


IP NETWORKING GUIDE FOR VIDEO AND AUDIO APPLICATIONS



Contents

1	Introduction	4	7.3	Communication details	22
2	Revision History	5	8	Switches	22
3	General Overview	6	8.1	Requirements	22
3.1	Bandwidth	6	8.1.1	Bandwidth	22
3.2	Latency	6	8.1.2	QoS	22
3.3	Packet Delay Variation	6	8.1.3	Spanning tree	22
3.4	Error Rate	7	8.1.4	PTP	22
3.5	Additional factors	7	8.1.5	Multicast	22
3.6	Physical Layer	7	8.1.6	Performance	22
3.6.1	Copper	7	8.2	Settings	23
3.6.2	Fiber	7	8.3	Limitations	23
3.6.3	SFP Connectors	8	8.4	Switch models	24
4	Multicast	8	9	Leased data connections	24
4.1	Overview	8	9.1	Link Types	24
4.2	Querier Election	9	9.1.1	Physical Links	24
4.3	Multicast Router port	9	9.1.2	Packet-switched Links	24
4.4	IGMP Report Flooding	10	9.1.3	Packet-routed Links	24
4.5	Unregistered multicast flooding	11	9.1.4	Analysis	24
4.6	IGMP Fast Leave	11	9.2	Bandwidth and QoS	25
4.6.1	Fast Leave – Tested Scenarios	11	9.3	Latency and Packet Delay Variation	25
4.7	Multicast Address Considerations	11	9.4	Measurement	25
4.8	IGMP Performance	12	9.5	Sample line characteristics	25
4.9	“Half IGMP” mode	12	9.5.1	LAN	25
4.10	Protocol Independent Multicast (PIM)	13	9.5.2	WAN	25
4.10.1	PIM Sparse Mode	13	9.6	Lawo device performance	25
5	PTP	14	10	Bandwidth examples	26
5.1	Overview	14	11	Specific Configurations	27
5.2	Profiles	14	11.1	Cisco SG300	27
5.3	PTP Timing Accuracy	17	11.2	Arista (PTP E2E)	35
5.4	Grandmaster models	18	11.3	Artel Video Systems ARG Quarra Switches	35
6	Quality of Service	18	12	Troubleshooting	38
6.1	Overview	18	12.1	Multicast	38
6.2	DiffServ	19	12.2	PTP	41
7	Ports and communication details	20			
7.1	Ports	20			
7.2	Commonly used multicast addresses	21			

13	FAQ.....	44
13.1	General.....	44
13.1.1	How much bandwidth is used? What connectivity do I need?	44
13.1.2	What cabling shall I use?	44
13.1.3	Can I combine multimode and singlemode SFPs?.....	44
13.1.4	What multicast scheme shall I use?.....	44
13.1.5	Do I need to worry about oversubscription? ...	44
13.2	Switches	45
13.2.1	Which switch types does Lawo support?	45
13.2.2	Which switch types does Lawo recommend?	45
13.2.3	Can you provide switch configs?	45
13.3	Redundancy	45
13.3.1	Does Lawo support SMPTE ST2022-7 redundancy?.....	45
13.4	PTP	45
13.4.1	What is the difference between the One-Step mode and the Two-Step mode in PTP?.....	45
13.4.2	Do I need PTP or can I work without it?	45
13.4.3	Which PTP Grandmasters do we work with?..	46
13.4.4	How many Grandmasters do I need in a system?	46
13.4.5	My Grandmaster does not support enough clients / how do I scale PTP?.....	46
13.5	Ravenna	46
13.5.1	What is the Payload?	46
13.5.2	Do I need to stick to the Payload Presets?.....	46
13.5.3	How does the Payload affect Latency?	47
13.5.4	What is the Network Packet Size?	47
13.5.5	What if the Network Packet Size is too big? ...	47
13.5.6	Calculating the Bandwidth of a Stream	47
13.5.7	How does the Network Bandwidth affect RAVENNA Streaming?	48
13.5.8	How does jitter affect RAVENNA streaming?	48
13.5.9	How long will it take for my audio stream to reach its destination?	48
13.5.10	How do I integrate Ravenna/AES67 and DANTE?	48
14	Glossary	49
15	Standards	50

1 Introduction

Transporting audio and video streams over IP networks creates requirements beyond the well-known settings for e.g. office networks. This guide outlines general aspects of networking as well as the specifics needed for the correct transport of real-time audio and video data over standard IP networks. Even though the guide uses Lawo devices as examples for specific settings, these are also applicable for equipment from other brands. The guide covers the following aspects:

- Generality of Service

2 Revision History

Revision	Author	Date	Changes
1.0	JKU, RS, Jo	2016-08-21	Initial Version
1.1	JKU	2016-09-06	Added section on relevant standards Added troubleshooting section
1.2	JKU	2016-09-07	Added FAQ section
1.3	JKU	2017-01-25	Clarifications on PTP, added section on bandwidths, added ARG switch config
1.4	JKU	2017-01-26	Enhancements on ARG switch config
1.5	JKU	2017-03-27	Added link to ARG manual, added remaining information from Ravenna networking guide in v1.0.3 (physical connections, SG300 switch, glossary), port usage information
1.6	JKU	2017-04-05	Clarification on fiber section, detailed descriptions on leased lines
1.7	JKU	2017-09-05	Added information on multicast addresses, changed ARG to Artel Video Systems ARG, adjusted PTP recommendations
1.8	JKU, Jo, GH	2018-01-10	Added section about PTP timing accuracy for the A__line, clarifications on “Limitations” section and “Switch models” section Added scenarios for fast leave Added section on PIM Added section on multicast address considerations
1.9	JKU, Jo, Tü	2018-07-15	Added section on WAN line types
1.10	JKU	2018-12-10	Fixed some typos, added Ravenna FAQs, maintenance on the “Switch models” section, clarifications on the “Limitations” section, added “Grandmaster models” section, extended FAQ section with input from SSE
1.11	JKU	2019-06-17	Added a section on which switches NOT to use, more concrete detail on Cisco switch models

3 General Overview

Network performance is defined by the following criteria:

- Bandwidth
- Latency
- Packet Delay Variation
- Error rate

3.1 Bandwidth

The bandwidth describes the amount of data that can be transported over the network per time unit. The usual unit is Mbit/s or Gbit/s.

When bandwidth is specified it is not always clear whether the bandwidth given is the net bitrate or the gross bitrate, but for Ethernet the data rates given are the net bitrates. For example, 100BaseT Ethernet provides 100Mbit/s usable data rate, while the amount of data on the physical connection is 125Mbit/s including the coding necessary to transport the data safely.

The bandwidth for a given technology is fixed. There is no way of increasing the bandwidth of e.g. 100Mbit/s Ethernet, except by changing to another technology such as 1000Mbit/s Ethernet (you can combine multiple links of a given technology to increase the available bandwidth, but that comes with other drawbacks).

For networks dedicated to media transport the bandwidth is determined by the bandwidth each media stream has and the amount of media stream that need to be transported over a single link. E.g. 1 HD-SDI signal has a bandwidth of 1.485Gbit/s. Encapsulated as SMPTE ST2022-6 the bandwidth increases to 1.57Gbit/s, so if you need to transport 3 HD video signals, the link needs to support at least 4.71Gbit/s. Leaving aside 5Gbit/s Ethernet, the choice in this case would be a 10Gbit/s Ethernet connection.

3.2 Latency

Latency describes the time that the information needs to travel from source to destination.

The most basic foundation is the speed of light (roughly 300'000'000 meters / second in vacuum); no information can travel faster. Depending on the medium used to transport the information, the time is longer. E.g. in optical fibers the index of refraction is 1.5, meaning that the light travels 1.5 times slower

than in a vacuum. That results in approximately 5 μ s of latency for every kilometer of fiber.

Of course latency is added for other elements in the data's path as well: packetizing of data, encoding, queueing in active network elements, etc.

For networks dedicated to media transport you want to minimize latency as latency ultimately translates into delay between the actual event and the representation of the event on screens or speakers. Since the distance between source and destination is usually fixed, only the processing on the path can be influenced, e.g. by limiting the amount of active network elements in the data path and limiting the processing in source and destination.

3.3 Packet Delay Variation

Packet delay variation ("PDV") is a measurement for the difference of the one-way, end-to-end delay of packets. Sometimes this is also referred to as "network jitter".

In an ideal network all packets would take the same time to travel from the source to the destination, but in real networks various factors cause this time to vary. When the packets are handled by an active network element such as a switch, the processing of the packets depends on the processing load of that switch. The load is mainly related to the number of concurrent packets to process and the complexity of the processing. Due to the "bursty" nature of data transport in Ethernet networks, the load varies quickly and thus influences the packet delay variation.

Packet delay variation can only be counteracted by adding buffers at the receiving end. The incoming data is first written to a buffer. Once the buffer is filled with an amount of data which can compensate the longest packet delay variation that you expect on the network (plus a little safety), the receiving device can start reading data from the buffer using a constant rate.

For networks dedicated to media transport you want to minimize packet delay variation, because the buffers you need to add in order to compensate for the packet delay variation add to the unwanted latency in the signal path. This is usually done by limiting the amount of active network elements in the data path and using techniques like Quality of Service (QoS) to prioritize the processing of data packets carrying real-time media over other traffic.

3.4 Error Rate

The error rate describes the amount of data that has been altered on the path from source to destination. In data networks this is usually related to corrupted packets (bit errors) or lost packets.

Bit errors are usually compensated by adding some redundancy to the data being transported, often in the form of error correction data transported on the physical transport layer in addition to the actual data. Only errors that cannot be corrected by these measures will be noticeable to the user of the network and will need to be handled by higher protocol layers.

Since bit errors happen randomly, the error rate describes a probability in the form of percentage of packets likely to be affected by errors.

Lost packets are usually caused by overloading one or multiple network elements in the data path. E.g. if two sources try to send 1Gbit/s each to a destination connected to the network with a 1Gbit/s connection only half of the packets can actually be forwarded, the rest of them needs to be dropped.

For networks dedicated to real-time media transport you want low error rates. Bit errors can be compensated by higher protocol layers, e.g. by adding redundancy to the data transported which allows reconstructing the original data, even if some packets are corrupted or lost on the transport.

Lost packets caused by overloaded paths in the network can only be handled by careful network design and management, including but not limited to correct choice of bandwidth and prioritization of traffic using Quality-of-Service.

Another method of dealing with lost data for video streams on the application layer is concealment: replacing missing data from previous video frames or from another part of the same frame. However, this is only the last resort and it is preferable to ensure that all packets are arriving at the destination.

Lost synchronization information can be compensated by the “flywheel” design (keeping the last known speed until new sync information is received).

3.5 Additional factors

Another factor to be observed is packet reordering: When data is transported over different network paths, it can happen that a packet that has been created first arrives at the destination after packets created later. In order to correctly process the data from the packets, these packets need to be sorted back into order before processing. Allowing for packets to be re-ordered also necessitates a larger buffer, causing more latency in processing.

In more complex networks it can also happen that a packet is duplicated and arrives twice at the destination. The device reading the data from the network and the protocols need to provide measures to identify and discard duplicate packets.

