannot be enabled on the port with MACsec-enabled.`
	If trunk is configured on the port:
	`Error configuring port A1: lldp-bypass cannot be enabled on the port with mesh or manual trunks configured.`
	If mesh is configured on the port:
	`lldp-bypass cannot be enabled on the port with mesh or manual trunks configured.`
	If Distributed Trunking is configured on the port:
	`lldp-bypass cannot be enabled on the port with mesh or manual trunks configured.`

When MACsec is enabled on the port:	If lldp-bypass is enabled on the port: `Cannot apply MACsec on the port A1 when lldp-bypass is enabled on that port.`
When learn-mode is configured on the port:	If lldp-bypass is enabled on the port: `A1: Cannot apply learn-mode on the port A1 when lldp-bypass is enabled on that port.`

Table Continued

Validation	Error/Warning/Prompt
When trunk, distributed trunk or mesh is configured on the port:	If lldp-bypass is enabled on the port: `Cannot apply mesh or manual trunks on the port A1 when lldp-bypass is enabled on that port.`
When MAC-lockdown is enabled on the port:	If lldp-bypass is enabled on the port: `Cannot apply MAC lock-enable on the port A1 when lldp-bypass is enabled on that port.`
Security Warning when enabling lldp-bypass on the port.	`Enabling lldp-bypass on the port may give access to any Aruba-AP that sends a special LLDP TLV without undergoing any authentication. This configuration may allow network access to the rogue devices that are capable of sending the special LLDP TLV` `Do you want to continue? [y/n]:`

Show commands

show port-access lldp-bypass clients

Syntax

`show port-access lldp-bypass clients`

Description

Displays the clients which bypassed the authentication.

Options

<PORT-LIST>

Show information for specified ports only.

Usage

`show port-access lldp-bypass clients [<PORT-LIST>]`

show port-access lldp-bypass clients

```
switch#show port-access lldp-bypass clients

Port Access lldp-bypass Client Status
Port      MAC Address
-------   ----------------
A1        000005-010203
A2        010203-040506
```

Stackable switch: show port-access lldp-bypass clients

```
switch(config)# show port-access lldp-bypass clients

Port Access lldp-bypass Client Status
Port     MAC Address
-------   -----------------
1/1       000005-010203
1/2       005056-bd7039
```

show port-access lldp-bypass clients A1

```
switch#show port-access lldp-bypass clients A1

Port Access lldp-bypass Client Status
Port     MAC Address
-------   -----------------
A1       000005-010203
```

Stackable switch: show port-access lldp-bypass clients 1/1

```
switch(config)# show port-access lldp-bypass clients 1/1

Port Access lldp-bypass Client Status
Port     MAC Address
-------   -----------------
1/1       000005-010203
```

show port-access lldp-bypass config

Syntax

`show port-access lldp-bypass config`

Description

Displays the lldp-bypass configuration applied on all switch ports.

show port-access lldp-bypass config

```
switch#show port-access lldp-bypass config

Port Access lldp-bypass Configuration
Port   Enabled
------  ----------
A1     Yes
A2     Yes
A3     No
A4     No
...
A24    No
F1     No
F2     No
F3     No
```

```
F24    No
```

Stackable switch: show port-access lldp-bypass config

```
switch(config)#show port-access lldp-bypass config

Port Access lldp-bypass Configuration
Port   Enabled
------  ----------
1/1    Yes
1/2    Yes
1/3    No
...
1/52   No
2/1    No

2/26   No
3/1    No

3/26   No
```

Error Log

Event	Message
CLIERR_CANNOT_ENABLE_LLDP_BYPASS_MAC_LOCKDOWN_ENABLED	lldp-bypass is not allowed on the port where MAC-lockdown is enabled. `lldp-bypass cannot be enabled on a port with MAC lock-enabled.`
CLIERR_MACLOCK_AND_LLDP_BYPASS	MAC-lockdown is not permitted on the port where is enabled lldp-bypass. `Cannot configure MAC lock-enable on the port A1 when lldp-bypass is enabled on that port.`
CLIERR_CANNOT_ENABLE_LLDP_BYPASS_MACSEC_ENABLED	lldp-bypass is not allowed on the port MACsec is configured. `lldp-bypass cannot be enabled on a port when MACsec is enabled.`
CLIERR_CANNOT_ENABLE_MACSEC_AS_LLDP_BYPASS_CONFIGURED	MACsec is not permitted on the port where is enabled lldp-bypass. `Cannot apply MACsec on the port A1 when lldp-bypass is enabled on that port.`

Table Continued

Event	Message
CLIERR_CANNOT_ENABLE_LEARN_MODE_CONFIGURED_LLDP_BYPASS	Port-security learn-mode configured is not permitted when lldp-bypass is enabled on the port. `A1: Cannot apply learn-mode on the port A1 when lldp-bypass is enabled on that port.`
CLIERR_LLDP_BYPASS_AND_LEARN_MODE_CONFIGURED	lldp-bypass is not permitted when port-security learn-mode is configured. `lldp-bypass cannot be enabled on a port when learn-mode is enabled.`
CLIERR_LLDP_BYPASS_AND_MESH_OR_MANUAL_TRUNK	Trunk/ mesh/Distributed Trunk is not permitted on the lldp-bypass enabled port. `Cannot apply mesh or manual trunks on the port A1 when lldp-bypass is enabled on that port.`
Existing Log:CLIERR_MESH_OR_MANUAL_TRUNK	lldp-bypass cannot be enabled for trunk/mesh/ Distributed Trunk ports. `lldp-bypass cannot be enabled on a port when mesh or manual trunks is enabled.`

Debug log

Comment	Message
Security warning to be displayed when lldp-bypass configuration is enabled on the port.	`Enabling lldp-bypass on the port may give access to any Aruba-AP that sends a special LLDP TLV without undergoing any authentication. This configuration may allow network access to the rogue devices that are capable of sending the special LLDP TLV` `Do you want to continue? [y/n]:`
When adding the Aruba-AP into the authorized client list.	`Will use the existing debug log:` `0000:00:24:25.07 PSEC mPORTSECMCtrl:added new SA 000005-000000 to authorized    addr list of port A1 for vlan 1.`
When removing the Aruba-AP from the authorized client list.	`Will use the existing debug log:` `0000:00:01:47.07 PSEC mPORTSECMCtrl:removed 000006-000000 from authorized addr list of port A1 for vlan 1 due to delete.`

Table Continued

Comment	Message
When Aruba-AP is detected on lldp-bypass enabled port:	`0000:00:13:57.64 PSEC mPORTSECMCtrl: Received PROFMGR_DEVICE_CONNECTED event for 40e3d6-c6d492 on port A1.`
When already connected Aruba-AP is disconnected/removed on lldp-bypass enabled port.	`0000:00:13:07.96 PSEC mPORTSECMCtrl: Received PROFMGR_DEVICE_DISCONNECTED event for 40e3d6-c6d492 on port A1.`

Net-service Overview

Net-service names are used as alias in defining ACL rules for defined lists. An alias of net-service will configure a list of hosts, networks, or subnets.

Extended ACL can have both source IP, destination IP and port number along with protocol in its ACE. An alias-based ACE for an extended ACL therefore allows the use of an alias of net-service protocol and destination port.

Limitations

- Alias-based ACE will not support access-control based on source port which is a limitation of the net-service command. The use of net-service will also restrict operators specified for port number to `equals` and `range`.
  - Operators `lt`, `gt`, `equal`,`negative`, and `range` for the source port in the ACL rule are not specified using the options available in net-service.
  - Operators `lt`, `gt`, `negative` are not specified for destination port using the options available in net-service.
  - Only the ACL will be affected when changes are made to an existing net-service. Either the rule must be reapplied to the ACL or the switch must be rebooted to affect the service.

netservice [tcp | udp | port]

Syntax

```
[no] netservice <NAME-STR> [tcp | udp | <PROTOCOL>]
 port <PORT-LIST>
```

Description

Configures net-service.

Parameters

protocol

IP protocol number.

Range: 0-255

TCP

Configure an alias for a TCP protocol.

UDP

Configure an alias for a UDP protocol.

port

Specify a single port or a list of noncontiguous port numbers, by entering up to six port numbers, separated by commas or range of ports.

Range: 0-65535

Example net-service tcp-service tcp 100

```
net-destination src-ip
```

```
      host 10.120.0.1
      host 10.91.1.1
      host 10.0.100.12

   net-destination destn-ip
      host 16.90.51.12
      host 10.93.24.1

   net-service tcp-service tcp 100
   ip access-list extended “acl1”
      permit alias src-ip alias destn-ip alias tcp-service
```

Net-destination overview

The use of net-destination and net-service helps reduce effort required to configure ACL rules.

Net-destination is a list of hosts, networks, or subnets that are used to configure an ACL rules.

There are two types of ACLs supported and configured on the switch:

- Standard
- Extended

Standard

The standard ACL can have an IP source or network in the ACE. Defining the alias-based ACE for standard ACL, only use an alias of net-destination for the source.

Example - standard

```
net-destination src-ip
      host 10.120.0.1
      host 10.91.1.1
      host 10.0.100.12

   net-destination destn-ip
      host 16.90.51.12
      host 10.93.24.1

   net-service tcp-service tcp 100
   ip access-list extended “acl1”
      permit alias src-ip alias destn-ip
         alias tcp-service
```

Extended

The extended ACL can have both source IP, destination IP and port number along with protocol in its ACE. Defining an alias-based ACE for an extended ACL can use an alias of net-destination for the source and destination and an alias of net-service for the protocol and destination port. Alias-based ACE will not support access-control based on source port which is a limitation of the net-service command. The use of net-service will also restrict the operators that can be specified for port number to `equals`and `range`.

Example - extended

```
HP-Switch-5406Rzl2(config)# ip access-list extended aext1
HP-Switch-5406Rzl2(config-ext-nacl)#
   permit tcp host 10.100.12.1 gt 23 16.90.0.0 /16 range 200 400
HP-Switch-5406Rzl2(config-ext-nacl)# exit
```

Limitations

- Limited to IPv4 addresses per syntax.
- Any changes made to an existing net-destination that is used by an ACL, will be applied on the ACL only when the rule is reapplied to it or when switch is rebooted.
- The number of entries for a single net-destination is limited. The number of net-destinations configurable on a switch is also limited.
- A considerable amount of memory (for global structures) will be allocated when alias-based ACEs are configured which may cause issues on a switch with low memory.
- The Host or Domain name cannot be specified as an entry in a net-destination.
- Application level gateway will not be supported as the existing ACL infra does not support ALG.
- SNMP support to configure and delete net-destination, net-service, and the alias-based rules will not be provided.
- The ‘invert’ and ‘range’ option have been deprecated as per ArubaOS-Switch 7.4 CLI Reference Guide and hence will not be supported. However, the functionality of ‘invert’ option can be achieved through the ‘deny’ rule.
- RADIUS server-based ACL application to interface/VLAN will not be supported for ACLs with alias-based rules.

net-destination host |position | network

Syntax

```
[no] netdestination <NAME-STR> [host <IP-ADDR>
 [position <NUM>] network <IP-ADDR/MASK-LENGTH>
 [position <NUM>]]
```

Description

Net-destination is a list of hosts, networks, or subnets that are used to configure an ACL rule.

Parameters

host

Configures a single IPv4 host.

network

An IPv4 subnet consisting of an IP address and netmask.

no

Removes any configured item in list or an entire net-destination.

position

Specifies the position of a host, network, or range in the net-destination. This optional parameter is specific to a net-destination and may only be used to sort entries in a list.

show net-destination

Syntax

```
show net-destination <NAME-STR>
```

Description

Show a host-specific net-destination.

Static IP Visibility

Only IP addresses assigned by the DHCP server are visible in RADIUS accounting on an ArubaOS-Switch. Visibility of statically assigned IP addresses in RADIUS accounting is available with a command that enables and disables static IP visibility for an authenticated client.

IP client-tracker

Syntax

```
ip client-tracker [trusted | untrusted]

no ip client-tracker [trusted | untrusted]
```

Description

Enables the visibility of statically and dynamically assigned IPv4 and IPv6 addresses for both authenticated and unauthenticated client.

The `no` form of this command disables the visibility of statically and dynamically assigned IPv4 and IPv6 addresses for both authenticated and unauthenticated client.

Command context

```
config
```

Parameters

`trusted`

Enables or disables the visibility of statically and dynamically assigned IPv4 and IPv6 addresses for authenticated clients.

`untrusted`

Enables or disables the visibility of statically and dynamically assigned IPv4 and IPv6 addresses for unauthenticated client.

Usage

- Enabling this feature will send ARP probes to the client at regular intervals. This interval is determined by setting `arp-age timeout`. By default `arp-age timeout` is 20 minutes however the default timeout can be changed by using the command `ip arp-age` *`<timeout value in minutes>`*.
  - The periodic ARP probe aids in detecting any change of IP addresses on end clients.
  - The periodic probe aids silent clients to stay connected to the switch. Silent clients do not send any traffic to the switch after authentication which causes de-authenticated after a logoff period of 5 minutes (default value) due to inactivity. The periodic ARP probe triggers the end client to send response packets and aids

the client in staying connected. The customer must manually configure the IP arp-age value to 1 minute with the command `ip arp-age 1` to avoid being de-authenticated.

- When the `ip client-tracker` command is executed more than once, it takes the last commands' behavior. For example when the command `ip client-tracker trusted` is run after the command `ip client-tracker`, the behavior will follow the last command, `ip client-tracker trusted`.
    - When the administrator tries to execute the `no` command that has not been configured (does not exist in running configuration), an error will appear.

Example

Show port-access client with multiple addresses.

```
switch#  show port-access clients

Port Access Client Status

Port    Client Name    MAC Address           IP Address      User Role Type  VLAN
----   -------------- -------------- ---------------------- ---------- ----- ----
1       005056bd3ff7   005056-bd3ff7  3ffe:501:ffff:100::5e                MAC  1
```

Example

Show the port-access IPv4 client.

```
Switch-Stack(config)# show port-access clients

 Port Access Client Status

  Port  Client Name   MAC Address   IP Address      User Role          Type  VLAN
  ----- ------------- ------------- --------------- ------------------ ----- ----
  1/3   000002b85001  000002-b85001 10.1.1.30                          MAC   10
```

Example

Show the port-access IPv6 client.

```
switch(config)# show port-access clients 22

Port Access Client Status

  Port  Client Name   MAC Address   IP Address      User Role          Type  VLAN
  ----- ------------- ------------- --------------- ------------------ ----- ----
  22    0000005daa34  000000-5daa34 n/a                                MAC   20
```

Example

Show the port-access client detail.

```
switch(config)# show port-access clients 22 detailed

Port Access Client Status Detail

  Client Base Details :
   Port             : 22                   Authentication Type : mac-based
   Client Status    : authenticated        Session Time        : 64 seconds
   Client Name      : 0000005daa34         Session Timeout     : 0 seconds
   MAC Address      : 000000-5daa34
   IP               : n/a

  Access Policy Details :
   COS Map          : Not Defined          In Limit Kbps       : Not Set
   Untagged VLAN    : 20                   Out Limit Kbps      : Not Set
```

```
   Tagged VLANs    : No Tagged VLANs
   Port Mode       : 1000FDx
   RADIUS ACL List : No Radius ACL List
   IPV6 Address    : 2000::10
```

Chapter 37
Network Out-of-Band Management (OOBM)

OOBM concepts

Management communications with a managed switch can be:

- In band—through the networked data ports of the switch
- Out of band—through a dedicated management port (or ports) separate from the data ports

Out-of-band ports have typically been serial console ports using DB-9 or specially wired 8-pin modular (RJ-style) connectors. Some recent switches have added networked OOBM ports.

Figure 173: *Switch management port identification*

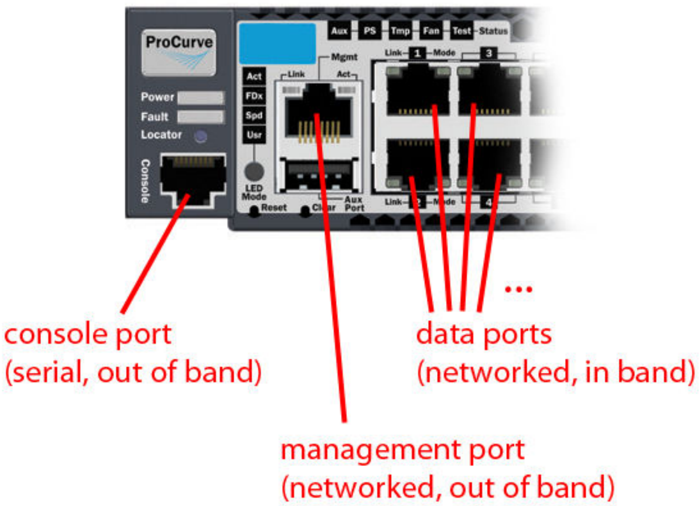

OOBM operates on a "management plane" that is separate from the "data plane" used by data traffic on the switch and by in-band management traffic. That separation means that OOBM can continue to function even during periods of traffic congestion, equipment malfunction, or attacks on the network. In addition, it can provide improved switch security: a properly configured switch can limit management access to the management port only, preventing malicious attempts to gain access via the data ports.

Network OOBM typically occurs on a management network that connects multiple switches. It has the added advantage that it can be done from a central location and does not require an individual physical cable from the management station to each switch's console port.

Table 28: *Switch management ports*

	In band	Out of band	
	Networked	**Directly connected**	**Networked**
Management interface	Command line (CLI), menu, Web	Command line (CLI), menu	Command line (CLI), menu
Communication plane	Data plane	Management plane	Management plane
Connection port	Any data port	Dedicated serial or USB console port	Dedicated networked management port
Connector type	Usually RJ-45; also CX4, SFP, SFP+, and XFP	DB9 serial, serial-wired 8-pin RJ	RJ-45
Advantages	Allows centralized management	Not affected by events on data network, shows boot sequence	Not affected by events on data network, allows centralized management, allows improved security
Disadvantages	Can be affected by events on data network; does not show boot sequence	Requires direct connection to console port (can be done via networked terminal server)	Does not show boot sequence

OOBM and switch applications

The table below shows the switch applications that are supported on the OOBM interface as well as on the data interfaces. In this list, some applications are client-only, some are server-only, and some are both.

