IP Primer

Everything you need to know

CALREC

IP Primer

Everything you need to know

This AoIP network primer offers a comprehensive overview of the technologies involved in implementing AoIP.

One does not need to read this guide from front to back to implement AoIP networks. AoIP networks can consist of a few switches and some AoIP devices.

Network configuration can be done using factory defaults or prebuilt configurations available from various sources.

On the other hand, AoIP networks can be extraordinarily complex ecosystems specifically designed to meet the demanding workflows, resilience, and elastic requirements of a broadcaster.

This guide details the underlying technologies that are required to implement AoIP in its simplest and most complex forms.

Some prior knowledge is assumed and networking fundamentals are not covered in this guide. We have however created a series of webinars for this exact purpose which can be found on our website:

https://calrec.com/calrec-sound-institute/

Contents

Calrec IP Primer
calrec.com

Terminology

ACL – Access Control List
BMCA – Best Master Clock Algorithm
bps – bits per second
CLI – Command Line Interface
COTS – Commercial off the Shelf
DHCP – Dynamic Host Configuration Protocol
DNS – Domain Name System
DNS-SD – Domain Name System – Service Discovery
DSCP – Differentiated services code point
Gbps – Gigabits per second
GMC – Grand Master Clock
GUI – Graphical User Interface
IB – In band
IGMP – Internet group membership protocol
IP – Internet Protocol
MAC – Media Access Control
Mbps – Megabits per second
mDNS – Multicast Domain Name System
Mpps – Millions of packets per second
MTU – Maximum transmission unit
NAT – Network Address Translation
NIC – Network Interface Card
OB – Outside Broadcast
OOB – Out of Band
OUI – Organizationally unique identifier
PDV – Packet delay variation
PIM – Protocol Independent Multicast
PTP – Precision Time Protocol
QoE – Quality of Experience
QoS – Quality of Service
RP – Rendezvous Point
RPF – Reverse Path Forward
RTP – Real-time transport protocol
SDN – Software defined networking
SPF – Shortest Path First
SSM – Source specific Multicast
SNMP – Simple Network Messaging Protocol
Tbps – Terabit per second
TCP – Transmission Control Protocol
UDP – User Datagram Protocol
UI – User Interface
VPN – Virtual Private Network
VRF – Virtual Routing and Forwarding

Introduction

This AoIP network guide offers a comprehensive overview of the technologies involved in implementing AoIP.

What is a network?

Fundamentally, a network is a tool that allows the exchange of data between devices. In the case of AoIP, the data is audio, and the devices are mixing desks, loudspeakers, microphone preamps etc. The network should be invisible to the end-user; for all the user needs to be concerned about is audio production.

The network should enable the workflow requirements of the broadcaster. Due to the variable nature of different broadcaster workflows, this often means that there is not a one size fits all network solution. Thankfully, IP technologies are flexible and can be adapted to many different workflows and use cases. AoIP Networks can enable your desired workflow without defining it.

Practically, networks are made up of network devices such as switches, routers, firewalls and others. In broadcast, layer 2 and layer 3 switches are the most common due to their scalability and performance characteristics.

The methods of passing data through a network vary between layer 2 and layer 3 technologies, and each has its pros and cons. This guide goes into these technologies in detail.

Benefits of AoIP Networking

The huge rise in demand for content and evolving consumption methods have fuelled demand for flexible and scalable infrastructures. IP infrastructures contain all the necessary characteristics needed to aid content creation in an environment where production requirements are often demanding, scalable and variable.

AoIP facilitates flexible workflows and provide opportunities for content creation that were not previously available, whilst also maintaining broadcaster's legacy workflows. It allows interoperable resource sharing that can increase agility and collaboration. It provides an infrastructure that can be used for both file based and live to air workflows.

AoIP networking is future proofed too; IP networks can scale upwards to supply higher levels of throughput to account for:

- Future production requirements
- Higher bandwidths required by the next generation of audio and video formats and standards.

Facilities that are IP native also have the added benefit of simplified interfacing with other facilities, whether they be studios, OB vans, remote locations or cloud environments.

Standards

For AoIP to be interoperable, the rules of transmission, discovery and connection must be standardised. This effectively means that all AoIP products speak the same language for all their operational interactions.

This section outlines the main standards that allow for efficient and interoperable AoIP.

AES67

AES67 defines the format for the transport of audio over IP using a collection of existing, proven and open protocols. The goal of AES67 is to provide the interoperable exchange of audio between different product vendors.

Amongst other things, AES67 defines audio format (sample rate, bit depth), packet format, clock usage, QoS markings, buffering requirements and addressing. Protocols referenced in AES67 include IP, UDP, RTP, SDP, IGMP, PTPv2.

AES67 mentions several options for device discovery but does not mandate any particular technology.

SMPTE ST2110

SMPTE's 2110 defines "elemental streaming". The elemental streaming methodology uses separate RTP streams for video, audio and metadata, making audio/metadata payload more accessible and independently route-able, while the RTP protocol takes care of time-alignment of the different streams. SMPTE ST2110 is popular within the broadcast industry because of inclusiveness of video, audio and data.

ST2110-30

ST2110-30 is the audio specific ST2110 standard. For the most part, ST2110-30 references AES67 for media transport, but adds more specific audio and packet formats that would grant an AoIP device different levels of compliance.

ST2022-7

Network redundancy is not mandatory under either AES67 or ST2110, but ST2110 does say that if redundancy is provided, then it should conform to ST2022-7.

Under ST2022-7, a media stream is duplicated, allowing receiving devices to take data on a packet-by-packet basis from multiple streams. Typically, there are two streams used for each media flow (although 2022-7 itself is not limited to two); devices output media over their "first" connection and duplicate that media over streams on their "second" connection.

There are several nomenclatures in circulation for names of the "first" and "secondary" networks including: primary/secondary, A/B, Red/Blue, Amber/Blue. In this guide we will refer to independent ST2022-7 networks as Amber/Blue as to be consistent with the JT-NMs TR1001-1.

Both Amber and Blue connections can pass over the same network hardware (and some devices even support ST2022-7 redundant streams over a single physical NIC interface), but for best redundancy, including against network failure, physically separate networks are often deployed. A receiving device is free to take the media from either the Amber or Blue stream, seamlessly on a packet-by-packet basis, which protects against packets being dropped or becoming corrupt as they traverse the network.

PTP

Precision Timing Protocol, specifically version 2 (with a formal name of IEEE-1588-2008), is a method of synchronising clocks across a network with sub microsecond accuracy.

PTP is an extremely important protocol for AoIP. Without it, exchange of audio is not possible. As AoIP networks scale, it is important to understand your PTP configuration and to ensure the network and AoIP devices are correctly configured for PTP.

ST2059-1

ST2059-1 describes the clock generation capabilities of an end device based on timing information contained within PTP messaging.

ST2059-2

ST2059-2 defines a standard set of PTP parameter sets to be supported on end devices. The ST2059-2 has been specified specifically for media applications.

AMWA NMOS IS04 & IS05

NMOS provides a centralised server approach allowing AoIP devices from different manufacturers to be managed in the same way, reducing the reliance on the individual and varied web-UIs that each device may serve.

IS-04 provides a discovery and advertisement mechanism for end devices to pass their stream configuration information to a central server.

IS-05 provides connection management, allowing users to interact with a single application that can present and allow management of connectivity in a familiar way that fits with broadcast workflows.

TR1001-1

TR1001-1 is a technical recommendation rather than a standard, produced by the JT-NM.

TR1001-1 offers a framework for an interoperable AoIP facility making recommendations on network configuration, device discovery and connection management. TR1001-1 references the standards above but goes beyond by recommending technologies to enable simplified operations to broadcasters.

Calrec

Calrec is committed to the standardisation of AoIP and Networking tools that enable its customers to meet their AoIP deployments regardless of scale or complexity.

All Calrec AoIP implementations are compliant to:

- AES67
- ST2110-30
- ST2022-7
- ST2059-1
- ST2059-2
- PTPv2
- NMOS IS-04
- NMOS IS-05

Summary

In this guide, all topics are covered to enable a user to understand and design AoIP networks that help broadcasters achieve their unique workflow goals.

Multicast

In an ST2110-30 environment, media streams must have mandatory support for multicast stream transmission.

Multicast is a method whereby packets are sent to multiple devices on a network at the same time. The packets themselves are transmitted once but are replicated as and when necessary by the network.

The following images show how multicast differs from unicast and broadcast transmission.

Unicast (fig 1): Data is sent from one host directly to another using the destination hosts IP address in its IP header.

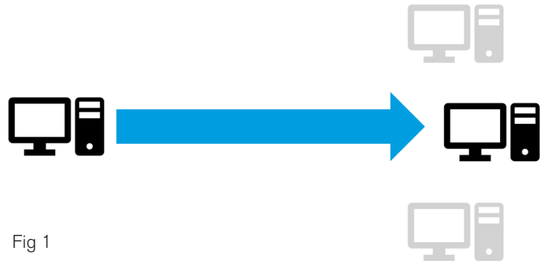

Fig 1

Broadcast (fig 2): Data is sent to all hosts on the local network or vLAN. The destination IP address in the IP header is the broadcast address of the subnet.

Multicast (fig 3): Data is sent to some hosts on the network but not all. Data is sent to a multicast IP address (G). Any hosts which wish to receive traffic that is sent to (G) register their interest with the network and the network dynamically replicates and forwards the traffic as required.

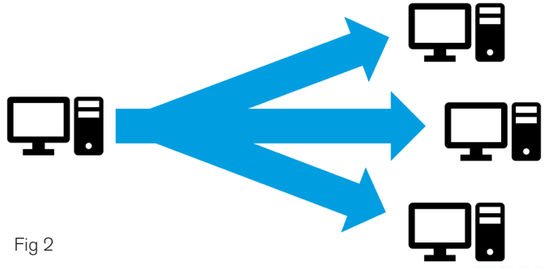

Fig 2

Addressing

Multicast packets are addressed from the reserved address space: 224.0.0.0/4 (224.0.0.0 – 239.255.255.255).

The Multicast IP address space can be further broken into the following three categories shows in fig 4. AES67 specifies the use of the administratively scoped (local) multicast address space (239.0.0.0/8)

The “well-known” multicast address range cannot be routed across layer 3 networks. This is of particular importance with devices that use mDNS (Bonjour) to advertise themselves and their resources. Devices that rely on mDNS should be placed within a single network or vLAN.

A multicast stream can be referred to as a group where any receiver of that steam can be referred to as members of that group.

Multicast MAC addresses always have the OUI: 01-00-5e.

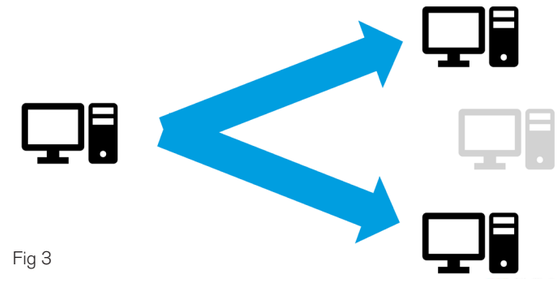

Fig 3

Start Address	End Address	Description
224.0.0.0	224.0.0.255	Reserved for special “well-known” multicast addresses
224.0.1.0	238.255.255.255	Globally scoped (internet-wide) multicast addresses
239.0.0.0	239.255.255.255	Administratively scoped (local) multicast addresses

Fig 4

The MAC Address ambiguity problem

The relationship between MAC addresses, NICs and IP addresses are slightly different with multicast compared to unicast. MAC addresses are burnt into a NIC so that every NIC has a unique MAC address. This guarantees that all devices are uniquely identifiable on a global scale. With multicast, rather than a NIC being associated with a MAC address, a multicast IP address is associated with a multicast MAC address.

When multicast was being developed, researchers were unable to secure enough MAC addresses to make a 1 to 1 mapping to multicast IP addresses. Therefore every multicast MAC address maps to 32 multicast IP addresses. For example, the following MAC address: 01-00-5e-01-01-01; corresponds to the following multicast IP addresses:

224.1.1.1	*224.129.1.1*
225.1.1.1	*225.129.1.1*
226.1.1.1	*226.129.1.1*
227.1.1.1	*227.129.1.1*
228.1.1.1	*228.129.1.1*
229.1.1.1	*229.129.1.1*
230.1.1.1	*230.129.1.1*
231.1.1.1	*231.129.1.1*
232.1.1.1	*232.129.1.1*
233.1.1.1	*233.129.1.1*
234.1.1.1	*234.129.1.1*
235.1.1.1	*235.129.1.1*
236.1.1.1	*236.129.1.1*
237.1.1.1	*237.129.1.1*
238.1.1.1	*238.129.1.1*
239.1.1.1	*239.129.1.1*

When a switch makes a forwarding decision based on the MAC address 01-00-5e-01-01-01, the switch will forward any of the available 32 multicast streams from the list to the destination. This is known as the MAC address ambiguity problem. This is a well understood limitation with multicast.

However, it is easy to avoid by administratively managing the available address space. When planning your multicast address space, ensure that the multicast IP addresses that are in use fit between two of the 32 IP addresses. This ensures that a switch will not unnecessarily forward multicast traffic from streams that contain the same MAC address but different IP content.

For a more technical view on multicast MAC and IP address mapping, see Appendix 1.

Internet Group Management Protocol (IGMP)

IGMP is a protocol that was designed to efficiently manage multicast on a local network; it is specifically referenced within AES67 as a method to initiate multicast communication on a network.

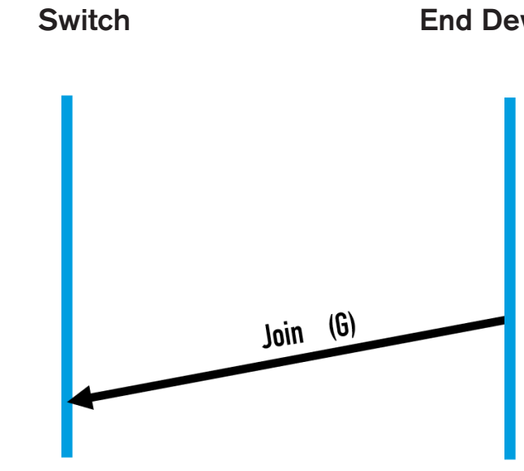

Fig 5

IGMP was initially developed in the early 1990s, long before AoIP at this scale was even conceptualized. IGMP was designed to request and deliver multicast content from and to the internet respectively.

An AoIP receiver signals to the network that it would like to receive a stream by sending an IGMP membership report for that stream. The informal name for an IGMP membership report is an IGMP join or simply just a join message. The group that wishes to be joined is referenced in the join message as (G) (fig 5).

When the switch receives the join message, it knows that any packets for the joined group (G) should be replicated (if necessary) and passed to the end device.

IGMP then maintains this connection by querying the AoIP device to ensure that the

device is still interested in receiving the stream. It does this by sending queries to the end device and the end device responds. Queries can be for all multicast groups or just a single group.

In this case, all groups are queried. The AoIP device responds with all groups that it is currently a member of (G1, G2, G3). (fig 6)

If the querier does not receive a response back from the AoIP device, it will time out the entry in its IGMP table and stop forwarding traffic to the device.

Usually, queries are sent every two minutes, but can be tuned to be shorter or longer depending on the switches' capabilities and the requirements for the network. If audio stops passing two minutes after connection, it may be that there is a problem with the querier.

Query messages are sent to all devices on the network. To avoid bursty protocol messaging, a query message contains a maximum response time. The maximum response time indicates to the AoIP device how long it can wait before it responds to its membership reports.

AoIP devices wait for a random time that is no longer than the maximum response time before responding. This ensures that all end devices do not respond with their membership reports at the same time causing bursty traffic patterns.

An AoIP device can signal to the switch that it no longer wishes to receive this stream by sending an IGMP leave message to the switch. This is also known as "IGMP Fast Leave". (fig7)

Querier

Historically, IGMP queriers were internet-facing routers. Their job was to maintain memberships on a local network and signal to the internet which groups are required on the local network.

Part of this mechanism also specifies that multicast should be flooded towards the querier. This is not desirable on AoIP networks.

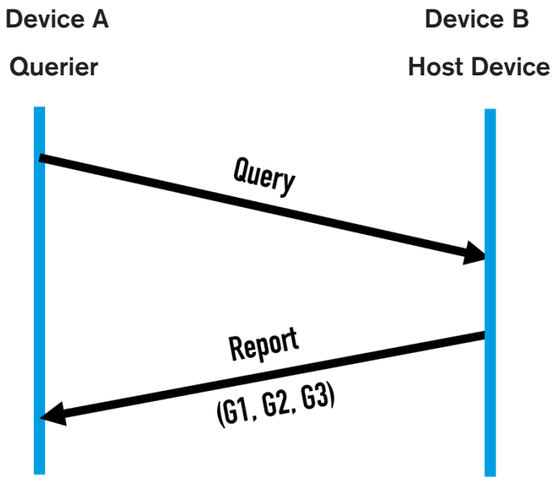

Fig 6

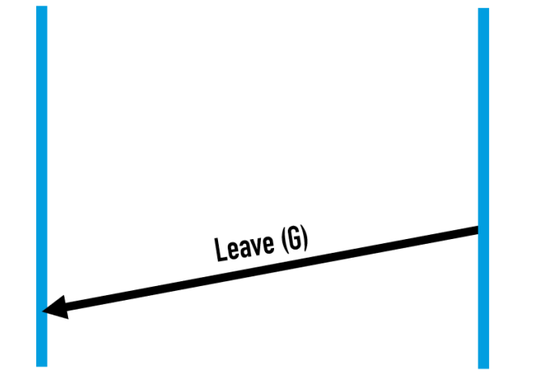

Fig 7

The port that is flooded with multicast on the switch is known as the Multicast Router port or the mrouter port. Switches will dynamically learn that a switch port is a mrouter port by the presence of query messages or multicast routing protocol messages on that port.