3.6 Physical Layer

The physical connections for Ethernet can either be copper cables or fiber (optical) cables.

3.6.1 Copper

Copper connections are established using twisted-pair cables, which are available in different categories. Gigabit Ethernet requires at least category 5e while category 6 cables (or higher) are recommended for longer distances.

The connectors that terminate the copper cables are called “RJ45”.

The maximum distance for Ethernet with copper cables is defined as 100m (328ft) in the Ethernet standard. However, the distance achievable depends on the quality of the cable, the quality of the connectors and the number of connectors. It is therefore recommended to plan for a cable length of no more than 70-80m.

If you need to cover more distance, active network elements can be used to refresh the signal (e.g. a switch), but the use of optical cables is highly recommended.


While there is the possibility to transport 10Gigabit Ethernet over copper cables, it is common to use fiber cables.

3.6.2 Fiber

Fiber cables exist of two varieties, multimode and single (or mono) mode, and come with a wide range of connectors. The multimode cables are suitable for distances up to 550m with 1Gbit/s and

up to 300m with 10Gbit/s. Single-mode fibers are suitable for distances up to 80km with 10Gbit/s or even higher.

Please be aware that any fiber connection longer than 10km will require individual tuning of the fibers and the lasers to cater for the exact transmission power levels.



3.6.3 SFP Connectors

Many devices today offer a choice of connector, by providing an “intermediate” connector called the “small form-factor pluggable” (SFP) or its enhanced version for higher data rates SFP+.

The device side of the SFP connector is standardized while the connector side can provide copper connections (on RJ45) or different types of fiber connections (with LC connectors).

When choosing an SFP, make sure pick the correct type:

- Data rate: 100Mb/s, 1Gb/s, 10Gb/s
- Connector: Copper with RJ45 connector, Fiber with LC connector
- Fiber type and distance: Multimode (up to 550m for 1Gbit/s, up to 300m for 10Gbit/s) vs. Single-mode (1km, 10km, 40km, 80km)

For short distances (~ up to 30m) SFP modules with direct attached fiber cables exist (AOC – Active Optical Cable).

While there are copper cables with SFP connectors we strongly advise against using them as they have proven to cause interoperability issues. Use AOCs instead. Some devices (especially switches) might require approved SFPs and will reject non-approved ones.

4 Multicast

4.1 Overview

The vast majority of communication in traditional IT networks runs as unicast, meaning that messages are always sent from a single source to a single destination. While this method is perfectly sufficient to transport video and audio signals between two devices it would mean doubling the amount of data being sent if one source were to send the same video to two destinations; so, this method does not scale if applied to a traditional broadcast infrastructure where a router is able to distribute a single source to as many outputs as the router is equipped with.

To alleviate this limitation, data transmission with a multicast addressing scheme is used, which allows sending data to a group of destinations without the need for the source to send the data multiple times. The actual distribution of data to the destinations is handled by the network elements (switches and routers) connecting the source with the destinations: the source sends the data once and the switch will duplicate the packet for each destination.

If the switch does not know which destination is interested in which source, it needs to duplicate the source packets on every port, thereby turning the multicast into a broadcast. This is very inefficient in terms of bandwidth usage and does not scale well.

To orchestrate multicast in a more efficient way, the Internet Group Management Protocol (IGMP) is used. A destination interested in a certain stream will send a “Join” message. Switches can listen to IGMP communication (“IGMP snooping”) and use this information to selectively duplicate the data stream to the port on which the destination is attached. The destination becomes a “member” in the requested “multicast group”. The switch keeps track of all members of a multicast group and sends periodic queries (“membership queries”) whether they are still interested – the switch acts as an “IGMP Querier”. If a destination does not reply in a specified time, the switch assumes that the destination is no longer interested and stops duplicating the packets. In more recent versions of IGMP, the destinations can also send “leave” messages, actively informing the switch about not wanting the multicast data anymore (see below).

IGMP Snooping and IGMP Querier are features not found on all switches but are needed for a correct and performant network

function in the realm of audio / video networking. The V__line switch has both of these features built-in.

When referring to multicast people usually talk about the layer 2 Ethernet multicast, limiting the data distribution to a single subnet. In larger setups, when the layer 2 domains get too large, a separation into several smaller networks might be desirable. Every network, usually kept in its own VLAN, needs its own IGMP Querier and IGMP snooping to be enabled. To transfer data between the different subnets, IP routing is necessary. This is done by a dedicated router or a routing-enabled switch. In case of media data using multicast, a more specialized form of routing is necessary: PIM. This feature can only be found in carrier-grade / enterprise-grade switches or dedicated routers.

In setups where a more rigid and deterministic control is necessary, other means of routing multicast data might be necessary to ensure correct audio and video distribution, such as an application layer-based control of the audio and video streams (sometimes termed an “Software Defined Network” or “SDN”).

Multicast uses a reserved range of IP addresses to identify a multicast group (224.0.0.0 to 239.255.255.255 for IPv4). Some of them are reserved for specific protocols or purposes (e.g. 224.0.0.1 for all systems on this subnet or 224.0.0.22 for IGMP messages). The range of 239.0.0.0 to 239.255.255.255 can be freely used.

4.2 Querier Election

There will be only one querier (one device managing multicast group information) per subnet / VLAN. This querier is elected between all the devices capable of fulfilling this task.

When such a device starts, it multicasts a general query to all other systems to the 224.0.0.1 address using its own (unicast) IP address as source address. When a device receives such a query it compares the source IP from the message with its own IP address. The device with the lowest IP address is elected to be the querier for the subnet.

All other devices start an internal timer which is reset every time they receive a general query from the querier. If those messages cease, a new querier election takes place after the timer expires.

In larger installations the querier role is usually taken by a multicast capable router, which handles the data exchange to and from other subnets.

4.3 Multicast Router port

Network elements need to differentiate between ports to which hosts are attached and ports to which other multicast-aware network elements are attached. Ports to which hosts are connected receive only those multicasts which they have explicitly requested using IGMP “join” messages. Ports to which routers are connected commonly receive all multicasts, so that they can forward the multicast traffic to other network segments.

Switches receiving general IGMP queries on a specific port assume by default, that this port is connected to a multicast router that is interested in all multicast traffic. This port is generally termed a multicast router (“mrouter”) port. Since the switch has no knowledge about the group membership of hosts in other subnets, it simply sends all the local multicast traffic to all mrouter ports (also referred to as “multicast data flooding”). In media applications, especially with high-bandwidth video streaming, this behavior might be undesirable, since it can easily lead to congestion on the mrouter ports.

Since the V__line contains its own switch, many switches will recognize the V__line as a multicast router and flood all multicast traffic to the V__line.

To counteract this behavior three options exist:

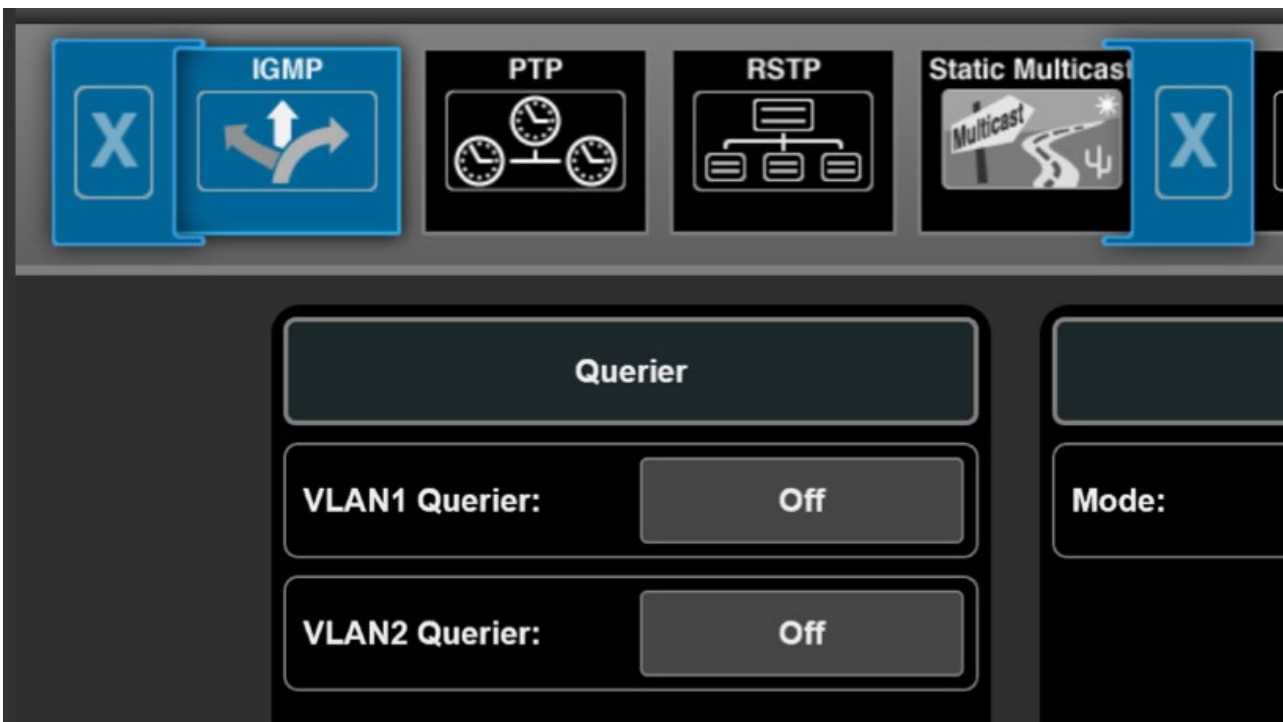
- disabling the mrouter port learning on the switch
- disabling the multicast data flooding on the switch (aka “report flooding”, see below)
- disabling the V__line’s IGMP querier function

Disabling the mrouter port learning alleviates the issue above, but also necessitates that you manually need to assign mrouter ports where you want the multicast traffic to be forwarded without explicit IGMP “join” messages.

Disabling the V__line’s IGMP querier function leads to the V__line no longer sending general IGMP queries and thus not being recognized as a multicast router anymore. This can be done without

negative side-effects, provided the network has another active IGMP querier (e.g. the switch).

For switches that support report flooding, we recommend using this functionality. For other switches we recommend disabling the V_line's IGMP querier function ("Settings" > "Switch" > "IGMP"):



4.4 IGMP Report Flooding

To overcome the restrictions that might occur from multicast router port learning and the multicast data flooding, an alternative port mode for inter-switch links is available on some switches: IGMP report flooding.

If a port is in this mode, automatic multicast router port learning is disabled and multicast data flooding is inhibited as well. A port will be included in a multicast group only, when IGMP snooping has detected an IGMP membership report for it. Whenever a locally attached node sends an IGMP membership report, it is reflected to all ports in IGMP Report Flooding Mode. This allows the neighboring network device to correctly detect memberships across switches without flooding multicast data unnecessarily.

This mode leads to slightly higher IGMP traffic and IGMP snooping workload in a layer 2 domain, but effectively avoids congestion on inter-switch links.

IGMP Report Flooding Mode is the preferred mode in a medium sized, flat hierarchy layer 2 topology. The V__line built-in switch supports IGMP Report Flooding as well.