Application	Inbound OOBM (server)	Outbound OOBM (client)	Inbound data plane (server)	Outbound data plane (client)
Telnet	yes	yes	yes	yes
SSH	yes		yes	[1]
SNMP	yes	yes	yes	yes
TFTP	yes	yes	yes	yes
HTTP	yes	[1]	yes	[1]
SNTP	[1]	yes	[1]	yes
TIMEP	[1]	yes	[1]	yes
RADIUS	[1]	yes	[1]	yes
TACACS	[1]	yes	[1]	yes

Table Continued

Application	Inbound OOBM (server)	Outbound OOBM (client)	Inbound data plane (server)	Outbound data plane (client)
DNS	[1]	yes	[1]	yes
Syslog	[1]	yes	[1]	yes
Ping	yes	yes	yes[4]	yes
Traceroute	yes[4]	yes	yes[4]	yes

[1] N/A = not applicable

[4] ***=Ping and Traceroute do not have explicit servers. Ping and Traceroute responses are sent by the host stack.

For applications that have servers, `oobm/data/both` options have been added to listen mode. There is now a `listen` keyword in the CLI commands to allow selection of those options. Default **value** is `both` for all servers.

Example

In a typical data center installation, top-of-rack switches connect servers to the data network, while the management ports of those switches connect to a physically and logically separate management network. This allows network administrators to manage the switches even if operation on the data network is disrupted.

In **Figure 174: Network OOBM in a data center** on page 688, the switches face the hot aisle of the data center, allowing easy connection to the network ports on the backs of the servers.

Figure 174: *Network OOBM in a data center*

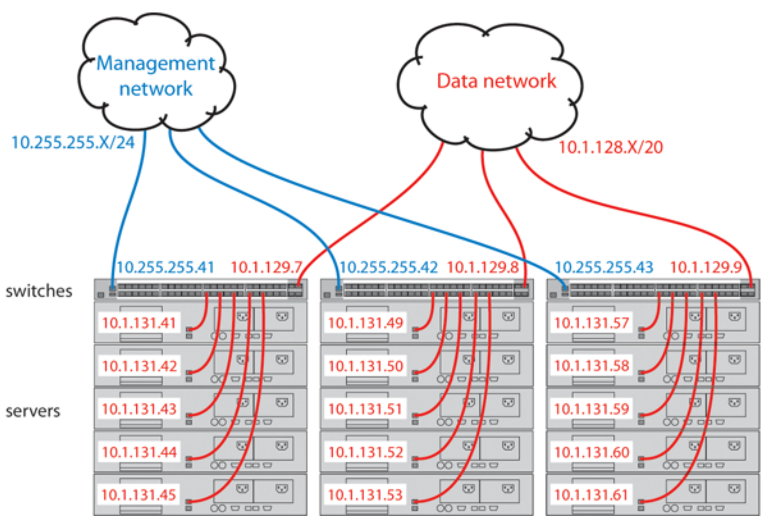

For even more control, the serial console ports of the switches can be connected to the management network through a serial console server (essentially, a networked serial switch), allowing the network administrators to view the CLI activity of each switch at boot time and to control the switches through the console ports (as well as through the management ports.)

OOBM Configuration

OOBM configuration commands can be issued from the global configuration context (config) or from a specific OOBM configuration context (oobm.)

Entering the OOBM configuration context from the general configuration context

Syntax

```
oobm
```

Enters the OOBM context from the general configuration context.

Example

```
switch(config)# oobm
HP Switch (oobm)#
```

Enabling and disabling OOBM

From the OOBM context:

Syntax

```
enable disable
```

From the general configuration context:

Syntax

```
oobm enable oobm disable
```

Enables or disables networked OOBM on the switch.

OOBM is not compatible with either a management VLAN or stacking. If you attempt to enable OOBM when a management VLAN is enabled or when stacking is enabled, the command will be rejected and you will receive an error message.

If an OOBM IP address exists and you disable OOBM, the OOBM IP address configuration is maintained. If you enable OOBM and there is a pre-existing OOBM IP address, it will be reinstated.

Network OOBM is enabled by default.

Examples

```
HP Switch (oobm)# enable
HP Switch (oobm)# disable
switch(config)# oobm enable
switch(config)# oobm disable
```

Enabling and disabling the OOBM port

The OOBM `interface` command enables or disables the OOBM interface (that is, the OOBM port, as opposed to the OOBM function.)

From the OOBM context:

Syntax

```
interface [enable | disable]
```

From the general configuration context:

Syntax

```
oobm interface [enable | disable]
```

Enables or disables the networked OOBM interface (port.)

Examples

```
HP Switch (oobm)# interface enable
switch(config)# oobm interface disable
```

Setting the OOBM port speed

The OOBM port operates at 10 Mbps or 100 Mbps, half or full duplex. These can be set explicitly or they can be automatically negotiated using the `auto` setting.

From the OOBM context:

Syntax

```
interface speed-duplex [10-half | 10-full | 100-half | 100-full | auto]
```

From the general configuration context:

Syntax

```
oobm interface speed-duplex [10-half | 10-full | 100-half | 100-full | auto]
```

Enables or disables the networked OOBM interface (port.) Available settings are:

`10-half`	10 Mbps, half-duplex
`10-full`	10-Mbps, full-duplex
`100-half`	100-Mbps, half-duplex
`100-full`	100-Mbps, full-duplex
`auto`	auto negotiate for speed and duplex

Example

```
HP Switch (oobm)# interface speed-duplex auto
```

Configuring an OOBM IPv4 address

Configuring an IPv4 address for the OOBM interface is similar to VLAN IP address configuration, but it is accomplished within the OOBM context.

From the OOBM context:

Syntax

```
[no] ip address [dhcp-bootp | ip-address/mask-length]
```

From the general configuration context:

Syntax

```
[no] oobm ip address [dhcp-bootp | ip-address/mask-length]
```

Configures an IPv4 address for the switch's OOBM interface.

You can configure an IPv4 address even when global OOBM is disabled; that address will become effective when OOBM is enabled.

Example

```
HP Switch (oobm)# ip address 10.1.1.17/24
```

Configuring an OOBM IPv4 default gateway

Configuring an IPv4 default gateway for the OOBM interface is similar to VLAN default gateway configuration, but it is accomplished within the OOBM context.

From the OOBM context:

Syntax

`[no] ip default-gateway` *`ip-address`*

From the general configuration context:

Syntax

`[no] oobm ip default-gateway` *`ip-address`*

Configures an IPv4 default gateway for the switch's OOBM interface.

Example

```
HP Switch (oobm)# ip default-gateway 10.1.1.1
```

Configuring an IPv6 default gateway for OOBM devices

An OOBM interface is used for managing devices from remote sites. OOBM devices must be given a default gateway responsible to maintain a network connection when these devices are placed in an IPv6 network enabled with RA suppression.

To configure and enable an IPv6 default gateway for OOBM interfaces, use the `oobm ipv6 default-gateway` command.

oobm ipv6 default-gateway

Syntax

`oobm ipv6 default-gateway` *`<IPV6-ADDR>`*

`no oobm ipv6 default-gateway`

Description

Configures the IPv6 default gateway address for OOBM interfaces.

The *no* form of the command deletes the default gateway. It is imparative that an IPv6 address is specified when the *no* form of the command is used.

Command context

`config`

Parameters

<IPV6-ADDR>

Specifies the IPv6 address when configuring the OOBM for a specific gateway.

Example

```
switch(config)# oobm ipv6 default-gateway 1001::1/64
```

oobm member ipv6 default-gateway

Syntax

`oobm member <MEMBER-ID> ipv6 default-gateway <IPV6-ADDR>`

`no oobm member <MEMBER-ID> ipv6 default-gateway`

Description

Configures the IPv6 default gateway address for an OOBM member using their unique identifier and the IPv6 address of the default gateway.

When *no* proceeds the command, the default gateway address is deleted.

Command context

`config`

Parameters

<MEMBER-ID>

Specifies the unique member-id which allows the OOBM device access to the IPv6 default-gateway.

<IPV6-ADDR>

Specifies the IPv6 address of the default gateway for a member OOBM interface.

Example

Configuring and deleting the OOBM member from a specific IPv6 gateway.

```
switch(config)# oobm member 1 ipv6 default-gateway 1001::1

switch(config)# no oobm member 1 ipv6 default-gateway
```

IPv6 default router preferences

The command `ipv6 nd ra router-preference {high | medium | low}` provides an extension to the Neighbor Discovery Router Advertisement messages for communicating default router preferences from routers to hosts. The extension improves the ability of hosts to pick the appropriate router for an off-link destination. In network topologies, where the host has multiple routers on its Default Router list, the choice of router for an off-link destination is critical for making the communication more efficient. For example, one router may provide much better performance than another for a destination while choosing a wrong router may result in failure to communicate.

`ipv6 nd ra router-preference`

Syntax

```
ipv6 nd ra router-preference {low | medium | high}

no ipv6 nd ra router-preference
```

Description

Sets the router-preference configuration for communicating default router preferences from routers to hosts. Improves the ability of hosts to pick the appropriate router for an off-link destination by providing options at the operator level which set the router preference value as low, medium, or high. Depending on the router preference value set, the host receives the value as part of the IPv6 neighbor discovery router advertisement and chooses the best router for communication.

The `no` form of this command removes the router-preference configuration.

Command context

`vlan`

Parameters

`low`

Specifies the router-preference value to low.

`medium`

Specifies the router-preference value to medium. Medium is the default router-preference value.

`high`

Specifies the router-preference value to high.

Usage

- When VRRP is configured, RA messages for Virtual IP are sent with configured router-preference values.
- This command complies with RFC 4191.

Example

Default router-preferences.

```
switch(vlan-1)# ipv6 nd ra router-preference
 low                    Set the Default Router Preference to low.
 medium                 Set the Default Router Preference to medium (default).
 high                   Set the Default Router Preference to high.

switch(vlan-1)# ipv6 nd ra router-preference high
```

OOBM `show` commands

The `show` commands for OOBM are similar to the analogous commands for the data plane. Note that you must always include the `oobm` parameter to see the information for the OOBM interface, regardless of the context. For instance, even from the OOBM context, the `show ip` command displays the IP configuration for the data plane; to see the IP configuration of the OOBM interface, you need to use `show oobm ip`.

Showing the global OOBM and OOBM port configuration

Syntax

```
show oobm
```

Summarizes OOBM configuration information. This command displays the global OOBM configuration (enabled or disabled), the OOBM interface status (up or down), and the port status (enabled/disabled, duplex, and speed.)

You can issue this command from any context

Example

```
HP Switch# show oobm

Global Configuration
  OOBM Enabled    : Yes
  OOBM Port Type   : 10/100TX
  OOBM Interface Status : Up
  OOBM Port    : Enabled
  OOBM Port Speed   : Auto
```

Showing OOBM IP configuration

Syntax

```
show oobm ip
```

Summarizes the IP configuration of the OOBM interface. This command displays the status of IPv4 (enabled/disabled), the IPv4 default gateway, and the IPv4 address configured for the interface.

You can issue this command from any context.

Example

```
HP Switch# show oobm ip
```

Showing OOBM ARP information

Syntax

```
show oobm arp
```

Summarizes the ARP table entries for the OOBM interface.

You can issue this command from any context.

Example

HP Switch# show oobm arp

`show oobm ipv6`

Syntax

`show oobm ipv6`

Description

Shows the IPv6 service status for OOBM interfaces.

Command context

`operator`

Example

Shows the IPv6 service status for OOBM interfaces.

```
switch# show oobm ipv6

Internet (IPv6) Service for OOBM Interface


  IPv6 Status     : Enabled
  IPv6 Default Gateway : 1000::2

                                                               Address   Intf
 Member IP Config  IP Address/Prefix Length                   Status    Status
 ------ ---------- ------------------------------------------ --------- ------
 Global manual     1000::1/64                                           Up
 Global autoconfig fe80::42a8:f0ff:fe9e:901/64                          Up
```

show oobm ipv6 (for stacked switches)

Syntax

`show oobm ipv6`

Description

Shows the OOBM IPv6 interface for a stacked switch.

Command context

`operator`

Example

Shows the OOBM IPv6 interface for a stacked switch.

```
stack-switch# show oobm ipv6

 Internet (IPv6) Service for OOBM Interface


  IPv6 Status     : Enabled
  IPv6 Default Gateway : 3000::2

                                                               Address   Intf
 Member IP Config  IP Address/Prefix Length                   Status    Status
 ------ ---------- ------------------------------------------ --------- ------
 Global manual     3000::1/64                                 Active    Down
 Global autoconfig fe80::42a8:f0ff:fe9b:a581/64               Active    Down
 1      manual     1000::1/64                                 Active    Down
 2      manual     2000::1/64                                 Active    Up
```

show oobm ipv6 member (for stacked switches)

Syntax

```
show oobm ipv6 member <MEMBER-ID>
```

Description

Shows the OOBM IPv6 service detail for a specific member.

Command context

`operator`

Example

Shows the OOBM IPv6 service detail for a specific member.

```
stack-switch# show oobm ipv6 member 2

 Internet (IPv6) Service for OOBM Interface


  IPv6 Status    : Enabled
  IPv6 Default Gateway : 1000::1

                                                                Address   Intf
 Member IP Config  IP Address/Prefix Length                     Status    Status
 ------ ---------- -------------------------------------------- --------- ------
 2      manual     1000::2/64                                   Active    Up
```

show oobm ip detail (for stacked switches)

Syntax

`show oobm ip detail`

Description

Shows the OOBM IP detail for a stacked switch.

Command context

`operator`

Example

Show the OOBM IP detail for a stacked switch.

```
stack-switch# show oobm ip detail

 Internet (IP) Service for OOBM Interface

Global Configuration

  IPv4 Status    : Enabled
  IPv6 Status    : Enabled

  IPv4 Default Gateway : 3.3.3.1
  IPv6 Default Gateway : 3000::1
```

```
 Origin     | IP Address/Prefix Length                    Status
 ---------- + ------------------------------------------- -----------
 manual     | 3.3.3.2/24                                  preferred
 manual     | 3000::2/64                                  preferred
 autoconfig | fe80::42a8:f0ff:fe9b:a581/64                preferred


Member 1

  IPv4 Status    : Enabled
  IPv6 Status    : Enabled

  IPv4 Default Gateway : 2.2.2.1
  IPv6 Default Gateway : 2000::1

 Origin     | IP Address/Prefix Length                    Status
 ---------- + ------------------------------------------- -----------
 manual     | 2.2.2.2/24                                  preferred
 manual     | 2000::2/64                                  preferred


Member 2

  IPv4 Status    : Enabled
  IPv6 Status    : Enabled

  IPv4 Default Gateway : 1.1.1.1
  IPv6 Default Gateway : 1000::1

 Origin     | IP Address/Prefix Length                    Status
 ---------- + ------------------------------------------- -----------
 manual     | 1.1.1.2/24                                  preferred
 manual     | 1000::2/64                                  preferred
```

Application server commands

Application servers (as described in OOBM and server applications in **OOBM concepts** on page 686) have added a `listen` keyword with `oobm|data|both` options to specify which interfaces are active.

Default value is `both` for all servers.

Syntax

```
telnet-server [listen {oobm | data | both}]
```

Syntax

```
ip ssh [listen {oobm | data | both}]
```

Syntax

```
snmp-server [listen {oobm | data | both}]
```

Syntax

```
tftp server [listen {oobm | data | both}]
```

Syntax

```
web-management [listen {oobm | data | both}]
```

In all cases, `show running-config` displays the server configurations.

Use the `no` form of the command to prevent the server from running on either interface.

Examples

Telnet: `no telnet-server`

SSH: `no ip ssh …`

SNMP: `no snmp-server …`

TFTP: `no tftp server`

HTTP: `no web-management …`

The `show servers` command shows the listen mode of the servers:

```
HP Switch# show servers
Server listen mode

Server   Listen mode
-----------------------------
Telnet        | both
Ssh          | both
Tftp         | both
Web-management | both
Snmp         | both
```

Application client commands

CLI commands for client applications have added the `oobm` keyword to allow you to specify that the outgoing request be issued from the OOBM interface. If you do not specify the `oobm` keyword, the request will be issued from the appropriate in-band data interface. Command syntax is:

Telnet:

```
telnet ip-address [oobm]
```

TFTP:

```
copy tftp ... ip-address filename... [oobm]
```

SNTP:

```
[no] sntp server priority priority ip-address [oobm] [version]
```

TIMEP:

```
[no] ip timep [dhcp | manual ip-address | [oobm]]
```

RADIUS:

```
[no] radius-server host ip-address [oobm]
```

TACACS+:

```
[no] tacacs-server host ip-address [oobm]
```

DNS:

```
[no] ip dns server-address priority priority ip-address [oobm]
```

Syslog:

```
[no] logging ip-address [[control-descr] | [oobm]]
```

Ping:

```
ping [...] [source [ip-address | vlan-id | oobm]]
```

Traceroute:

`traceroute [...][source [ip-address|vlan-id|oobm]]` *Management and Configuration Guide*

Example

Figure 175: Example data center on page 699 shows setup and use of network OOBM using the commands described above.

Assume that the figure below describes how you want to set up your data center.

Figure 175: *Example data center*

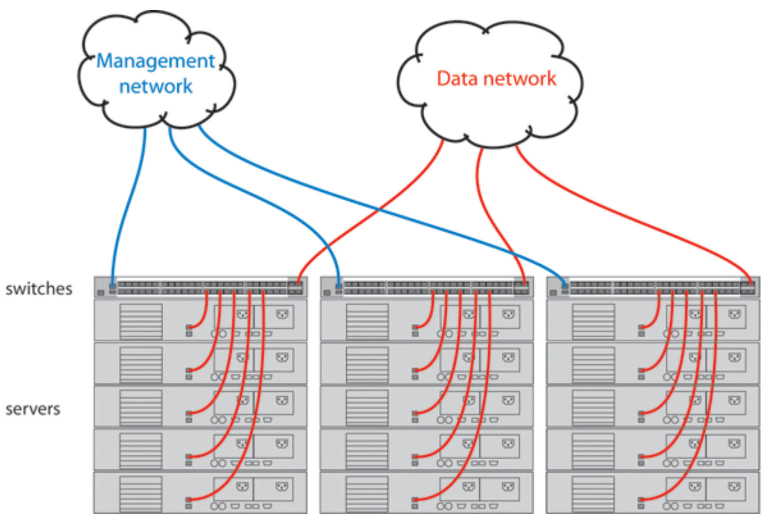

Assume that you are configuring the switch in the left-hand rack to communicate on both the data and management networks. You might do the following:

- Configure an IP address on the data network.
- Verify that out-of-band management is enabled. (It is enabled by default.)
- Configure an IP address on the management network.
- Verify that the switch can communicate on both networks.