For example, if each of the three end devices was to produce 500Mbps of multicast, the combined total of 1.5Gbps would be flooded towards the querier. This link may become saturated if this there is not enough bandwidth available on the link.

Fig 8a

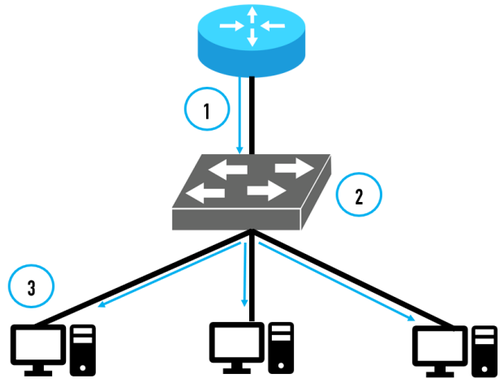

Fig 8b

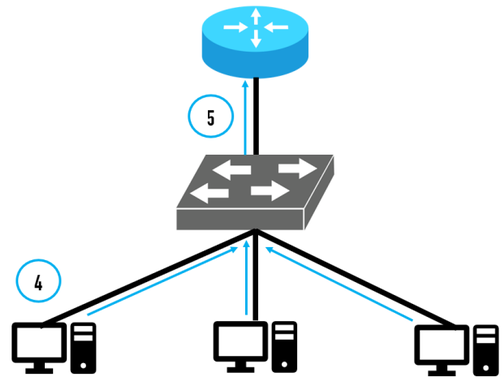

Fig 8c

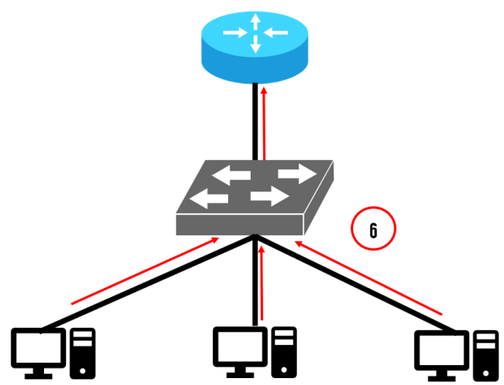

Fig 8 (a, b & c) shows how the switch floods multicast towards the querier.

Fig 8a illustrates how the router sends a query message onto the network (1). The switch detects a query message on the port connected to the router (2) and it marks this port as a 'mrouter' port. The switch forwards on query messages to the rest of network (3).

Fig 8b shows how the hosts respond with membership reports (4). Membership reports are then passed to the router as the switch had previously marked the port as an mrouter (5).

Any streams that are produced by the host devices, are also flooded to the mrouter (6) in Fig 8c.

Querier Election

Only one device on a network can be the querier at any one time. If there are multiple configured queriers on a network, they will elect a single device to be the active querier while the others remain idle. In the event of a querier failure, the idle queriers will begin the election process again. The election process is decided by the IP address of the querier. The lower the IP address, the higher the priority of becoming querier on the network.

It is recommended that all switches on a network are configured to be queriers and that the election process takes place. This is so that in the event of a failure of the active querier, another device will take over to maintain the multicast memberships across the network.

Membership Report Flooding

In AoIP, it is not desirable to flood multicast groups unnecessarily given the bandwidths involved. Some switch manufacturers have implemented a feature called "multicast report flooding". This feature only floods the IGMP Membership Reports but not the actual multicast. The multicast content will only be forwarded if it is specifically requested by an upstream device. This avoids unnecessary bandwidth consumption on inter-switch links.

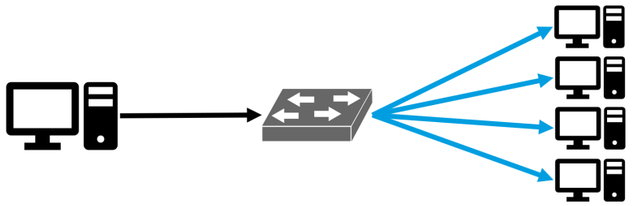

Fig 9

IGMP Snooping

A typical unmanaged or unconfigured switch will broadcast multicast packets out of all ports that they were not received on. This very much contradicts the very purpose of IGMP; it does not allow for efficient transport of multicast across a network. (fig 9)

IGMP snooping is a method of constraining multicast flooding. A switch that is enabled with IGMP snooping will listen (or snoop) on the conversations between the querier and the end device. When an IGMP snooping switch hears an IGMP Membership report, it does the following:

1. Adds the port number and Multicast MAC address to its MAC table
2. Adds the port number and the Multicast IP to its IGMP table

This allows the switch to create a map that ties together multicast groups with the switchports they are required to be delivered to.

It is recommended that IGMP snooping is turned on for all switches in a network to ensure multicast flooding does not take place.

Unknown Multicast Flooding

Another feature that some switches have is the ability to disable unknown multicast flooding. That is, if a multicast packet arrives at the switch and the switch has no members for that group, the switch will drop the packet rather than flood it. It is recommended that unknown multicast flooding is disabled where possible.

AES67 – IGMP senders' clause

AES67 specifies that senders (as well as receivers) should send an IGMP membership report for the group that it is being sent to. This is to ensure that a switch always has a member of the group and thus avoid multicast flooding. This is a rudimentary way of avoiding this problem, but it is not always desirable.

For example, some networks may use SDN or orchestration tools that make decisions based off IGMP membership. IGMP states that only receivers should join the group so this AES67 clause may have unintended consequences on some networks.

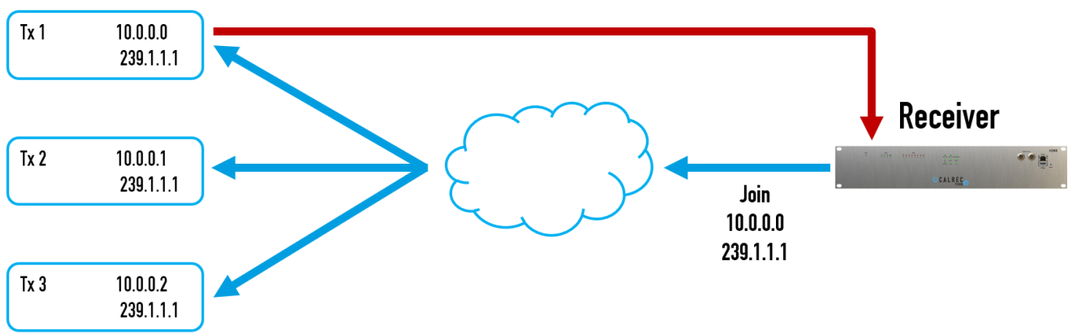

Fig 10

IGMPv2 vs IGMPv3

There are two major revisions of IGMP in circulation today: IGMPv2 and IGMPv3. IGMPv3 is backwards compatible with IGMPv2.

Devices that run IGMPv3 are required to support and communicate with IGMPv2 devices. IGMPv3 adds extra features to IGMPv2. If a device is configured to operate at IGMPv3 but is required to "downgrade" to IGMPv2, there may be a small timing delay while this process occurs.

A comparison table for IGMPv2 and IGMPv3 can be found in Appendix 6.

IGMPv3 – Source Specific Multicast (SSM)

IGMPv3 supports a feature that allows a receiver to signal to the network which device it would like to receive a stream from. This is known as source-specific multicast (SSM). Common notation when sending an SSM IGMPv3 membership report is (S, G) where S = the source IP address and G = the group IP address.

For example, a receiver could send a Membership report for 239.1.1.1 but only if that group is produced by 10.0.0.0. (fig 10)

SSM is considered a security feature to ensure that the multicast received is only received from a trusted sender and that the accidental creation of a stream on a pre-existing group does not affect the reception of existing receivers.

IGMPv3 works on the basis that a group can be received from a set of IP addresses (INCLUDE mode) or not from a set of IP addresses (EXCLUDE mode). If a receiver does not wish to use SSM but does wish to use IGMPv3, it would send the membership report with the instruction EXLUDE: NONE.

Therefore the stream will be received from all sources, excluding none.

Multicast Routing

The discussion of multicast so far has been with reference to layer 2 networks. In larger AoIP deployments it is becoming increasingly common to use layer 3 multicast routing to manage multicast packets across a network. There are numerous benefits to this:

- Further efficiencies compared to IGMP
- Great visibility
- Finer control over packet flow
- Load balancing and resilience options
- More accessible for SDN control

The most common protocol for routing multicast is Protocol Independent Multicast (PIM). PIM uses the unicast routing table from a switch to engineer a reverse path forward (RPF) from source to destination. A Shortest Path Tree (SPT) is created from source to destination and packets are routed from source to destination. PIM has several variations including:

- PIM-SM (Sparse Mode)
- PIM-DM (Dense Mode)
- PIM-SSM (Source Specific Multicast)

Some PIM protocols rely on the concept of a Rendezvous Point (RP). The Rendezvous point is a single point on the network where senders register their multicast groups. When a receiver requests a group, the request is forwarded towards the RP at which point the multicast group can be transmitted towards the receiver.

This is an overly simplistic view of the PIM protocol. Mechanisms are in place that ensure that bandwidth is not unnecessarily consumed on a network (for example if a sender has no active receivers, the multicast is not unnecessarily sent to the RP). The specifics on how PIM forwards traffic between networks varies between the PIM variants.

Practically, the RP is usually implemented using a virtual loopback adapter on the network, so its availability is not dictated by any physical link.

Precision Time Protocol (PTP – IEEE1588)

In a broadcast environment, it is essential that all digital audio devices remain in sync.

This was previously achieved in the SDI/AES world by distributing baseband synchronization signals throughout a facility on independent infrastructure. With the move to IP comes a different approach to synchronization.

In an IP environment, a protocol called PTP is used to sync devices across an IP network, removing the need for separate synchronization infrastructure.

AES67 and SMPTE ST2110 mandate the use of the 2008 revision of IEEE1588, more commonly referred to as PTPv2. Devices that require PTPv1 can co-exist on a network that runs PTPv2 devices. In this chapter whenever the term PTP is used, it refers to PTPv2.

A PTP grand master clock (GMC) is used to originate timing data that all devices on the network derive their time from. The GMC may be a dedicated device, a network device or an audio device on the network. (fig 11)

PTP devices belong to a single domain. All devices that can communicate with each other and negotiate master/slave relationships must be in the same domain.

Two devices on different domains can co-exist on a network but they will not exchange timing information. The domain of a device is usually configurable through the control UI for that device. Calrec Connect is used to manage TP settings of Calrec's AoIP devices.

PTP uses multicast to transmit packets between clocks on a network. However, unicast operation is becoming more popular for security and efficiency reasons.

PTP Accuracy

SMPTE ST2059-2 states that it is acceptable for slave devices to achieve an offset from the master of 1us or less. In practice, PTP can achieve much tighter accuracy than this in a finely tuned environment.

The offset from master is a reading that is usually presented to the user from within the control UI of a device. The offset from master is what the time is on the slave device, compared to the time on the master device. An offset from master of 1us, would mean the time on the slave device is +/- 1us from the master.

To understand why 1us is acceptable, we can dissect the sample time, in this example at 48kHz. 1 sample in time is:

$$\frac{1\text{second}}{48000} \approx 0.00002\text{seconds/20us}$$

It is generally acceptable to achieve synchronization to +/- 5% of the sample period. 5% of 20us is 1us.

Practically, it is possible to achieve an offset from master of +/- 100ns on a tightly controlled PTP network. With increasing sample rates, higher accuracy is required for AoIP devices.

PTP Domains

PTP devices are said to be included within a single domain as specified by the domain number in the PTP devices configuration. Only devices that are in the same domain will exchange timing information. PTP domains can be between 0 and 255. The BMCA (as discussed later) will not take part between two devices where the domain is different.

Network Clock Types

When working with PTP on a network, there are three options on how to treat PTP packets as they pass through network devices. In this section we will go through the pros and cons of each. It is worth noting, that all options can be mixed and matched on a single network.

Grandmaster Clocks

The grandmaster clock (GMC) is the source of time for the entire IP network. There may be many grandmaster capable clocks located on a network but only one may be actively providing the network with timing information at any given time. The grandmaster is often a dedicated device on the network but may also be a network or audio device. Dedicated grandmaster devices often have the benefit of an increased feature set such as the ability to generate baseband sync or get its own time from GPS.

Fig 11

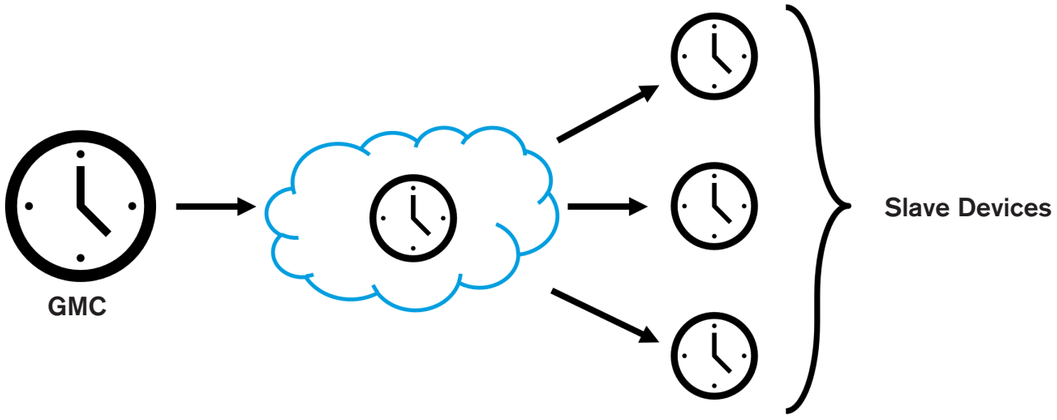

Ordinary Clocks

An ordinary clock only has a single NIC exposed to the PTP network. This NIC will provide timing to the network when it is in the Master state or receive timing from the network when in a Slave state.

Transparent Clocks

Transparent clocks run on switches and offer a substantial step upwards in performance compared to non-PTP aware switches. Transparent clocks inspect the timing information inside of a PTP packet when the packet is received by the switch. It then does the same thing before the packet is transmitted onwards. This way, the switch can calculate the residency time (the amount of time the packet spends inside of a switch):

egress time-ingress time = residency time

This residency time is added as a correction time into the PTP packets. When the AoIP device receives the PTP messages, it also receives all correction information from the network. It can then compensate its calculations to derive its own time.

This means that PTP packets are less affected by a noisy environment as there is a degree of natural compensation for latency and jitter through the network.

The GMC is still responsible for communicating with every device on the network and all devices continue to receive each others' messages. This may limit the stability of the PTP network as the network is scaled upwards. (fig 12)

Pros

- Less reliant on network performance and QoS
- Can generally achieve better accuracy than non PTP aware systems

Fig 12

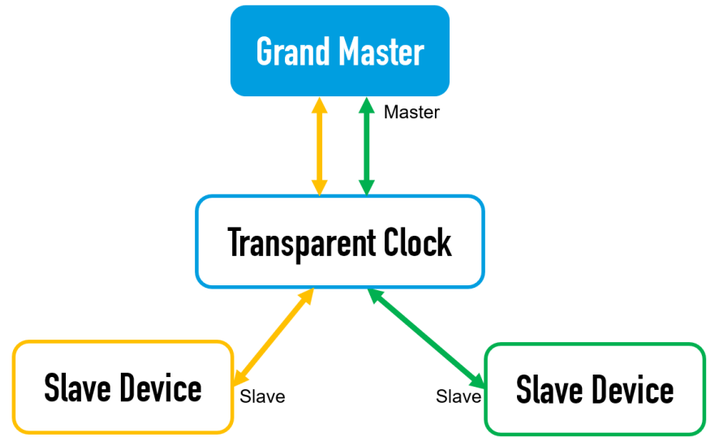

Cons

- Network infrastructure is generally more expensive
- Limited Scalability

Boundary Clocks

Boundary clocks offer further performance and scalability improvements over transparent clocks. Switches that can run as boundary clocks will slave to the GMC and then in turn, be a Master to downstream devices.

This forms a hierarchal Master/Slave relationship throughout the network. PTP messages that are generated by the GMC are not forwarded to Slaves.

Instead, the switch generates new PTP messages and forms its own downstream Master/Slave relationship. This means that a boundary clock will only ever have a single Slave port.

One of the main benefits to this methodology is that Slave devices no longer receive PTP messages from other slave devices and the GMC does not have to directly communicate with all slave devices. This enables IP networks to be much more scalable and reliable compared to non-PTP aware networks and transparent clock networks.

Usually, a switch will intercept PTP messages outside of the QoS or ingress queuing mechanics that a switch may have in place. This means that PTP processing is expedited in the switch's hardware and software which in turn enables highly accurate synchronization across the network.