4.5 Unregistered multicast flooding

The default behavior of most switches with multicast traffic not being registered by IGMP join or membership reports, is to deliver the traffic to all ports, thus turning it into a broadcast. This is not acceptable when predictive and well-controlled bandwidth allocation is needed. Media applications must be capable of properly speaking IGMP to achieve stable multicast network operation.

Use the switch configuration to turn off unregistered multicast flooding (this feature is also referred to as “Unregistered Multicast Filtering”).

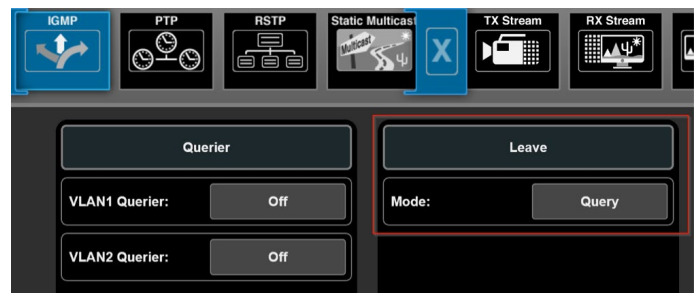
4.6 IGMP Fast Leave

IGMP Fast Leave was introduced with version 2 of IGMP. It refers to the capability of a host to explicitly announce that it is no longer interested in a specific multicast. The host sends out an IGMP “leave” message and the network immediately stops the delivery of multicast traffic. In version 1 of IGMP there was no leave message and a timeout was required to determine that a host no longer wanted to consume a specific multicast.

The advantage of fast leave is that the multicast traffic ceases immediately and does not consume any more network resources. It allows quickly re-using the bandwidth for another stream.

Since the implementations of the fast leave feature are different across device vendors, our recommendation is to verify correct functionality of fast leave in a concrete system setup before using it in production. Alternatively disable fast leave in the complete system (sources, switches and destinations).

For the V__line the respective settings can be found in “Settings” > “Switch” > “IGMP”:



4.6.1 Fast Leave – Tested Scenarios

Scenario	Result
V__lines (SW 1.4.0.106) connected to Arista 7150 (SW 4.20.3), Layer2, Immediate Leave activated on all devices	OK

4.7 Multicast Address Considerations

In the early days of the internet the address block 224.0.0.0/4 was reserved for multicast and designated “class D”.

This range has been split into multiple smaller ranges as outlined in RFC 5771, e.g. the range of 224.0.0.0 to 224.0.0.255 was reserved for use in local subnets and is used today in routing protocols such as OSPF.

The range of 239.0.0.0/8 has been reserved for private use within an organization (RFC 2365) and is commonly used and recommended for A/V streaming.

One additional thing to be aware of:

In the layer-2 domain, every multicast address is translated into a destination MAC address. As of design, the first half of the MAC address for multicast IP addresses is fixed to 01:00:5E (by OUI / IEEE; RFC 1112). The other half of the MAC address (the remaining 3 octets or 24 bits), will be derived from the last 24 bits of the multicast IP address.

Of these last 24 bits, the least significant bit of the first octet is always set to “1”, leaving only 23 bits which are then taken from the IPv4 multicast address (RFC 7042). This causes different multi-

cast addresses to be translated into the same MAC address: “1000 0000” and “0000 0000” for the octet are the same; resulting in e.g. 228.0.1.1 and 233.128.1.1 being translated into the same MAC address.

This behavior needs to be considered when creating a multicast address schema in order to avoid collisions.

4.8 IGMP Performance

The performance needed depends on the system size and expectations of the users. Usually switches process IGMP requests serially, so if 100 Joins messages need to be processed and each Join request takes 10ms, the 100 requests take a total of 1 second.

Looking at the emerging standards such as VSF TR-03 / SMPTE ST-2110 the goal is to separate the essence streams into video, audio and metadata. Assuming each audio stream carries 4 channels, the representation of each SDI signal results in 1 video stream, 4 audio streams (4x 4ch) and 1 metadata stream (VANC data, such as captions), a total of 6 streams in need to separate handling with IGMP.

Equating this to a normal baseband router of 500x500 signals, assuming that 25% of the signals need to be switched simultaneously at peak times and that the users' expectation is that switching signals takes 5 frames at maximum, this results in:

- 125 signals to be switched (25% of 500 signals)
- 750 resulting streams (6 streams per signal * 125 signals)
- Total time available for switching: 100ms (assuming 50fps)
- Many of the signals switched will require a “Leave” and a “Join” command, effectively doubling the number to 1500 requests for 750 streams

With 1500 requests to be processed in 100ms that leaves 67µs per request, if the requests are in fact processed sequentially.

The total time needed to establish a new essence connection needs to also take into account how much time the source and the destination devices need to create, provide and consume the new stream.

Looking at those numbers and being aware that the latency of IGMP message processing increases when using a network with

multiple cascaded switches, it becomes apparent that solely using IGMP to control the essence flows inside a media network will only be practical for limited size networks. Other control mechanisms will be needed to allow building bigger systems.

4.9 “Half IGMP” mode

As outlined above the total time taken to establish a new multicast data stream is composed of the time it takes for the IGMP message processing, the time for the source to setup the new multicast stream and the time for the destination to consume the stream.

Some source devices, such as the Lawo V__line, allow to “force” sending the multicast data to the switch, despite no one having requested the specific multicast data yet. This reduces the amount of time to establish the multicast data stream, since the source is already sending.

This behavior has no adverse side effects, provided the switch is configured correctly, as the switch will discard the incoming multicast data stream until a destination actually requests it.

To activate half IGMP mode and force the sending of streams irrespective of IGMP “join” messages, set the Stream Mode to the respective interface(s) in “Settings” > “TX Stream” > “RAW SDP” (or “J2K SDP”, “Ravenna SDP”, et. al.):



This can be done on a stream-by-stream basis.

4.10 Protocol Independent Multicast (PIM)

PIM is used to route multicast traffic across layer 3. It does not use its own routing tables but relies on information provided by routing protocols. That means that routing needs to be established, e.g. via OSPF or BGP protocols, in order for PIM to work.

There are four variants of PIM:

- Sparse Mode
- Dense Mode
- Bidirectional PIM
- PIM Source-Specific Multicast

Sparse Mode is the most commonly used mode in broadcast currently.

4.10.1 PIM Sparse Mode

PIM uses the concept of a Rendezvous Point (RP) to collect information about available multicasts in the network. This Rendezvous Point is a service running on one or multiple routers in the network. It will be informed about any multicasts present in the network and will be asked for the location of a multicast if someone wants to receive it:

As soon as a source host provides a multicast, the router to which the host is connected will inform the RP about this multicast using a PIM Register message. A destination host that is interested in the multicast, sends an IGMP Join message. The router to which the destination host is connected, sends a PIM Join message to the RP. The RP will start forwarding the multicast traffic to the destination host via the destination router – that data flows through the RP (the path in the network is referred to as “Root Path Tree” or “shared distribution tree”). If the routing path through the RP is not the optimal path for the underlying network, the destination router, having learned the source address of the multicast, creates the best route and subsequently start receiving the multicast via the new route (the new path through the network is referred to as “Source Path Tree”).

The changing of paths might lead to interruptions in the data flow and should be avoided by network design and configuration.

In a spine/leaf topology the RP is usually placed as an anycast RP on all spine switches to provide redundancy.

5 PTP

5.1 Overview

Precision Time Protocol (PTP) is a way of synchronizing clocks within a computer network. Correctly implemented it can achieve a clock accuracy in the sub-microsecond range and is suitable to synchronize media streams. Version 2 is applicable today and has been standardized as IEEE1588-2008.

PTP time uses the same epoch Unix time uses (00:00 on 1970-01-01), but is based on International Atomic Time which is not adjusted for leap seconds.

PTP uses a master-slave approach in which all master-capable devices elect the best master, called the grandmaster, following a common algorithm (“best master clock algorithm”, “BMCA”). To ensure accurate time at the client, the delay caused by the time it takes the packets to travel between master and slave is measured and compensated for in a continuous adjustment process.

PTP usually uses multicast to distribute time information, even though version 2 of the standard extended the protocol by an option for unicast transmission.

Generally, PTP can be run on any network. Switches that are unaware of PTP treat it as regular multicast traffic. This allows PTP to function, but decreases the clock accuracy since the switch itself introduces variable delays in packet processing which are not accounted for. Using switches without support for PTP should be avoided.

Some switches support PTP in order to increase clock accuracy. They can work in two modes: transparent and boundary clock.

The switches running in transparent PTP mode measure the time the PTP packets spend travelling through the switch and add a compensation value into the packets to account for that time. This increases the accuracy of PTP time. In this mode all slaves communicate directly with the master, which – depending on the amount of slaves – might cause a substantial amount of load on the master.

In boundary clock mode the switch participates in the master election process and presents itself as PTP master on any switch port that does not have “better” master attached to it. This mode

alleviates much of the processing load from the grandmaster (now only the switch request time directly from the grandmaster) and thus allows for bigger networks with more PTP clients.

With many participants in the synchronization process the amount of messages exchanged can increase substantially when using multicast without boundary clock, especially when using SMPTE ST2059-2 timing with 8 delay request messages per second: every PTP slave will see every other PTP slaves’ delay request messages and delay response messages from the grandmaster as they are all subscribed to the same multicast addresses. To avoid this, the delay requests and the corresponding delay responses from the master can be exchanged in unicast mode. This communication scheme is then often called “hybrid” mode. If the master is able to handle the amount of concurrent client requests from the slaves, this can be an alternative to a switch with boundary clock.

The V__line has been successfully tested as grandmaster with approximately 200 V__lines acting as PTP slaves in hybrid mode.

PTP will be used as a complement and later a replacement of video reference and word clock signals. The accuracy that can be achieved with a well-tuned system is perfectly sufficient for phase-accurate audio and video applications.

If the switch implements boundary clock mode correctly, we recommend using boundary clock. For smaller networks transparent clock is a valid option as well. For medium sized networks the hybrid mode can be used.

5.2 Profiles

AES-R16-2016 proposes the PTP settings for interoperability between the existing profiles. The first option includes the default Peer-to-Peer Profile as specified in IEEE1588-2008, the Media Profile as specified in AES67 and the SMPTE Profile as specified in ST2059-2.

The second option only includes the latter two (AES67, ST2059-2).

The following tables outline the parameter values for both proposals:

Parameter	Proposed Value	Comment
Domain	0	
Announce Interval	1	= 2 seconds
Announce Receipt Timeout	3	
Sync Interval	-1	= 0.5 seconds
Min Delay Request Interval	0	= 1 second

(Interoperability between IEEE1588-2008 Default Profile, AES67 Media Profile and SMPTE ST2059-2 Profile)

Parameter	Proposed Value	Comment
Domain	0	
Announce Interval	0	= 1 seconds
Announce Receipt Timeout	3	
Sync Interval	-3	= 0.125 seconds
Min Delay Request Interval	-3	= 0.125 seconds

(Interoperability between AES67 Media Profile and SMPTE ST2059-2 Profile)

Parameter	Proposed Value	Comment
Domain	127	
Announce Interval	-2	= 0.25 seconds
Announce Receipt Timeout	3	
Sync Interval	-3	= 0.125 seconds
Min Delay Request Interval	-3	= 0.125 seconds

(SMPTE ST2059-2 Profile)

Please keep in mind that in order for PTP to work correctly, all PTP masters AND slaves must use the same values.