The CLI commands that follow would accomplish those tasks. (The first time through the process you might easily make the omission shown near the end of the example.)

```
Switch 41# config
Switch 41(config)# vlan 1
Switch 41(vlan-1)# ip address 10.1.129.7/20        Set up IP address on data
network.
Switch 41(vlan-1)# end                                   Exit back to manager
context.
Switch 41# show oobm                                Look at default OOBM
```

```
configuration.

Global Configuration
  OOBM Enabled   : Yes
  OOBM Port Type  : 10/100TX
  OOBM Interface Status : Up                              Defaults look
appropriate.
  OOBM Port   : Enabled
  OOBM Port Speed  : Auto

Switch 41# config
Switch 41(config)# oobm                                      Go to OOBM context and
Switch 41(oobm)# ip address 10.255.255.41/24                     add IP address and
Switch 41(oobm)# ip default-gateway 10.255.255.1                     default
gateway.
Switch 41(oobm)# end                                   Exit back to manager
context.
Switch 41# ping 10.1.131.44              Ping server in this rack (on data
network.)
10.1.131.44 is alive, time = 19 ms
Switch 41# ping 10.1.131.51                            Ping server in adjacent
rack.
10.1.131.51 is alive, time = 15 ms
Switch 41# ping 10.255.255.42                          Ping switch in adjacent
rack.
The destination address is unreachable.        Oops! It’s on the management
network.
Switch 41# ping source oobm 10.255.255.42             Go through the management
port
10.255.255.42 is alive, time = 2 ms                               and it works
fine.
Switch 41#
```

Chapter 38 Websites

Networking Websites

Hewlett Packard Enterprise Networking Information Library

www.hpe.com/networking/resourcefinder

Hewlett Packard Enterprise Networking Software

www.hpe.com/networking/software

Hewlett Packard Enterprise Networking website

www.hpe.com/info/networking

Hewlett Packard Enterprise My Networking website

www.hpe.com/networking/support

Hewlett Packard Enterprise My Networking Portal

www.hpe.com/networking/mynetworking

Hewlett Packard Enterprise Networking Warranty

www.hpe.com/networking/warranty

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

For additional websites, see **Support and other resources**.

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  http://www.hpe.com/assistance
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  http://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center

  www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  www.hpe.com/support/downloads

  Software Depot

  www.hpe.com/support/softwaredepot
- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates
- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

  www.hpe.com/support/AccessToSupportMaterials

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

HPE ProLiant and x86 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Overview of chassis management redundancy

Some HPE switches provide high availability through the use of hot-swappable, redundant management modules. In the event of a failure on the active management module, management module redundancy allows a quick and unattended transition from the active management module to the standby management module. The standby management module now becomes the active management module. Management module redundancy keeps the switch operating and reduces network downtime.

The advantages of redundant management are:

- Maintaining switch operation if a hardware failure occurs on the active management module
- Minimizing restart time caused by the failure of a management module
- Hotswapping a failed management module with no downtime

Nonstop switching with redundant management modules

Beginning with software version K.15.01, you can use either nonstop switching or warm-standby redundant management.

The advantages of nonstop switching are:

- Quick, seamless transition to the standby management module; no reboot is necessary
- Switching of packets continues without interruption

How the management modules interact

When the switch boots up, the management modules run selftest to decide which is the active module and which is the standby module. The module that becomes active finishes booting and then brings up the interface modules and ports.

If you are using nonstop switching mode, the standby management module is synced continuously with the active management module so that all features and config files are the same on both management modules. The standby management module is ready to become the active management module. If the active management module fails or if there is a manual switchover, switching continues without interruption.

If you are using warm-standby mode, the standby module boots to a certain point, syncs basic files such as the config and security files, and finishes booting only if the active management module fails or you choose to change which module is the active module.

The two management modules communicate by sending heartbeats back and forth.

About using redundant management

The CLI commands for redundant management are shown at the beginning of the chapter. Additionally, other commands are affected by redundant management.

Transition from no redundancy to nonstop switching

While the switch is transitioning from no redundancy mode to nonstop switching mode, no configuration changes are allowed. The management modules are syncing information during the transition period.

About setting the rapid switchover stale timer

After a failover has occurred, use the rapid switchover stale timer to set the amount of time that you want route and neighbor table entries to be re-added to the FIB on the active management module.

Layer 3 applications and protocols rely on existing routing information in the FIB. They restart and operate as if the switch performed a quick reset.

When a failover occurs, the interface modules and the fabric modules continue forwarding Layer 3 traffic based on the information in the FIB. The transitioning standby management module marks all routes in the FIB as "stale". The routing protocols restart, reestablish their neighbors and reconverge. As a route is added in again, the route's stale designation is removed. After the rapid switchover stale timer expires, the remaining stale route entries are removed. Multicast flows are also removed; the multicast application re-adds the flows after failover completes.

About directing the standby module to become active

To make the standby management module become the active management module, use the `redundancy switchover` command. The switch will switchover after all files have finished synchronizing.

In nonstop switching mode:

- The switchover occurs quickly and seamlessly; no reboot is needed.
- There is no interruption in switching operations.

In warm-standby mode:

- The switchover may take several minutes if there have been recent configuration file changes or if you have downloaded a new operating system.
- The standby module finishes booting and becomes the active module.

The formerly active module becomes the standby module if it passes selftest.

Preferred management module

NOTE

ArubaOS-54XXR platform only

On a 54xxR chassis with the ability to support two management modules, one of the management modules can be persistently set as Active management module through a redundancy `preferred-management-module` command.

- A `preferred-active-module` may be configured by command to retain its role as an Active module under both hard power cycle and soft reboot scenarios.
- In the case of a failure of a `preferred-active-module`, the Standby management module takes over and becomes new Active management module.
- The `preferred-active-module` is seen as configured in the `show running config` command.

redundancy active-management

Syntax

```
redundancy active-management {management-module1 | management-module2 | standby}
```

Description

Specifies the management module that will become active at the next boot.

Command context

`config`

Parameters

`management-module` *`<1>`* `|` *`<2>`*

Specifies the redundant management module to enable.

`standby`

Specifies the standby management module.

Restrictions

- The `redundancy active-management` command will fail if the other module is in a failed state or if VSF is enabled.

Usage

`redundancy rapid-switchover <Seconds>`

Specifies the rapid-switchover timer in seconds for L3 Hitless forwarding when used as a parameter for redundancy.

`management-module`

The selected active management module will continue as the active management module on boot until the user issues the command `redundancy active-management` which then changes the selected module.

`redundancy switchover`

Instructs the switch to immediately switchover to the standby management module.

Example

Show redundancy for management module.

```
Switch(config)# show redundancy

Configured Mode: Nonstop Switching
Current Mode   : Nonstop Switching

Rapid Switchover Stale Timer : 90
Failovers     : 0
Last Failover :
Preferred Active Management: management-module1

Slot Module Description                       State    SW Version    Boot Image
---- ---------------------------------------- -------- ------------- ----------
MM1  HP J9827A Management Module 5400Rzl2     Active   KB.16.02.0014 Primary
MM2  HP J9827A Management Module 5400Rzl2     Standby  KB.16.02.0014 Primary
```

`redundancy preferred-active-management`

Syntax

```
redundancy preferred-active-management {management-module1 | management-module2}

no redundancy preferred-active-management {management-module1 | management-module2}
```

Description

Configures priority to specified management-module to be active management module across multiple boots. Command does not change management module roles instantaneously and would require a reboot for configuration to come into effect.

The `no` form of this command removes the configuration for the preferred-active-management module.

> **NOTE:** The command `redundancy preferred-active-management` requires a reboot for the configuration to change.

Command context

`config`

Parameters

`management-module` ***`<1>`*** **`|`** ***`<2>`***

Configures selected management module as an active management module across multiple reboots.

Restrictions

- The command `redundancy preferred-active-management` is mutually exclusive with VSF. The preferred-management-module must be disabled to enable VSF on an ArubaOS-switch.
- When a preferred module taking over as an Active module fails, a fail-over to another management module will occur. User intervention must rectify the cause for the failure or if a chassis reboot is attempted, the attempt may fail. In the case of a boot, there will be an extra redundancy switchover.

Example

Configuring preferred active management module.

```
switch(config)# redundancy preferred-active-management management-module1

switch(config)# show running config

Running configuration

; J9851A Configuration Editor; Created on release #KB.16.03.0000x
; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
hostname "HP-Switch-5412Rzl2"
module B type j9993a
module G type j9987a
redundancy preferred-active-management management-module1
snmp-server community "public" unrestricted
oobm
   ip address dhcp-bootp
   exit
vlan 1
   name "DEFAULT_VLAN"
   untagged B1-B8,G1-G24
   ip address dhcp-bootp
   exit
```

CLI warnings in response to `preferred-active-management` command

CLI Warnings	Scenario
Active-management cannot be set when preferred-active-management is enabled. Unconfigure before attempting to set active-management.	The command `active-management` is mutually exclusive with existing `redundancy active-management` command.
Preferred-active-management configuration overrides active-management configuration.	If `redundancy active-management` is set and user attempts to configure `preferred-active-management` a warning message displays when configuring `redundancy preferred-active-management`. NOTE: Once set, there is no way to unconfigure active-management setting.
VSF cannot be enabled when preferred-active-management is enabled. Un-configure before attempting to enable VSF.	The command `preferred-active-management` is mutually exclusive with VSF configuration. NOTE: No warning/error message would be displayed upon ports being configured as part of VSF.

show redundancy

Syntax

`show redundancy`

Description

Displays the status of the management and fabric modules.

Command context

`config`

Examples

Shows the redundancy of the switch for an active management module.

```
switch(config)# show redundancy

Configured Mode: Nonstop Switching
Current Mode   : Nonstop Switching

Rapid Switchover Stale Timer : 90
Failovers     : 0
Last Failover :
Preferred Active Management: management-module1

Slot Module Description                        State    SW Version    Boot Image
---- ---------------------------------------- -------- ------------- ----------
```

```
MM1   HP J9827A Management Module 5400Rzl2      Active   KB.16.02.0014 Primary
MM2   HP J9827A Management Module 5400Rzl2      Standby  KB.16.02.0014 Primary
```

Determining active module

Both management modules run selftest routines to determine which module becomes the active management module and which becomes the standby management module. The module that was last active in the chassis is given precedence and becomes the "active" module. This module is the one that is booted going forward. If a module fails selftest and is unable to communicate with the other module, it does not take control as the management module. The other management module takes control and becomes the active module.

If both modules fail selftest, the fault LED flashes and neither module is operational.

NOTE

You are not allowed to switchover to a management module that is not in standby mode. The module must have passed selftest and be in standby mode.

The entire boot decision process works as follows:

1. If there is only one management module, that is the active management module.
2. If one module is already booted and operational, a newly inserted module or the other management module booting always becomes the standby module. The standby module does not become active unless a switchover occurs.
3. If there are two management modules and one fails selftest, the one that passes selftest becomes the active management module.
4. If only one of two modules was ever booted in the chassis, that module is given precedence.
5. The module that was active on the last boot becomes the active management module. This guarantees that the active module has the latest configuration data.
6. If both management modules have previously booted in this chassis and were "active" the last time booted, the module that booted most recently becomes the active management module.
7. If none of the above conditions are applicable, the module in the lowest slot becomes the active management module.

Diagram of the decision process

Figure 176: *Active module decision flow chart at boot*

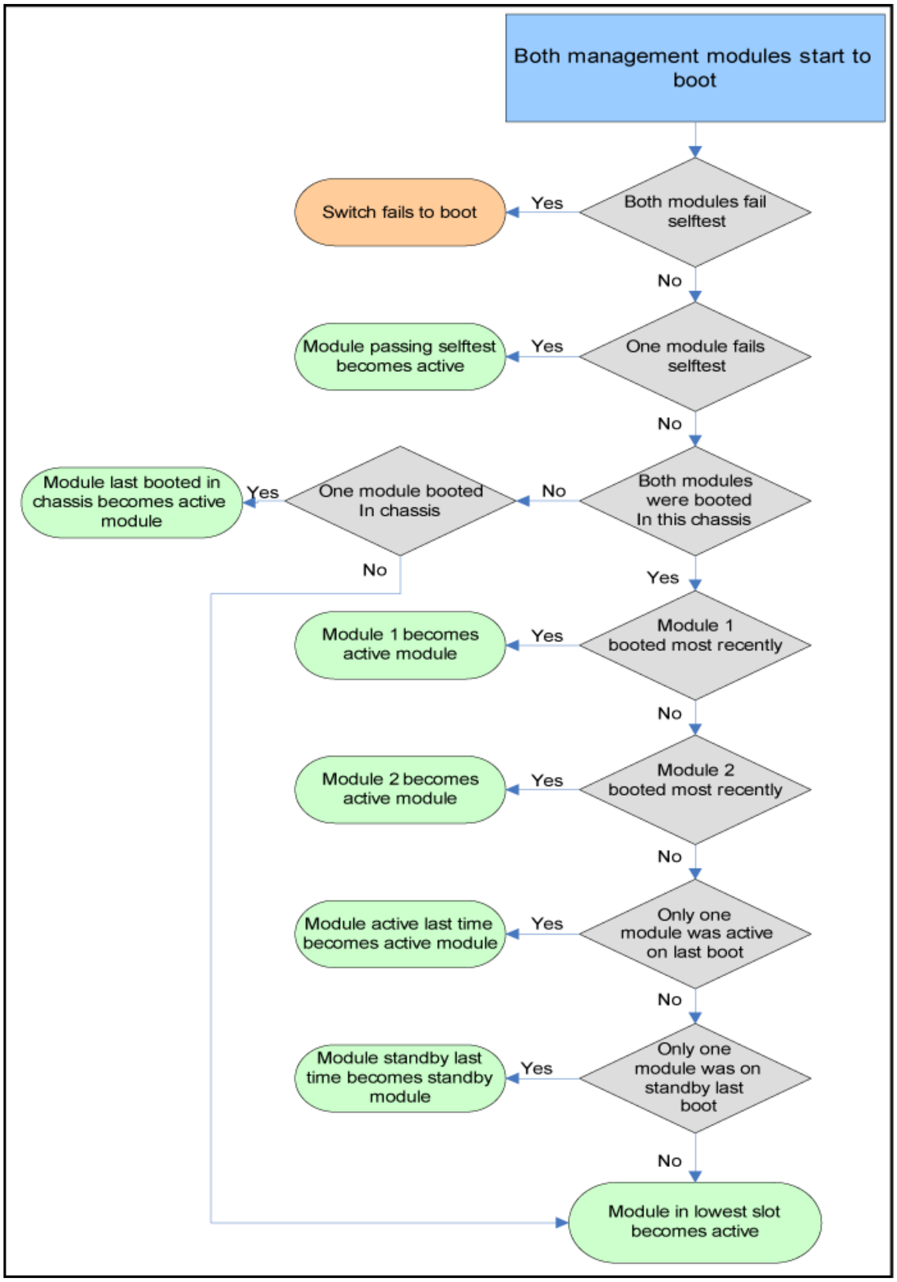

Hotswapping management modules

Hotswapping out the active management module

1. On the management module to be hotswapped out, press the **MM Shutdown** button. It is located between the **Module Operation** and **Component Status** LEDs. (See **Figure 177: The MM Shutdown button** on page 712.)

   Figure 177: *The MM Shutdown button*

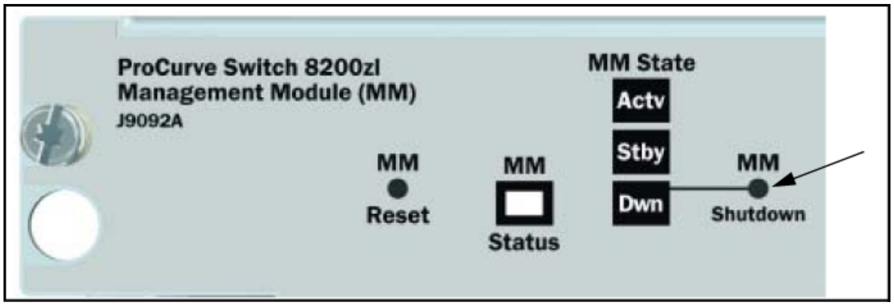

2. The **Dwn** LED to the right of the **MM Shutdown** button begins flashing green. File synchronization will complete before shutdown occurs.
3. The standby module takes control and the switchover occurs. It is now the active management module.
4. The **Dwn** LED on the management module being hotswapped out turns green and all other LEDs go out when it is OK to remove the module.
5. The module being hotswapped out goes into offline mode. In the offline mode, the module cannot take over when the active module fails over.

NOTE

If you remove the active management module without pressing the **MM Shutdown** button, any files that may have been in the process of synchronizing will not finish synchronizing to the standby module and all file transfer is aborted.

Management module switchover

Events that cause a switchover

There are a number of events that can cause the active management module to switchover to the standby management module when management module redundancy is enabled:

- The active management module crashes
- The standby management module does not receive a heartbeat from the active management module
- The `redundancy switchover` command is executed
- The **MM Reset** button on the active management module is pressed
- The **MM Shutdown** button on the active management module is pressed
- The `boot` or `boot active` command is executed
- The `reload` command is executed
- There is a hardware failure on the active management module

In all of these cases, the standby management module takes control and performs the actual switchover. The reason for the switchover is entered in log messages on the newly active management module and to any configured Syslog servers.

What happens when switchover occurs

When a switchover occurs, the features that support nonstop switching continue to operate in an uninterrupted manner. See **Nonstop switching features** on page 739 for a list of the supported features.

The features that do not support nonstop switching perform as if the switch had just finished booting; however, no actual boot time occurs.

NOTE

When meshing configuration changes are made on a redundant management system, you must execute `write mem` and then the `boot system` command to boot **both** management modules for the changes to be activated.

Meshing is not supported by nonstop switching.

NOTE

If the switch is a querier and a failover occurs, the querier continues to be the same on the standby management module; no new querier election process occurs on the standby management module.

When switchover will not occur

There are some events for which a switchover is not triggered:

- When a `boot system` command is executed
- When the **Clear** button on the System Support module is pressed
- When management module redundancy is disabled, unless there is a hardware failure and the system is rebooted.

When a management module crashes while the other management module is rebooting

If the uncommon situation occurs where the active management module (MM1) is trying to reboot and the standby management module (MM2) also crashes, the switch attempts to recover from the crash and eventually the standby management module becomes the active management module if it passes self-test. However, traffic can be disrupted for as long as five minutes before the newly active management module (MM2) has finished rebooting.