Boundary clocks can have a lot of parameters than need to be configured for correct operation. Often, each individual interface on a switch will have its own settings for packet generation including the intervals in which packets should be sent. It is essential that each link has the same packet interval settings at either side, but they can differ from link to link. This allows some devices to operate with a stricter packet timing than others depending on their capabilities.

The amount of slave devices that a boundary clock can support should be discussed with the switch vendor (fig 13).

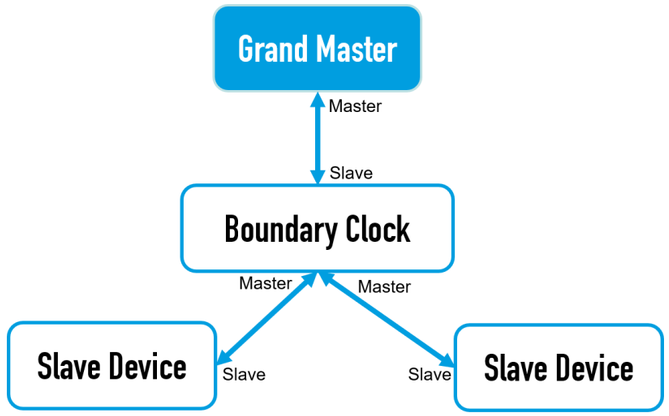

Fig 13

Pros
- Highly scalable
- Switches can become GMC
- Less reliance on QoS
- Higher accuracy than non PTP aware systems
- Added flexibility by adjusting packet intervals on a device-by-device basis

Cons
- Network infrastructure is generally more expensive
- High operational overhead

Non PTP Aware Devices

PTP devices can be connected with off the shelf networking devices. PTP uses multicast to distribute the packets between devices and therefore any network devices that have basic multicast functionality can pass PTP timing data between devices. The switch is not aware that the packets it is transporting are PTP packets and do not react to any of the timing data found within the PTP packets.

It is essential that non-PTP aware networks have adequate QoS strategies in place to give PTP packets priority over other traffic such as media and control. QoS strategies are discussed in the following chapter.

The quality of the timing data that is derived from PTP packets is entirely dependent on the network's ability to provide deterministic and acceptable latency and jitter performance to PTP packets.

If there are inefficiencies in the QoS strategy or there is a sudden burst of traffic, loss of sync accuracy may be unavoidable. The natural delay and jitter of a non-PTP aware switch is enough to see a large drop in performance compared to transparent and boundary clocks. Bad synchronization is often analogous to bad audio.

Because of the way PTP messaging works, the GMC is required to receive PTP messages from all Slave devices. This can put excessive load on a GMC device.

How many Slave devices a GMC device can support should be discussed with the vendor of the GMC. There is no protection against this when there is no PTP awareness on a network.

As PTP messaging is done via multicast, all devices will receive each others' PTP messages. This adds operational and bandwidth overhead to the network.

Pros
- Network infrastructure usually cheaper
- Less operational overhead

Cons
- Limited scalability
- Reliant on an acceptable and deterministic network environment
- Potentially vulnerable if QoS is not robust

Clock Type Summary

When deciding on a network infrastructure, special attention should be taken to PTP performance requirements. The scale or capacity requirements for the network may be a decider in which PTP technology to adopt. PTP awareness on the network improves performance and scalability but it comes at budgetary cost and becomes potentially more operationally intensive to implement.

To highlight the performance difference between a device that derives its time from a non PTP aware switch with a GMC attached to the same switch and a device that is directly connected to a boundary clock, see the following graph (fig 14).

PTP Port States

Each NIC or switch port on a PTP network is represented by a port state.

The port state indicates what role the NIC or switch port is performing on the PTP network. There are three main ports states that are common on PTP networks:

- Master
- Slave
- Passive

Master

Devices in the Master state provide timing information to other devices on the network. Some boundary clocks have a feature known as "Master-Only".

This ensures that the port is unable to Slave from a downstream device. This is a safety feature and should be applied to end device ports where possible. This will stop any misconfigured or accidently configured AoIP devices from causing a BMCA event on the network.

Slave

Slave devices receive timing information from Master devices and calculate their clocks based on this timing information.

Passive

Devices in Passive do not provide or receive timing information from the PTP network. It is common in ST2022-7 networks where the Amber NIC will Slave to a Master, the Blue NIC will go into the PASSIVE state. This is because devices cannot slave from two separate NICs.

Best Master Clock Algorithm (BMCA)

The BMCA is the part of the PTP protocol that decides which NICs and switchports should operate as a specific PTP port state. Every clock on a network advertises characteristics about itself.

By default, every clock assumes it is the Master until it hears about a better Master, at which point it will then stop advertising itself as the Master and instead advertise the clock that it has heard as a better Master.

A better Master is decided by these clock characteristics in this order:

1. Priority1
2. Class
3. Accuracy
4. Variance
5. Priority2
6. UUID

The class, accuracy, variance and UUID characteristics are usually burned into a device and are set based upon the operating mode of the clock.

The Priority1 and Priority2 values are user-definable and allow network designers to influence which clock becomes Master on the network.

The lower the value of the priority1 values, the more likely it is to become the Master. If there is a tie breaker, the GMC uses the next characteristic in the list to decide which clock is the better Master. Priority values can be between 0 and 255.

PTP Announce Messages

The characteristics of the clock are transported inside of announce messages. The rate at which announce messages are sent is called the announce interval. AoIP devices also allow the configuration of an announce timeout value.

This is the amount of announce messages that a Slave can miss from a Master before it believes that the Master is no longer operable. If the announce message timeout is reached, the Slave device will begin the BMCA process again in order to find the next suitable Master on the network.

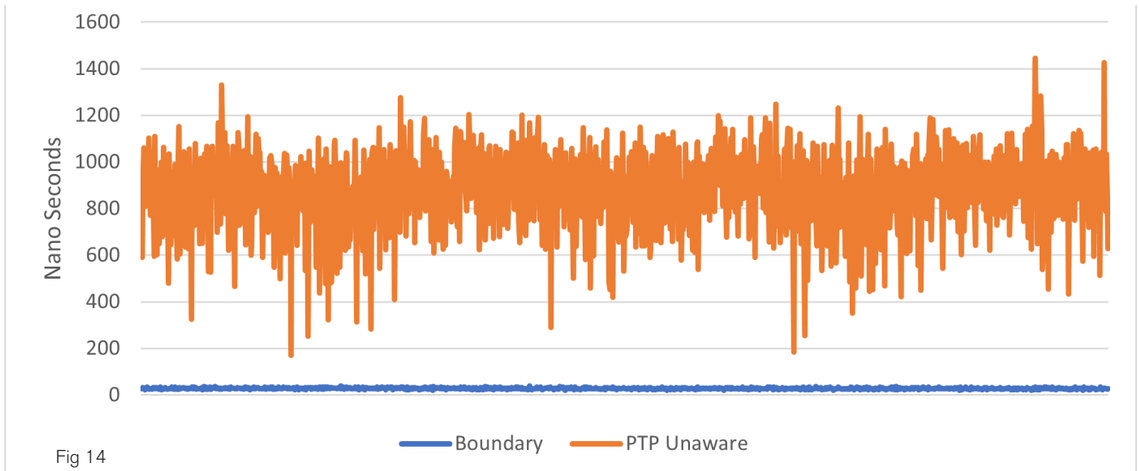

Fig 14

PTP Messaging

The way timing data is passed between Master and Slave devices is through a series of messages that calculate the offset from Master and the delay on the link. PTP has two operating modes:

- One Step
- Two Step

All devices on a network should be configured to either be one step or two step. The difference is the number of messages required to be sent from the Master to the Slave. One step requires one, two step requires two. The four messages required for two step are as follows:

1. Sync
2. Follow Up
3. Delay Request
4. Delay Response

For two step mode, the initial messaging from Master to Slave has timing information generated by the sync message which is transported to the Slave within the follow up message. In one step mode, there is no follow up message. Timing data is generated and transported within the sync message.

Fig 15 shows the flow of messages between the Master and Slave devices.

The slave device uses the timestamps found in these messages to derive its own time. The subsequent offset from master and mean path delay values are often presented to the user in the end devices UI.

These are continuously updated as new timing data is received from the master.

The interval between the sync and delay request messages are usually definable by the user or network designer.

PTP Message Intervals

PTP messages are sent at predefined intervals. The unit to measure the intervals can be variable between vendors. The descriptions of the various configurable message intervals and their units are detailed below:

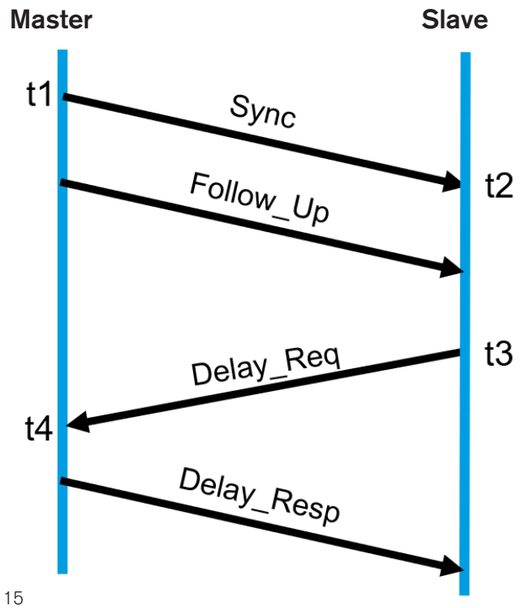

Fig 15

Announce Message Interval

- The amount of time between consecutive Announce messages.

Announce Receipt Timeout

- The amount of Announce messages that can be missed before the Master is considered not available.
- This is measured in multiples of the Announce message interval. For example, if the Announce Message interval is 1 second, and the Announce Message Timeout is 3, then the Timeout period would be 3 seconds.

Sync Message Interval

- The amount of time between consecutive sync messages

Minimum Delay Request Interval

- For a Slave device, this is the initial interval to send delay request messages until it receives a delay response from the Master
- For a Master device, this is the interval to add to its delay response messages to tell the Slave devices which rate to send. This allows the Master to control the amount of delay request messages that is sent to it.

Interval Units

Outside of the Announce Receipt Timeout, the other intervals can all be entered into an AoIP device UI in several different units. Calrec Connect uses $\log_2$ notation. Other notations are seconds and packets per second (pps). See Appendix 2 for a conversion table.

AES67 Media Default	SMPTE ST2059-2 Default
Domain = 0 Announce Interval = 2s Announce Timeout = 3 Sync Interval = 125ms Del_Req Interval = 125ms	Domain = 127 Announce Interval = 250ms Announce Timeout = 3 Sync Interval = 125ms Del_Req Interval = 125ms

Fig 16

PTP Profiles

To standardise the many different PTP parameters that are available, AES67 and SMPTE 2059-2 provide recommended default values and specific ranges that should be allowed.

Fig 16 shows the recommended default values for the AES and ST2059-2 PTP profiles.

The benefits of these profiles are the relatively short intervals between packets, specifically the announces message interval of 250ms in the ST2059-2 profile.

Devices using the ST2059-2 profile will react quicker to BMCA events on their network as they process announce messages 8 times quicker than devices using the AES67 profile.

Intervals must be set consistently between all MASTER/SLAVE pairs.

Hydra2 Considerations

When integrating AoIP and Hydra2 into a hybrid environment, timing information must be shared between the two domains.

AoIP products use PTP for timing information whereas the Hydra2 protocol relies on baseband timing information.

For this reason, the source of PTP and the source of baseband time must be tied together. This can be achieved by using a GMC to produce a baseband signal for the Hydra cores or by clocking both the hydra cores and the GMC from a baseband sync generator (as illustrated in fig 17).

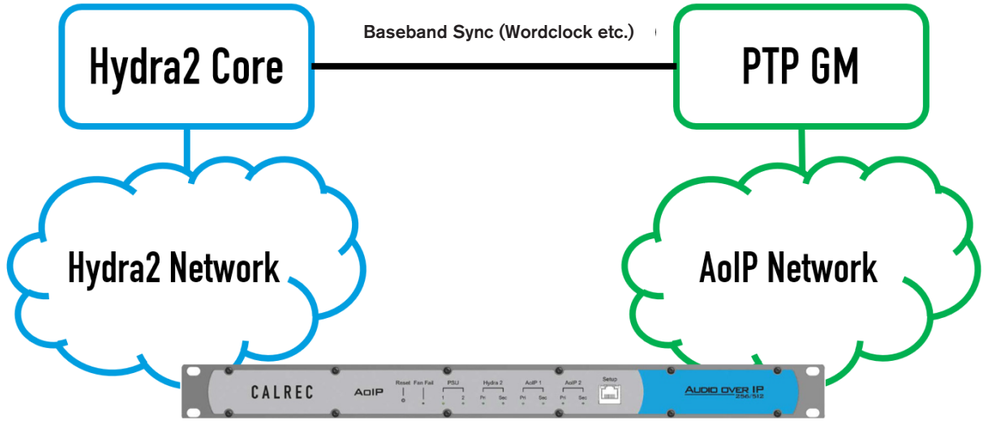

Fig 17

Bandwidth

Bandwidth is a term that has some crossover between audio engineering and networking engineering technologies.

In network engineering, you would be referring to the amount of data that can be produced or consumed by a device on an IP network. As this guide is networking focus, we refer to bandwidth in the networking context.

Audio bandwidth requirements are generally a lot less than with video. It is common for audio-only devices to have a single or a set of 1Gbps NICs. In larger capacity devices such as Calrec's ImPulse core, 10Gbps options are also available. In contrast, video equipment will have a minimum requirement of 10Gbps, and it is not uncommon to see devices with 25/40/100Gbps NICs.

Protocol overhead

Bandwidth is not just generated by audio data. It is also generated by control data and protocol overheads. Control data is generally low bandwidth and not prioritised by QoS strategies. Control data can even be abstracted to a separate management network to ensure there is no unnecessary bandwidth consumption on the media network. Protocol overhead however is something that needs to be accounted for. Whenever we create an audio packet it is split into several different parts:

1. Payload (the actual audio in the packet, can vary in size)
2. RTP Header (12 Bytes – may vary with presence of CSRCs or extended headers)
3. UDP Header (8 Bytes)
4. IP Header (20 Bytes)
5. Ethernet Header and Tailer (26 Bytes – including preamble)

For every packet of audio that is produced, an extra 66 Bytes of data is produced to aid in the transportation of the audio across the network. The more channels of audio that are required, the more streams are required, the more packets are required, the more protocol overhead becomes something to consider and account for.

What affects Bandwidth?

Bandwidth planning plays a large part in successful network design. Ensuring there are no bottlenecks in your network requires a good understanding of how bandwidth is generated and what can affect bandwidth consumption. The following sections explain the various components that are used to produce AES67 and ST2110 streams and how they affect bandwidth usage.

Packet Time

The packet time is the amount of audio in time that is placed in a packet. Common values are 125us and 1ms. When running at 48kHz, a 125us packet would contain 6 samples per audio channel where as a 1ms packet would contain 48 samples per audio channel.

The larger the packet time, the lower the bandwidth that is consumed due to less packets needing to be transmitted. Assuming a 24bit codec and a single audio channel, we can multiply the number of samples by the bit depth to find the payload size:

125us:

6 samples×24 =144 bits

1ms:

48 samples×24 =1152 bits

A single 125us contains an audio payload that is 144 bits and a 1ms audio payload contains 1152 bits. This does not mean that the 1ms stream consumes more bandwidth than the 125us stream. To pass one second of audio, the 125us stream needs to send 8000 packets whereas the 1ms only needs to send 1000.

If each audio packet has 528 bits of protocol header, the bits transported in a single second is as follows:

125us:

(528+144)×8000 = 5376000bps /5.3Mbps

1ms:

(528+1152)×1000 = 680000bps/1.68Mbps

Over the course of a second, the amount of data produced by a 1ms stream is less than 125us microseconds. This is because fewer packets are being generated and therefore the effects of protocol overhead are less pronounced. Although the 1ms stream operates with a lower bandwidth, this is at the cost of a larger latency in its transmission path. This will be discussed in the latency chapter.

Channel Count

The number of channels of audio to be transported over a given stream will change the audio payload size within a packet. For example, the payload of the 1ms stream in the previous example was 1152 bits. This was for a stream which was only transporting a single audio channel. If we increase the channel count to 8, the new audio payload size would be 9216 bits. This ultimately affects the number of bits per second (bps) that are transported across the network.

Codec

The codec used affects the amount of data contained within an audio payload. For example, in the packet time example, the 125us stream contained 6 samples using a 24bit codec. This resulted in a payload size of 144 bits. If a 16bit codec was used, this would reduce the payload to 96 bits. This would reduce the bps produced by this stream.

Sample Rate

The sample rate affects how many samples are played out during a second. At 48kHz, 48,000 samples a second are played out. At 96kHz, 96,000 samples are played out. Using a 1ms packet time, a 48kHz stream will put 48 samples of data into an audio packet. A 96kHz stream will double the number of samples (96) contained within a packet. This often means that to comply with the MTU (discussed next), the channel count must be halved.

Maximum Transmission Unit (MTU)

The MTU is the maximum amount of data that can be encapsulated by any given transport protocol. In AoIP, it is common to reference the MTU with regards to ethernet. The maximum amount of data that ethernet can encapsulate is 1500 bytes. This means 1460 bytes of audio can be placed inside of an ethernet frame (subtracting the IP, UDP and RTP protocol headers from the MTU). The usable data allowance for audio is called the payload.