As the new standard SMPTE ST2110 specifies the use of the SMPTE ST2959-2 profile, we recommend the use of this profile as well. If integrating with (older) audio devices, we recommend using the interoperability profile between AES67 and SMPTE ST2059-2.

Since the ST2059-2 profile uses an “aggressive” timing, the use of a high performance grand master, hybrid mode and / or boundary clock is recommended for larger systems. For a more detailed test of PTP scaling read the SMPTE Presentation “Large scale PTP: How big can it get?” by Nicholas Ciarleglio (Arista), Thomas Edwards (FOX) and Robert Welch (Arista).

For the V__line the PTP values can be set in the “Settings” > “Switch” > “PTP” menu:



For the A__line these values can be set in the “PTP” settings:

PTP Properties

Domain	127
Prio1	128
Prio2	128
Announce interval	-2 (1/4 sec)
Sync interval	-3 (1/8 sec)
Slave only	<input checked="" type="checkbox"/>
Delay mechanism	E2E
DSCP	56

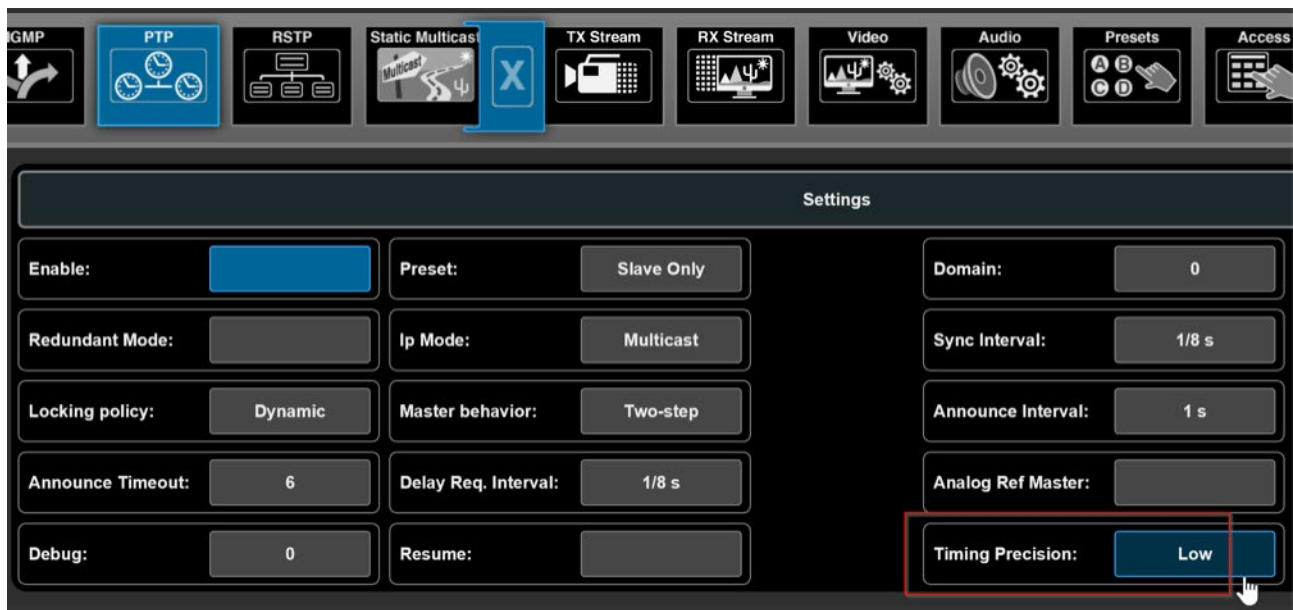
Changes will not take effect until after the device is rebooted.

5.3 PTP Timing Accuracy

The time distributed in the network needs to be accurate enough to synchronize essence streams including phase information. Therefore, it is essential that the PTP clients do not deviate too much from the PTP grandmaster's time.

With a good grandmaster and a PTP-aware network you should be able to reach an accuracy of $\pm 1\mu\text{s}$. An accuracy of $\pm 2\mu\text{s}$ is still considered sufficient.

The V__lines allow to measure and display these offsets and can be set to respect either accuracy definition ($\pm 1\mu\text{s}$ = Timing Precision High; $\pm 2\mu\text{s}$ = Timing Precision Low) or even create custom one. These settings can be found in "Settings" > "Switch" > "PTP":



Values of $\pm 5\mu\text{s}$ offset have been tested successfully between 2 V__lines without impact to the video on the output. Results with other devices may vary.

6 Quality of Service

6.1 Overview

By default, all packets are treated equally in a switch: whichever packet comes first is processed first. However, when it comes to requirements of audio / video networking, not all packets are equal. Packets transporting e.g. PTP time information are very sensitive to variations in latency, so they should be transported with the minimum amount of packet delay variation possible.


The ability to treat packets identified by certain criteria differently is referred to as Quality of Service. It allows a switch to decide to process packets with e.g. PTP information before processing packets with e.g. FTP data.

In order for the switch to process packets with different priorities, a switch has queues on the output ports which can be served according to different algorithms. Many switches have a total of 4 or 8 different queues per port and use a fixed ascending priority, the queue with the highest number having the highest and the one with the lowest number having the lowest priority. This means that the switch processes all packets from queue 8 first, then all packets from queue 7 and so on (“strict priority queuing”). This algorithm is simple and straightforward but might lead to “starving” of the lower priority queues.

Other switches assign a “weight” to the queues (e.g. 25% to queue 7, 15% to queue 6 and 10% for each of the lower queues) which makes sure that higher priority queues are served with preference, but lower priority queues don't starve. On enterprise level switches it is even possible to further tune the behavior of QoS according to an application's individual needs. Some switches also allow a combination of both queuing methods, assigning strict priority to some queues and weights to others.

It is good practice to leave the highest priority queue for essential network traffic, which is needed for the network itself to function. PTP is a traffic class that should run with highest priority in the network due its sensitiveness to packet delay variation. Other traffic that needs priority should be distributed to the queues below. Any traffic assigned to the lowest priority queue is transported as “best effort”.

In the A__line the timing precision is displayed as a color-coded graph on the “Ravenna” page.



Green means the offset is below 1µs, yellow means between 1µs and 5µs and red means above 5µs.

1µs translates into $\pm 5\%$ of media a clock period (at 48 kHz), which is the allowed phase tolerance for digital outputs specified in AES11-2009. 5 µs translates into $\pm 25\%$ of a media clock period (at 48 kHz), which is the required maximum tolerance for digital inputs of connected downstream equipment. This usually won't be noticeable with a common slave device. A value beyond 5µs may indicate a potential problem in certain critical applications. External MAD1 signals might not work together with the network-based signals. This has been chosen to become a red indication.

However, if there is only analogue audio signal exchange or if a connected audio device derives its own synchronization from the network device, timing precision even worse than 5 µs won't produce audible artifacts.

Noticeable issues, such as sample slip, may arise if the offset is larger than 10 µs, representing more than $\pm 50\%$ of a sample period.

5.4 Grandmaster models

The following PTP Grandmasters have successfully been used in Lawo deployments:

- Tektronix SPG8000A
- LAW0 V__link4/remote4 (no absolute time reference)
- Meinberg M1000 / M3000 (with the IMS-HPS100 card)
- Evertz 5601 MSC
- Trilogy Mentor

It is important to understand that the whole concept of QoS is no remedy for too little bandwidth on a link: If you try to transport 2GB of data through a 1GB link, packets will be dropped, irrespective of QoS. QoS will just help the switch to decide which packets to drop in case of bandwidth contention.

In common IT applications QoS is used to prioritize VoIP telephony data over the other data transported on the same network since VoIP data is very sensitive to latency and packet loss. The same applies to audio and video data in networks: it is sensitive to latency and packet loss and – in the case of video – often very bandwidth-intensive.

In order for the switch to sort the packets into the different queues, the switch must know the intended priority for each packet. This is achieved by marking the packets and evaluating these marks to decide on the correct queue (a process called “mapping”). One way to perform the marking is called Differentiated Services (“DiffServ”).

6.2 DiffServ

With DiffServ packets can be marked with a value for packet classification. This value is called the Differentiated Services Code Point (DSCP). It allows for 64 different values and is assigned to the IP packets. As a rough approximation you can say that the higher the value, the more important the packets is. However, in standard practice only a few of the available 64 values are actually used.

An important thing to understand is that these values, if they are set by the device sending the data, are often considered a “wish” concerning the processing priority and can be redefined on the network path to the destination; that means that setting a specific DSCP value does not mean that the packets are treated preferentially along the complete network path.

A switch evaluates the DSCP values in the IP packets, classifies traffic into different categories and maps them to its internal queues according to a user defined policy. In order for the complete path to provide the same (preferential) handling of packets, the complete path needs to be under the users’ control.

Lawo devices use different values to distinguish packets:

- All PTP packets are marked either with “CS7” (56) or “EF” (46) [AES67 uses “EF”]
- All packets containing audio / video data are either marked as “EF” (46) or “AF41” (34) [AES67 uses “AF41”]
- All other packets are marked with “BE” (0)

The switches processing these packets should be setup to handle the packets according to those priorities and treat the remaining, unmarked traffic as “best effort” (DSCP value = 0).

When leasing network lines on which you have no direct influence on the DSCP evaluation and processing, talk to the provider to ensure correct processing and check the line prior to usage (e.g. using the V__line sounding feature). During those tests transport audio and video streams in parallel to e.g. FTP traffic to verify that the audio video streams will indeed receive preferential treatment. No audio or video dropouts should occur even when the leased line is fully utilized by the streams and the FTP transfer.

7 Ports and communication details

7.1 Ports

Lawo devices offer their services on different ports. The following is a list of network ports used by Lawo devices and their purpose. If you need to access a specific service, make sure that the access to the relevant port is available through all of the network in between (firewalls, etc.).

Device	Port	Service
V__pro8 / V__link4 / V__remote4	80	Web UI
	9000	Ember+
A__mic8 / A__digital8	21	telnet
	22	SSH
	80	Web UI (Landing Page)
	5060	SIP
	8050	Web UI (Audio Device Settings, GPI/O, Ember+ Web Access)
	8081	Web UI (Streaming Parameters)
	9000	Ember+
	9009	Ember+

7.2 Commonly used multicast addresses

The following tables lists some commonly used multicast addresses and the default multicast addresses of some Lawo devices.

Multicast Address	Description
224.0.0.1 / ff02::1	Addresses all hosts in the subnet
224.0.0.2 / ff02::2	Addresses all routers in the subnet
224.0.0.251 / ff0x::fb	Multicast DNS (mDNS)
224.0.1.129-132	PTP version 1
224.0.1.129 / ff0x::181	PTP version 2
239.255.255.255	SAP Announcements (for administrative scope 239.0.0.0/8)
239.0-5.x.y	V__line Ravenna Audio Streams (x = third octet from device's unicast address, y = fourth octet of device's unicast address)
239.16-19.x.y	V__line J2K Video Streams (x = third octet from device's unicast address, y = fourth octet of device's unicast address)
239.25-28.x.y	V__line RAW Video Streams (x = third octet from device's unicast address, y = fourth octet of device's unicast address)
239.68-71.x.y	V__line MJPEG Video Streams (x = third octet from device's unicast address, y = fourth octet of device's unicast address)
239.w.x.y	A__line Ravenna Audio Streams (w = Index of Stream, x = third octet from device's unicast address, y = fourth octet of device's unicast address)

The stream addresses used in the Lawo devices can be changed in the respective devices' configuration to match an existing multi-cast scheme. Please consult the devices' manuals for details.