Hotswapping out the active management module

You can hotswap out the active management module and have switch operations taken over by the standby management module by following the correct shutdown procedure on the active module using the **MM Shutdown** button. When the **MM Shutdown** button is pressed, any file synchronization in progress completes before the shutdown begins, and then a graceful shutdown of that management module occurs.

When the standby module is not available

If you have disabled management module redundancy with the `no redundancy management-module` command, or the standby module failed selftest, the **Dwn** LED does not turn green to indicate it is OK to hotswap out the active management module.

NOTE

If you remove the active management module without pressing the **MM Shutdown** button, any files that may have been in the process of synchronizing will not finish synchronizing to the standby module and all file transfer is aborted.

Hotswapping in a management module

If another management module is hotswapped in while there is an active management module booted up, the newly hotswapped management module becomes the standby module.

No negotiating is needed as to which module becomes the active management module, because there is already a functioning active management module. However, the following conditions must be met to determine if the hotswapped module can become a standby management module:

- The hotswapped module must pass selftest
- Management module redundancy is not administratively disabled (using the `no redundancy management-module` command.) If the active management module's config file has redundancy administratively disabled, the hotswapped management module goes into "offline" mode.

In nonstop switching mode—The active management module's files and features are synced with the standby management module. Heartbeats are sent back and forth, and the standby management module is ready to quickly take over in the event of a switchover or a failure on the active management module.

In warm-standby mode—The standby management module partially boots up and heartbeats are sent back and forth with the active management module.

Software version mismatch between active and hotswapped module

If the software version in the hotswapped module does not match the software version in the active module, the following occurs:

Procedure

1. The active module sends the primary and secondary images in flash to the hotswapped module.
2. The module that was hotswapped in then reboots if necessary to primary or secondary flash, whichever matches (if it does not already match.)
3. After the hotswapped management module finishes booting, it is sent the config and other critical files from the active management module.
4. The hotswapped management module goes into standby mode and is ready to take over in case of a switchover.

NOTE

After the

`boot standby`

command is executed, if the software versions on the active management module and the standby management module are not compatible, the standby module does not sync with the active management module. The standby module then enters warm-standby redundancy mode.

Other software version mismatch conditions

The following steps describe the behavior that may when a new software image is installed in secondary flash of the AMM and a `redundancy switchover` command is executed.

1. A new software image, K.15.04.0002 containing ROM upgrade K.15.12 is installed in secondary flash of the AMM/MM1.
2. The AMM/MM1 automatically syncs the images to the secondary flash in the SMM/MM2. Now both AMM/MM1 and SMM/MM2 have identical software and ROM in secondary flash.
3. The SMM/MM2 is booted from secondary. It boots into the new K.15.04.0002 software version. The new ROM is applied and the SMM/ MM2 reboots.
4. After the SMM/MM2 finishes rebooting, it reconnects to the AMM/MM1 and prepares to take the standby role by rebooting.
5. However, the AMM/MM1 is running software version K.15.03.0008 in its primary flash, and the SMM/MM2 is running software version K.15.04.0002 in its secondary flash, so the SMM/MM2 pauses its reboot because of the software mismatch.
6. If a `redundancy switchover` command is executed, the AMM/MM1 will give control to the SMM/MM2, which can then finish booting and become the new AMM/MM2. This is the warm-start behavior.
7. The SMM/MM1 (former AMM/MM1) reboots, but unless the reboot is executed from secondary flash, it reboots into primary flash, which contains the older software version K.15.03.0008 with no ROM upgrade.
8. If the SMM/MM1 is forced to boot from secondary before executing the `redundancy switchover` command, it will boot into the new K.15.04.0002 software and upgrade the ROM. After the reboot that occurs with the ROM upgrade, the SMM/MM1 connects to the new AMM/MM2 and takes the standby role.

About turning off redundant management

Disable management module redundancy with two modules present

To troubleshoot a suspect management module, you may want to operate the switch with redundant management disabled by entering this command:

```
switch# no redundancy management-module
```

After executing this command, the second management module will not boot into standby mode—it is offline and no longer receives configuration file changes from the active module. The active management module updates its config file with the information that redundancy is disabled.

NOTE

Even if redundancy has been disabled, the specified management module becomes the active management module at the next system boot if you use the `redundancy active-management` command. You are warned that you may not be using current configurations. See **Setting the active management module for next boot** on page 723.

The second management module is enabled as the active management module in the event of a hardware failure of the first management module.

Figure 178: Results of disabling redundancy on page 715 shows that redundant management was disabled.

Figure 178: *Results of disabling redundancy*

```
HP Switch(config)# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active
management module. Do you want to continue [y/n]? y

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                          Status   SW Version    Boot Image
---- ---------------------------------------- -------- ------------ ----------
MM1  HP Switch J9092A Management Module 8200zl Offline  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled
s
```

Disable management module redundancy with only one module present

If you disable redundancy when there is only one management module in the switch, and then you insert a second management module, the second module never goes into standby mode. You must re-enable redundant management using this command:

```
switch# redundancy management-module
```

The currently active module remains active on boot (assuming no selftest failure) unless you make the newly inserted management module active using this command:

```
switch# redundancy active-management standby
```

The standby management module becomes the active management module.

Active management module commands

Viewing modules

The `show modules` command displays information about all the modules in the switch, as well as additional component information for the following:

- System Support Modules (SSM)—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, "J" number, and serial number (when available)

Syncing commands

The following CLI commands can be executed during initial syncing between the active management module and the standby management module, which occurs when the standby module is inserted or after a reboot of the system. All other CLI commands will not be executed until after the initial syncing completes.

During initial syncing, no SNMP set requests are executed, except the SNMP request for ping.

Operator commands		
dir	menu	traceroute6
enable	ping	dbgstack
exit	ping6	wireless-services
link-test	show	services
logout	traceroute	

Manager commands		
boot system	copy running-config	page
boot active	copy startup-config	print
boot standby	copy event-log	redo
configure	copy core-dump	reload
copy command-output	recopy	repeat
copy config tftp	display	task-monitor
copy config xmodem	end	telnet
copy crash-data	getMIB	terminal
copy crash-log	kill	walkMIB
copy flash tftp	licenses	write-terminal
copy flash xmodem	log	redundancy

Using the WebAgent for redundant management

The WebAgent can be used to display information about the active and standby management modules.

Online Help is available for the WebAgent, which you can open by clicking on the question mark (?) in the upper right corner of any of the WebAgent screens. An example redundancy screen is shown in **Figure 179: Example of redundancy screen in the WebAgent** on page 717.

To access the redundancy information in the WebAgent:

1. In the WebAgent navigation panel, click System.
2. Click Redundancy. The following screen displays.

Figure 179: *Example of redundancy screen in the WebAgent*

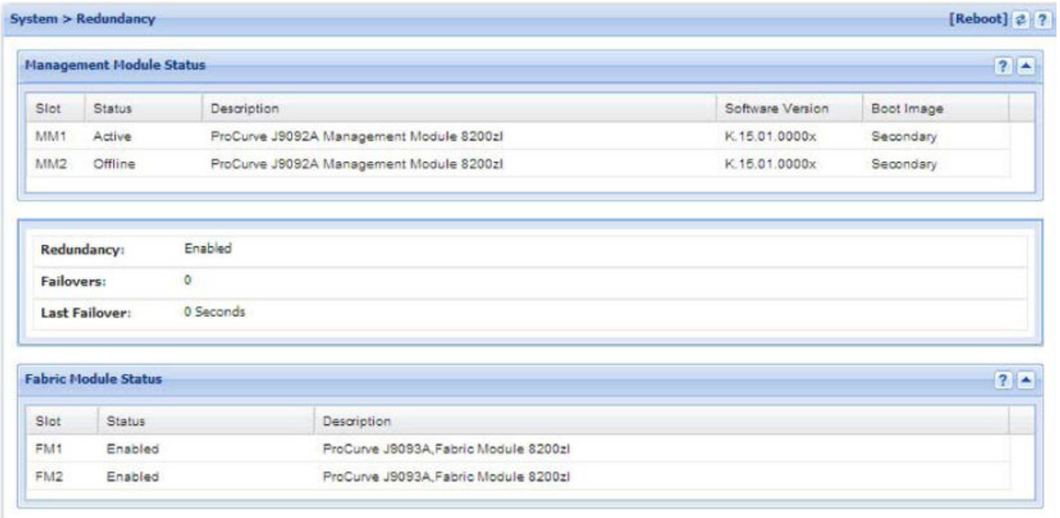

Enabling or disabling redundant management

There are two modes for management module redundancy—warm standby mode (the default) and Nonstop switching mode. In warm-standby mode, the active management module does not sync continuously with the standby management module. The standby management module boots to a certain point, syncs basic files, and only finishes booting if the active management module fails or you choose to change which module is the active management module. The transition is not seamless or immediate.

In Nonstop switching mode, the standby management module is synced continuously with the active management module so that all features and config files are the same on both management modules. The standby management module is ready to become the active management module. The transition is quick and seamless; switching continues without interruption.

Syntax

```
[no] redundancy management-module [nonstop-switching]
```

Allows enabling or disabling of redundant management. The current active module continues to be the active module on boot unless you use the `redundancy active-management` command to enable redundant behavior.

(Default: Warm-standby redundancy mode)

The `nonstop-switching` parameter sets the redundancy mode to Nonstop switching.

You are prompted with "All configuration files and software images on the off-line management module will be overwritten with the data from the current active management module. During initial syncing from active to standby management module configuration changes are disallowed. Do you want to continue [y/n]?"

When the `nonstop-switching` option is **not** selected, the switch enters warm-standby redundancy mode.

You are prompted with "All configuration files and software images on the off-line management module will be overwritten with the data from the current active management module. Do you want to continue [y/n]?"

The `no` version of the command disables redundant management. You are prompted with this message: "The other management module may reboot and it will no longer be used for system redundancy, except in the case of a hardware failure of the active management module. Do you want to continue [y/n]?".

Example

The `redundancy management-module` command in **Figure 180: Enabling warm-standby redundancy** on page 718 shows **warm-standby redundant management** being **enabled**. The `show redundancy` command displays "**Mgmt Redundancy**" as **warm-standby redundancy enabled** . Management Module 1 (MM1) is the active management module and Management Module 2 (MM2) is the standby management module.

Figure 180: *Enabling warm-standby redundancy*

```
HP Switch(config)# redundancy management-module
All configuration files and software images on the off-line management
module will be overwritten with the data from the current active
management module. Do you want to continue [y/n]?  y

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Warm-standby redundancy enabled      <-- Redundancy enabled
Rapid Switchover Stale Timer : 1

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                           Status   SW Version   Boot Image
---- ---------------------------------------- -------- ----------- ----------
MM1  HP Switch J9092A Management Module 8200zl Active   K.15.01.000x Secondary
MM2  HP Switch J9092A Management Module 8200zl Standby  K.15.01.000x Secondary

FM1  HP Switch J9093A,Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A,Fabric Module 8200zl     Enabled
```

The `redundancy management-module` command in Figure **Figure 181: Enabling nonstop-switching redundancy** on page 718 shows Non-stop switching redundant management being **enabled.** The `show redundancy` command displays "**Mgmt Redundancy**" as **Nonstop switching enabled** . Management Module 1 (MM1) is the standby management module and Management Module 2 (MM2) is the active management module.

Figure 181: *Enabling nonstop-switching redundancy*

```
HP Switch(config)# redundancy management-module nonstop-switching
All configuration files and software images on the off-line management module
will be overwritten with the data from the current active management module.
During initial syncing from active to standby management module configuration
changes are disallowed. Do you want to continue [y/n]?  y

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching enabled      <-- Redundancy enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                           Status   SW Version    Boot Image
---- ---------------------------------------- -------- ------------ ---------
MM1  HP Switch J9092A Management Module 8200zl Standby  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled
```

The `no` version of the `redundancy management-module` command is used to disable management module redundancy on the switch, as seen in **Figure 182: Disabling redundancy** on page 719. The `show redundancy` command displays "**Mgmt Redundancy**" as **Nonstop switching disabled** . The standby management module in slot MM1 is now offline. The management module in slot MM2 remains the active management module.

NOTE

Hewlett Packard Enterprise recommends that you leave management module redundancy enabled. If the active management module has a hardware failure, the standby module may take over and may have an old configuration since file synchronization has not occurred when management module redundancy was disabled.

The `no redundancy management-module` command allows you to shut down a management module that is not functioning correctly without physically removing the module. If you want to remove the module, first perform the shutdown procedure as explained in **Hotswapping out the active management module** on page 712 and then remove the module.

Figure 182: *Disabling redundancy*

```
HP Switch(config)# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active management
module. Do you want to continue[y/n]?  y

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching disabled            Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 1
Last Failover : Tue Mar 19 12:42:31 2009

Slot Module Description                             Status   SW Version   Boot Image
---- ---------------------------------------- -------- ------------ ---------
MM1  HP Switch J9092A Management Module 8200zl Offline  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled
```

The `redundancy management-module` command shows Nonstop switching redundant management being enabled. The `show redundancy` command displays “Mgmt Redundancy” as Nonstop switching enabled. Management Module 1 (MM1) is the standby management module and Management Module 2 (MM2) is the active management module.

Example

Enabling non-stop switching redundancy.

```
switch# redundancy management-module nonstop-switching
All configuration files and software images on the off-line management module
will be overwritten with the data from the current active management module.
During initial syncing from active to standby management module configuration
changes are disallowed. Do you want to continue [y/n]? y
switch# show redundancy
Settings
--------
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0
Statistics
----------
Failovers : 0
Last Failover :
Slot Module    Description                Status   SW Version    Boot Image
---- ----------------------------------------   ----------    ------------
MM1  HP J9092A Management Module 8200zl  Standby  K.15.01.000x  Primary
```

```
MM2  HP J9092A Management Module 8200zl  Active   K.15.01.000x  Primary
FM1  HP J9093A Fabric Module 8200zl      Enabled
FM2  HP J9093A Fabric Module 8200zl      Enabled
```

The no version of the redundancy management-module command is used to disable management module redundancy on the switch, as seen in Figure 7-4. The show redundancy command displays "Mgmt Redundancy" as Nonstop switching disabled. The standby management module in slot MM1 is now offline. The management module in slot MM2 remains the active management module.

NOTE Hewlett Packard Enterprise recommends that you leave management module redundancy enabled. If the active management module has a hardware failure, the standby module may take over and may have an old configuration since file synchronization has not occurred when management module redundancy was disabled.

The `no redundancy management-module` command allows you to shut down a management module that is not functioning correctly without physically removing the module. If you want to remove the module, first perform the shutdown procedure as explained in "Hotswapping Out the Active Management Module" on page 7-25, and then remove the module.

Example

Disabling redundancy:

```
switch# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active management
module. Do you want to continue[y/n]? y
switch# show redundancy
Settings
--------
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0
Statistics
----------
Failovers : 1
Last Failover : Tue Mar 19 12:42:31 2009
Slot   Module     Description               Status   SW Version   Boot Image
----   ---------  ------------------------  -------  -----------  ------------
MM1    HP J9092A  Management Module 8200zl  Offline  K.15.01.000x  Primary
MM2    HP J9092A  Management Module 8200zl  Active   K.15.01.000x  Primary
FM1    HP J9093A  Fabric Module 8200zl      Enabled
FM2    HP J9093A  Fabric Module 8200zl      Enabled
```

Transitioning from `no redundancy` to nonstop switching

While the switch is transitioning from no redundancy mode to Nonstop switching mode, no configuration changes are allowed. The management modules are syncing information during the transition period.

Setting the Rapid Switchover Stale Timer

Use the Rapid Switchover Stale Timer to set the amount of time that you want route and neighbor table entries to be re-added to the Forwarding Information Base on the active management module after a failover has occurred.

Layer 3 applications and protocols rely on existing routing information in the FIB. They restart and operate as if the switch performed a quick reset.

When a failover occurs, the interface modules and the fabric modules continue forwarding Layer 3 traffic based on the information in the FIB. The transitioning standby management module marks all routes in the FIB as "stale". The routing protocols restart, reestablish their neighbors and reconverge. As the routes are added in again, the route's stale designation is removed. After the Rapid Switchover Stale Timer expires, the remaining stale route entries are removed. Multicast flows are also removed; the multicast application re-adds the flows after failover completes.

Syntax

```
redundancy rapid-switchover <0-36000>
```

Allows configuration of a timer (in seconds) for Layer 3 forwarding of packets when Nonstop switching is configured for redundancy. After failover, the route and neighbor entries in the Forwarding Information Base (FIB) on the active management module are marked as stale. As new routes are added, the stale flag is reset. This continues for the number of seconds indicated by the timer, after which all remaining stale entries (entries not re-added) are removed.

A setting of zero indicates that no Layer 3 Nonstop switching behavior is wanted.

When the switch fails over, the FIB entries and corresponding hardware entries are removed. Default: 90 seconds

To display information about stale FIB routes, enter the `show tech route stale` command. The VLAN ID and IP route are shown, as well as other information used only for technical support.

Directing the standby module to become active

To make the standby management module become the active management module, use the `redundancy switchover` command. The switch will switchover after all files have finished synchronizing.

In Nonstop switching mode:

- The switchover occurs quickly and seamlessly. No reboot is needed.
- There is no interruption in switching operations.

In warm-standby mode:

- The switchover may take a couple of minutes if there have been recent configuration file changes or if you have downloaded a new operating system.
- The standby module finishes booting and becomes the active module.

The formerly active module becomes the standby module if it passes selftest.

Syntax

```
redundancy switchover
```

Causes a switchover to the standby module.

For Nonstop switching, the warning displays:`A nonstop switching failover will occur; L2 operations will not be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module’s console interface. Do you want to continue [y/n]?`

In warm-standby mode the warning displays:`A warm failover will occur; all networking operations will be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module’s console interface. Do you want to continue [y/n]?`

If management module redundancy has been disabled, or there is no standby module, or the standby module is not in standby mode, this message displays:`The other management module does not exist or is not in standby mode` An example of the

```
redundancy switchover
```

command when the switch is in Nonstop switching mode is shown in the example below.