The MTU affects the maximum amount of audio of any given format into an ethernet frame. The format of the audio (Packet Time, Codec, Sample Rate, Channel count) must produce a payload that fits within the MTU of an ethernet frame. Some networks and network devices may be able to support Jumbo frames which allows an ethernet payload of up to 9216 bytes.

Over-subscription

Over-subscription refers to when the bandwidth produced by a device or by several devices exceeds the available bandwidth of a single link. This is not likely to occur if you have several devices plugged into a single switch for two reasons:

1. Most AoIP devices are designed so that they cannot produce or receive a number of streams that exceed its bandwidth own capabilities
2. Most COTS network switches are non-blocking which means they are designed with enough throughput to forward enough packets to account for all its ports and their respective bandwidths.

Over-subscription is much more likely to occur when connecting two switches together. For example, fig 18 shows three devices all producing 600Mbps of data and this is being received by devices on the other side of a inter switch link.

The three devices producing 600Mbps each are aggregated and sent from SW1 to SW2. The total aggregated bandwidth is the sum of all bandwidths that to be transported across that link. In this case 18Gbps.

Special attention is required to ensure that any links within a network cannot be compromised by being oversubscribed.

The method of doing this may mean installing high-capacity links to account for peak traffic levels or deploying an SDN solution that reserves bandwidth on a link.

In fig 18 there would be severe packet loss which would cause distorted audio at the receivers. It must be ensured that the total upstream bandwidth is greater or equal to the downstream bandwidth.

Not only does the format of audio matter when calculating a networks bandwidth requirement, it is also necessary to calculate the number of streams required and how they are distributed throughout the network.

Planning the distribution of bandwidth is essential for designing scalable and reliable IP networks. We discuss specific network design techniques in the network design chapter.

Bandwidth Examples

For convenience, some example bandwidths for various audio formats can be found in Appendix 3. These numbers consider inter-packet gaps which the previous examples don't which makes these examples more accurate references.

Recommendations

In general, it is advised to not exceed 60% of the available bandwidth on the link. Staying below 60% offers the following benefits:

- 40% free bandwidth for control/ user traffic which may be bursty or unpredictable
- Room to add ad-hoc audio streams if necessary
- Reduces the load on network infrastructure and thus ensuring deterministic performance

Despite this recommendation, it is perfectly acceptable to exceed this value if the network performance is known and trusted. For example, in an SDN environment, bandwidth can be guaranteed on a stream-by-stream basis. This allows confidence in using a much higher percentage of a link's bandwidth.

Fig 18

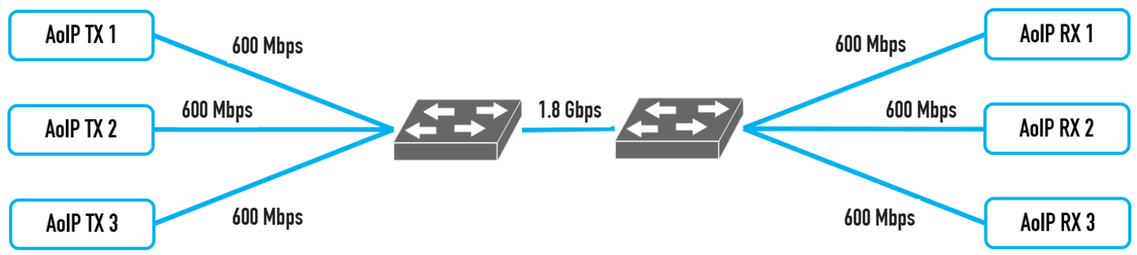

Latency/Jitter

Latency and jitter are normal properties of any audio network, regardless of transportation format.

Analogue audio, for example, has a natural latency. Baseband signals do not travel at an infinite speed, they are restricted by the speed of light.

Jitter is the rate of change in the amount of latency within a transmission path. For example, packets across a network may not be delivered at the exact same interval as they were produced. Packet 14 may be delivered before packet 9 and 10. Jitter describes the variance in arrival time.

Latency

Latency on the most part is deterministic. There are several parts of an AoIP network where delay may be introduced:

- Packetisation
- Switch residency time (the time a packet spends in switch queues)
- In flight delay (time it takes to get from one end of a cable to another)

Packetisation

Latency is introduced during the encapsulation process that takes place on a sender. A packet cannot be transmitted until its payload has been filled. The amount of time it takes for a payload to fill is equal to the packet time. Therefore, the larger the packet time, the larger the latency across the network. If live audio monitoring is required, it is recommended to use a lower packet time such as 125us.

Switch residency time

When a packet is received by a switch, it is queued before it is transmitted onwards. The length of time that a packet spends in a queue may be affected by:

1. The amount of traffic that is currently queued
2. The QoS polices on the network

A good rule of thumb is that at worst case 100us of delay is added for every switch hop. This value will vary significantly depending on the switch and its configuration.

In-flight Delay

The amount of time a packet spends on the cable can be ignored for the most part. In a local network, the delay is negligible. The speed at which a signal passes through a cable is the speed of light. Delays may be introduced on much longer cable runs, specifically for WAN links. However, these are again negligible compared to the other delays in the transmission path.

Jitter

Jitter is introduced into a network when the amount of delay that packets experience varies over time. In a perfect world, jitter would equal 0, meaning that the network latency is understood and guaranteed.

Jitter on an IP network is also known as “Packet delay variation” or PDV. There are multiple reasons why jitter might occur in a network:

1. The amount of data the network processes varies
2. Bursty control traffic may put the network under a temporary strain
3. QoS is not configured or is incorrectly configured

To absorb jitter, AoIP devices contain buffers so that receivers are able to store packets for a defined amount of time before playing the audio out. This means that there is room for error in case a packet arrives early or late.

Link Offset

The link offset is the size of the buffer allocated to a receiver, commonly configured in time.

The size of the link offset is definable by the user and can be tailored to the performance of the network. If for example a receiver was created with a link offset of 20ms, the receiver would receive packets for 20ms before it played out the audio.

This means that any one packet can have an instantaneous latency of 20ms, and the link offset buffer will “absorb” the delay. If WAN links or networks with no deterministic performance are in the transmission path, larger link offsets are required.

Theoretically, the minimum link offset will be the packet time + the measured network delay. This will allow audio to pass with the minimum amount of buffering although it will not be able to tolerate any amount of jitter. If a single packet it delayed, there is a high chance it will not arrive in time to be played out.

Practically, the link offset should be set to two times the packet time. This allows for the packet to be filled and for transmission over the network, while accounting for any unexpected jitter.

A good idea is to start with a high link offset and reduce until audio problems begin. Then increase until audio problems disappear. This ensures you have adequately allowed for any jitter on the network.

QoS

QoS is an essential component of AoIP networks.

QoS stands for Quality of Service and is an umbrella term for several tools, technologies and protocols that give network designers finer control over the priority and network performance for different types of traffic.

Using QoS tools, a network designer can give priority to some traffic over others. In this chapter, we will discuss what these technologies are and how they should be applied to AoIP networks.

It is worth noting what QoS is not. QoS cannot guarantee protection against all network conditions. If a network experiences unprecedented or unexpected demand, QoS can only prioritise traffic, it cannot guarantee that lower priority traffic will be passed. The event where lower priority traffic is dropped due to higher priority traffic using all available bandwidths is known as queue starvation.

QoS is also not a magic wand. If some traffic is grouped together as priority traffic, this intrinsically lowers the priority of other traffic.

QoS Components

There are several different sub technologies that make up an overall QoS feature set within a networking device. Each of these sub technologies affect how individual packets are prioritised or dropped through a device. The following points give a brief overview of these technologies.

Admission Control

Admission control is as much of a security feature as it is a QoS feature. Admission control simply permits or denies traffic the ability to enter a networking device. This can be achieved by using Access Control Lists (ACL), firewalls or port security using MAC address learning. By blocking unrequired traffic at ingest, bandwidth can be preserved throughout the network.

Classification

Classification is the process of identifying the priority of traffic. This is commonly achieved by inspecting the DSCP field on the IP header. DSCP is discussed in the next section.

Marking

Marking is when a switch actively changes the QoS of a packet in order to force classification on it. This technique is also known as remarking or remapping which is discussed in the following sections.

Traffic Policing

Traffic policing involves analysing the data rate of a particular port or traffic class and making a decision on whether to drop that traffic based on it exceeding a particular threshold. You can employ traffic policing to ensure that low priority traffic does not exceed a specific data rate in order to protect higher priority traffic.

Traffic Shaping

Traffic shaping is a less aggressive form of traffic policing. Rather than allowing the data rate to exceed a threshold and then begin dropping traffic, traffic shaping schedules packets for transmission in such a way that lowers the data rate. Traffic shaping delays some of the packets to keep the data rate below the drop threshold to prevent bursty traffic exceeding that given threshold.

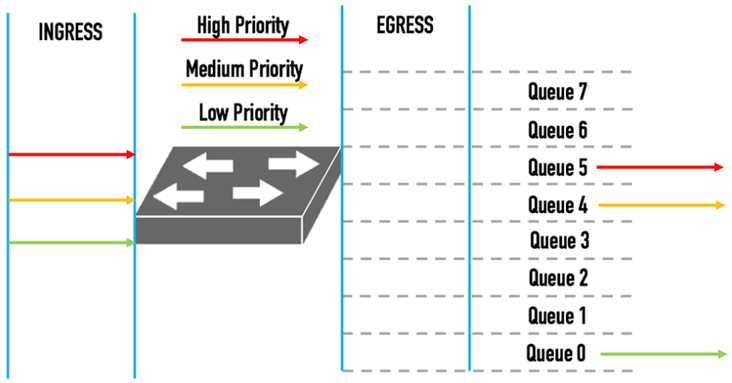

Fig 19

Egress Queuing

Before traffic is transmitted outbound from a network device, it is queued in a buffer. There are often several different buffers which are given different priorities for transmission. Packets are placed in specific buffers depending on their classification (fig 19).

There are multiple packet selection algorithms that can be found within a network device's feature set.

Differentiated Services Code Point (DSCP)

DSCP is a 6-bit value that is found within the Differentiated services (Diff Serv) field of an IP packets header. This 6-bit value can be set to tell the network about its priority. When a network receives an IP packet, it inspects the Diff Serv field and can then process the packet based on its priority.

AES67

AES67 states that PTP, Media and Management traffic should contain the following DSCP values in their diff serv headers:

- PTP: 46: Expedited Forwarding
- Media: 34: Assured Forwarding 41
- Discovery/Management: 0: DF0 Best Effort

Using these values will allow a network to make the best forwarding decision for the packet type. Some AoIP devices will allow you to move away from these values to accommodate a QoS policy that has not been designed around AES67.

Although AES67 specifies the DSCP values that should be used, it does not go as far as mandating a specific QoS policy for network switches to adopt; this would be too restrictive given the enormous number of options available to network designers.

A QoS policy is defined by a network engineer and it is a set of instructions for the switch to follow when it receives a packet of a certain priority. QoS strategies will be discussed later in this chapter.

Trust Boundaries /DSCP Remapping

Trust boundaries are defined by the network designer and specify whether a switch should trust the DSCP value on traffic as it ingresses into the switch. It is plausible that a device on the network uses DSCP: 46 for media traffic which goes against the AES67 standard. This would mean that media traffic from this device would be processed with the same priority as PTP traffic. This may have a performance impact on the PTP protocol across the network.

If this scenario occurs there are two options:

- Do not trust DSCP markings and assign this traffic to the default queue (usually the lowest priority)
- Re-mark the DSCP field and pass the packets to the appropriate queue

By not trusting the port that packets are coming in on, all packets will be assigned to the default queue. The default queue is usually queue 0 with the lowest priority. If it is a single stream with an incorrect DSCP marking, this option may not be wise as all traffic on this link will now be treated the same. If there are other streams or PTP packets on this ingress port, then forwarding performance could be affected.

Re-marking the DSCP involves changing the DSCP value in the packet before forwarding.

In this instance, the switch could be programmed to change 46 to 34 to prioritise all packets with the same priority as the rest of the media. The difficulty with this strategy is that this would also remark PTP packets (which also use DSCP46) so PTP on that link would also be treated the same as media packets.

Switches have different capabilities when it comes to QoS. It is not uncommon for a switch to use further logic such as checking the DSCP value and the UDP port of a packet to remark it. This would allow a stream with DSCP 46 to be identified independently of PTP traffic and remarked to DSCP 34.

QoS policies

QoS policies come in many different shapes and sizes and can be determined by the network's capabilities as much as the DSCP requirements. There are three areas that should be addressed with implementing an AoIP QoS strategy:

- PTP
- Media
- Discovery/Management

PTP

PTP is relatively low bandwidth compared to media but does not tolerate large amounts of latency and jitter. It is therefore advised that PTP should take priority over other network traffic.

Using DSCP 46, PTP should be mapped to a strict priority queue. When a strict priority queue is configured, the oldest packet in the queue gets transmitted before anything else.

Because there are generally fewer than 100 packets per second on a congested non PTP aware network, there is little to no risk of queue starvation. There are always plenty of gaps between the PTP packets for the switch to play out media packets.

Media

Media traffic is higher bandwidth than PTP although it does tolerate latency and jitter better than PTP traffic. Therefore, media is generally assigned to a medium priority queue.

There are several algorithms that can be used to pick which order packets are taken from non-strict queues such as round robin, weighted round robin and dedicated bandwidths. If a network link is designed to run at around 60% of its total bandwidth, the QoS choice algorithm generally has less significance. Round robin is usually default on most switches.

Discovery/Management

Discovery and management traffic is usually low bandwidth and can handle latency, jitter and packet loss exceptionally well. Given the importance of this traffic compared to PTP and media, it should generally be assigned to the lowest priority queue.

Some discovery and management traffic is sent using the TCP protocol. TCP has an in-built mechanism for retransmission if packets are not delivered to their destination. For this reason, adding this traffic to the lowest priority queue increases its chances of being dropped if the switch comes under unpredicted load. This is desirable as it is understood that the data will be retransmitted.

One QoS strategy for management queues is to limit the amount of bandwidth that can be used. For example, 10Mbps of bandwidth could be reserved for use on the lowest priority queue. This means that if any bursty or unpredictable traffic exceeds this threshold, it will be dropped to protect the other packets (media and PTP) on the network.

Switches

When deploying AoIP, at the centre of everything is the network. Networks are made up of switches that interact with each other and end devices to transport media. Switches can either be managed or unmanaged.

For AES67 and ST2110 purposes, managed switches are required. Managed switches are provided with the ability to manipulate and configure the feature set of the switch to provide the best performance to the AoIP network. Managed switches will present a UI or a CLI to allow the user to configure the switches' features.

In this chapter we will review the essential features that switches must have to successfully operate in an AoIP environment.

Port Count

The number of available ports a switch has determines the number of devices that can be plugged into it.

Switches are available in lots of different shapes and sizes. A switch that is used to aggregate many devices may have many ports to do so. A switch that is used to aggregate a few switches, does not necessarily need to have the same port count.

All of the following must be considered when deciding on the port requirements for a switch:

- AoIP devices
- Servers
  - *Calrec Connect*
  - *NMOS*
  - *Misc*
- Control and diagnostic PCs
- PTP GMCs
- Inter switch links.

It is wise to plan for future capacity when implementing an AoIP network. For an AoIP network to scale, special attention is needed at the start of a project to understand where you might need to expand in the future. A 12-port switch may be adequate for today's use case, but what about tomorrow's?

Throughput

A switches' throughput is its maximum capability to pass packets between ports on the switch. Generally, most switches that are considered for AoIP projects are designed so that if every link on the switch was run at full capacity, the internal switching would be able to cope with this demand. These switches are known as non-blocking switches.

Throughput can be presented in two different units (both show up on switch data sheets):

- Gbps or Tbps (Gigabit or Terabit per second)
- Mpps (millions of packets per second)

Throughput is the amount of link usage that is intended to be used, calculated and cross referenced with the data sheet. To calculate the amount of throughput (measured in bandwidth) you require:

Port Speed×Number of ports×2 = required throughput

For example, if you are designing a network and you want to check how much throughput an 8 port 1G switch would need to be non-blocking you would make the following calculation:

1×8×2 = 16Gbps

The reason you multiply by two is to account for Rx and Tx traffic.

Calculating throughput for switches measured in Mpps is a little more difficult. But generally, network designers will use the worst-case scenario of a 125us packet time, which produces 8000 packets a second. Take the intended number of streams you think will transmitted per device, multiply by 8000 and then apply the same logic as above.

Uplink Speed

Most switches will have a set of standard ports and maybe one or two ports that are intended as uplink ports. Uplink ports usually operate at a higher bandwidth than standard ports. For a 1G native switch, it is common to see multiple 2.5G or 10G uplink ports.

The reason uplink ports operate at higher speeds than standard switch ports is because they are responsible for passing data to/from many end devices. If for example a switch has 8 x 1G ports and 1 x 10G port, the 10G port can be used to transport data from the 8 x 1G ports onto another switch or network.

It is important to avoid bottlenecks when bandwidth planning. Uplink ports should be able to okay to use the maximum port speeds when doing calculations, but for more realistic estimates, it may be advisable to use the 60% guidance, or methodically calculate the actual intended bandwidth usage for each device.