7.3 Communication details

The device discovery mechanisms offered by e.g. the A__mic8 make use of the Bonjour protocol and thus only work in a single network segment. An mDNS repeater can be used to exchange this information across network segments. Please consult your network administrator for details.

Some devices or combinations of devices might require direct connections or have additional restrictions (such as a flat layer 2 network). Please consult the devices' manuals for details.

To make use of all devices' capabilities, more protocols (e.g. NTP, Syslog, etc.) might need to be allowed through the firewalls. The exact setup depends on your network topology; consult with a Lawo representative and your network administrator for details.

8 Switches

8.1 Requirements

8.1.1 Bandwidth

The switch needs to support the bandwidth needed for the respective application. Lawo recommends

- 1GB Ethernet for J2K video applications and for pure audio applications
- 10GB Ethernet for video applications and mixed audio / video applications, including 4K using VC-2 compression
- 40GB Ethernet for the V__matrix, also supporting uncompressed 4K video applications

8.1.2 QoS

The switch must support DiffServ (RFC2474) and traffic prioritization according to IEEE802.1p.

8.1.3 Spanning tree

The switch must support Rapid Spanning Tree (RSTP) or Multiple Spanning Tree (MSTP) as defined in IEEE802.1w or IEEE802.1s respectively.

8.1.4 PTP

The switch should support PTP as defined in IEEE1588-2008 with at least E2E Transparent Mode, preferably also Boundary Clock mode.

The switch shall also support settings for all necessary PTP parameters to comply with the SMPTE ST2059-2 and AES-R16-2016 profiles.

8.1.5 Multicast

The switch must support multicast traffic (RFC 1112)

The switch must support multicast forwarding

The switch must support IGMPv2 (RFC 2236)

The switch must support IGMP snooping (RFC 4541)

The switch should support report flooding or the manual configuration of mrouter ports

8.1.6 Performance

The switch must meet the following performance characteristics:

- Non-blocking: The switch must use a non-blocking architecture, meaning that the switch internal forwarding capacity is equal to the capacity of a single port times the number of ports (e.g. a 48 port 10GbE switch needs $48 \times 10\text{Gb} \times 2$ (full duplex) = 960Gb/s switching capacity)
- Support for enough multicast groups: The switch must support as many multicast groups as you intend to send streams plus another 20 groups for management traffic (PTP, IGMP messages, etc). If you intend to use TR03 / ST2110 compliant setups, bear in mind that each signal will be transported as separate audio, video and data streams. 1024 multicast groups are a reasonable starting point for smaller setups.
- The switch may not lose IGMP messages, even when operating under high load
- IGMP processing time: the faster, the better. This directly correlates with the number of program switches that a system will be able to perform per time unit. See the "IGMP Performance" section above for an example.

- Port to Port latency (10Gb): 5µs
- PTP accuracy for PTP-aware switches: less than 1µs

8.2 Settings

These settings should be applied to switches in order to make them work correctly:

- Multicast / IGMP

Enable IGMP Snooping

Disable IGMP Fast Leave (aka Immediate Leave) unless the complete system has been tested and proven to work with this setting

Enable IGMP Querier and set IGMP Querier version to 2

Set the querier IP address to the IP of the switch (if needed)

Enable report flooding or manually define the mrouter ports (or disable the querier in the V__line)

Filter all unregistered / unknown multicast traffic

- QoS

Enable accepting (“trusting”) DSCP values on the ingress interfaces

Enable correct mapping of DSCP values to internal (IEE802.1p) CoS queues

- Miscellaneous

Disable power management on all ports (Green Ethernet, Energy Efficient Ethernet)

Disable Jumbo Frames

- PTP

Enable PTP

Set PTP to Boundary Clock mode (if available and proven to work correctly) or E2E Transparent mode as a fallback

- Spanning Tree

Enable Rapid Spanning Tree

Set all ports that are connected to devices (as opposed to other switches) to “Portfast” with BPDU Guard enabled (Note: the V__link4/V__remote4 contains a switch. The respective ports should NOT be set to Port fast).

Some of these settings might need to be applied to individual ports or VLANs. Consult your switch manual for details.

8.3 Limitations

- The following features are not supported at this time, either because they are known not to work or because they have not been sufficiently tested:
- Link aggregation: the available algorithms in the switch chips do not lend themselves to high bitrate multicast stream distribution. We thus recommend to not use link aggregation or limit yourself to using it in a failover mode.
- PIM routing for video streams: the V__link4/remote4 currently uses the same source IP and multicast destination IP. This mode isn't well supported with PIM in a redundant (SMPTE ST 2022-7) setup. Other devices are OK.
- IGMP version 3: The V__link4/remote4 only supports IGMP v2. V__matrix, PowerCore and A__lines support IGMPv3.
- Network Address Translation (NAT): NAT has been used in the context of integrations with Nevision's iPath, but is considered non-standard, requires careful system design and an adjustment of the multicast addresses inside the SDPs. Use with caution.
- Tagged VLANs: The switch inside the V__link4/remote4 does not support tagged VLANs.

8.4 Switch models

The following switches have successfully been used in LAWO deployments:

- Cisco SG300 [discontinued]: Audio, non PTP
- Cisco SG350: Audio, non-PTP
- Cisco Catalyst 3750X [discontinued]: Audio, non-PTP
- Cisco Catalyst 3650: Audio, non PTP, Commentary System
- Cisco Catalyst 9300: Audio, PTP
- Cisco Nexus 3000 Series: Audio, Video, selected models PTP (not compatible with Cisco DCNM at this time)
- Cisco Nexus 9000 Series: Audio, Video, PTP (boundary clock; observe compatibility with IP Fabric for Media; 9272Q, 9236C, 9336FX2, 9500R, 92160YC, 93180YC and 93240YC)Cisco Industrial Ethernet 4010: Audio, PTP (E2E, Boundary Clock)
- Artel Video Systems ARG Quarra (formerly ARG, formerly “Stage-box”): Audio, PTP (E2E, Boundary)
- Arista Networks 7150: Video, PTP (E2E, Boundary), Report Flooding
- Arista Networks 7280: Video, PTP (E2E), Report Flooding
- Arista Networks 7500 Series: Video, PTP (E2E), Report Flooding

The following switches should not be used:

- Arista Networks 7280QR-C36: uses multiple internal chips with an interconnect which might lead to bandwidth limitations on multicast
- Arista Networks 7280QRA-C36S: uses multiple internal chips with an interconnect which might lead to bandwidth limitations on multicast

9 Leased data connections

9.1 Link Types

Leased WAN connections can be established using different link types:

- Physical Links
- Packet-switched Links
- Packets-routed Links

9.1.1 Physical Links

Physical links exist as dark fiber and gray fiber. Dark fiber is a passive optical fiber cable with lengths up to 80km; In contrast to dark fibers, grey fibers come with additional optical or electrical elements that combine signals onto a single fiber (Wavelength-division multiplexing, WDM) or add timing and level regeneration. These additional elements connecting the individual fibers can add latency and jitter to the signal, so that the parameters need to be assessed before using such a connection in production.

9.1.2 Packet-switched Links

Use a technique such as Multiprotocol Label Switching (MPLS). A virtual point-to-point connection is established. These connections can be set up conforming to defined bandwidth and QoS parameters, but the physical infrastructure underneath is shared.

9.1.3 Packet-routed Links

Access to the network is established using techniques such as Digital Subscriber Line (DSL) and uses IP packet routing. The path is shared and no guarantees for bandwidth or QoS are given. Multicast traffic forwarding is usually not available.

9.1.4 Analysis

Broadly speaking the physical links provide the best performance at the highest cost, while the packet-routed links provide the least performance for the lowest cost. Choosing the right WAN connection for your application depends both on budget and requirements, which are also determined by the capabilities of the devices you finally connect to the link.

9.2 Bandwidth and QoS

First, determine how much traffic you will need to transport at peak times, e.g. 4 J2K compressed video streams with a data rate of 100Mbit/s each, equaling a total net data rate of 400Mbit/s. Consider that there is overhead incurred by the packaging of the audio / video data into e.g. SMPTE ST2022-6. If you cannot measure the final data rate, calculate with approximately 5% overhead.

Data streams originating from the V__line are marked “Expedite Forwarding” (EF; see above). Ask your provider how much “EF” traffic he can guarantee on the line (e.g. 90% of a 1Gbit/s line). The remaining traffic is handled as “best effort” and comes with no guarantees in regards to latency, packet delay variation and packet loss.

If the bandwidth guarantee is below 100%, ask your provider about how “bursty” traffic is handled.

Choose the bandwidth of the connection based on the comparison of the needed bandwidth and the bandwidth guarantee of the provider.

9.3 Latency and Packet Delay Variation

Query the packet delay variation (“PDV”, aka “jitter”) and the packet loss rate for the “EF” traffic on the line. Packet loss is often qualified as “Packet Loss Ratio” describing the number of packet losses per time unit, e.g. one packet loss per 10 days. Be aware that “bursty” (as opposed to evenly distributed) packet loss has more impact on the resulting video quality and can lead to a degraded experience even on lines rated with relatively low loss.

When renting two lines to use for Seamless Protection Switching (“SPS”, aka hitless merge), the latency difference (“offset”) must be below the maximum data buffer size including a reserve to compensate for the Packet Delay Variation on the lines.

9.4 Measurement

Once you have the leased line available and a V__line unit connected to each end, you can use the V__line “Sounding” feature to measure the line characteristics. For details see the V__line manual.

9.5 Sample line characteristics

9.5.1 LAN

For LANs, the error rate should be below 10⁻⁸ (Copper) 10⁻¹² (Fiber). PDV is expected to be below 10µsec.

9.5.2 WAN

For WANs, the error rate should be below 10⁻³. PDV is expected to be below 20ms.

9.6 Lawo device performance

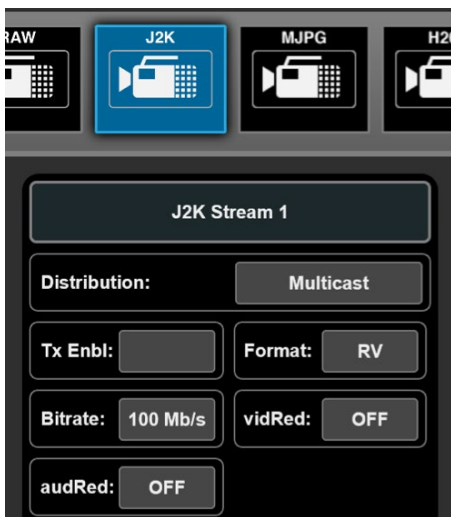
The V__line has been tested successfully with 30ms of PDV using RAW video (SMPTE ST-2022-6) and 150ms of PDV using J2K-encoded video (2x2 mode; currently only 8ms for 4x4 mode).