Example

`Redundancy switchover` command when in nonstop switching mode.

```
switch# redundancy switchover
A nonstop switching failover will occur; L2 operations will not be interrupted.
This management module will now reboot and will become the standby
module! You will need to use the other management module's console interface.
Do you want to continue [y/n]? y
This management module will now boot from the primary image and will
become the standby module! You will need to used the other management module’s
console interface. Do you want to continue [y/n]? y
ROM information:
Build directory: /sw/rom/build/bmrom(t2g)
Build date: Oct 15 2009
Build time: 08:24:27
Build version: K.15.01
Build number: 13040
Select profile (primary):
Booting Primary Software Image...
...
Standby Console>
```

Directing the standby module to become active

Syntax

`redundancy switchover`

Causes a switchover to the standby module.

For nonstop switching, the warning displays: "A nonstop switching failover will occur; L2 operations will not be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?"

In warm-standby mode the warning displays: "A warm failover will occur; all networking operations will be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?"

If management module redundancy has been disabled, or if there is no standby module, or if the standby module is not in standby mode, this message displays:

```
The other management module does not exist or is not in standby mode
```

Example

Figure 183: The redundancy switchover command when in nonstop switching mode on page 723 shows an example of the `redundancy switchover` command when the switch is in **nonstop switching** mode.

Figure 183: *The* `redundancy switchover` *command when in nonstop switching mode*

```
HP Switch(config)# redundancy switchover
A nonstop switching failover will occur; L2 operations will not be inter-
rupted. This management module will now reboot and will become the standby
module! You will need to use the other management module's console inter-
face. Do you want to continue [y/n]? y
This management module will now boot from the primary image and will
become the standby module! You will need to used the other management mod-
ule's console interface. Do you want to continue [y/n]? y

ROM information:
   Build directory: /sw/rom/build/bmrom(t2g)
   Build date:      Oct 15 2009
   Build time:      08:24:27
   Build version:   K.15.01
   Build number:    13040
Select profile (primary):

Booting Primary Software Image...
.
.
.
Standby Console>
```

Setting the active management module for next boot

Syntax

`redundancy active-management [management-module1 | management-module2 | standby]`

The specified module becomes the active management module at the next system boot. This message displays:`On the next system boot, the module specified will become active.`

This command does not take effect if the standby management module has failed selftest.

`management-module1`	Configures management-module 1 as the active management module for the next system boot.
`management-module2`	Configures management-module 2 as the active management module for the next system boot.
`standby`	Configures the current standby module as the active management module for the next system boot if management module redundancy is enabled. If redundancy is disabled, it becomes enabled as a standby module at the next boot or failover event.

If the specified management module is not there or is in failed mode, this message displays:

```
The specified module is not present or is in failed state.
```

Example

Figure 184: Setting a management module to be active on the next boot on page 724 shows an example of setting **management module 2** to be the **active management module**.

Figure 184: *Setting a management module to be active on the next boot*

```
HP Switch(config)# redundancy active-management management-module2

On the next system boot, the management-module2 will become active.

HP Switch(config)# boot system
(boot occurs...)

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop Switching enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                       Status  SW Version   Boot Image
---- ---------------------------------------- ------  ------------ ----------
MM1  HP Switch J9092A Management Module 8200zl Standby K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Active  K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled
```

If management module redundancy has been disabled and you specify the standby module with the `active-management` command, upon rebooting, the offline module becomes the standby module. The state of redundancy (enabled or disabled) is based on the value in the configuration file in the offline (now standby) module. The configuration files have not been synchronized if management module redundancy has been disabled. An example of making the offline management module become the standby management module when

redundancy is disabled is shown in **Figure 185: Showing the results of switching to standby module when redundancy is disabled** on page 725.

Figure 185: *Showing the results of switching to standby module when redundancy is disabled*

```
HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                        Status   SW Version   Boot Image
---- ----------------------------------------  -------- ------------ ---------
MM1  HP Switch J9092A Management Module 8200zl Active    K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Offline   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled


HP Switch(config)# redundancy active-management standby
On the next system boot, the standby will become active.
Redundancy and Synchronization have been disabled, so it will
not have current configurations.

HP Switch(config)# boot
The other management module is not in standby mode and this command will
not cause a switchover. System will reboot from primary image.
Do you want to continue [y/n]? y

(After system reboots...)

HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                        Status   SW Version   Boot Image
---- ----------------------------------------  -------- ------------ ---------
MM1  HP Switch J9092A Management Module 8200zl Standby   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Active    K.15.01.000x Primary
```

Nonstop switching disabled

When redundancy is disabled and the redundancy active-management standby command is executed, the offline MM becomes the active MM.

Resetting the management module

The **MM Reset** button, shown in **Figure 186: The MM Reset button on the management module** on page 726, found on each management module reboots its management module. If the management module is active and management module redundancy is enabled, switchover occurs. The standby management module is notified immediately. It then takes over and becomes the active management module. If the **MM Reset** button is pressed on the standby management module, that module reboots but no other switch operations are affected. The active management module remains in control.

If management module redundancy is disabled, the active management module reboots and remains in control, as long as it passes selftest.

Hewlett Packard Enterprise does not recommend using the **MM Reset** button to trigger a switchover. Files being copied over at the time of the reset will be aborted.

Figure 186: *The MM Reset button on the management module*

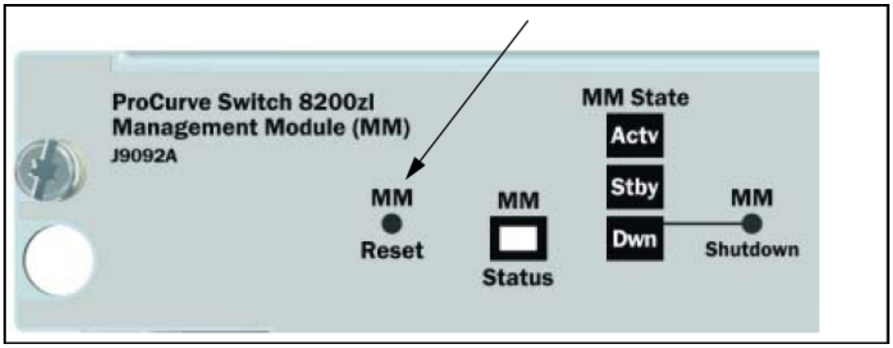

Viewing management information

Syntax

```
show modules [details]
```

Displays information about the installed modules, including:

- The slot in which the module is installed
- The module description
- The serial number
- The status
- Core dump
- Model Version

Additionally, the part number (J number) and serial number of the chassis is displayed.

Example

```
Status and Counters - Module Information
Chassis: 8212zl J9091A         Serial Number:   LP713BX004
Allow V1 Modules:    Yes
                                   Core   Mod
 Slot Module Description                    Serial Number  Status   Dump  Ver
 ---- -------------------------------------- -------------- -------- ----- ---
MM1  HP J9092A Management Module 8200zl     sg844bp012     Active   NO    1
SSM  HP J9095A System Support Module        SG911BZ00N
FM1  HP J9093A Fabric Module 8200zl         SG911BQ015     Enabled  -     1
FM2  HP J9093A Fabric Module 8200zl         SG911BQ04T     Enabled  -     1
A    HP J9536A 20p GT PoE+/2p SFP+ v2 zl... SG0607T124     Up       YES   2
B    HP        Enh Svs v2 zl Module                        Up       YES   2
C    HP J8702A 24p Gig-T zl Module                         Up       NO    1
D    HP J9840A Adv Svs v2 zl Module         ID3ZG6N008     Up       YES   2
E    HP J8705A Gig-T/SFP zl Module                         Up       NO    1
F    HP J9857A Adv Svs v2 zl Module         SG2ZFNX166     Up       YES   2
G    HP J8708A 4p 10G CX4 zl Module                        Up       NO    1
H    HP J9154A Services zl Module           SG811GG01N     Up       NO    1
I    HP J9051A Wireless Edge Services zl... SG660ZB095     Up       NO    1
J    HP J9545A ONE Adv Svs zl Module        SG9604P933     Up       NO    1
```

```
K     HP J9051A Wireless Edge Services zl... 1111           Up         NO     1
L     HP J9154A Services zl Module           SG811GG01M     Up         NO     1
```

Viewing information about the management and fabric modules

The `show redundancy` command displays information about the management and fabric modules. It displays the flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot.

Example

Figure 187: *`show redundancy` command*

```
HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers      : 0
Last Failover :

Slot Module Description                          Status   SW Version   Boot Image
---- ---------------------------------------- -------- ------------ ----------
MM1  HP J9092A Management Module 8200zl        Standby  K.15.01.000x Primary
MM2  HP J9092A Management Module 8200zl        Active   K.15.01.000x Secondary

FM1  HP J9093A Fabric Module 8200zl            Enabled
FM2  HP J9093A Fabric Module 8200zl            Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.

Viewing information about the redundancy role of each management module

The `show redundancy` command with the `detail` option displays information about the redundancy role of each management module, as well as statistical information such as how long the module has been up.

Example

Figure 188: *`show redundancy detail` command*

```
HP Switch(config)# show redundancy detail

 Redundancy Information:

  Slot Role     Card Up Since        Role Since           Redundancy State
  ---- -------- -------------------- -------------------- -----------------
  1    Active   11/11/09 23:40:22    11/04/09 23:33:15    Active
  2    Standby  11/11/09 23:40:24    11/04/09 23:33:15    Nonstop switching

 Fail-Over Log:

  Slot Role     Time                 Reason
  ---- -------- -------------------- ---------------
  2    Standby  11/01/09 10:16:04    Standby Reset
  2    Active   11/02/09 17:46:03    Hot Swap
  1    Standby  11/03/09 15:39:06    Standby Reset
  1    Active   11/04/09 09:25:39    Switchover
```

Viewing which software version is in each flash image

The `show flash` command displays which software version is in each flash image. The **Default Boot** field displays which flash image will be used for the next boot.

Example

Figure 189: *`show flash` command*

```
HP Switch(config)# show flash
Image              Size(Bytes)    Date    Version
-----              ----------   --------  -------------------
Primary Image    : 7463821     09/05/09 K.15.00.0001
Secondary Image  : 7463821     09/05/09 K.15.00.0001

Boot Rom Version: K.15.07
Default Boot    : Primary
```

Will boot from primary flash on the next boot.

Viewing system software image information for both management modules

The `show version` command displays system software image information for both management modules, as well as which module is the active management module and which is the standby management module. The Boot Image field displays which flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot. The output of the `show version` command when redundancy is enabled is shown in **Figure 190: show version command when redundancy is enabled** on page 728.

Example

Figure 190: *`show version` command when redundancy is enabled*

```
HP Switch(config)# show version
Management Module 1: Standby
Image stamp:     /sw/code/build/btm(t2g)
                 Mar  5 2009 13:20:59
                 K.15.01.0001
                 351
Boot Image:      Primary

Management Module 2: Active
Image stamp:     /sw/code/build/btm(t2g)
                 Mar  5 2009 13:20:59
                 K.15.01.0001
                 351
Boot Image:      Primary
```

Both management modules were booted from primary flash.

When redundancy is disabled, the output of the `show version` command changes, as shown in **Figure 191: show version command when redundancy is disabled** on page 728.

Example

Figure 191: *`show version` command when redundancy is disabled*

```
HP Switch(config)# show version
Management Module 1:  Redundancy and Synchronization has been disabled;
                      enable with the 'redundancy' command.

Management Module 2: Active
Image stamp:     /sw/code/build/btm(t2g)
                 Mar  5 2009 13:20:59
                 K.15.01.0001
                 351
Boot Image:      Primary
```

Viewing the status of the switch and its management modules

The `show logging` command displays the status of the switch and its management modules. See **Displaying module events** on page 736. To show log messages in reverse chronological order (most recent messages displayed first), enter `show log -r`.

Example

Figure 192: *`show log` command output*

```
HP Switch(config)# show logging
 Keys:   W=Warning   I=Information
         M=Major     D=Debug E=Error
---- Event Log listing: Events Since Boot ----
I 10/28/09 21:45:42 00061 system: AM1: -----------

I 10/28/09 21:45:42 00062 system: AM1: Mgmt Module 1 went down without saving cr
            ash information
M 10/28/09 21:45:42 03002 system:  AM1: System reboot due to Reset Switch
I 10/28/09 21:45:42 02759 chassis: AM1: Savepower LED timer is OFF.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot A configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot B configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot C configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot D configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot E configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot F configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot G configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot H configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot I configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot J configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot K configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot L configured ON.
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 1 inserted
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 2 inserted
I 10/28/09 21:45:43 00092 dhcp: AM1: Enabling Auto Image Config Download via
            DHCP and turning off auto-tftp if enabled
I 10/28/09 21:45:43 00690 udpf: AM1: DHCP relay agent feature enabled
I 10/28/09 21:45:43 02637 srcip: AM1: TACACS admin policy is 'outgoing interface
            '
I 10/28/09 21:45:43 02638 srcip: AM1: TACACS oper policy is 'outgoing interface'
```

AM1 = Active management module in slot 1
AM2 = Active management module in slot 2
SM1 = Standby management module in slot 1
SM2 = Standby management module in slot 2

Standby management module commands

The standby management module, by design, has very little console capability. You can use three commands— `show flash`, `show version`, and `show redundancy`. The `show redundancy` command displays when a management module is in standby mode.

Viewing redundancy status on the standby module

Use the `show redundancy` command to display redundancy status on the standby module, as shown in **Figure 193: show redundancy command for standby module** on page 730. This command displays the flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot.

Example

Figure 193: *`show redundancy` command for standby module*

```
Standby Console> show redundancy

Settings
--------
Mgmt Redundancy : Nonstop Switching Enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 1
Last Failover : Mon Sep 26 09:50:40 2009


Slot Module Description                           Status   SW Version   Boot Image
---- --------------------------------------------- -------- ------------ ----------
MM1  HP Switch J9092A Management Module 8200zl Active   K.15.01.0001 Secondary
MM2  HP Switch J9092A Management Module 8200zl Standby  K.15.01.0001 Primary

FM1  HP Switch J9093A Fabric Module 8200zl      Enabled
FM2  HP Switch J9093A Fabric Module 8200zl      Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.

Viewing the flash information on the standby module

Use the `show flash` command to display the flash information on the standby module, as shown in **Figure 194: show flash command for standby module** on page 730. The **Default Boot** field displays which **flash image** will be used for the next boot.

Example

Figure 194: *`show flash` command for standby module*

```
Standby Console> show flash
Image              Size(Bytes)   Date    Version
-----              ----------  --------- ------------------
Primary Image    : 7493854    09/21/09 K.15.00.0001
Secondary Image  : 7463821    09/05/09 K.15.00.0001

Boot Rom Version: K.15.07
Default Boot    : Primary
```

Will boot from primary flash on the next boot.

Viewing the version information on the standby module

Use the `show version` command to display the version information on the standby module, as shown in **Figure 195: show version command for standby module** on page 730. The **Boot Image** field displays which flash image was last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot. Unlike executing the `show version` command on an active management module, this command shows only the running version of software on the standby management module.

Example

Figure 195: *`show version` command for standby module*

```
Standby Console> show version
Image stamp:     /sw/code/build/btm(t2g)
                 Mar 21 2009 15:03:31
                 K.15.01.0001
                 1617
Boot Image:      Primary
```

Was booted from primary flash.

Setting the default flash for boot

You can set which flash image to boot from as the default image on boot by using this command:

Syntax

```
boot set-default flash [primary | secondary]
```

Sets the flash image to boot from on the next boot.

`primary`	Boots the primary flash image.
`secondary`	Boots the secondary flash image.

Example

Figure 196: boot set-default command defaulting to secondary flash on page 731 shows an example of the output when the command is used to set the boot default to secondary flash.

Figure 196: *`boot set-default` command defaulting to secondary flash*

```
HP Switch(config)# show flash
Image             Size(Bytes)   Date   Version
-----             ----------  -------- -------------------
Primary Image    : 7463821    11/05/09 K.15.01.0001
Secondary Image : 7463821    11/05/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot     : Primary


HP Switch(config)# boot set-default flash secondary
This command changes the location of the default boot. This
command will change the default flash image to boot from
secondary. Hereafter, 'reload' and 'boot' commands will boot
from secondary. Do you want to continue [y/n]? y

HP Switch(config)# show flash
Image             Size(Bytes)   Date   Version
-----             ----------  -------- -------------------
Primary Image    : 7463821    03/05/09 K.15.01.0001
Secondary Image : 7463821    03/05/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot     : Secondary
```

Booting the active management module from the current default flash

Use the `reload` command to boot the active management module from the current default flash (You can change the default flash with the `boot set-default` command. See **Setting the default flash for boot** on page 731.) Switchover occurs if redundancy is enabled and the standby management module is in standby mode. If redundancy is disabled or the standby management module is not present, the `reload` command boots the system.

NOTE

The `reload` command is a "warm" reboot; it skips the Power on Self Test routine.

Syntax

```
reload <cr>
```

Boots (warm reboot) the active management module. Switchover to the standby management module occurs if `management module` redundancy is enabled. If redundancy is disabled or if there is no standby management module, the `reload` command boots the system.

NOTE

If the `running config` file is different from the stored config file, you are prompted to save the config file. The `reload at/after` versions of this command do not display a prompt to save configuration file changes: the changes are lost on the scheduled reload.

Example

Figure 197: *`reload` command with redundancy enabled*

```
HP Switch(config)# reload
This command will cause a switchover to the other management module
which may not be running the same software image and configurations.
Do you want to continue [y/n]? y


(Boots….)


HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop Switching Enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 1
Last Failover : Mon April 30 09:10:11 2009

Slot Module Description                          Status   SW Version   Boot Image
---- ---------------------------------------- -------- ------------ ----------
MM1  HP Switch J9092A Management Module 8200zl Active   K.15.01.0001 Primary
MM2  HP Switch J9092A Management Module 8200zl Standby  K.15.01.0001 Primary
```

`boot` command

In redundant management systems, the `boot` or `boot active` command causes a switchover to the standby management module as long as the standby module is in standby mode. This message displays:

```
This management module will now reboot and will become the
standby module! You will need to use the other management
module's console interface. Do you want to continue [y/n]?
```

If you select **y**, switchover is initiated by the standby management module, which becomes the active management module after boot completes.

If the standby module is not in standby mode (for example, it is in failed mode or offline mode), switchover to the standby module does not occur. The system is rebooted and this message displays:

```
The other management module is not in standby mode and this
command will not cause a switchover, but will reboot the
system, do you want to continue [y/n]?
```

If the other management module is not present in the switch, the system simply reboots.

The `boot` command has these options.