Connectivity

Some switches will have fixed connectivity, usually for RJ45 connections. Most 1G switches that are applicable for AoIP networks are fixed copper. However, it is becoming increasingly common to deploy switches which have SFP or QSFP cages, so the user has full control of the connectivity for each switch port.

One thing to be aware of, is some switch manufacturers will only support SFPs that have been approved for use by themselves. This is always worth checking with the switch vendor, as SFPs supplied by other vendors may not be compatible with the switch.

Spanning Tree

Spanning Tree is a feature that is found on almost all switches. The spanning tree protocol is used to detect frame loops in a layer 2 network.

By default, switch manufacturers will have this protocol enabled by default. In general, this should never be turned off and switches should not be used where spanning tree is not a supported protocol.

It is generally advised to leave spanning tree settings to their default values unless there is good reason to change them. In layer 2 environments, if a new switch is detected on the network, spanning tree can converge which may lead to temporary disruption to audio.

Energy Saving

Some switches have a feature that enables energy saving which can diminish performance of the switch. This feature may be named slightly differently for vendor to vendor. Energy saving on a switch should be turned off.

IGMP

Switches should have basic multicast functionality including IGMP snooping and the ability to be an IGMP querier. See the multicast chapter for more information.

QoS

Switches should have adequate QoS functionality with at least 4 queues. See the QoS Chapter for more details.

PTP

Switches do not necessarily need to be PTP aware for use on AoIP networks, but it is often desirable. For more information on PTP awareness and its requirements see the PTP and Network Design chapters.

Layer 2 vs Layer 3

Switches can operate at layer 2 (forward layer 2 frames) or layer 3 (route IP packets). The protocols used to manage packets across these boundaries offer different paths to deployment and should be considered carefully.

Generally, layer 2 networks scale vertically. To expand a layer 2 network, you add more devices to the network and assign them IP addresses from a subnet. If the IP address space is depleted for that network, then all devices on the network may need to be configured with a larger subnet to continue adding devices.

Any device that is on a layer 2 subnet will be able to broadcast packets to any other device on that network. This is seen as a vulnerability on high performance networks because the larger the subnets grow, the more broadcast messages each device must listen to and throw away. This means that the larger a Layer2 network grows, the more inefficient and noisier it becomes.

However, layer 2 is generally considered easier to work with and requires less specialist knowledge to implement. Protocols such as mDNS are entirely dependent on layer 2 networks to operate.

Layer 3 networks scale horizontally. To add more devices to the network, you can add another network with a different subnet and route traffic to that new subnet.

Fig 20

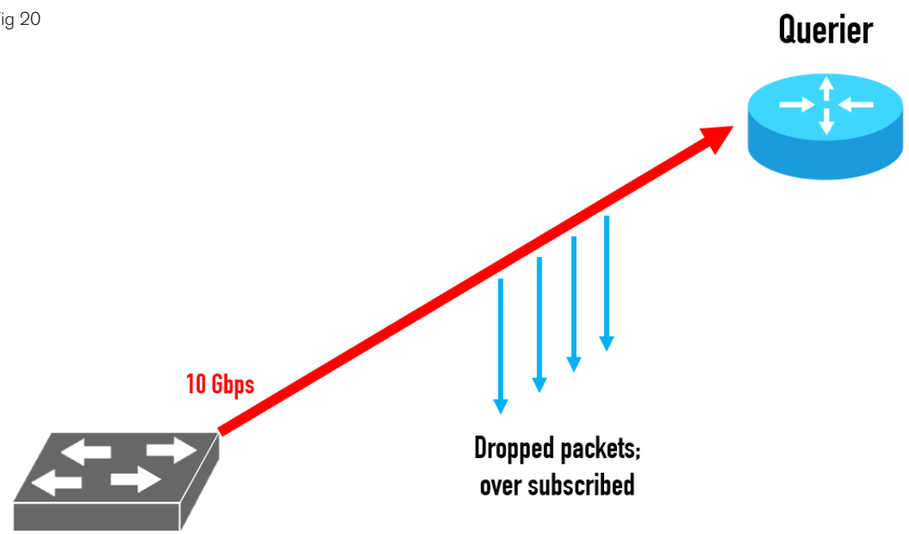

Deploying a second studio for example would involve creating a second AoIP network and routing between them. This means that when the second network is implemented, the first remains unchanged which is a highly desirable deployment strategy.

A major benefit of layer 3 technologies is the ability to avoid flooding of unknown or undesired multicast. In ST2110 environments, it is common for audio and video to co-exist on the same network. Commonly, audio devices will have their own switches due to 1G connectivity as opposed to 10/25/40g on video end devices.

If a video stream with bandwidth 10Gbps (a nominal value in the range of 4k video) is unintentionally exposed to a layer 2 network through accidental mispatching, protocol or device failure, or malicious means, then that 10Gbps will be flooded towards the querier.

This may mean that the link to the querier is now oversubscribed and there will be audio failure due to lost packets (fig 20).

Strategic configuration of layer 3 multicast routing and network/subnet boundaries can ensure that this flooding does not occur at all or if it needs to, it is only temporary.

Layer 3 end points on a network also typically have more featureful security, monitoring and SDN capabilities which can offer strong protection against multicast flooding.

Implementing AoIP networks at layer 3 requires more specialist skills, especially with reference to multicast routing. Network infrastructure must also be capable of routing at layer 3 which generally comes at a monetary cost.

Layer 2 Pros and Cons	
Pros	**Cons**
▪ IGMP is very easy to administrate ▪ mDNS discovery ▪ Works well if the network is undersubscribed	▪ Multiple switches may cause IGMP querying complications ▪ Broadcast/failure domains ▪ Limited scalability

Layer 3 Pros and Cons	
Pros	**Cons**
▪ PIM creates trees that ensure multicast is never flooded ▪ Smaller and segmented failure domains ▪ Scales horizontally well ▪ Better traffic visibility ▪ Flood protection	▪ Not as easy to configure or administrate ▪ Cannot pass some discovery traffic between networks ▪ Requires network hardware capable of unicast and multicast routing ▪ Increased feature set -> increased cost

Network Design

The goal of understanding IP technologies is to enable us to design scalable and highly available AoIP networks that meet the resiliency standards we have come to expect during the previous generations of broadcast technologies.

In an ideal world, IP networks should be invisible to the end user as they are perfectly designed to meet their workflows. IP networks should enable the desired workflow of the broadcaster without defining or compromising operations.

Network design and implementation can be as easy as plugging AoIP devices into some network switches and pressing play. Network design and implementation that is deterministic, resilient, cost effective and scalable requires more thought and experience.

This chapter brings together some of the concepts discussed in previous chapters as well as Calrec's extensive experience in implementing AoIP networks in broadcast environments.

Although there are some definitive dos and don'ts, much of this chapter simply presents best practices. Best practices may or may not be applicable for any given AoIP deployment. It may be that a best practice for a multi-studio environment is a multi-switch layer 3 deployment. This does not mean that the workflow required cannot be achieved with a single switch layer 2 deployment or even that it is more desirable to do so.

Benefits of good network design

Good network design can:

- Improve performance
- Increase manageability
- Allow you to scale more efficiently in the future
- Decrease implementation time
- Decrease capital and operational expense

Performance Benchmarks

Ensuring the performance benchmarks of an AoIP network is essential. Ensuring that the bandwidth is available where and when you need it is of the upmost importance. Some network designs will achieve this by ensuring that the bandwidth capabilities of the network are way higher than the bandwidth that it is possible to produce.

Some network designs will rely on other technologies to reserve bandwidth per source across the network.

Both solutions avoid oversubscription but use vastly different methods. Either way, the network has been designed to meet the performance requirements of the workflow.

Unconfigured or misconfigured networks can generally not meet the vital performance benchmarks that are required.

Manageability

It is important that operations teams are considered during the network design phase. Networking implementations can be easier to work with if certain parameters are designed to be human readable and memorable. For example, Amber and Blue vLANs may be configured as:

- Amber: vLAN: 654
- Blue: vLAN: 129

Or they could be:

- Amber: vLAN: 100
- Blue: vLAN: 200

Which configuration is easier to remember? vLAN 100 and 200. It is generally easier to maintain, and fault find IP systems when the numbering system of various components ties into the physical location or practical application of the device. This concept also applies to IP addressing, which we will cover in a following section.

Scalability

The network design phase is the best time to plan for future expansion, whether that is the addition of ad hoc AoIP devices, or the longer-term integrations of more facilities. IP can adapt and scale to meet the requirements of tomorrow and it is made much easier by planning for your future needs.

For example, an IO box cannot be added to an AoIP network that has no switchports available. This dilemma would make a broadcaster's intended workflow impossible to achieve.

We discuss future proofing in a following section. An IP network's scalability is enabled by its design as much as it is defined by the network's capabilities.

Implementation Time

Using network design best practices, the risk of unknowns that could affect stability or performance of an IP network are reduced.

By considering device placement, PTP awareness, multicast performance, addressing structure, and bandwidth management, you can remove a lot of risk from the installation and implementation phase of a project.

Reducing Expenditure

With reduced implementation time comes less operational expenditure on integrating and testing the platform. Over time, these savings increase, especially if the alternative would be a last minute network redesign due to lack of preparation.

By thoroughly assessing your needs and the networking devices that you intend to deploy you can be confident that you are not over specifying the hardware you need, and this can reduce capital expenditure at the start of a project.

There are also opportunities to economise by exploring alternative workflows.

Identifying your needs

The first step in network design is to identify the needs of your workflow. The following sections discuss various aspects of identifying your needs.

Port Count

One of the most fundamental things to identify is that you have enough network ports on your network switches to accommodate your network devices. This will often mean examining what baseband audio requirements you have and how many AoIP devices are required to achieve the required channel count.

It may not be okay to assume that each device requires a single port on a switch. For example, Calrec's Modular IO boxes have the option for two NICs to increase the capacity of the device. Special attention should be taken to ensure that the port count is counted independently of the device count.

Geography

Understanding the port count is important but a second consideration is to understand how these audio channels are distributed through the facility.

For example, you may identify a requirement for 300 audio inputs. 100 of these come from the studio floor, 100 from the control room, and 100 from the machine room. The questions that are posed in the scenario are:

- Is a switch required in each area?
- If both switches connect back to the machine room, does the bandwidth of the inter switch links need to be higher than the AoIP devices?
- What is the most efficient and scalable way to run cables?
- Should spare cables be run for future expansion?

Unfortunately, there are no stock answers for these questions. The answers will be entirely dependent on the layout of the broadcast facility and the workflows that are to be met.

Bandwidth

Once you have an idea on the number of ports that are needed and where they will be positioned geographically, it important that the links between the switches do not present a bottle neck.

An example of this was given in the bandwidth chapter. It is essential that the available bandwidth between the switches is equal or greater than the bandwidth of the end devices that are connected to them.

The 60% rule may also be applied here, where no link should exceed 60% of its full capacity.

This is more applicable on layer 2 IGMP networks compared to higher performance layer 3 and SDN environments.

If abiding by the 60% rule, it could be said that 60% of the uplink bandwidth must be equal to or greater than the combined downstream bandwidth as shown in fig 21 below.

This methodology of calculating bandwidth usage should be considered between all switches in both directions.

Calrec recommends abiding by the 60% rule as a matter of precaution but can be altered as appropriate to a given implementation.

PTP

Deciding on whether go with PTP aware switches comes down to several factors:

- The amount of AoIP devices
- The amount of switch hops
- The supported number of slave devices as defined by the GMC
- QoS features available on the network
- Expected load on the network
- Planned expansion for the future

Due to the number of variables, the suitability of PTP aware switches should be decided on a case-by-case basis. Below are some rough guidelines that can be followed to help with decision-making.

- For deterministic, reliable and accurate PTP performance, boundary clocks should be used.
- As the number of end devices on a system increase, consider the inclusion of boundary clocks running on switches to manage the workload of the GMC and to remove inter-device PTP traffic.
- If the AoIP devices are lower capacity devices (ie. they produce and consume less than 100Mbps of media) then the number of devices may be increased without using a boundary clock as the load on the network is mitigated.

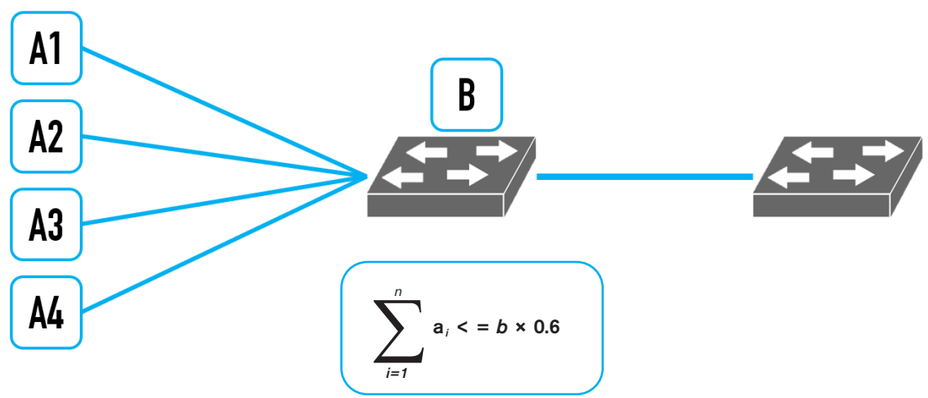

Fig 21

- If there are multiple switch hops needed between the GMC and the slaves, transparent or boundary clocks should be considered.
- If the switches can place PTP traffic into a strict QoS queue, this improves the performance of non-PTP aware switches.
- If the network is likely to expand over time, it is recommended that boundary clocks are put on the network from the start to ensure that any expansion does not add load to the existing PTP architecture.

Note that the guidance provided above is flexible and may have different levels of effectiveness depending on the variables stated at the start of this section.

Identifying the capabilities of AoIP devices

AoIP devices are generally fully interoperable, but there are areas that are not covered by standards which can lead to unforeseen challenges. The two main resources to get an idea of interoperability are:

1. JT-NM Tested Event results

   The JT-NM publish test results from their test events. This gives broadcasters a level of confidence of compliance to ST2110, AMWA NMOS and TR1001-1 for IP products.

2. AES67 PICS

   AES67-2018 contains a compliance checklist called PICS that is to be completed for all products that claim AES67 interoperability. This can be obtained by contacting the vendor of the product or may be published publicly.

Future proofing

A major part of network design is ensuring that the network can scale to meet the needs of your future workflows.

This may mean scaling down but often it means scaling up. Installing an infrastructure that is incapable of scaling upwards can reduce the initial capital expenditure of a project, but can also lead to complications when implementing new workflows in the future. These complications will often require additional capital and operational expenditure to put right.

There are two scenarios that cover all bases when considering future proofing your network design:

1. The future is understood.
2. The future is not understood.

In the following sections, we will discuss these two scenarios and what solutions are available for both.

When the future is understood

Maybe you are implementing a multi-studio facility in stages. You know that you will start by implementing one AoIP network, but this AOIP network will grow to ensure connectivity to the full multi-studio facility in the future.

There will certainly be unknowns and unexpected challenges to the plan along the road, but in general you have good visibility as to what levels of capacity need to be accounted for in the initial design.

One approach would be to install a high-capacity backbone (known as a spine) in a centralized location. From this spine, multiple leaf switches can be attached with each leaf switch aggregating devices from a single location. This increases capital expenditure at the start of a project, as you are installing a spine and a leaf switch where you could maybe work without the spine.

However, by installing this infrastructure from the start you can prove the deterministic performance of the AoIP, PTP and multicast.

With the network secured and trusted, scaling the network is as simple as adding extra leaf switches to the topology. This allows for agile scaling of the AoIP network where the capital and operational expenditure is limited to the costs of the new leaf networks. The more leaves that you add, the more benefit you get from future proofing the design early.

When the future is not understood

Broadcast workflows can be volatile. A facility may have a last-minute booking with special requirements that the facility must quickly adapt to.

IP networks can scale upwards and downwards to meet demand with relative ease, but there is a basic requirement that the static part of the network has the capacity and capability to do so.

When provisioning resources for an AoIP network, it may be worth adding an extra 20% capacity to account for any unknowns that might appear during deployment or operation.

For example, if it has been determined that a switch must have 10 ports to satisfy the basic requirements of the project, purchasing a switch that has 12 ports allows for any unknowns or for future expansion.

Furthermore, it may be worth ensuring that switches have spare uplink ports (which are usually a higher bandwidth than the standard ports) so you are able to ensure connectivity to another switch or network in the future.

The 20% extra applies to many areas of network design including:

- Port count
- Bandwidth usage
- AoIP channel count
- Subnetting

20% is used for guidance only, if the network designer understands that there may be significant unknowns or is keen to be able to react to expansion opportunities this value can be increased to any value as necessary.

IP Addressing Strategies

Successful planning of the IP address space can have a significant impact on the manageability of an AoIP network. AoIP networks can have hundreds if not thousands of IP addresses across the entire network, and structuring them in a way that they are more memorable can bring huge efficiencies. In this section, different approaches to IP addressing will be discussed.

Dynamic Host Configuration Protocol (DHCP)

DHCP is a protocol that allows an external entity to apply an IP address, subnet, gateway, DNS servers or other items of configuration to a device. DHCP servers can be configured on some switches or may be present as a standalone server.