10 Bandwidth examples

The following table lists some bandwidth requirements depending on the essence transported. These numbers are rounded to provide estimation guidelines.

Essence	Bandwidth (including overhead)	Packets per second
Audio, Linear PCM, 16bit, 16 channels, 48kHz	14Mbit/s	2000 packets/sec
Audio, Linear PCM, 16bit, 8 channels, 48kHz	7Mbit/s	1500 packets/sec
Audio, Linear PCM, 24bit, 16 channels, 48kHz	20Mbit/s	2000 packets/sec
Audio, AM824, 16 channels, 48kHz,	27Mbit/s	3000 packets/sec
Audio, Linear PCM, 24bit, 32 channels, 48kHz	53Mbit/s	6000 packets/sec
Video SD (270Mbit/s), SMPTE 2022-6	286Mbit/s	25000 packets/sec
Video, HD (1.485Gbit/s), SMPTE 2022-6	1.57Gbit/s	135000 packets/sec
Video, 3G (2.970Gbit/s), SMPTE 2022-6	3.15Gbit/s	270000 packet/sec
Video, HD (1.485Gbit/s), SMPTE 2110-20	1.1Gbit/s	108000 packets/sec
Video, 3G (2.970Gbit/s), SMPTE 2011-20	2.2Gbit/s	216000 packets/sec

Note: Bandwidth requirements for J2K-encoded video depends on the actual picture; the upper limit can be defined in the V_line settings (Settings > TX Stream > J2K):



Please note that the bitrate is the total bitrate including audio and VANC data. If you add more audio channels, but do not increase the overall bitrate, the bitrate capacity for video decreases.

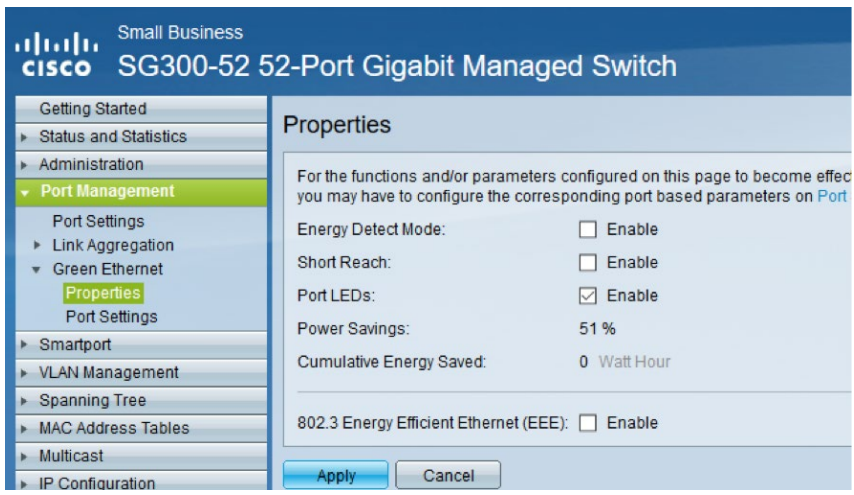
11 Specific Configurations

11.1 Cisco SG300

Navigate to Port Management > Port Settings and disable Jumbo Frames:



Navigate to Port Management > Green Ethernet > Properties and disable Energy Efficient Ethernet:



Navigate to Multicast > Multicast Router Port and set all ports to None:

Small Business
CISCO SG300-52 52-Port Gigabit Managed Switch

Getting Started
 ▶ Status and Statistics
 ▶ Administration
 ▶ Port Management
 ▶ Smartport
 ▶ VLAN Management
 ▶ Spanning Tree
 ▶ MAC Address Tables
 ▼ **Multicast**
 Properties
 MAC Group Address
 IP Multicast Group Address
 IGMP Snooping
 MLD Snooping
 IGMP/MLD IP Multicast Group
Multicast Router Port
 Forward All
 Unregistered Multicast
 ▶ IP Configuration
 ▶ Security
 ▶ Access Control
 ▶ Quality of Service
 ▶ SNMP

Multicast Router Port

Filter: VLAN ID equals to AND IP Version equals to

Port	GE1	GE2	GE3	GE4	GE5	GE6	GE7	GE8	GE9
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Dynamic	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Port	GE25	GE26	GE27	GE28	GE29	GE30	GE31	GE32	GE33
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Dynamic	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Port	GE49	GE50	GE51	GE52
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Dynamic	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Navigate to Multicast > Forward All and set all ports to None:

The screenshot shows the configuration page for the Cisco Small Business SG300-52 52-Port Gigabit Managed Switch. The left sidebar contains a navigation menu with the following items: Getting Started, Status and Statistics, Administration, Port Management, Smartport, VLAN Management, Spanning Tree, MAC Address Tables, Multicast (expanded), Properties, MAC Group Address, IP Multicast Group Address, IGMP Snooping, MLD Snooping, IGMP/MLD IP Multicast Group, Multicast Router Port, Forward All (highlighted), Unregistered Multicast, IP Configuration, Security, and Access Control.

The main content area is titled "Forward All" and features a filter: "Filter: VLAN ID equals to 101 AND Interface Type equals to". Below the filter are three tables of radio button options for different ports. In all cases, the "None" option is selected.

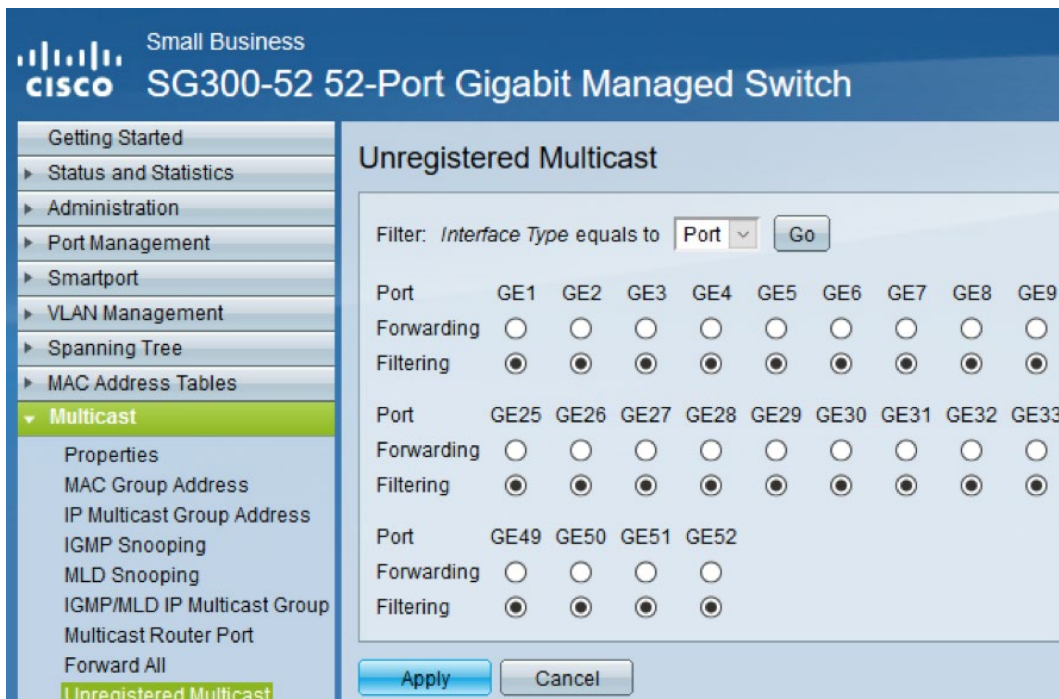
Port	GE1	GE2	GE3	GE4	GE5	GE6	GE7	GE8	GE9
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Port	GE25	GE26	GE27	GE28	GE29	GE30	GE31	GE32	GE33
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

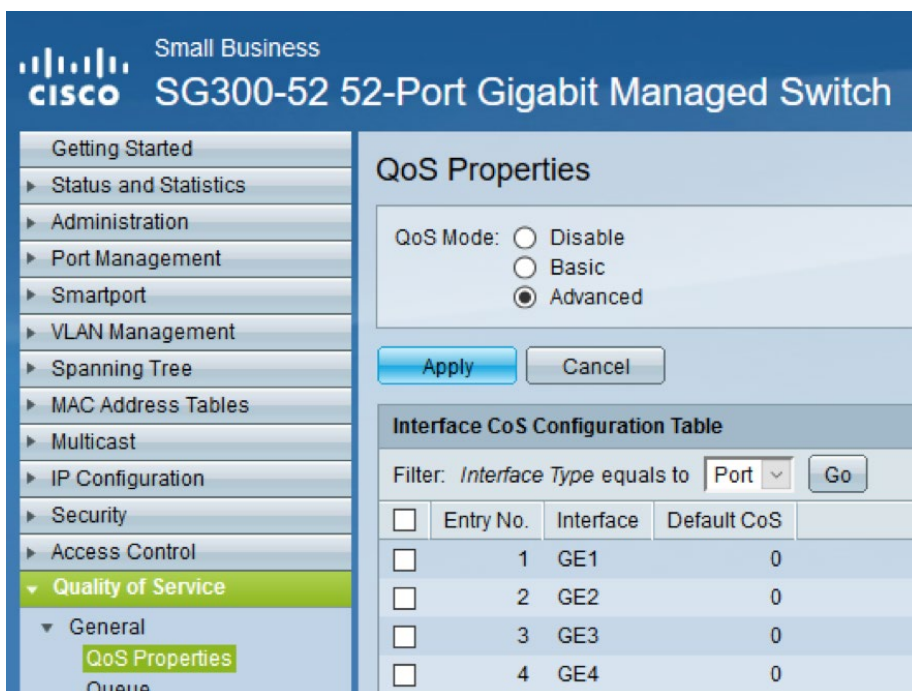
Port	GE49	GE50	GE51	GE52
Static	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
None	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

At the bottom of the configuration area are "Apply" and "Cancel" buttons.