Command	Action
`boot cr`	Reboots the active management module from the flash image that is specified for the default boot. This can be changed with the `boot set-default flash` command. You can select which image to boot from during the boot process itself. (See **Figure 198: The management module rebooting, showing boot profiles to select** on page 733.) The switch will switchover to the standby management module. NOTE: This is changed from always booting from primary flash. You are prompted with a message, which indicates the flash being booted from.
`boot active`	Boots the active management module. The switch starts to boot from the default flash image. You can select which image to boot from during the boot process itself. (See **Figure 198: The management module rebooting, showing boot profiles to select** on page 733.) The switch will switchover to the standby management module. If a second management module is not present in the switch, the system is rebooted.
`boot standby`	Boots the standby management module. The switch does not switchover.If the standby module is not present, this message displays: "The other management module is not present."
`boot system [flash [primary | secondary]]`	Boots both the active and standby management modules. You can specify the flash image to boot from.
`boot set-default flash {primary | secondary}`	Sets the default flash for the next boot to primary or secondary. You see this message:"This command changes the location of the default boot. This command will change the default flash image to boot from flash chosen>. Hereafter, 'reload' and 'boot' commands will boot from flash chosen>. Do you want to continue [y/n]?"

You can select a **boot profile** during the reboot process, as shown in **Figure 198: The management module rebooting, showing boot profiles to select** on page 733. If you make no selection, the boot defaults to the image displayed as the default choice (shown in parentheses.)

Figure 198: *The management module rebooting, showing boot profiles to select*

```
Boot Profiles:

0. Monitor ROM Console
1. Primary Software Image
2. Secondary Software Image

Select profile(primary): 2

Booting Secondary Software Image...
```

An example of the `boot` command with the **default flash** set to **secondary** is shown in **Figure 199: Showing boot command with default flash set to secondary** on page 734.

Figure 199: *Showing boot command with default flash set to secondary*

```
HP Switch(config)# boot set-default flash secondary
This command changes the location of the default boot. This command will
change the default flash image to boot from secondary image. Hereafter,
'reload' and 'boot' commands will boot from secondary image. Do you want
to continue [y/n]? y

HP Switch(config)# show flash
Image             Size(Bytes)   Date    Version
-----             ----------  -------- --------------------
Primary Image    : 7476770    11/01/09 K.15.01.0001
Secondary Image  : 7476770    11/01/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot    : Secondary

HP Switch(config)# boot
This management module will now reboot from secondary and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]?
```

CAUTION

For a given reboot, the switch automatically reboots from the `startup-config` file assigned to the flash (primary or secondary) being used for the current reboot. The `startup-default` command can be used to set a boot configuration policy. This means that both the flash image and one of the three configuration files can be specified as the default boot policy.

`Boot` and `reload` commands with OSPFv2 or OSPFv3 enabled

It is now possible to gracefully shut down OSPFv2 or OSPFv3 routing on switches without losing packets that are in transit. OSPF neighbors are informed that the router should not be used for forwarding traffic, which allows for maintenance on the switch without interrupting traffic in the network. There is no effect on the saved switch configuration

Prior to a switch shutdown, the CLI/SNMP `reload` command or the CLI `boot` command is executed to initiate the sending of OSPF "empty Hello list" messages on the interfaces that are part of the OSPF routing configuration. After a small delay (approximately 2 seconds) that allows the messages to be transmitted on all applicable interfaces the `boot` or `reload` command continues.

Modules operating in nonstop mode

When a switch is in standalone mode and OSPF routing is enabled, the "empty Hello list" is transmitted whenever the `boot` or `reload` command is executed.

When the switch is operating in nonstop switching mode (redundant), and a single module is being reloaded or booted, the standby module notifies neighboring switches of the management module failover. If the failover fails, the "empty Hello list" is transmitted before the switch is rebooted.

When a switch is operating with multiple management modules in warm standby mode, the "empty Hello list" is sent when a `reload` or `boot` command is executed. The standby management module sends out OSPF Hello packets after becoming the active management module.

Additional commands affected by redundant management

The other existing commands operate with redundant management as shown below.

Command	Action
`auto-tftp`	If a new image is downloaded using `auto-tftp`, the active management module downloads the new software version to both the active and standby modules. Rebooting after the `auto-tftp` completes reboots the entire system.
`banner`	The banner will not been seen on the standby module, only the active module.
`chassislocate`	If the management module performs a switchover, the LED does not remain lit.
`clear`	The `clear crypto` command causes public keys to be deleted from both modules when the second module is in standby mode.
`console`	Console settings, such as mode, flow-control, and baud-rate, are the same on both management modules. There cannot be individual settings for each management module.
`copy`	Files are automatically sync'd from the active management module to the standby management module.When no parameter is specified with the `copy crash-data` or `copy crash-log` command, files from all modules (management and interface) are concatenated. NOTE: If redundancy is disabled or the standby module failed selftest, the `copy` command affects only the active management module.
`copy core-dump [mm | standby | flash | xmodem | usb` *`filename`*`]`	The `copy core-dump standby flash` command copies the standby management module's coredump to the active management module's flash. The destination file is fixed as dumpM1.cor or dumpM2.cor, depending on which module is the standby management module.The `copy core-dump [mm | standby | flash | xmodem | usb` *`<filename>`*`]` command copies the core file of the active management module or the standby management module to a USB flash drive or to an xmodem host.
`core-dump management-module`	Enables or disables a core dump on a management module.
`crypto`	Authentication files for ssh or the https server are copied to the standby management module. The `clear crypto` command deletes the public keys from both modules when the second module is in standby mode.
`erase flash`	Erases the software version on the active and standby modules. If redundancy has been disabled, or if the standby module has not passed selftest, the flash is not erased on the standby module.

Table Continued

Command	Action
`erase config`	Erases the config file on the active and standby modules. If redundancy has been disabled, or if the standby module has not passed selftest, the config file is not erased on the standby module.
`erase startup-config`	Affects both modules if the second module is in standby mode. If redundancy has been disabled, or if the standby module has not passed selftest, the `startup-config` file is not erased on the standby module.
`fastboot`	When fastboot is enabled, this information is saved to the standby management module when the config files are sync'd. The fastboot value is used during the next boot on both modules.
`front-panel-securityfactory-reset password-clear password-recovery`	This command and its options affect only the active management module.
`kill`	Does not affect the console on the standby module.
`log`	Log messages from a formerly active management module are available on the current active management module after a switchover.
`password (set` or `clear)`	Affects only the active management module until a switchover occurs, at which time it affects the new active module.
`startup-default`	Affects both modules. The config file is immediately sent to the standby module and also becomes the default on that module when the next boot occurs.
`update`	Affects only the active module. The standby may become the active module when the updated active module is booted.
`write`	A `write memory` updates the config file in flash on the active module. The file is then sync'd to the standby module.

Displaying module events

Viewing log events

The log file displays messages about the activities and status of the management modules. Enter this command to display the messages:

Syntax

`show logging [-a, -b, -r, -s, -t, -m, -e, -p, -w, -i, -d, command, filter, *option-str*]`

Displays log events.

The event messages are tagged with the management module state and the management module slot (AM1 or AM2, SM1 or SM2.) Synchronization is maintained by syncing the standby management module log events with the active management module. In this way, events are available for both management modules. Only the active management module events are shown unless you select the `-s` option. This option works like the `-a` option, except that the events for both the active management module and standby management module are displayed.

Example

Figure 200: *Log file listing*

```
HP Switch(config)# show logging
 Keys:   W=Warning   I=Information
         M=Major     D=Debug E=Error
---- Event Log listing: Events Since Boot ----
I 10/28/09 21:45:42 00061 system:  AM1: ----------------------------------------

I 10/28/09 21:45:42 00062 system:  AM1: Mgmt Module 1 went down without saving
           crash information
M 10/28/09 21:45:42 03002 system:  AM1: System reboot due to Reset Switch
I 10/28/09 21:45:42 02759 chassis: AM1: Savepower LED timer is OFF.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot A configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot B configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot C configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot D configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot E configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot F configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot G configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot H configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot I configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot J configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot K configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot L configured ON.
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 1 inserted
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 2 inserted
I 10/28/09 21:45:43 00092 dhcp: AM1: Enabling Auto Image Config Download via
            DHCP and turning off auto-tftp if enabled
I 10/28/09 21:45:43 00690 udpf: AM1: DHCP relay agent feature enabled
I 10/28/09 21:45:43 02637 srcip: AM1: TACACS admin policy is 'outgoing interface
            '
I 10/28/09 21:45:43 02638 srcip: AM1: TACACS oper policy is 'outgoing interface'
```

AM1 = Active management module in slot 1
AM2 = Active management module in slot 2
SM1 = Standby management module in slot 1
SM2 = Standby management module in slot 2

Copying crash file information to another file

Crash logs for all modules are always available on the active management module. You can use the `copy crash-log` and `copy crash-data` commands to copy the information to a file of your choice.

Syntax

`copy crash-log [`*`slot-id`*` | mm] tftp `*`ip-address filename`*

Copies the crash logs of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

`slot-id`	Retrieves the crash log from the module in the specified slot.
`mm`	Retrieves the crash logs from both management modules and concatenates them.

Syntax

`copy crash-data [`*`slot-id`*` | mm] tftp `*`ip-address filename`*

Copies the crash data of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

`slot-id`	Retrieves the crash data from the module in the specified slot.
`mm`	Retrieves the crash data from both management modules and concatenates them.

Viewing saved crash information

Syntax

`show boot-history`

Displays the system boot log.

Example

Figure 201: *The system boot log file*

```
HP Switch Switch 8200zl$ show boot-history

Mgmt Module 1 -- Saved Crash Information (most recent first):
=============================================================
Mgmt Module 1 in Active Mode went down:  11/07/09 14:48:36
Operator warm reload from CONSOLE session.

Mgmt Module 1 in Active Mode went down:  11/07/09 11:43:10
Operator cold reboot from CONSOLE session.


Mgmt Module 2 -- Saved Crash Information (most recent first):
=============================================================
  No Saved Crash Information
```

Enabling and disabling fabric modules

The fabric modules can be enabled or disabled even if they are not present in the switch. You cannot disable both fabric modules at the same time; one must be enabled.

Use this command to enable or disable the redundant fabric modules. Disabling one fabric module reduces the overall switching capacity of the series switches. On some networks where network utilization is less than 50%, you may not notice any degradation of performance.

Syntax

`redundancy fabric-module [1 | 2] [enable | disable]`

Allows enabling or disabling of fabric modules. (You cannot have both fabric modules disabled at the same time.)

Default: Both fabric modules are enabled.

NOTE

The redundant fabric modules do not support nonstop switching.

Example

Figure 202: *Disabling a fabric module*

```
HP Switch(config)# redundancy fabric-module 2 disable
HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                        Status   SW Version   Boot Image
---- ---------------------------------------- -------- ----------   ----------
MM1  HP Switch J9092A Management Module 8200zl Active   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl Standby  K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Disabled
```

Nonstop switching features

Nonstop switching features are synced at initialization of the standby management module.

802.1X and Web/MAC authenticationMAC Lockout/ LockdownACLs/Qos PoliciesPower over EthernetPort SecurityDHCP SnoopingDynamic ARP ProtectionDynamic IP Lockdown	Spanning Tree (MSTP)GVRPLoop ProtectionLACPSyslogUDLDVirus ThrottlingLLDP

Nonstop switching with VRRP

When Nonstop VRRP is enabled, VRRP continues to operate in its current state when a failover from the AMM to the SMM occurs. This provides an additional layer of redundancy in a switched network. VRRP state information is maintained between MMs so that VRRP operations resume immediately after failover from the AMM to SMM. Because of this quick resumption of operations there is no failover to the backup VRRP router in the network. The Master VRRP router continues to be active and operate as is.

The command for enabling Nonstop mode for VRRP must be executed in VRRP context.

Syntax

```
(vrrp#) [no] nonstop
```

Enabling Nonstop VRRP allows the VRRP router to retain control of IP addresses when the AMM fails over. The VRRP Backup router does not take control of the virtual IP addresses on the network.

The no version of the command disables Nonstop VRRP.

When Nonstop behavior is disabled, failure of the AMM on the VRRP Master results in the VRRP Backup router taking control of the virtual IP addresses on the network.

The commands must be executed in VRRP context.

NOTE

Before this command is executed, the command `redundancy management nonstop-switching` should be configured. Any prerequisites required for VRRP configuration commands, such as IP routing being enabled, remain as required prerequisites.

Default: Disabled

Example

Example of enabling nonstop switching for VRRP and then displaying the output

This example shows nonstop VRRP being enabled. The `show vrrp config` command output displays the enabled status (see bold line below.)

```
switch(vlan-10-vrid-1)# nonstop
switch(vlan-10-vrid-1)# show vrrp config

VRRP Global Configuration Information

VRRP Enabled  : Yes
Traps Enabled : Yes
Virtual Routers Respond to Ping Requests : Yes
VRRP Nonstop Enabled: Yes

VRRP Virtual Router Configuration Information

Vlan ID : 10
Virtual Router ID : 1

Administrative Status [Disabled] : Enabled
Mode [Uninitialized] : Backup
Priority [100] : 150
Advertisement Interval [1] : 1
Preempt Mode [True] : True
Preempt delay time : 0
Respond to Virtual IP Ping Requests [Yes] : Yes
Primary IP Address : Lowest

IP Address      Subnet Mask
--------------- ---------------
10.0.202.87     255.255.0.0
```

Example nonstop routing configuration

Example of configuring the owner routing switch

```
HP Switch C(config#) ip routing
HP Switch C(config#) router vrrp
HP Switch C(vrrp)# enable
HP Switch C(vrrp)# vlan 201
HP Switch C(vlan-201)# untag a1-a10
HP Switch C(vlan-201)# ip address 20.0.0.1/24
HP Switch C(vlan-201)# vrrp vrid 1
HP Switch C(vlan-201-vrid-1)# owner
HP Switch C(vlan-201-vrid-1)# virtual-ip-address 20.0.0.1/24
HP Switch C(vlan-201-vrid-1)# enable
```

Example of configuring the backup routing switch

```
HP Switch D(config#) ip routing
HP Switch D(config#) router vrrp
HP Switch D(vrrp)# enable
HP Switch D(vrrp)# vlan 201
HP Switch D(vlan-201)# untag a1-a10
HP Switch D(vlan-201)# ip address 20.0.0.2/24
HP Switch D(vlan-201)# vrrp vrid 1
```

```
HP Switch D(vlan-201-vrid-1)# backup
HP Switch D(vlan-201-vrid-1)# virtual-ip-address 2.1.1.1/24
HP Switch D(vlan-201-vrid-1)# enable
```

The configuration is shown graphically in **Figure 203: Example of nonstop routing configuration** on page 741.

Figure 203: *Example of nonstop routing configuration*

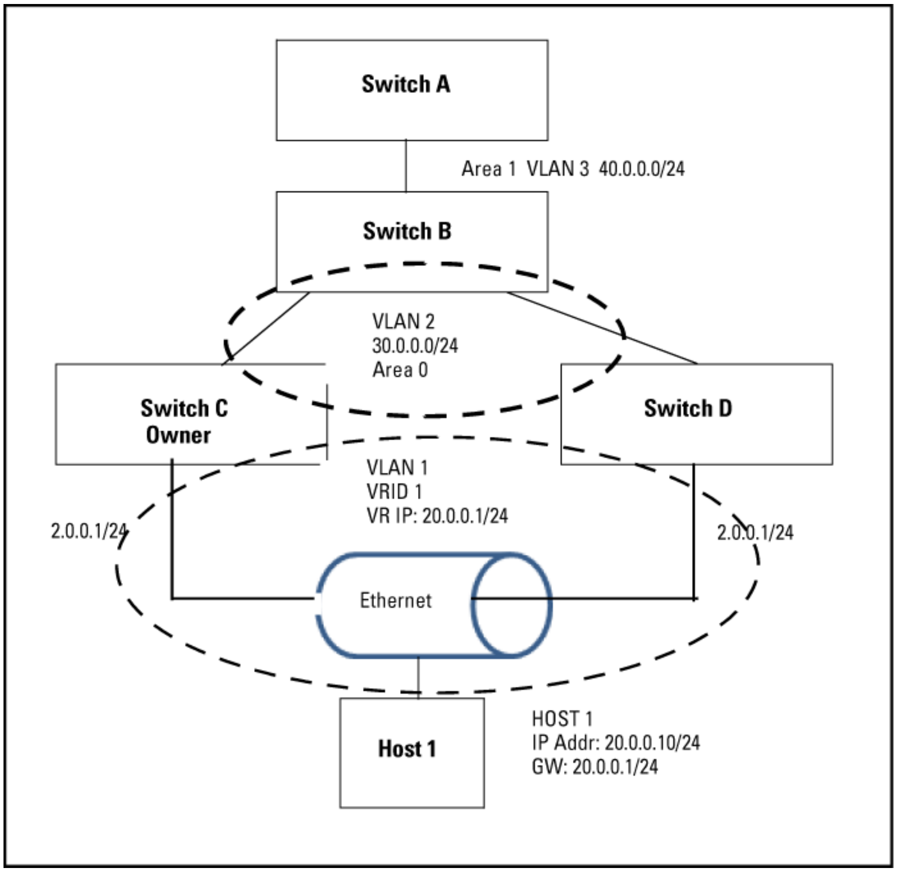

Nonstop forwarding with RIP

On a Nonstop RIP router, the traffic does not get re-routed when the MM fails over. A request packet is sent on failover that asks for the router's peers to send routing updates to the requesting router. There is no loss of routed traffic.

Nonstop forwarding with OSPFv2 and OSPFv3

On a Nonstop OSPFv2 router, failover of a MM does not result in the OSPF v2 router being removed from the OSPFv2 domain. A restart request is sent by the Nonstop OSPFv2 router to the neighboring OSPFv2 routers, after which the graceful restart process begins. This behavior applies to OSPFv3 as well.

A graceful restart allows an OSPF routing switch to stay on the forwarding path while being restarted. The routing switch sends "grace LSAs" that notify its neighbors that it intends to perform a graceful restart. During the configurable grace period, the restarting switch's neighbors continue to announce the routing switch in their LSAs as long as the network topology remains unchanged. The neighbors run in "helper mode" while the routing switch restarts.

Graceful restart will fail under these conditions:

- There is a topology change during the graceful restart period. The helper switches exit helper mode and adjacencies are lost until the restarting switch rebuilds the adjacencies.
- The neighbor switches do not support helper mode.