The benefit of using DHCP on AoIP networks is that network engineers only manage the pool of IP addresses that the DHCP server serves to clients, and not the individual IP addresses of the AoIP devices.

This allows for the automated configuration of AoIP devices. DHCP is gaining popularity within the industry with bodies such as the JT-NM promoting its use as documented in TR1001-1.

One potential issue with relying heavily on DHCP is the lack of clarity on IP addresses that have been assigned to individual devices. If there are no service discovery mechanisms in place on the network, it may be difficult to understand what IP addresses have been allocated.

For example, if NMOS IS-04 has been adopted as the service discovery mechanism, then a user would be able to resolve the IP addresses of all IS-04 enabled devices on the network using an NMOS controller.

If there is no mechanism in place, and there is no physical way to read back the IP address of a device, then access to that device for operational or diagnostical purposes may be difficult.

DHCP leases IP addresses to its clients which are bound to the device for a particular time before the client must renew its license. There are mechanisms in place with modern DHCP implementations that devices retain their IP addresses if they are temporarily unavailable on the network.

Although this is made possible by design, it cannot always be guaranteed and it is possible that an AoIP devices' IP address may change.

Static Addressing

An alternative to DHCP is to deploy static addressing across an AoIP network. Although this comes at a higher operational cost, it allows fine grain control of IP addresses and allows the implementation of custom logic which is specific to the broadcaster's requirements.

For an example addressing strategy see Appendix 4.

Subnet sizes

Typically, a 255.255.255.0 (/24) subnet size will be adequate for most AoIP networks. There are 254 available IP addresses in this range, and even if only a fraction of the available addresses are used, there is still plenty of room to grow.

A /24 also has the added benefit of having a clear boundary between the network bits and host bits which makes the subnet easy to work with.

Larger subnet sizes such as 255.255.0.0 (/16) may be easier to work with as there are far more addresses to use but the cost of this is increased address space wastage and being tied into a more vulnerable layer 2 design.

A comparison table of subnet sizes can be found in Appendix 5.

Layer 2 vulnerabilities

A vLAN is a layer 2 domain, which means that if a device sends a broadcast onto a vLAN, all other devices on the vLAN will also receive that message. This wastes CPU time because every device will receive the message even though it was not intended for that particular device.

The number of devices on a vLAN is dictated by the subnet associated with that vLAN. For example, a 255.255.255.252 (/30) subnet would have 2 devices whereas a 255.255.255.0 (/24) could have 254. The larger the subnet, the more potential there is for increased broadcast traffic on a network.

As well as being broadcast domains, vLANs can also be described as layer 2 failure domains. There are several layer 2 vulnerabilities that vLANs are suspectable to:

- Broadcast Storms
- Spanning tree converging events
- ARP Spoofing/Man in the Middle attacks
- Multicast flooding

Broadcast storms can be caused by unintentional loops being created within a network topology.

Spanning tree will usually disable a port if a loop is detected in order to stop a broadcast storm from occurring, but is also likely to cause interruption on the network.

This would be unacceptable if it happened on air. There are various ways someone with malicious intent could attack a layer 2 network which may not exist if the network was segmented into smaller and more manageable segments.

The solution is to decrease the size of the subnets in use on the network and use Layer 3 routing technologies to guide traffic between the vLANs or networks. This comes at a higher cost and requires more specialist networking knowledge to achieve.

Designing for Resilience

Resilience is at the heart of a broadcaster's needs; recovery in the event of a disaster should always be considered when designing AoIP networks.

This may mean different things for different broadcasters. Manual changeover to a separate baseband topology with an approximate recovery time of two minutes may be perfectly acceptable for one broadcaster, but not for another.

When it comes to adding resilience to an infrastructure, over-provisioning resources is a good idea. Having the capacity to continue operations on your network in the event of a partial network failure can ensure swift and automatic recovery times.

This can be as simple as adding extra links between switches and aggregating them, or it could mean deploying multiple spine switches to ensure paths remain open in the event of a spine switch failure.

Each network design has its own opportunities for added resilience. The example topologies at the end of this chapter work through these various opportunities in more detail.

ST2022-7

SMPTE ST2022-7 is a standard that allows ST2110 streams to be transmitted redundantly across discrete AoIP networks.

An AoIP device transmits two streams with the same audio content onto two different networks. The receivers of those streams then "pick" packets from each network to rebuild the audio.

This is known as *hitless packet merging.*

There are several different nomenclatures when describing the two networks used in a ST2022-7 solution. A and B, Red and Blue, Amber and Blue, and Primary and Secondary are all in regular use.

For the purposes of this document, we will use Amber and Blue to maintain consistency with TR1001-1.

There are several strategies to account for ST2022-7 within a network design:

- Implement one network and do not use a second
  - » *This is a perfectly valid strategy, albeit to be discouraged. ST2022-7 is an option that is highly recommended to protect audio content*
- Implement vLANs within a single switch to isolate Amber and Blue traffic
  - » *This allows cable and packet redundancy but may contain single points of failure within the network*
- Implement two completely isolated networks
  - » *This allows cable, packet and network hardware redundancy*
- Implement a single "purple" network which accepts streams from amber and blue networks and processes them independently
  - » *Purple networks are usually implemented in Layer3 routed multicast environments and make use of routing protocols to isolate traffic across a shared infrastructure.*

The option that is most desirable will depend on the requirements of the project.

Both Amber and Blue networks should be equal in performance. ST2022-7 specifies that AoIP devices should be able to handle acceptable amounts of variance between the networks, but a good network design will mitigate any variances in performance.

Typically, the same hardware running the same configurations is used on both sides of the network.

Alternatively, having different vendors for each side of the network diversifies the AoIP network in case of a vendor-specific failure on the network.

This would however mean maintaining two sets of equipment that may have significant operational differences and incompatibilities.

Both networks should have their own IGMP queriers and IGMP snooping should be enabled globally.

PTP Considerations

A challenge that is unique to implementing ST2022-7 is how to distribute PTP effectively between both Amber and Blue networks.

How PTP can be distributed may be influenced by:

- Network connectivity on the GMC
- The network's PTP Awareness
- Switch functionality

How PTP is distributed between both sides of a ST2022-7 network will dictate how the PTP BMCA converges. End devices can end up in several different states depending on what the configuration of the network is.

Some PTP states you may come across on a ST2022-7 network are demonstrated in the following examples. The examples assume boundary clocks are used throughout the network.

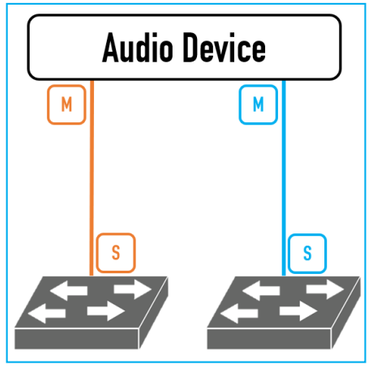

Amber: MASTER - Blue: MASTER

The audio device is the master device on the network and is sending announce and synchronization messages to the network on both of its ports.

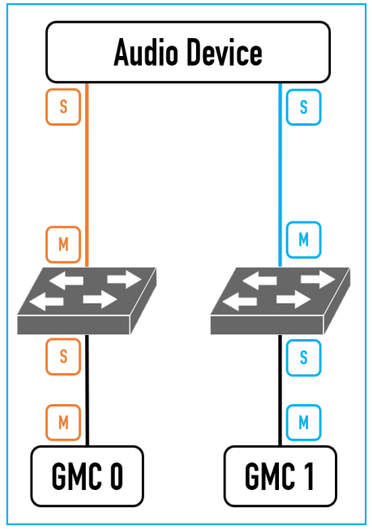

Amber: SLAVE - Blue: SLAVE

This may occur if the PTP implementation on the end device treats each NIC as its own discrete PTP clock. This allows both NICs to SLAVE to the MASTER. Devices that operate like this often have internal logic to choose which of the two PTP clocks on the device is the source for the audio clock synchronisation.

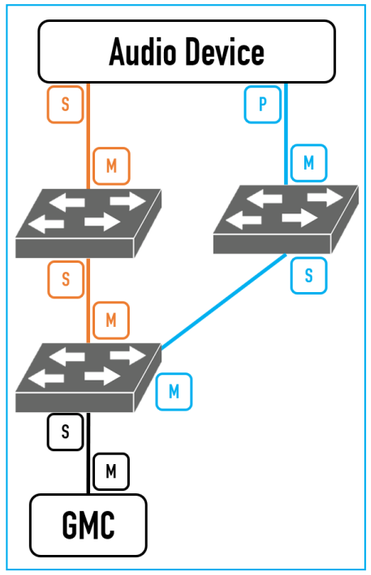

Amber: SLAVE - Blue: PASSIVE

This may occur if the PTP implementation of the end device treats both NICs as a contributor to a single PTP clock. The PTP clock can only SLAVE from one NIC. The NIC that is SLAVE is topologically closer to the GMC than the NIC that is PASSIVE. The NIC that is PASSIVE is also connected to a MASTER but does not update its PTP clock with timing information. These are the ideal port states for devices that operate this way.

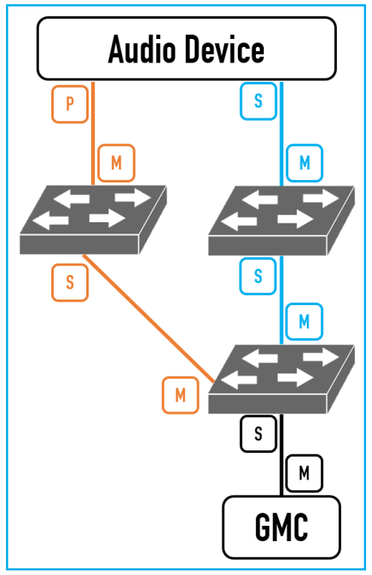

Amber: PASSIVE - Blue: SLAVE

As b efore however this time the GMC is topologically closer to the Blue NIC than the Amber.

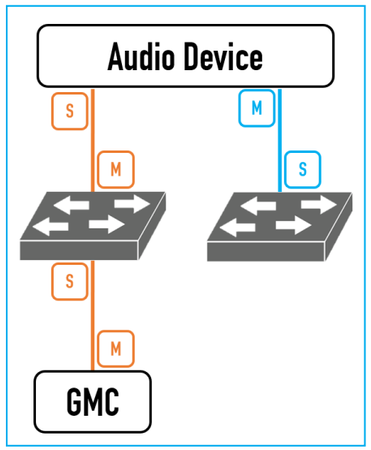

Amber: SLAVE - Blue: MASTER

In this situation, the Blue NIC is topologically closer to the GMC than the nearest PTP port on the network . Therefore the end device effectively runs as a boundary clock. It is generally advisable to try and avoid this configuration as adding and removing audio devices from the network can cause wider BMCA events.

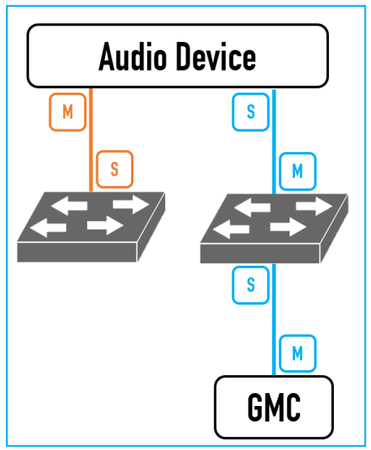

Amber: MASTER - Blue: SLAVE

As above however this time the GMC is topologically closer to the Blue NIC than the Amber.

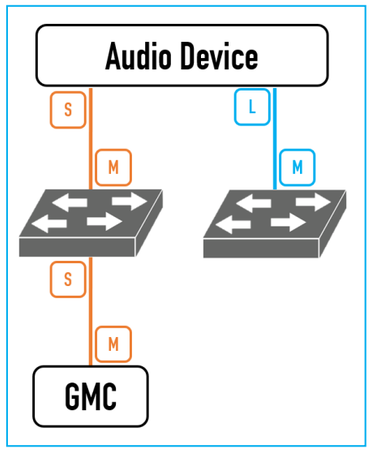

Amber: SLAVE - Blue: LISTENING

When a ST2022-7 device is set to “Slave Only”, and the device uses both NICs as a source for the PTP clock, one of the ports will try to become MASTER. However, this is blocked due to the Slave Only setting. Therefore one of the ports will go LISTENING. On a failover, the LISTENING NIC would go SLAVE but may take take a significant amount of time to regain the current time as the LISTENING NIC is not taking part in time synchronization. A ST2022-7 NIC pair that goes SLAVE/ LISTENING is not recommended.

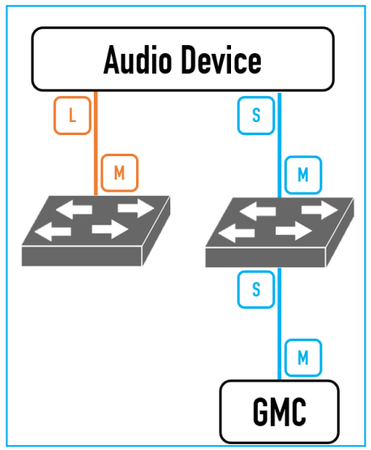

Amber: LISTENING - Blue: BLUE

As above however this time the GMC is topologically closer to the Blue NIC than the Amber.

When working with PTP on a ST2022-7 network, there are two approaches:

1. Using a separate GMC on each network
2. Using the same GMC on both networks

Although these are presented as two different approaches, option 2 is a natural progression from option 1. The only difference is the point in which timing data is shared between the Amber and the Blue network.

Separate GMCs

When separate GMCs are used on the Amber and Blue network, only one of them will act as the GMC at abt one time while the other may slave from the it depending on the PTP implementations of the AoIP devices. In this setup, it would also be common for both GMCs to reference GPS to provide consistent timing data to both Amber and Blue networks.

In this topology, the AoIP devices will generally bridge the Amber and Blue networks by passing PTP timing information between them (subject to the AoIP devices’ PTP implementation.

It is common for a pair of NICs to SLAVE to the Amber network and then be MASTER to the Blue (and vice versa if the GMC were placed on the Blue side). This is because the AoIP devices may receive information about the GMC on one NIC and transmit that information the other. In this topology, the AoIP devices effectively run as boundary clocks between the Amber and Blue networks; this generally isn’t advisable, as a failed or removed/reintroduced audio device may have wider implications on the BMCA.

Same GMCs

An alternative approach would be to use the same GMCs on both sides of the network. This involves passing PTP timing data at the GMC or network level, rather than relying on AoIP devices to bridge the two networks.

This is generally the desired option as it allows the PTP BMCA to converge irrespective of the AoIP devices abilities or configuration.

If a GMC has two NICs that can serve the Amber and Blue networks concurrently, this is the most operationally efficient way to provide the same GMC to both sides.

If a GMC only has a single NIC then it is the responsibility of the network designer to provide a solution to share the PTP data between the networks. Common solutions are:

1. Implementation of a dedicated infrastructure to provide multiple paths to each network from redundant GMCs
2. Adding a network link between the Amber and Blue network
3. Using a GMC that has two NICs to separately feed the Amber and the Blue network.

It is essential that only PTP data is shared between the networks when a dedicated PTP link is used. The presence of media traffic that was meant for the Amber network on the Blue network may cause problems for some AoIP devices (and vice versa). There are several methods for achieving this traffic isolation:

- The use of ACLs to permit PTP traffic but deny other traffic
- Dedicate vLANs specifically for PTP distribution
- Use routed switchports for PTP distribution
  - » A useful trick is to not assign IP addresses to these ports so they do not influence any routing protocols that may be active on the network

Dedicated Grand Master clocks may have the ability to be “Master Only” devices. This would force the port state on the backup grandmaster to be passive rather than slave.

This allows the backup GMC to continue to rely on its own GPS hardware to derive time rather than relying on the active GMC. Some switch manufacturers also provide the ability to lock certain ports into a “Master Only” state.

This allows the network designer to influence the BMCA at a topological level to prevent end devices from converging timing data between Amber and Blue networks.

PTP BMCA and GPS Design

If a single GMC is deployed on an AoIP network, then the only point of concern is the user definable Priority1 value for each clock.

Ensuring that the Priority1 value is set lower than all other devices guarantees which device will become GMC on the network.

When working with redundant GMCs consideration can be given to more intricate tuning of the BMCA. Take following example in Fig 22 where there are two GMCs are on the same network; the active on the Amber network and the backup on the Blue network:

Both GMCs have independent GPS connections for full redundancy. A common approach in this situation is to apply the same Priority1 value to both GMCs. This forces the BMCA to use Clock Class, Accuracy and Variance as tie breakers before the next user-definable tie break in Priority2.

This allows the PTP BMCA to dynamically choose which GMC should be active on the network based on the quality and performance of the timing information that is produced by each GMC. The Priority 2 value can then be used as a user-definable tiebreaker if performance is equal between the clocks.

Dual NIC Reachability Considerations

As standard, all Calrec AoIP hardware is supplied with a pair of NICs for all media applications. On ImPulse and Type R, control applications can also be mapped to dual NICs. For such functions as management and control, it may be necessary to add a default gateway to the NICs in order to allow the application to communicate with a remote subnet.

This causes an issue, as generally you cannot add multiple default gateways to the same device because it causes a conflict in the devices routing table. Using static routes rather than default gateways provides finer control over the direction traffic will flow to subnet devices.