Navigate to Multicast > Unregistered Multicast and set all ports to Filtering:



In Quality of Service > General > Properties set the QoS mode to Advanced:



Navigate to Quality of Service > General > Queue and set all queues to Strict Priority:

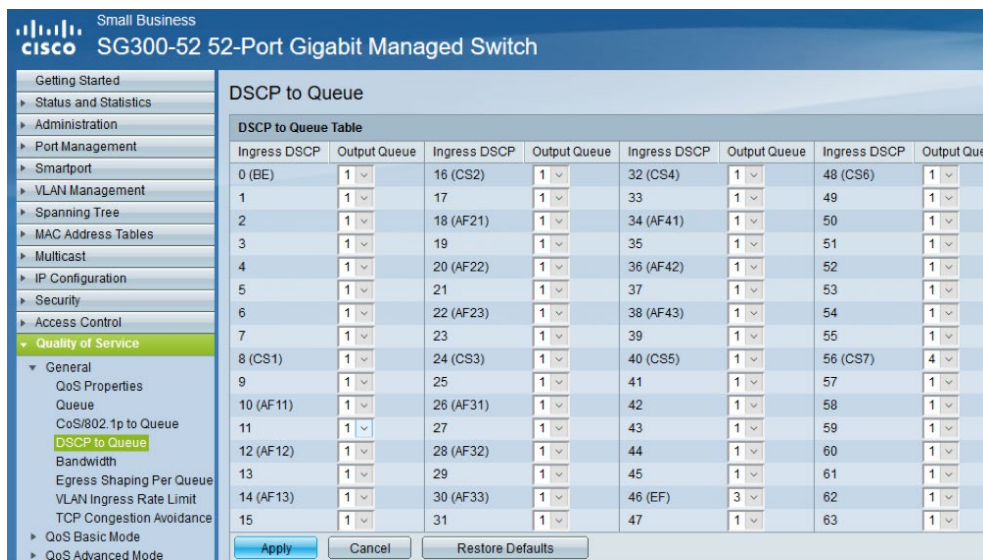
The screenshot shows the configuration page for the Queue on a Cisco Small Business SG300-52 52-Port Gigabit Managed Switch. The left sidebar contains a navigation menu with the following items: Getting Started, Status and Statistics, Administration, Port Management, Smartport, VLAN Management, Spanning Tree, MAC Address Tables, Multicast, IP Configuration, Security, Access Control, Quality of Service (expanded), General, QoS Properties, Queue (highlighted), and CoS/802.1p to Queue. The main content area is titled "Queue" and contains a "Queue Table" section. The table has four columns: Queue, Scheduling Method, WRR, WRR Weight, and % of WRR Bandwidth. The Scheduling Method column has sub-columns for Strict Priority and WRR. Queue 1 has Strict Priority selected (radio button checked) and WRR Weight of 1. Queue 2 has Strict Priority selected and WRR Weight of 2. Queue 3 has Strict Priority selected and WRR Weight of 4. Queue 4 has Strict Priority selected and WRR Weight of 8. Below the table are "Apply" and "Cancel" buttons. A note below the buttons states: "Queue 1 has the lowest priority, queue 4 has the highest priority."

Queue	Scheduling Method		WRR Weight	% of WRR Bandwidth
	Strict Priority	WRR		
1	<input checked="" type="radio"/>	<input type="radio"/>	1	
2	<input checked="" type="radio"/>	<input type="radio"/>	2	
3	<input checked="" type="radio"/>	<input type="radio"/>	4	
4	<input checked="" type="radio"/>	<input type="radio"/>	8	

Apply Cancel

Queue 1 has the lowest priority, queue 4 has the highest priority.

In Quality of Service > General > DSCP to Queue set DSCP 56 (CS7; PTP packets) to queue 4 (highest priority), DSCP 46 (EF, essence streams) to queue 3 and everything else to queue 1 (lowest priority):

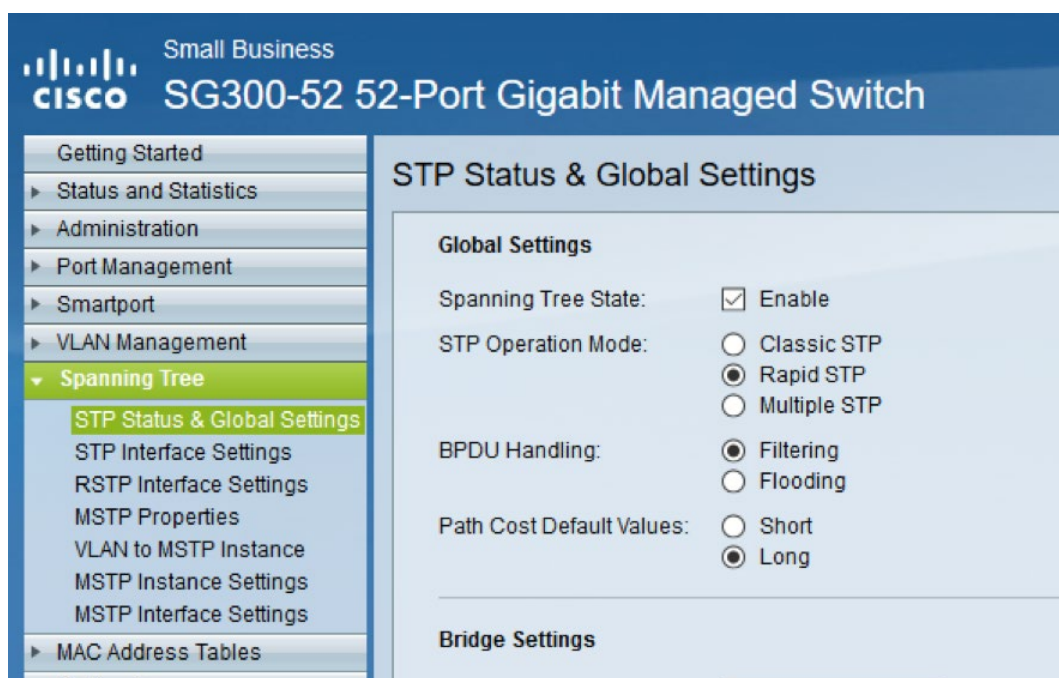


Navigate to Quality of Service > QoS Advanced Mode > Global Settings and set the Trust Mode to DSCP:

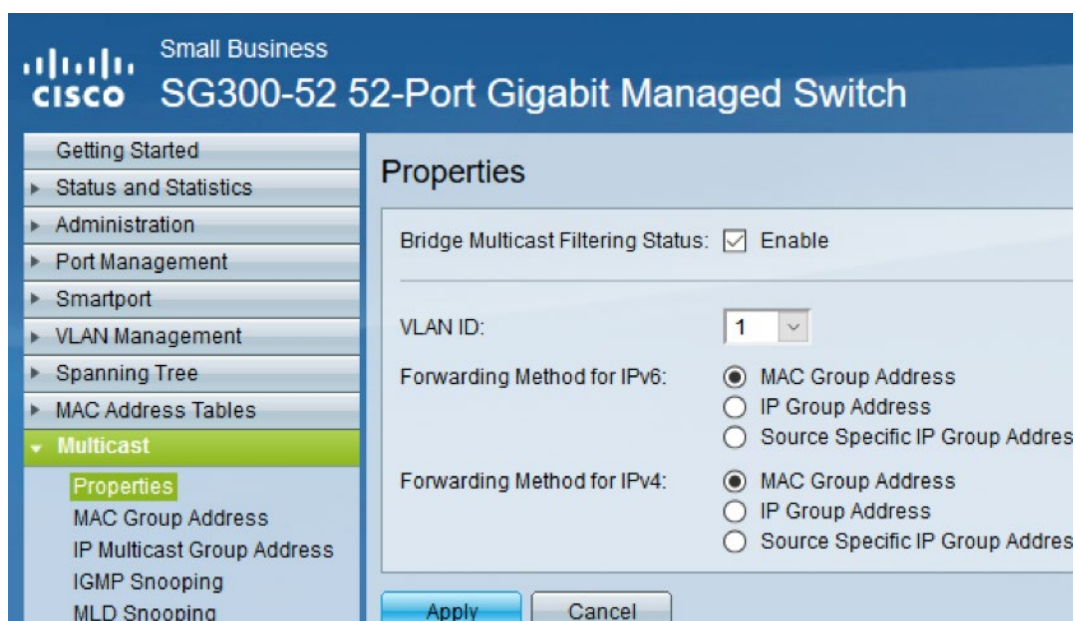


Remember to save your configuration!

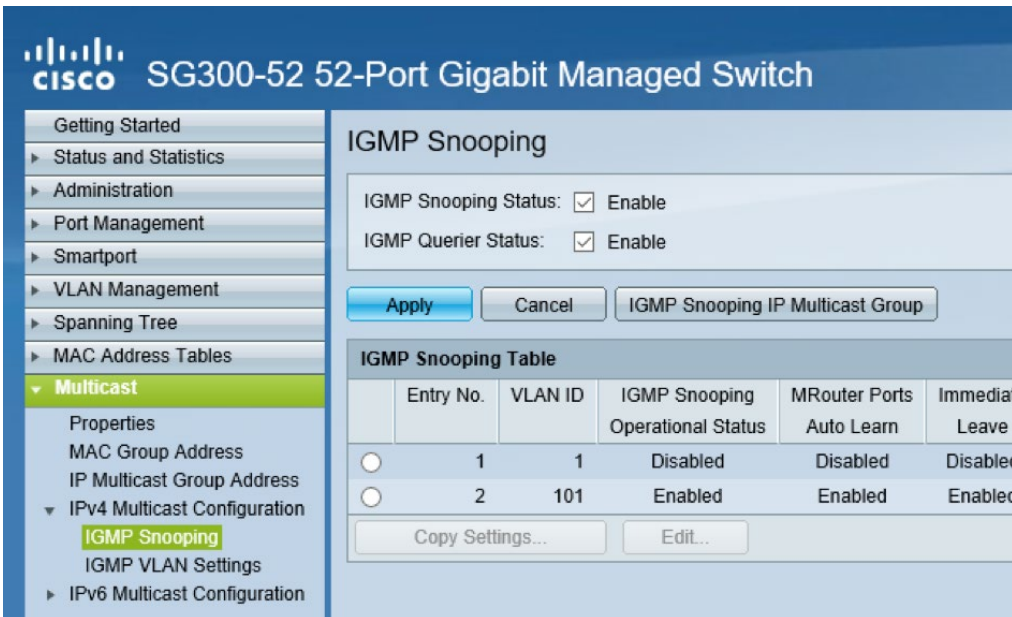
In Spanning Tree > STP Status and Global Settings enable Spanning Tree, set the mode to Rapid Spanning Tree and set the BPDU Handling to Filtering:



Navigate to Multicast > Properties and enable Bridge Multicast Filtering Status:

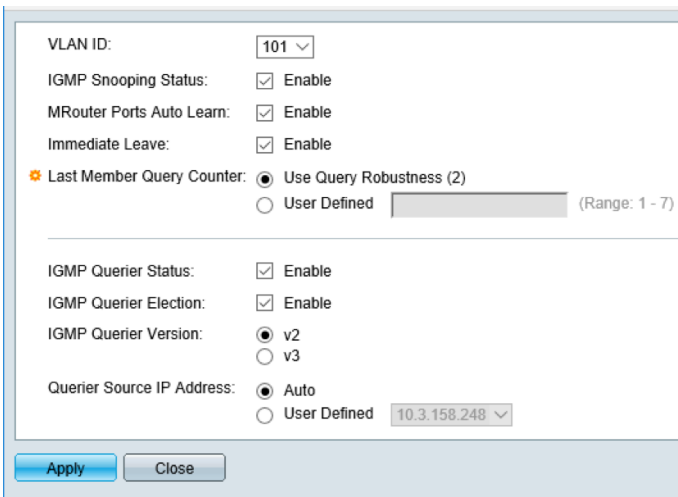


In Multicast > IPv4 Multicast Configuration > IGMP Snooping
enable IGMP Snooping Status and IGMP Querier Status:



Choose to edit the settings for the relevant VLAN. Enable the IGMP Snooping Status, enable the MRouter Ports Auto Learn (for audio networks only), enable Immediate Leave (if certain that all devices implement it correctly).