NOTE

- Configure router-id or IPv4 loopback address for OSPFv3 Non-Stop Forwarding to work on the switch.
- For OSPF nonstop switching to work without traffic loss (after switchover), redundancy rapid-switchover time should be greater than(~10secs) restart interval time of OSPF.

For more information on OSPFv2 and OSPFv3 graceful restart, see RFC 3623 and RFC 5187.

Enabling nonstop forwarding for OSPFv2

The routing switch must be in ospf context when enabling Nonstop forwarding for OSPFv2. To enable Nonstop forwarding, enter this command.

Syntax

```
(ospf)# [no] nonstop
```

Enables nonstop forwarding for OSPFv2.

The no version of the command disables nonstop forwarding.

The commands must be executed in `ospf` context.

Default: Disabled

Example of enabling nonstop forwarding for OSPFv2

```
switch(ospf)# nonstop
```

Configuring restart parameters for OSPFv2

Syntax

```
(ospf)# [no] restart interval 1-1800 [strict-lsa-checking]
```

Specify the graceful restart timeout interval in seconds.

The no version of the command sets the restart parameters to the default values.

Default: Disabled

`interval 1-1800`

The graceful restart timeout interval (grace period) in seconds. Default: 120 seconds

`strict-lsa-checking`

Used in OSFPv2 context to enable or disable strict LSA operation in a network segment for a neighboring router that is attempting a graceful restart. When enabled, this operation halts Helper mode support if a change in LSAs (topology change) is detected during the neighbor's restart period.

The no form of this command disables strict LSA operation.

Default: Strict LSA operation enabled

Viewing OSPFv2 nonstop forwarding information

To display the status of Nonstop forwarding information, enter the `show ip ospf general` command.

Example of output showing status of nonstop forwarding for OSPFv2

```
switch# show ip ospf general

OSPF General Status

OSPF protocol   :enabled
Router ID       :10.10.10.80
.
.
.
Nonstop forwarding : Enabled
Graceful Restart Interval : 500
Graceful Restart Helper Mode : Enabled
.
.
.
```

Enabling nonstop forwarding for OSPFv3

The routing switch must be in `ospf3` context when enabling Nonstop forwarding for OSPFv3. To enable nonstop forwarding, enter this command.

Syntax

```
(ospf3)# [no] nonstop
```

Enables nonstop forwarding for OSPFv3.

The no version of the command disables nonstop forwarding.

The commands must be executed in `ospf3` context.

Default: Disabled

Example of enabling nonstop forwarding for OSPFv3

```
switch(ospf3)# nonstop
```

Configuring restart parameters for OSPFv3

Syntax

```
(ospf3)# [no] restart interval 1-1800 [strict-lsa-checking]
```

Specify the graceful restart timeout interval in seconds.

The no version of the command sets the restart parameters to the default values. Default: Disabled

`interval 1-1800`

The graceful restart timeout interval (grace period) in seconds. Default: 120 seconds

`strict-lsa-checking`

Used in OSFPv3 context to enable or disable strict LSA operation in a network segment for a neighboring router that is attempting a graceful restart. When enabled, this operation halts Helper mode support if a change in LSAs (topology change) is detected during the neighbor's restart period.

The no form of this command disables strict LSA operation.

Default: Strict LSA operation enabled

Viewing OSPFv3 nonstop forwarding information

To display the status of Nonstop forwarding information, enter the `show ipv6 ospf3 general` command.

Example of output showing status of nonstop forwarding for OSPFv3

```
switch# show ipv6 ospf3 general

OSPFv3 General Status

 OSPFv3 protocol  :enabled
 Router ID      :10.10.10.80
.
.
.
 Nonstop forwarding : Enabled
 Graceful Restart Interval : 500
 Graceful Restart Helper Mode : Enabled
.
.
.
```

About downloading a new software version

File synchronization after downloading

After downloading a new software version to either the primary or secondary flash of the active management module, the software version is immediately copied to the corresponding flash (primary or secondary) of the standby module, unless the standby module failed selftest or redundancy was disabled with the `no redundancy management-module` command.

The configuration files, including which configuration file to use for that flash image, are synchronized. For example, if the active management module is using `config1`, the standby module is also synchronized to use `config1`.

Table 29: *Example of upgrading software version K.15.01.0003 to version K.15.01.0004*

	Newer code to secondary flash		New code to primary flash	
	Active MM	Standby MM	Active MM	Standby MM
Software version downloaded to Primary flash image	K.15.01.0003	K.15.01.0003	K.15.01.0004	K.15.01.0004
Software version downloaded to Secondary flash image	K.15.01.0004	K.15.01.0004	K.15.01.0003	K.15.01.0003

After installing the new software to the active management module, wait a few minutes, and then verify that the standby management module has been synchronized with the new software as well (use the `show flash` command.) If the default flash for boot is set correctly, you can start the standby management module on the new software by executing the `boot standby` command. This does not interrupt current switch operations yet. After the standby management module has rebooted and is ready for takeover in standby mode (you can verify this using the `show redundancy` command.) you can now switch over to the management module running the newer software with this command:

```
HP Switch# redundancy switchover
```

This causes a switchover to the management module that received the new software version, which becomes the active management module. This method incurs the least amount of network downtime for booting. If downtime is not an issue, use the `boot system` command. Both management modules are then running the new software version.

Potential software version mismatches after downloading

When a new software version is downloaded to the active management module, it is immediately copied to the corresponding flash (primary or secondary) in the standby management module, unless redundancy has been disabled. If the standby management module is rebooted, it will be running a different software version than the active management module. You can direct the standby module to boot from the non-corresponding flash image that has a different software version during the actual reboot process of the standby module when the prompt to select the **Boot Profile** appears, as shown in **Figure 204: Booting the standby management module to secondary flash** on page 745.

Figure 204: *Booting the standby management module to secondary flash*

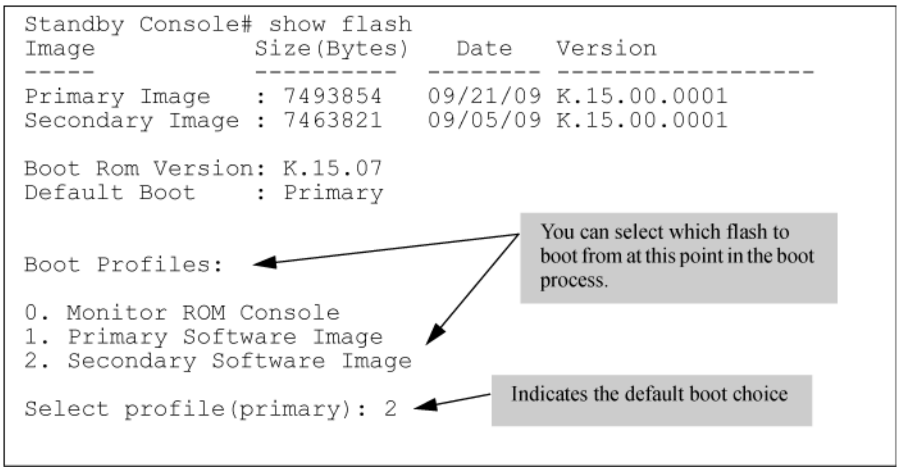

CAUTION

If you have booted one module out of primary flash and one module out of secondary flash, and the secondary flash is running a prior software version because the latest version was never copied over from the primary flash, you will have a software version mismatch. The configuration file may not work with that software version.

The standby module enters warm-standby redundancy mode and boots to a certain point, syncs basic files such as the config and security files, and finishes booting only if the active management module fails or you choose to change which module is the active module..

Additionally, if a switchover occurs, or if you reboot to make the standby module become the active module, any configuration file changes made may not work on the active module if it has a different software version from the standby module.

When you enter the `show redundancy` command and a software version mismatch exists, a warning message is displayed, as shown at the bottom of **Figure 205: Example of a software version mismatch between the active and standby modules** on page 746.

Figure 205: *Example of a software version mismatch between the active and standby modules*

```
HP Switch(config)# show version
Management Module 1: Active
Image stamp:    /sw/code/build/btm(t2g)
                Mar 15 2007 12:28:32
                K.15.01.0001
                64
Boot Image:     Primary


Management Module 2: Standby
Image stamp:    /sw/code/build/btm(t2g)
                Mar 21 2007 14:24:38
                K.15.01.0002
                789
Boot Image:     Secondary


HP Switch(config)# show redundancy

Settings
--------
Mgmt Redundancy : Warm-standby redundancy enabled
Rapid Switchover Stale Timer : 0

Statistics
----------
Failovers     : 0
Last Failover :

Slot Module Description                        Status   SW Version   Boot Image
---- ---------------------------------------- -------- ----------   ----------
MM1  HP Switch J9092A Management Module 8200zl Active   K.15.01.0001 Primary
MM2  HP Switch J9092A Management Module 8200zl Standby  K.15.01.0002 Secondary

FM1  HP Switch J9093A Fabric Module 8200zl     Enabled
FM2  HP Switch J9093A Fabric Module 8200zl     Enabled


Warning: Standby module is running a different software version and may be using
a different configuration file. Configuration changes on active management
module may not take effect on a failover.
```

Mismatch exists

Downloading a software version serially if the management module is corrupted

If the software version on a management module becomes corrupted, you may need to do a serial download to restore the affected module. The non-corrupted management module becomes the active module. You can then use the serial port on the corrupted management module to download a new software version. When the corrupted module is rebooted, the software version in the corrupted module is immediately overwritten by the software version in the active management module. Both management modules should now operate on the same software version.

Unsupported zl modules

ZL modules/controllers that do not support the nonstop switching feature include the following:

- HPE ONE Services zl Module (J9289A)
- HPE Threat Management Services zl Module (J9155A)
- HPE Threat Management Services zl Module with 1-year IDS/IPS subscription (J9156A)
- HPE Wireless Edge Services zl Module (J9051A) and Redundant Wireless Services zl Module (J9052A)
- HPE MSM765zl Mobility Controller (J9370A)

During a nonstop switching failover, unsupported modules will not failover seamlessly to the standby module. A nonstop switching failover causes a forced reboot on these modules. After rebooting, these modules then sync with the newly active management module and begin operation again. Module traffic is disconnected until the module completes the reboot process.

Hot swapping of management modules

Use the MM Shutdown button on the front of the management module before removal. The Shutdown button ensures that the management module is shut down properly. If nonstop switching is enabled, using the Shutdown button prior to removal ensures failover to the standby module will be successful.

Rapid routing switchover and stale timer

With K.15.01.0031, nonstop switching supports only Layer 2 functions on the switch. During a failover, traffic routed through the switch at Layer 3 will see an interruption. When a failover from active to standby occurs, the routing table is "frozen." All routes that existed at the time of the failover are marked as "stale." While dynamic routing protocols running at the time may act as if the routing protocol has been restarted and rebuilds the table, the switch on which the failover occurred continues to rout traffic using the ‘stale routes.'

The "stale timer" begins counting when the switchover occurs. When the "stale timer" expires, any routes that are still marked as stale are purged from the routing table. Because of the nature of rapid routing switchover, if there are multiple simultaneous failures, network loops could occur or traffic could flow through unpredictable paths.

Use caution if setting the "rapid-switchover" timer higher than the default. To disable "rapid routing switchover" and to ensure that all routing is based on the most current routing protocol information, set the "rapid-switchover" timer to 0.

Task Usage Reporting

The task usage reporting feature provides the ability to collect and display CPU usage data (with a refresh rate of 5 seconds) of running tasks on the switch. It includes the following commands:

- `process-tracking`

  : Use this command to enable/disable the task-usage collecting capability for a specific module on the switch.
- `show cpu process`

  : Use this command to display task-usage statistics for a specific module.

Help text

process-tracking help

Usage: `[no] process-tracking [slot[SLOT-LIST] [<time>]] [<time>]`

Description: Enable/disable module process-tracking functionality.

show cpu help

Usage: `show cpu [<CHASSIS_MIN_CPU_UTIL_INDEX-CHASSIS_MAX_CPU_UTIL_INDEX>]`

```
[slot <SLOT-LIST> [<CHASSIS_MIN_CPU_UTIL_INDEX-
CHASSIS_MODULE_MAX_CPU_UTIL_INDEX>] ]
```

```
[process [[slot <SLOT-LIST>] [refresh <iterations>]]
```

```
[refresh <iterations>]
```

Description: Show average CPU utilization over the last 1, 5, and 60 seconds, or the number of seconds specified.

Use the 'slot' argument to display CPU utilization for the specified modules, rather than the chassis CPU.

Use the 'process' argument to display module process usages.

show cpu process help

Usage: `show cpu process [slot [SLOT-LIST][refresh <iterations>]]`

`[refresh <iterations>]`

Description: Display module process usage.

Command tab

process-tracking

process-tracking <tab>

`slot`

Enable/disable process-tracking for a module

`INTEGER`

Specify time to track value between 1 second to 30 seconds

`<cr>`

process-tracking slot <tab>

`SLOT-ID-RANGE`

Enter an alphabetic device slot identifier or slot range

process-tracking slot A

`INTEGER`

Specify time to track value between 1 second to 30 seconds

`<cr>`

process-tracking slot A 10 <tab>

`<cr>`

process-tracking 10 <tab>

`<cr>`

show cpu process

show cpu <tab>

`process`

Display process usage

`slot`

Display module CPU statistics

`<1-300>`

Time (in seconds) over which to average CPU utilization

`<cr>`

show cpu process <tab>

refresh

Number of times to refresh process usage display

slot

Display module process usage

<cr>

show cpu process refresh <tab>

INTEGER

Enter an integer number

show cpu process refresh 10 <tab>

<cr>

show cpu process slot <tab>

SLOT-ID-RANGE

Enter an alphabetic device slot identifier or slot range

show cpu process slot A <tab>

refresh

Number of times to refresh process usage display

<cr>

show cpu process slot A refresh <tab>

INTEGER

Enter an integer number

show cpu process slot A refresh 10 <tab>

<cr>

Command ouput

show cpu process

HPE-5406zl# show cpu process

```
                             | Recent |  %  | Time Since|   Times   |   Max
Process Name        | Priority | Time   | CPU | Last Ran  |    Ran    |  Time
--------------------+----------+--------+-----+-----------+----------+-------
           Idle-1 |      226 |  10 s  |  41 |     57 us |   380986 |  69 us
           Idle-3 |        1 |   5 s  |  20 |     52 us |   761665 |  55 us
           Idle-0 |      226 |   8 s  |  33 |     19 us |   380867 |  66 us
Sessions & I/O-24 |      171 | 926 ms |   3 |      1 ms |      150 | 335 ms
```

show cpu process slot <SLOT-LIST>

HPE-5406zl# show cpu process slot A

slot a:

```
                             | Recent |  %  | Time Since|   Times   |   Max
```

```
Process Name       | Priority | Time    | CPU | Last Ran   |    Ran   |  Time
--------------------+----------+--------+-----+-----------+----------+-------
System Services-2 |      156 | 253 ms |   2 |    767 ms |       12 |  35 ms
           Idle-3 |        1 |   3 s  |  28 |     13 ms |   101309 | 150 us
  Hardware Mgmt-2 |      192 | 282 ms |   2 |    303 us |       44 |  12 ms
           Idle-1 |      226 |   6 s  |  55 |     13 ms |    50793 | 233 us
           Idle-0 |      226 |   1 s  |   9 |     14 ms |    50633 | 106 us
```

Smart Rate is a new technology designed to enable higher port link speeds on legacy cabling where an Ethernet RJ45 port type can link at 1Gbps, 2.5Gbps, 5Gbps, or 10Gbps.

When situations occur where a network link establishes at a lower than expected speed (or not at all) due to marginal or bad cabling, the Smart Rate port technology allows administrators to triage cabling issues and determine root causes of lower than expected performance.

Smart Rate Technology is available on the following products:

- Switch 5400R v3 zl2 modules (J9991A, J9995A)
- Switch 5400R chassis switch bundles (JL002A)

Show Smart Rate port

Syntax

`show interface` *`PORT-LIST`* `smartrate`

Displays port diagnostics on a Smart Rate port.

Unlinked Smart Rate port

```
show interface C5 smartrate

Status and Counters - Smart Rate information for Port C5

Model          : 0x03a1
Chip          : 0xb4b3
Firmware (major)       : 0x0002
Firmware (minor)       : 0x0003
Firmware (candidate)      : 0x0005
Firmware (provision)      : 0x0001

      Chan1    Chan2    Chan3    Chan4 (in db)
Current SNR  9.000000  6.700000  3.500000  9.200000
Minimum SNR  9.000000  6.700000  3.500000  9.200000

CRC8 errors:     0
LDPC errors:     0
LDPC 1 iteration:    27620089
LDPC 2 iterations:    954117
LDPC 3 iterations:    0
LDPC 4 iterations:    0
LDPC 5 iterations:    0
LDPC 6 iterations:    0
LDPC 7 iterations:    0
LDPC 8 iterations:    0

23  Number of fast retrains requested by Local Device.
32  Number of fast retrains requested by Link Partner.
```

```
150 Accumulated time (ms) spent in fast retrain since last AN.
9  Number of RFI Training Link Recovery Events since last AN.
3  Number of Link Recover Events since last AN.

Established link speed       : 5000Mbps
Number of attempts to establish link   : 5
Uptime since link was last established (ms) : 5099

Local port advertised speeds

1000Mbps 2500Mbps  5000Mbps  10Gbps
 No   Yes   Yes   No

Link partner speed capability

1000Mbps  2500Mbps  5000Mbps  10Gbps
 Yes   Yes   Yes   No

Link Partner matching vendor: Yes
```

Smart Rate port that is linked at 1Gbps

```
show interface C5 smartrate

Status and Counters - Smart Rate information for Port C5

Model          : 0x03a1
Chip           : 0xb4b3
Firmware (major)     : 0x0002
Firmware (minor)     : 0x0003
Firmware (candidate)    : 0x0005
Firmware (provision)    : 0x0001

Established link speed         :1000Mbps
Number of attempts to establish link    :5
Uptime since link was last established (ms)  : 5099

Local port advertised speeds

1000Mbps  2500Mbps  5000Mbps  10Gbps
 No     No     No     No

Link partner speed capability

1000Mbps  2500Mbps  5000Mbps  10Gbps
 Yes    Yes    Yes    Yes

Link Partner matching vendor: Yes
```

Rate-Limiting — GMB features when Fast-Connect SmartRate ports are configured

When Rate-Limiting or Guaranteed Minimum Bandwidths are configured for 5Gbps ports, the granularity of percentage-based rates for the 5Gbps speed is in steps of 2%. For example, a 1% rate-limit for a 5Gbps port will function as a 2% limit while a 5% limit will function as a 6% limit.