For example, Amber and Blue networks could be using 172.16.10.0/24 and 172.16.20.0/24 respectively (illustrated in fig 23). However, within the /24 block, the address space could be further divided into multiple /30 networks so that every NIC is effectively on its own network. Each device should have a static route applied to the applicable /24. This allows the AoIP device to communicate with all devices in the dedicated range.

A default gateway may also be needed to ensure that the AoIP device can communicate with any other subnet, to enable management and control from a remote or dedicated management network.

Management Networks

Without discovery, control and management traffic, AoIP traffic would not be able to flow between senders and receivers. Different AoIP solutions have different requirements for discovery, control and management data, and the route this data takes through a network may or may not be definable by the protocol or the product.

From this point the terms discovery, control and management will all be handled under the same umbrella and will just be referred to as management traffic.

If management traffic is to be sent to the Amber and Blue networks, this is considered in band (IB) management traffic. If the management data has its own physical connections, then it can be placed on its own network. This is known as out of band (OOB) management traffic.

In Band Management (IB)

Generally, the requirements of Amber and Blue media networks for AES67 and ST2110 go far beyond what is needed for management traffic. It is useful to have an understanding whether the management traffic is unicast or multicast, which may help when planning an address space for the network.

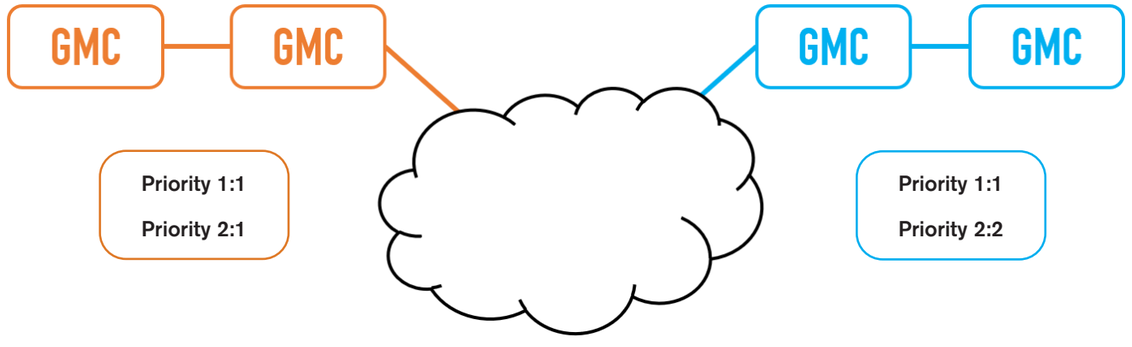

Fig 22

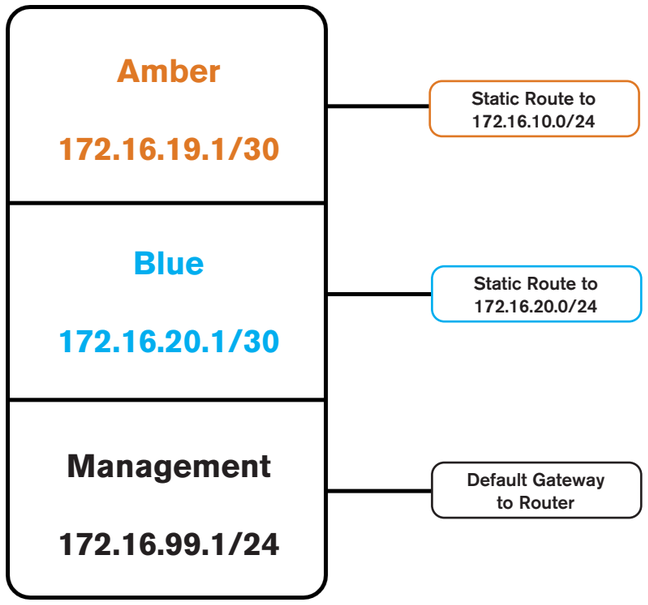

Fig 23

Networks that contain In Band management traffic should protect PTP and media traffic by ensuring that management traffic is not prioritised for transmission by the network. This was discussed in detail in the QoS chapter.

A risk of deploying in band management is that if the network topology develops a fault, management traffic may also be compromised. This can significantly reduce the availability of network applications and may hinder the ability to resolve the fault.

Out of Band Management (OOB)

Out of Band management networks have several key benefits over in band management networks, although these benefits come at the capital cost of implementing a dedicated management infrastructure.

By moving as much management traffic away from the media networks, there is a natural protection from adverse network conditions caused by management devices and protocols. This also allows added flexibility as out of band management networks can be integrated easily into existing enterprise environments.

Out of Band management networks can also offer viable paths to the media network if the media network is in a fault condition. This allows the network to remain managed even if the media network has been compromised and allows for more rapid maintenance and fault finding.

Commonly, the OOB links will be connected using layer 3 interfaces or virtual routing and forwarding systems (VRF) to provide sufficient fault isolation between the networks.

Users generally access management user interfaces through PCs connected to the management networks. PCs that are connected to the internet pose a vulnerability if connected directly to a media network. In the event of an OS failure due to technology reasons or otherwise, it is safer for the failure to happen in a dedicated management environment rather than in a media environment. Placing PCs out of band wherever possible achieves this.

Some AoIP equipment is only manageable over its media network NICs. A hybrid architecture could be implemented where the management network is able to exchange management traffic with a media network.

By routing packets between management and media vLANs or switches, the network designer can make use of common switch features to secure access to the media network. Some switch features which may be used to secure a management link include:

- ACL or Firewalls
- Port Shaping or Policing

Software Defined Networking (SDN)

SDN is a branch of networking technologies that abstracts the control plane from the data plane. Traditionally, the mechanisms that physically pass packets from one port to another (data plane) are tightly coupled to the protocols that pass packets from one port to another (control plane). SDN removes the control plane from inside the switch and offloads it to a central location.

It is common that a SDN controller (also known as a broadcast controller) takes control inputs (such as routing requests from a control panel) and translates them into instructions to the network. The SDN controller dictates packet flow through the network using custom forwarding rules which are created based on the routing instruction from the controller. This process is also known as flow orchestration.

SDN technologies are often aware of a network's bandwidth utilization and topology. SDN Controllers can protect against oversubscription of network links by load balancing traffic or rejecting routing requests that would cause oversubscription on a network. This is also how SDN provides superior resilience in case of network component failure, as alternative paths can be securely and deterministically utilized.

Security

Securing AoIP networks is as important as any other topic in this guide. Content is a broadcaster's most important commodity and having it compromised can have significant implications.

Security is best implemented at every level so no single system's failure can compromise operations. There is a joint responsibility from all members of a broadcasting team to ensure that best practice security processes are in place and followed. This section describes some strategies for ensuring that a network is secure from accidental or malicious vulnerabilities.

Physical Access

Ensuring that physical access to a network is carefully controlled is one of the most basic ways to secure a network. By placing mission-critical network hardware behind doors with restricted access will help prevent any unauthorised access to the network. If someone with malicious intent gains physical access to a network switch, it may be relatively simple to disrupt.

Cabling should be undamaged, tidied and well labelled. This will prevent accidental removal of cables due to broken connectors or being pulled due to accidental contact. Labelling cables helps protect against personnel misidentifying cables and removing incorrect links. Cables that must be run in vulnerable places such as across a floor, should be well protected.

Switch Credentials

Usernames and passwords should be set on all network infrastructure in accordance with best practices. This often means minimum character limits and other enforcements. On some switches, it may also be possible to allow access through digital keys, which means only those who have a particular key can access hardware and there is no potential for a password to be stolen, guessed or cracked.

Switchports

Unused switchports should be turned off when not in use. This is a function that is available on most switches. By turning off unused ports, any device that is plugged into the switch cannot affect its operation without authorised personnel enabling the port.

Network Separation

Using vLANs to separate different devices from each other may make passing multicast and audio between them more operationally difficult, but there are several security benefits too. For example, if two studios are separated into their own vLANs, then the consequences of a compromise in one Studio will not affect the other. This logic can be taken further by saying each IO device gets its own vLAN or network. There is no specific guidance to this other than the smaller the subnet, the less impact a compromise will have on that subnet.

Monitor the Network and its Performance

Protocols such as SNMP provide methods of collecting information from all corners of a network. This can be diagnostic or performance related, and can be an early warning system for potential compromises.

Understanding the base performance levels of your network can help you identify any issues quickly. If the load or performance of the network changes unexpectedly it could be a sign that something is wrong and may require intervention.

Example Topologies

Previous sections have discussed a variety of variables and challenges to consider when designing AoIP networks. The following sections will present some practical examples with a discussion on the pros and cons of each.

Single Switch

This network demonstrates a simple network topology consisting of a single switch and several AoIP devices. The number of AoIP devices is limited by the number of ports on the switch. Amber and Blue will be implemented on their own vLANs to ensure logical separation. If ST2022-7 is not required by the project, no vLANs are required. Depending on the number of AoIP devices and desired performance, the switch may or may not be PTP aware.

- Media Network A
- Media Network B
- Connect A
- Connect B

Pros	Cons
Operationally simple to configure	Single point of failure on the network
Options for physical expansion using uplink ports	Potential single point of failure for PTP GMC (if network is not capable)
Good for isolated AoIP "islands"	PTP workload on an AoIP device could be a vulnerability as the network scales

Dual Switch with dedicated PTP GMC

This topology improves on some of the weaknesses of the single switch design. Having dedicated independent network devices for Amber and Blue networks removes the single point of failure from the network. Adding dedicated master clock devices removes workload from AoIP devices, and allows devices designed to perform as GMCs to take over. Dedicated GMCs often have additional feature sets including GPS and baseband sync connectivity which may be a project requirement. A dedicated PTP link is installed between the Amber and Blue switches to allow a communication path for timing information. The Amber and Blue switches have the same multicast configuration and offer identical performance to both sides of the network. Depending on desired scale and performance, PTP awareness remains optional.

- Media Network A
- Media Network B
- Connect A
- Connect B
- PTP Link

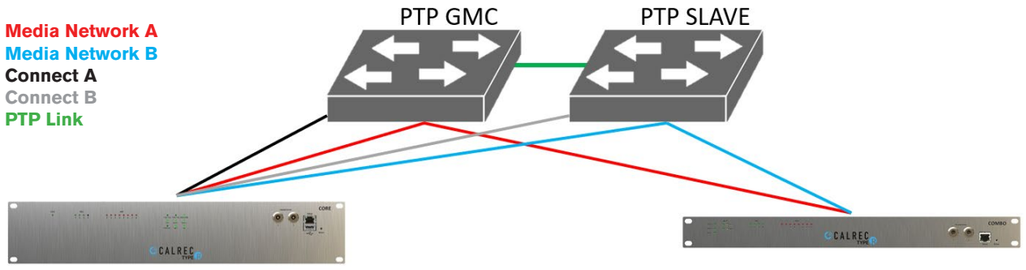

Pros	Cons
Operationally simple to configure	More configuration required due to increased switch count
Options for physical expansion using uplink ports	
Good for isolated AoIP "islands"	

Dual switch: multi-studio environment

Topology wise, this is the same as the previous example. However, as this is a multi-studio environment, there are more things to consider. With two studios in operation there is twice the amount of content, which potentially doubles the failure domain of the network. In addition, if studio1 and studio 2 need to exchange audio this could dramatically change the approach to switch configuration.

If Studio 1 and Studio 2 do need to exchange audio:
- Having both studios in the same vLAN makes audio exchange operationally easier, but ties both studios into the same broadcast domain
- Both studios can still be isolated by their own vLANs with multicast routers between them. This will often require a dedicated routing protocol and some networking expertise

If Studio 1 and Studio 2 do not need to exchange audio:
- vLANs could be used to isolated them. This ensures that each studio is its own broadcast domain
- The switch must be capable of passing PTP between vLANs. This could be through static multicast routing or through a boundary clock mechanism
- Each vLAN should contain an IGMP querier and have Snooping enabled

This example shows how a broadcaster's workflow defines the technology deployed within an IP network. The two options as presented above both look the same physically, but the underlying configuration introduces significant implications. Given the increased device count and increased consequence of network disturbance, it would be advisable to consider PTP awareness on the network at this point. This will add resilience and accuracy to the network now that it has begun to scale.

- Media Network A
- Media Network B
- Connect A
- Connect B
- PTP Link

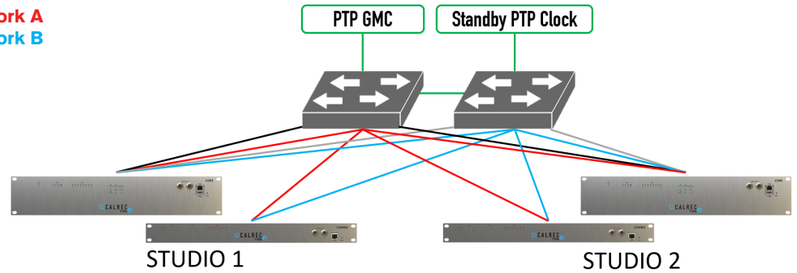

Both Studios share the same vLANs:

Pros	Cons
Operationally simple to configure	More configuration required due to increased switch count
Options for physical expansion using uplink ports	Two studios share the same broadcast domain, both are equally susceptible to broadcast storms or Spanning tree events
Configuration is the same as a single studio environment	Studios must be in proximity for cabling

Isolating each studio with vLANs

Pros	Cons
Each studio has its own broadcast domain; gives natural resilience to layer 2 network issues	Specialist configuration required
Provides a scalable template for further studio deployment	Switch must be capable of passing PTP between vLANs
	Multicast routing required if audio is to be shared between studios
	Studios must be in proximity for cabling

Four switches: Multi-Studio Environment

The limiting factor in the previous example is geography. Both studios must be within cable run distance of the switches. There would likely be multiple devices that all need a ST2022-7 pair of cables running to the switches. Furthermore, it is increasingly likely that the switch will not have the port count to accommodate multiple studios. The solution is to ensure that each studio has its own switch. The options for sharing vLANs or having independent vLANs still apply, but the implications of sharing vLANs changes.

In a shared vLAN topology, multicast from one studio will be flooded to the other. Placement of the IGMP querier should be understood and the election process should be enabled in case of interswitch failure. To prevent the unnecessary flooding of multicast, switches that are equipped with a multicast report flooding feature should be evaluated.

Due to the increased amount of traffic between the switches, it is advisable to use dedicated uplink ports to connect the two studios together. Often, this will mean having a 10Gbps link between the switches and 1Gbps connections to the end devices.

Due to the increased switch count and the likely increased AoIP device count, PTP awareness on the network may be required.

- Media Network A
- Media Network B
- Connect A
- Connect B
- PTP Link

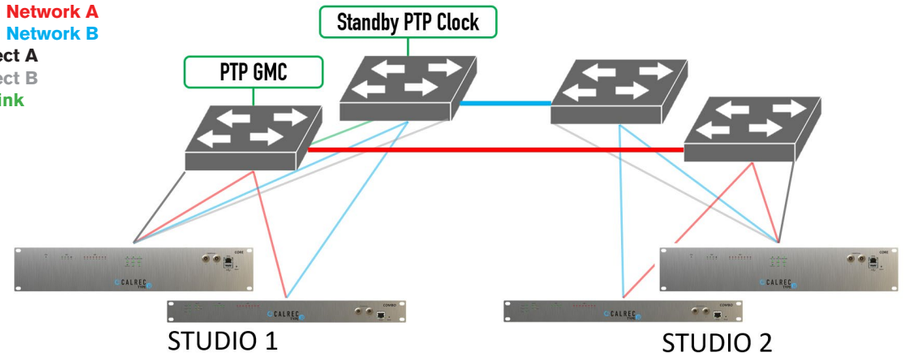

Both Studios share the same vLANs:

Pros	Cons
Operationally simple to configure	More configuration required due to increased switch count
Options for physical expansion using uplink ports	Two studios share the same broadcast domain, both are equally susceptible to broadcast storms or Spanning tree events
Configuration is the same as a single studio environment	Multicast flooding between Studios, higher bandwidth link or specific switch feature set is required

Isolating each studio with vLANs

Pros	Cons
Each studio has its own broadcast domain; gives natural resilience to layer 2 network issues	Specialist configuration required
Provides a scalable template for further studio deployment	Switch must be capable of passing PTP between vLANs
No multicast flooding	Multicast routing required if audio is to be shared between studios
Potential to not need 10Gbps between studios	

Spine and Leaf

Spine and leaf networks consist of a dedicated pair of spine switches which are used to aggregate leaf switches together. Commonly, spine and leaf networks are implemented using layer 3 routing, but can also use layer 2 switching.

The spine switch usually has a higher native bandwidth than the leaf switches. If the leaf switches at 1Gbps native with 10Gbps uplinks, then the spine should be 10Gbps native to accommodate the 10Gbps uplink connectivity from the leaf switches.

Spine and leaf networks are relatively easy to scale. As the network grows, the number of leaf switches can be increased to provide access to the network. If each leaf represents a studio, all studios are equidistant from each other providing deterministic latency.