The Guaranteed Minimum Bandwith profiles will show the same behavior. For example on an 8-queue system, the actual default servicing profile will be 2%, 4%, 30%, 10%, 10%, 10%, 16%, and 20%. The CLI and SNMP values for these ports will show what the customer configured, but the actual hardware results will be in steps of 2%.

This limitation only applies to 5Gbps ports. Ports running at 2.5Gbps have the same 1% granularity as all previously-offered port speeds.

Error messages

When the `show interface PORT-LIST smartrate` command is run on a non-Smart Rate port, the command will fail with an error message similar to the following: `Port A1: This command is only applicable to Smart Rate ports.`

When the `show interface PORT-LIST smartrate` command is run on a Smart Rate port, but is unable to retrieve all results the command will fail with an error message similar to the following: `Port A1: This command did not complete successfully. Please try again.`

Speed-duplex

Syntax

```
interface PORT-LIST speed-duplex
```

Options

auto	Auto-negotiate link parameters.
auto-1000	1000 Mbps only, auto-negotiate link parameters.
auto-2500	2500 Mbps only, auto-negotiate link parameters.
auto-5000	5000 Mbps only, auto-negotiate link parameters.
auto-2500-5000	2500 or 5000 Mbps only, auto-negotiate link parameters.
auto-10g	10 Gbps only, auto-negotiate link parameters.

Limitations on 5Gbps ports

For 5Gbps ports, when the customer has Rate-Limiting or Guaranteed Minimum Bandwidths configured, the granularity of percentage-based rates for the 5Gbps speed is in steps of 2%. For example a 1% rate-limit for a 5Gbps port will function as a 2% limit, a 5% limit will function as a 6% limit. The Guaranteed Minimum Bandwith profiles will show the same behavior. On an 8-queue system, the actual default servicing profile will be 2% 4% 30% 10% 10% 10% 16% 20%. The CLI and SNMP values for these ports will show what the customer configured, but the actual hardware results will be in steps of 2%.

NOTE

This limitation only applies to the 5Gbps ports. Ports running at 2.5Gbps have a 1% granularity in port speeds.

Error messages

- On ports that do not support the respective speed-duplex option, the command will fail with an error message similar to the following:
  - Value `auto-10` is not applicable to port E1.
- The following speed-duplex options are not available on switch platforms that do not have Smart Rate ports.

- auto-2500
- auto-5000
- auto-2500-5000

Introduction

The HPE Networking 6th Generation Switch ASIC based module creates compatibility between v2 and v3 blades on the 5400R Chassis Switches. When the 5400R Chassis Switch platform detects a mix of v2 and v3 blades, the v3 feature will default the platform to v2 behavior. The default behavior is v2.

The compatibility mode of v2 and v3 modules are controlled by configuration. When the compatibility mode is disabled, v2 modules in the system will be disabled.

Commands

Configuration commands enable/disable the 5400R Chassis Switch platform v2/v3 interoperability.

Configuration setup

Syntax

```
[no] allow-v2-modules
```

Enables support for V2 modules. When enabled, V3 modules will operate in V2-compatiblity mode. When disabled, V3 modules will have full functionality and the ports of any V2 modules will be non-operable. Enabling the V2 module support erases the current configuration of the device and reboots the device. Whereas, disabling the V2 module support clears all V2 module specific configuration from startup configuration and reboots the device.

allow-v2-modules

Enable support for V2 modules.

Enabled/Disabled state

When V2 compatibility mode is disabled from an enabled state, the below message is displayed for user input.

```
HP-5406Rzl2(config)# no allow-v2-modules
This command will disable all V2 modules and reboot the switch.
Continue (y/n)?
```

When V2 compatibility mode is enabled from disabled state, the below message is displayed for user input.

```
HP-5406Rzl2(config)# allow-v2-modules
This command will erase the current configuration of the switch and reboot it.
Continue (y/n)?
```

V3 to V2 compatibility

On an Aruba 5400R switch loaded, when switching from v3 only mode to v2 compatibility mode, the V2 specific module entries and related configuration options are retained. The switch erases the entire configuration when moving from V3 only mode to V2 mode.

allow-v2-modules

Syntax

`allow-v2-modules`

Description

Enables the use of v2 modules in the switch. Before enabling v2 support, the command determines whether any of the v3-native features are configured. This command is active only in v3–native mode.

```
V3 modules cannot be changed to v2-compatibility mode when there are v3-specific
configuration settings. Use the command 'show running-config v3-specific' to
display the settings that must be changed before enabling v2 module support.
```

Unconfigure all v3-only features before moving to compatibility mode.

If the v3-native configuration is not present, the device reboots with the non-v3 configuration and issues the following message:`This command will save the running configuration and reboot the system with all V3 modules operating in v2-compatibility mode.Continue (y/n)?`

Options

erase

Erases the entire configuration and reboots the switch with either the custom default configuration (if it is available) or the factory default configuration. The command with the erase option is active only in v3–native mode.

Usage

`allow-v2-modules erase`

More Information

To list the v3-native configurations present in the current configuration, see **show running-config v3-specific** on page 756.

`show running-config v3-specific`

Syntax

`show running-config v3-specific`

Description

Provides a list of V3-native configurations present in the current configuration.

```
show running-config v3-specific

vsf
   enable domain 40
   member 1
      type "J9850A" mac-address 645106-8a0400
      priority 200
      link 1 1/A24
```

```
      link 1 name "I-Link1_1"
      exit
   member 2
      type "J9850A" mac-address 40a8f0-9e6600
      priority 150
      link 1 2/A24
      link 1 name "I-Link2_1"
      exit
   exit
oobm
   vsf member 1
      ip address dhcp-bootp
      exit
   vsf member 2
      ip address dhcp-bootp
      exit
   exit

HP-VSF-Switch#
```

Show commands

The `show module` command shows the configuration status of allowed V2 modules. The output will be available only for the 5400R Chassis Switches.

Show system

Syntax

```
show system
```

Show system output

```
Status and Counters - General System Information

System Name         : HP-5406Rzl2
System Contact      :
System Location     :
Allow V2 Modules    : Yes
MAC Age Time (sec)  : 300
Time Zone           : 0
Daylight Time Rule  : None
Software revision   : KB.15.16.0000x       Base MAC Addr       : 40a8f0-9d6f00
ROM Version         : KB.15.Z1.0012       Serial Number       : SG44G490FZ
                   Opacity Shields     : Not Installed

Up Time             : 7 mins                    Memory    - Total    : 763,846,656
CPU Util (%)        : 0                               Free     : 646,757,632

IP Mgmt   - Pkts Rx : 106                       Packet    - Total    : 6750
     Pkts Tx : 111                        Buffers     Free     : 4830
                   Lowest   : 4828
                   Missed   : 0
```

Show system information

Syntax

```
show system information
```

Show system information output

```
Status and Counters - General System Information

System Name         : HP-5406Rzl2
System Contact      :
System Location     :
Allow V2 Modules    : Yes
MAC Age Time (sec)  : 300
Time Zone           : 0
Daylight Time Rule  : None
Software revision   : KB.15.16.0000x       Base MAC Addr      : 40a8f0-9d6f00
ROM Version         : KB.15.Z1.0012       Serial Number      : SG44G490FZ
                   Opacity Shields    : Not Installed

Up Time             : 7 mins         Memory   - Total   : 763,846,656
CPU Util (%)        : 0           Free    : 646,757,632

IP Mgmt  -  Pkts Rx : 106          Packet   - Total   : 6750
     Pkts Tx : 111          Buffers    Free    : 4830
                   Lowest  : 4828
                   Missed  : 0
```

Show running configuration

The `show running-config` command shows the entry when disabled. This output will be available on 5400R Chassis Switches only.

Syntax

```
show running-config
```

Show running configuration output

```
; J9850A Configuration Editor; Created on release #KB.15.16.0000x
; Ver #05:18.7f.ff.3f.ef:4d
hostname "HP-5406Rzl2"
module A type j9987a
module F type j9993a
no allow-v2-modules
snmp-server community "public" unrestricted
oobm
  ip address dhcp-bootp
  exit
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,F1-F8
  ip address dhcp-bootp
  exit
```

Event logging

Table 30: *Interoperability messages*

Event		Message
When switch is rebooting after a change in the interoperability mode.	V2/V3 interoperability message	`M 05/23/14 05:50:15 00064 system: Rebooting the device because the module compatibility mode has changed.`
	V1/V2 Interoperability message for reference	`M 05/23/14 05:50:15 00064 system: Rebooting for interOperabilityMode change`
	Modified V1/V2 interoperability message (same as V2/V3 message)	`M 05/23/14 05:50:15 00064 system: : Rebooting the device because the module compatibility mode has changed.`

Version 2 — version 3 blade compatibility on the 5400R switch

Allow V2 command

The CLI commands `allow-v2-modules` and `[no] allow-v2-modules` enable the configuration for compatibility of V3 and V2 modules to operate simultaneously. Disabling compatibility will disallow V3 and V2 modules from operating simultaneously, allowing only V3 modules to operate.

Syntax

```
[no] allow-v2-modules
```

Enable/disable support for V2 modules.

Validation rules

Validation	Error/Warning/Prompt
Compatibility Mode enabled – ‘no allow-v2-modules’	Prompt: ‘All V2 modules will be disabled. Continue [y/n] ?’
Compatibility Mode enabled – ‘allow-v2-modules’	No prompt
Compatibility Mode disabled – ‘no allow-v2-modules’	No prompt
Compatibility Mode disabled – ‘allow-v2-modules’	Prompt: ‘This will erase the configuration and reboot the switch. Continue [y/n] ?’

Show commands

Syntax

```
show system
```

Enable/disable support for V2 modules.

Show system

```
Status and Counters - General System Information
System Name         : HP-5406zl2
System Contact      :
System Location     :
Allow V2 Modules    : Yes
```

Event Log

Event	Message
Compatibility Mode disabled – ‘allow-v2-modules’	Rebooting for interOperabilityMode change

Appendix D
MAC Address Management

Overview

The switch assigns MAC addresses in these areas:

- For management functions, one Base MAC address is assigned to the default VLAN (VID = 1.) (All VLANs on the switches covered in this guide use the same MAC address.)
- For internal switch operations: One MAC address per port.

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.

NOTE

The switch’s base MAC address is also printed on a label affixed to the switch.

Determining MAC addresses

`Use the CLI` to view the switch's port MAC addresses in hexadecimal format.

`Use the menu interface` to view the switch's base MAC address and the MAC address assigned to any VLAN you have configured on the switch. (The same MAC address is assigned to VLAN1 and all other VLANs configured on the switch.)

NOTE

The switch's base MAC address is used for the default VLAN (VID =1) that is always available on the switch. This is true for dynamic VLANs as well; the base MAC address is the same across all VLANs.

Viewing the MAC addresses of connected devices

Syntax

```
show mac-address [<PORT-LIST> | mac-addr | vlan vid]
```

Lists the MAC addresses of the devices the switch has detected, along with the number of the specific port on which each MAC address was detected.

[<*PORT-LIST*>]	Lists the MAC addresses of the devices the switch has detected, on the specified ports.
[*mac-addr*]	Lists the port on which the switch detects the specified MAC address.Returns the following message if the specified MAC address is not detected on any port in the switch: `MAC address  mac-addr  not found.`
[vlan *vid*]	Lists the MAC addresses of the devices the switch has detected on ports belonging to the specified VLAN, along with the number of the specific port on which each MAC address was detected.

Viewing the switch's MAC address assignments for VLANs configured on the switch

The Management Address Information screen lists the MAC addresses for:

Procedure

1. Base switch (default VLAN; VID=1)
2. Any additional VLANs configured on the switch.

Also, the Base MAC address appears on a label on the back of the switch.

NOTE

The Base MAC address is used by the first (default) VLAN in the switch. This is usually the VLAN named "DEFAULT_VLAN" unless the name has been changed (by using the VLAN Names screen.) On the switches covered in this guide, the VID (VLAN identification number) for the default VLAN is always "1,"

and cannot be changed

.

- From the Main Menu, select
- **1. Status and Counters**
- **2. Switch Management Address Information**
- If the switch has only the default VLAN, the following screen appears. If the switch has multiple static VLANs, each is listed with its address data.

Figure 206: *Example of the Management Address Information screen*

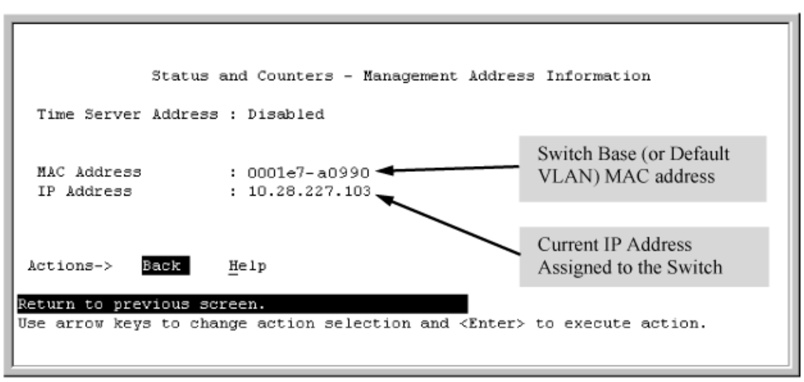

Viewing the port and VLAN MAC addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the spanning-tree protocol. Using the `walkmib` command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.

Table 31: *Switch series' and their MAC address allocations*

Switch series	MAC address allocation
8212zl	The switch allots 24 MAC addresses per slot. For a given slot, if a four-port module is installed, the switch uses the first four MAC addresses in the allotment for that slot, and the remaining 18 MAC addresses are unused.If a 24-port module is installed, the switch uses the first 24 MAC addresses in the allotment, and so on.
All Models	The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the `walkmib` listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)

NOTE

This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

Example

A 8212zl switch with the following module configuration shows MAC address assignments similar to those shown in **Figure 207: Example of Port MAC address assignments on a switch** on page 764:

- A 4-port module in slot A, a 24-port module in slot C, and no modules in slots B and D
- Two non-default VLANs configured

Figure 207: *Example of Port MAC address assignments on a switch*

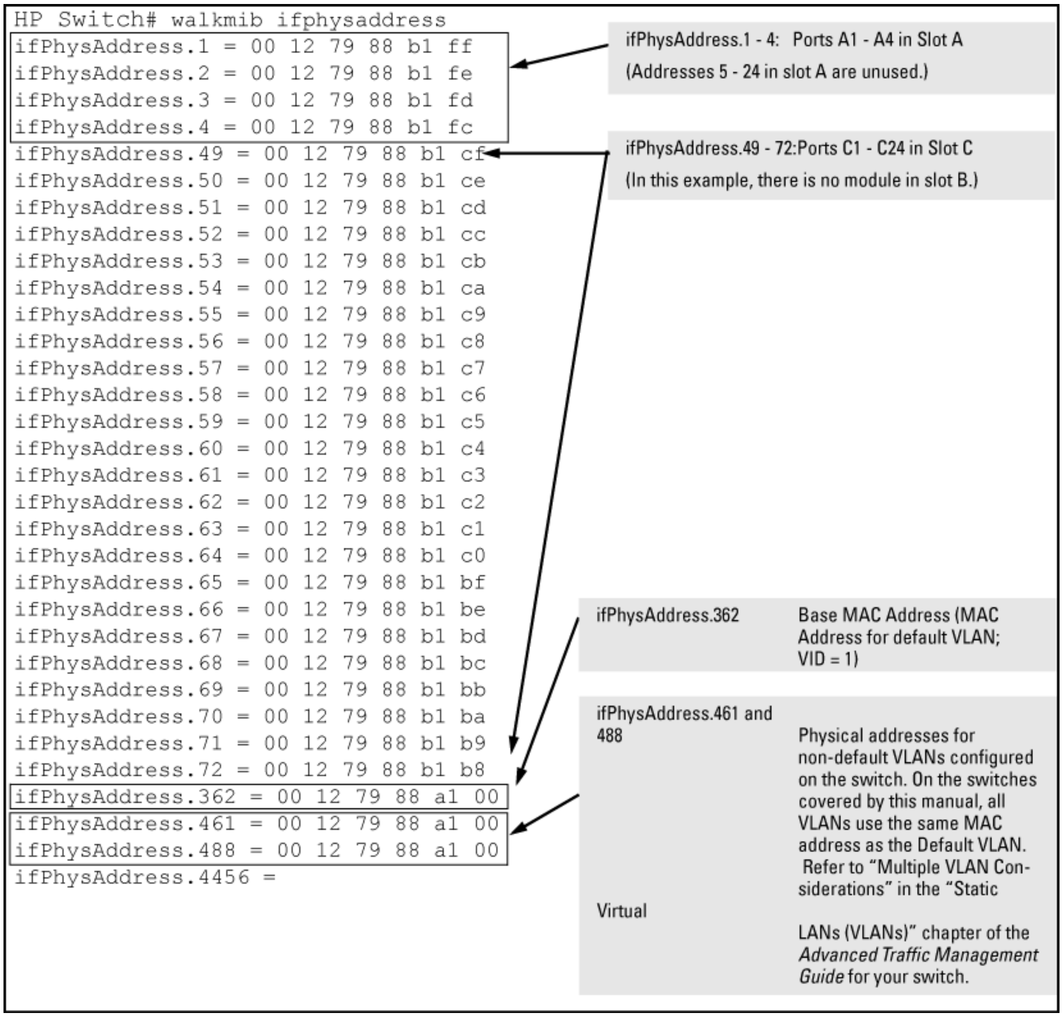

NOTE

When configuring an "out" monitor on a VLAN or an interface to a remote mirror, the mirrored packet will always be untagged when the original packet arrives on a zlV2 module, or a 3800 series.

When configuring a "both" monitor on an interface to a remote mirror, tags will not be present in the mirrored packet in these specific situations:

- For an interface monitor, packets transmitted by the monitored port that originally arrived on a zlV2 module or 3800 port.
- For a VLAN monitor, packets routed onto the monitored VLAN that originally arrived on a zlV2 module or 3800 port.

1. If the switch is at the CLI Operator level, use the `enable` command to enter the Manager level of the CLI.
2. Enter the following command to display the MAC address for each port on the switch:

```
HP Switch# walkmib ifPhysAddress
```

3. (The above command is not case-sensitive.)