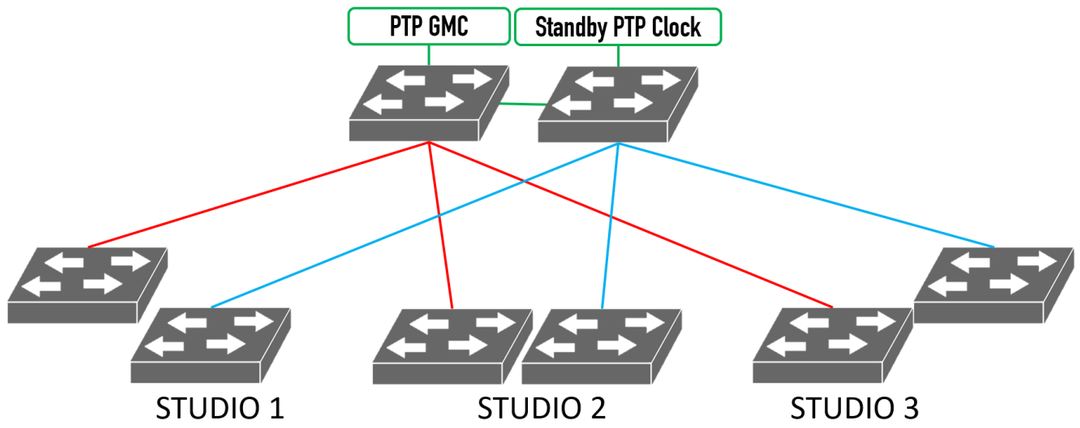

- **Media Network A**
- **Media Network B**
- **Connect A**
- **Connect B**
- **PTP Link**

Pros	Cons
Highly scalable: Studios can be added by adding leaf switches	Specialist configuration required if using layer 3
All studios are the same distance from each other in terms of switch hops	Initial configuration period can be longer
If layer3 implemented: each studio is layer 2 isolated	
Cost per studio is reduced due to less barriers to implementation	

FAQs

Which network design is the best?

Unfortunately, there is no single answer. The efficiency and appropriateness of a network falls entirely down to its use case. A complex spine leaf, layer 3 routed with sophisticated SDN control network may be a fantastic solution.

But if it's over-specified for a project, then it's a poor design due to unnecessary expenditure. Likewise, if a network is under-specified for its given deployment, it is also a poor design.

What network switches does Calrec support, and can Calrec provide configurations?

The switch choice for any project is completely defined by the scale and sophistication of the workflow. It is impossible to give a single answer. We are happy to have open discussions with our customers and make assessments based on their requirements.

Calrec can provide guidance on switch configuration but are unable to provide switch configurations for each use case. There are many variations of the same switch config which may be applicable for one customer but not for another.

Is Calrec's AoIP implementation compatible with Dante?

Yes.

For AES67 Mode, Dante is fully reliant on SAP advertisements to make its stream configuration data available to other devices as well as make other devices stream configuration data available to Dante controller.

Dante does not make manual stream configuration available in its ecosystem at this level. A third-party tool called RAV2SAP can be used to advertise Calrec's SDPs onto a Dante network via SAP.

It is worth noting that in Dante AES67 mode no ST2022-7 redundancy is available so we would recommend against this solution in a resilient broadcast environment.

For ST2110 Mode, Dante allows stream creation from SDP files manually. This allows for stream configuration data to be exchanged without the use of an external conversion tool. To use Dante ST2110 mode, a Dante Domain Manager (DDM) license must be purchased.

What are the networking implications of AoIP over a WAN or to the cloud?

There are many routes to passing audio over WAN links or into the cloud.

Leased lines are common as you can hold a service provider accountable for the safe and timely transport of data. A QoS strategy that has been developed locally may be able to be extended over the WAN depending on the service providers' offerings.

AoIP across the internet is more unpredictable. There is no guarantee of packet delivery due to the volatile nature of internet traffic. There is no control of QoS so there are limited options to improve service. VPN technology is a common method to ensure privacy and manageability of audio sources and destinations.

Cloud vendors often have solutions for expedited access to their networks. This often means establishing a direct connection to their network through edge locations.

One of the challenges with AoIP in a WAN setting is multicast. Most WAN networks are much more efficient when operating with unicast. NAT solutions can convert multicast to unicast and back again for transport over such networks.

What discovery mechanisms are available in Calrec Connect?

Currently CCP (Calrec Control Protocol) and NMOS IS-04 discovery mechanisms are supported in Calrec Connect. Users can input SDP information manually to receive streams from a network.

How much bandwidth should I provision on my network?

As a rule of thumb, audio devices will almost always only require 1Gbps connections. This is enough to produce 512 audio channels depending on packet and audio format.

Links between switches may need to be higher if there are a lot of aggregated devices on each side of the link. 10Gbps inter-switch links will usually be enough to handle the load.

The bandwidth section of this guide explains how to make much more accurate predictions of the bandwidth you will need to account for.

Can I avoid network oversubscription without an SDN solution?

Careful planning and network design can help avoid network oversubscription but may not be a guarantee. Network services such as QoS can guarantee the bandwidth is reserved for a given link, but they cannot control what bandwidth is produced and placed on that link.

Oversubscription can be avoided if the bandwidth between switches is higher than the maximum amount of bandwidth that can be generated by the end devices.

What is Calrec's default stream format?

Apart from the IP Gateway, all Calrec products ship without any configured streams. This gives the broadcaster full and unlimited control of the content that is provided to the network. When users create a stream they are presented with all the configuration parameters required. For convenience some of the fields are prefilled with:

- Channels: 8
- Sample Rate: 48000
- Codec: L24

The IP Gateway product also uses these defaults. The default IP addressing on an IP Gateway is based on the Hydra ID of the IO box.

What switches do Calrec support?

Calrec are committed to supporting any COTS switch that has the applicable feature set and performance for a given use case. A list of switch requirements and tested switches can be found in the switches chapter of this guide.

Do we need PTP-aware switches?

There are many factors in deciding if PTP awareness is required on an AoIP network. Guidelines on PTP awareness are given in the Network Design Chapter.

What is the difference between PTPv1 and PTPv2?

It is a common misconception that PTPv1 is less accurate than PTPv2. Both protocols aim for sub-microsecond accuracy.

PTPv1 and PTPv2 devices can co-exist on the same network, although PTPv2 adds extra features to PTPv1, including unicast communication and more flexible interval rates.

Transparent clocks are only available in PTPv2.

Appendix

1: Multicast MAC and IP address spaces

Mac addresses are 48 bits long (fig 24). The first 24 bits is the organizationally unique identifier. The OUI for multicast is 01:00:5E. The researchers who developed multicast had to share their address space with another user, so the remaining 43 bits were split in half by setting the MSB of the device ID to 0. Therefore there are only 23 unique bits available.

IP addresses are 32 bits long. The multicast address space reserved by IANA is 224.0.0.0/4. Therefore there are 28 unique bits available for multicast IP addresses.

There is a 5-bit difference between available MAC addresses and IP addresses; 2^5= 32. Therefore for every MAC address there are 32 correlating IP addresses.

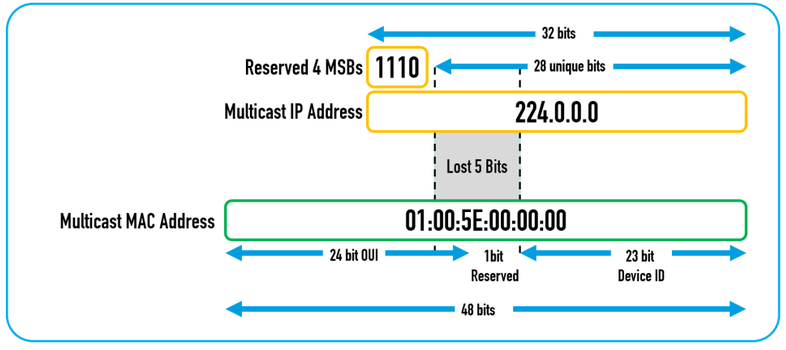

Fig 23

2: PTP interval unit conversion table

log_2	seconds	pps
-3	125ms	8
-2	250ms	4
-1	500ms	2
0	1s	1
1	2s	1/2
2	4s	1/4
3	8s	1/8

Calrec IP Primer
calrec.com

3: Bandwidth Examples

24-bit audio, 48kHz sample rate

Channels per stream	125us Packet Time		1ms Packet Time	
	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)
1	58	6.1 Mbps	184	1.8 Mbps
2	76	7.3 Mbps	328	2.9 Mbps
4	112	9.6 Mbps	616	5.2 Mbps
6	148	11.9 Mbps	904	7.5 Mbps
8	184	14.2 Mbps	1192	9.8 Mbps
10	220	16.5 Mbps	1480	12.1 Mbps
12	256	18.8 Mbps	N/A	N/A
16	328	23.4 Mbps	N/A	N/A
24	472	32.6 Mbps	N/A	N/A
32	616	41.9 Mbps	N/A	N/A
40	760	51.1 Mbps	N/A	N/A
48	904	60.3 Mbps	N/A	N/A
64	1192	78.7 Mbps	N/A	N/A
80	1480	97.2 Mbps	N/A	N/A

24-bit audio, 96kHz sample rate

Channels per stream	125us Packet Time		1ms Packet Time	
	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)
1	76	7.3 Mbps	328	2.9 Mbps
2	112	9.6 Mbps	616	5.2 Mbps
4	184	14.2 Mbps	1192	9.8 Mbps
6	256	18.8 Mbps	N/A	N/A
8	328	23.4 Mbps	N/A	N/A
10	400	28.0 Mbps	N/A	N/A
12	472	32.6 Mbps	N/A	N/A
16	616	41.9 Mbps	N/A	N/A
24	904	60.3 Mbps	N/A	N/A
32	1192	78.7 Mbps	N/A	N/A
40	1480	97.2 Mbps	N/A	N/A

32-bit audio, 48kHz sample rate

Channels per stream	125us Packet Time		1ms Packet Time	
	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)
1	64	6.5 Mbps	232	2.2 Mbps
2	88	8.1 Mbps	424	3.7 Mbps
4	136	11.1 Mbps	808	6.8 Mbps
6	184	14.2 Mbps	1192	9.8 Mbps
8	232	17.3 Mbps	N/A	N/A
10	280	20.4 Mbps	N/A	N/A
12	328	23.4 Mbps	N/A	N/A
16	424	29.6 Mbps	N/A	N/A
24	616	41.9 Mbps	N/A	N/A
32	808	54.1 Mbps	N/A	N/A
40	1000	66.4 Mbps	N/A	N/A
48	1192	78.7 Mbps	N/A	N/A
64	1192	78.7 Mbps	N/A	N/A
80	1480	97.2 Mbps	N/A	N/A

32-bit audio, 96kHz sample rate

Channels per stream	125us Packet Time		1ms Packet Time	
	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)	Layer3 IP Packet Size (Bytes)	Stream Bandwidth (Mega Bits per Second)
1	88	8.1 Mbps	424	3.7 Mbps
2	136	11.1 Mbps	808	6.8 Mbps
4	232	17.3 Mbps	N/A	N/A
6	328	23.4 Mbps	N/A	N/A
8	424	29.6 Mbps	N/A	N/A
10	520	35.7 Mbps	N/A	N/A
12	616	41.9 Mbps	N/A	N/A
16	808	54.1 Mbps	N/A	N/A
24	1192	78.7 Mbps	N/A	N/A
32	1192	78.7 Mbps	N/A	N/A
40	1480	97.2 Mbps	N/A	N/A
48	904	60.3 Mbps	N/A	N/A
64	1192	78.7 Mbps	N/A	N/A
80	1480	97.2 Mbps	N/A	N/A

4: Static IP Addressing Strategy

The following steps provide an option for managing IP addresses in an AoIP network. This can be used as a basic template, but the specifics may vary depending on requirements. The following strategy isn't an exact science and can be interpreted freely or ignored.

1. Choose a private address block to use within your network from the follow:

a. 10.0.0.0/8
b. 172.16.0.0/12
c. 192.168.0.0/16

These address ranges will not be the final subnets used for the AoIP networks. They just represent the usable address space for the entire network.

2. Assign a number to each network or vLAN. This number must be between 1 and 4096 as this will become the 802.1Q vLAN ID for each vLAN. As an example:

a. 10: Media Red
b. 20: Media Blue
c. 99: Management and Control

3. Decide on a subnet size for each network. This will be decided based on the following:

a. How many devices will be on the network plus any that a provisioned for future expansion
b. n this example, we will use 255.255.255.0 (/24)

4. Assign a range of numbers for each logical group of devices on the network. For example:

a. 1-30: Analogue audio devices
b. 60-70: playout devices
c. 150-170: Servers
d. 250-254: Network switches

The range should be big enough to account for all devices that fall into that logical group plus any amount that has been provisioned for future expansion.

5. Assign a unique number to each device from the applicable range, as decided in step 4. For example:

a. 1: Analogue device 1
b. 2: Analogue device 2
c. 3: Analogue device 3
d. 61: Playout device 1
e. 62: Playout device 2

6. Taking the framework as built in the previous 5 steps, create an IP addressing structure for all of the network devices. For example:

NICs:

192.168.10.1 /24

192.168. – Private Address block (step 1)
10. – vLAN/Subnet ID (step 2)
1 – Device ID (step 5)
/24 – Subnet mask (step 3)

Multicast:

239.10.1.1

239. – Local Multicast Range
10. – vLAN/Subnet ID
1 – Device ID
1 – Stream Number

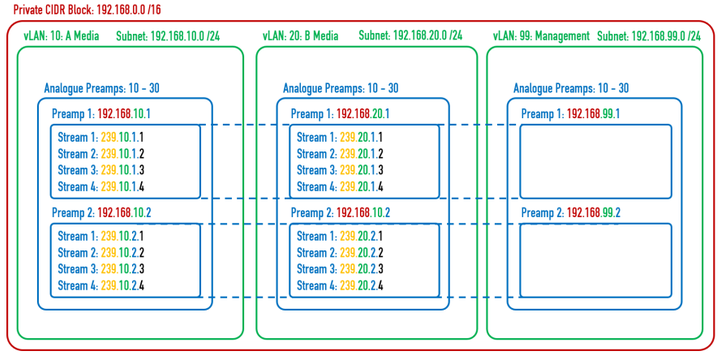

Fig 23

By taking a structured approach to IP addressing, it becomes easier for operations teams to translate IP addresses into physical devices, logical device groupings, vLANs and stream content.

As you build up an addressing schema, how the IP addressing relates to the physical network becomes more visible. Fig 24 shows how this IP addressing schema may look when documented.

In this IP addressing schema, a multicast stream can be identified by its second octet, with red streams belonging to vLAN 10 and blue streams belonging to vLAN 20. A device can be identified by the 4th octet of its IP address regardless of which vLAN the NIC is allocated to.

Addressing strategies can take many different forms. Studio numbers could be incorporated into the IP address so a device can be uniquely identified as a member of a particular studio. Operationally, there is little benefit to audio operators.

However, operations and maintenance teams will benefit hugely from a memorable IP address schema.

5. Subnet Size Comparison

Dotted Decimal Notation	Slash notation	Number of Hosts
255.255.0.0	/16	65534
255.255.128.0	/17	32766
255.255.192.0	/18	16382
225.255.224.0	/19	8190
255.255.240.0	/20	4094
255.255.248.0	/21	2046
255.255.252.0	/22	1022
255.255.254.0	/23	510
255.255.255.0	/24	254
255.255.255.128	/25	126
255.255.255.192	/26	62
255.255.255.224	/27	30
255.255.255.240	/28	14
255.255.255.248	/29	6
255.255.255.252	/30	2

6: IGMPv2/IGMPv3 Comparison

	IGMPv2	IGMP v3
General Query Address	224.0.0.1	
Group Specific Query Address	Group that is being queried	
Membership report	Group that is to be joined	224.0.0.22
Leave Message	224.0.0.2	224.0.0.22

References

JT-NM:
https://www.jt-nm.org/

AES67:
https://www.aes.org/publications/standards/search.cfm?docID=96

SMPTE ST2110-10:
https://ieeexplore.ieee.org/document/8165974

SMPTE ST2110-30:
https://ieeexplore.ieee.org/document/8167392

SMPTE ST2059-2:
https://ieeexplore.ieee.org/document/7291608

SMPTE ST2022-7:
https://ieeexplore.ieee.org/document/8716822

IEEE1588-2019: PTPv2:
https://ieeexplore.ieee.org/document/9120376

AMWA NMOS IS-04:
https://specs.amwa.tv/is-04/

AMWA NMOS IS-05:
https://specs.amwa.tv/is-05/

TR1001-1:2020:
https://static.jt-nm.org/documents/JT-NM_TR-1001-1_2020_v1.1.pdf

Ravenna:
https://www.ravenna-network.com/

RFC 3550: RTP: A Transport Protocol for Real-Time Applications
https://tools.ietf.org/html/rfc3550

RFC 2974: Session Announcement Protocol:
https://tools.ietf.org/html/rfc2974

RFC 4566: Session Description Protocol:
https://tools.ietf.org/html/rfc4566

RFC 2236: IGMPv2:
https://tools.ietf.org/html/rfc2236

RFC 3376: IGMPv3:
https://tools.ietf.org/html/rfc3376

Calrec Audio Ltd
Nutclough Mill
Victoria Road
Hebden Bridge
West Yorkshire
England UK
HX7 8EZ

Tel +44 (0)1422 842159
Fax +44(0)1422 845244
Email enquiries@calrec.com

calrec.com