Enable IGMP Querier Status and IGMP Querier Election. Set IGMP Querier Version to 2 and set the Querier Source IP Address to Auto:



11.2 Arista (PTP E2E)

Note: replace the bold values by the correct values for your infrastructure. These commands need to be issued in the correct context of the switch OS.

The following configuration is valid for all Arista switches tested.

```
hostname arista7150
username admin secret admin
interface management 1
    ip address 192.168.2.1/24
interface Vlan 1
    ip address 192.168.1.1/24
ip igmp snooping
ip igmp snooping vlan 1
ip igmp snooping report-flooding
ip igmp snooping vlan 1 report-flooding
ip igmp snooping report-flooding switch-port Ethernet 1-22
ip igmp snooping querier
ip igmp snooping vlan 1 querier
ip igmp snooping querier address 192.168.1.1
ip igmp snooping vlan 1 querier address 192.168.1.1
ip igmp snooping querier version 2
ip igmp snooping vlan 1 querier version 2
ip igmp snooping querier query-interval 30
ip igmp snooping vlan 1 querier query-interval 30
no ip igmp snooping vlan 1 immediate-leave
ptp mode p2pttransparent
ptp source ip 192.168.1.1
interface Ethernet 1-24
    ptp enable
```

11.3 Artel Video Systems ARG Quarra Switches

The Artel Video Systems ARG Quarra switches (formerly ARG) have only been tested for AES67 audio by Lawo, but are reported to support video with a SMPTE 2059-2 PTP timing.

Navigate to Configuration > IPMC > IGMP Snooping > Basic Configuration and set parameters as follows:

IGMP Snooping Configuration

Global Configuration	
Snooping Enabled	<input checked="" type="checkbox"/>
Unregistered IPMCv4 Flooding Enabled	<input type="checkbox"/>
IGMP SSM Range	232.0.0.0 / 8
Leave Proxy Enabled	<input checked="" type="checkbox"/>
Proxy Enabled	<input checked="" type="checkbox"/>

Port Related Configuration

Port	Router Port	Fast Leave	Throttling
*	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
2	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
3	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
4	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
5	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
6	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
7	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
8	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
9	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
10	<input type="checkbox"/>	<input type="checkbox"/>	unlimited

Navigate to Configuration > IPMC > IGMP Snooping > VLAN Configuration and set parameters as follows:

IGMP Snooping VLAN Configuration

Refresh

Start from VLAN 1 with 20 entries per page.

Delete	VLAN ID	Snooping Enabled	Querier Election	Querier Address	Compatibility	PRI	RV	QI (sec)	QRI (0.1 sec)	LLQI (0.1 sec)	URI (sec)
Delete	1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	0.0.0.0	Forced IGMPV2	0	2	125	100	10	1

Navigate to Configuration > QoS > Port Classification and set the parameters as follows:

QoS Ingress Port Classification

Port	CoS	DPL	PCP	DEI	Tag Class.	DSCP Based	Address Mode
*	<>	<>	<>	<>		<input checked="" type="checkbox"/>	<>
1	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
2	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
3	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
4	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
5	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
6	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
7	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
8	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
9	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source
10	0	0	0	0	Disabled	<input checked="" type="checkbox"/>	Source

Navigate to Configuration > QoS > DSCP-based QoS and set highest queue (7) for the DSCP value used for PTP, the second highest queue (6) for the RTP audio streams.

If using the ARG switch as a PTP clock, navigate to Configuration > PTP and choose to "Add new PTP clock". Set the parameters as follows:

PTP External Clock Mode

One_PPS_Mode	Output
External Enable	False
Adjust Method	LTC frequency
Clock Frequency	1

PTP Clock Configuration

Delete	Clock Instance	Device Type	Profile
Delete	0	E2eTransp	No Profile

Add New PTP Clock Save Reset

For more details, check the complete Setup Guide and switch manual:
<http://www.artel.com/media-transport-products/arg/quarra-ntp-10-gbps-ethernet-switch>

Click the "0" under Clock Instance and set the parameters as follows (don't forget to adjust the DSCP values):

PTP Clock's Configuration and Status

Clock Type and Profile

Clock Instance	Device Type	Profile	Apply Profile Defaults
0	E2eTransp	No Profile	n/a

Port Enable and Configuration

Port Enable										Configuration
1	2	3	4	5	6	7	8	9	10	Ports Configuration
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

Local Clock Current Time

PTP Time	Clock Adjustment method	Synchronize to System Clock
1970-01-01T04:31:51+00:00 374,943,740	Internal Timer	Synchronize to System Clock

Clock Current DataSet

stpRm	Offset From Master	Mean Path Delay
0	0.000,000,000	0.000,000,000

Clock Parent DataSet

Parent Port ID	Port	PStat	Var	Rate	GrandMaster ID	GrandMaster Clock Quality	Pri1	Pri2
00:50:c2:ff:fe:39:e9:f0	0	False	0	0	00:50:c2:ff:fe:39:e9:f0	Cl:251 Ac:Unknwn Va:65535	128	128

Clock Default DataSet

ClockId	Device Type	2 Step Flag	Ports	Clock Identity	Dom	Clock Quality	
0	E2eTransp	False	10	00:50:c2:ff:fe:39:e9:f0	0	Cl:251 Ac:Unknwn Va:65535	
Pri1	Pri2	Protocol	One-Way	VLAN Tag Enable	VID	PCP	DSCP
128	128	IPv4Multi	False	False	1	0	46

Clock Time Properties DataSet

UtcOffset	Valid	leap59	leap61	Time Trac	Freq Trac	ntp Time Scale	Time Source
0	False	False	False	False	False	True	160

Filter Parameters

Filter Type	Delay Filter	Period	Dist
Basic	6	1	2

Servo Parameters

Display	P-enable	I-enable	D-enable	'P' constant	'I' constant	'D' constant
False	True	True	True	3	80	40

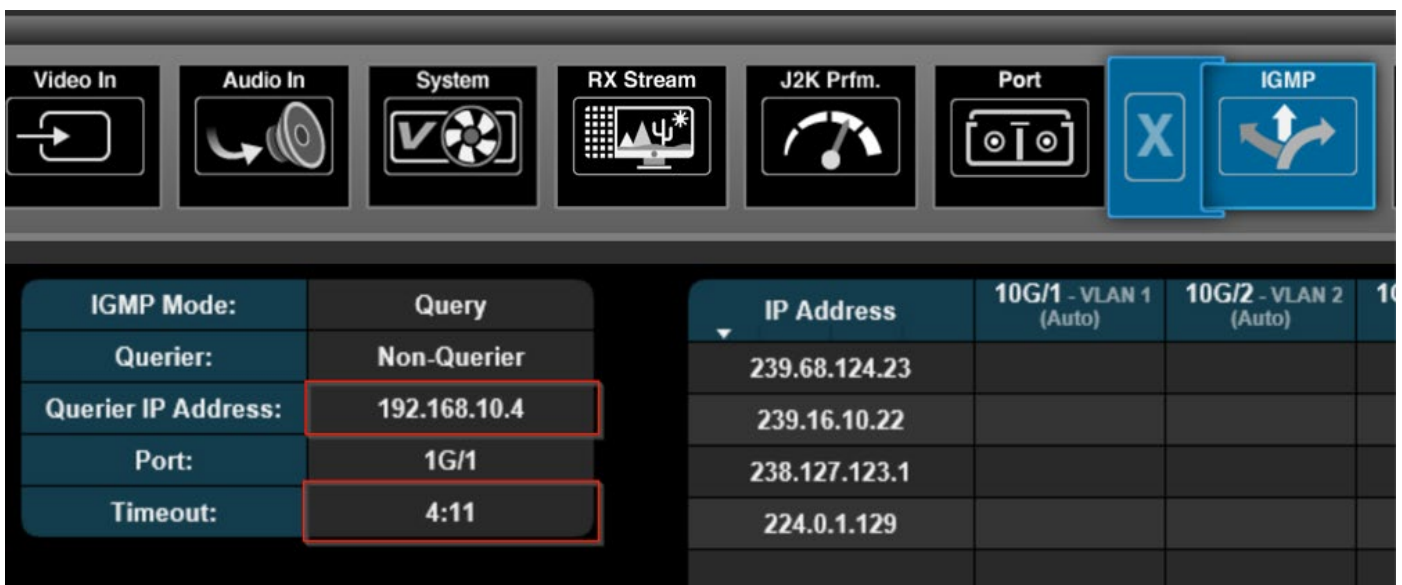
12 Troubleshooting

If you are having trouble establishing video streams the following tips can help you to track down the issue.

12.1 Multicast

Check if the IGMP Querier is setup correctly and recognized by the devices.

Navigate to the Web UI of the receiving V__line. Under “Status” > “Switch” > “IGMP” verify that the IP address of the querier is correctly reflected and that the timeout never actually reaches 0:00:



The screenshot shows the Web UI interface for the IGMP configuration. The top navigation bar includes tabs for Video In, Audio In, System, RX Stream, J2K Prfm., Port, and IGMP. The IGMP tab is selected, and a sub-tab with a blue 'X' icon is also visible. Below the navigation bar, there are two main sections. The left section displays the IGMP Mode and Query settings, with the Querier IP Address highlighted in a red box. The right section displays a table of IP addresses for different VLANs.

IGMP Mode:	Query	IP Address	10G/1 - VLAN 1 (Auto)	10G/2 - VLAN 2 (Auto)	10G/3 - VLAN 3 (Auto)
Querier:	Non-Querier	239.68.124.23			
Querier IP Address:	192.168.10.4	239.16.10.22			
Port:	1G/1	238.127.123.1			
Timeout:	4:11	224.0.1.129			

The IP address for the querier should either be the switch (if setup accordingly) or the IP address of the V__line fulfilling that role, but remember that the IGMP querier with the lowest IP address wins the querier election.

If you expect to see the switch, but instead see a V__line, try switching off the querier functionality of that V__line (“Settings” > “Switch” > “IGMP” > “Querier”).

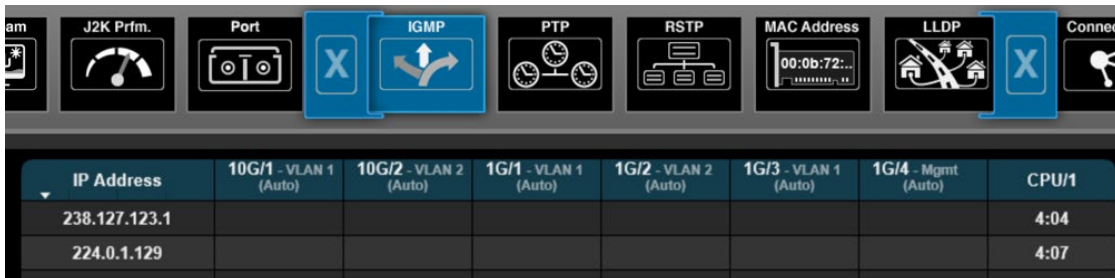
If you do not see a querier on the destination device, this might point to general issues with multicast between source and destination.

If the querier timeout reaches 0:00, the multicast streams will break down, as the source device assumes no one is interested in

the stream anymore. The V__lines will re-establish the stream, but there will be an interruption in the video.

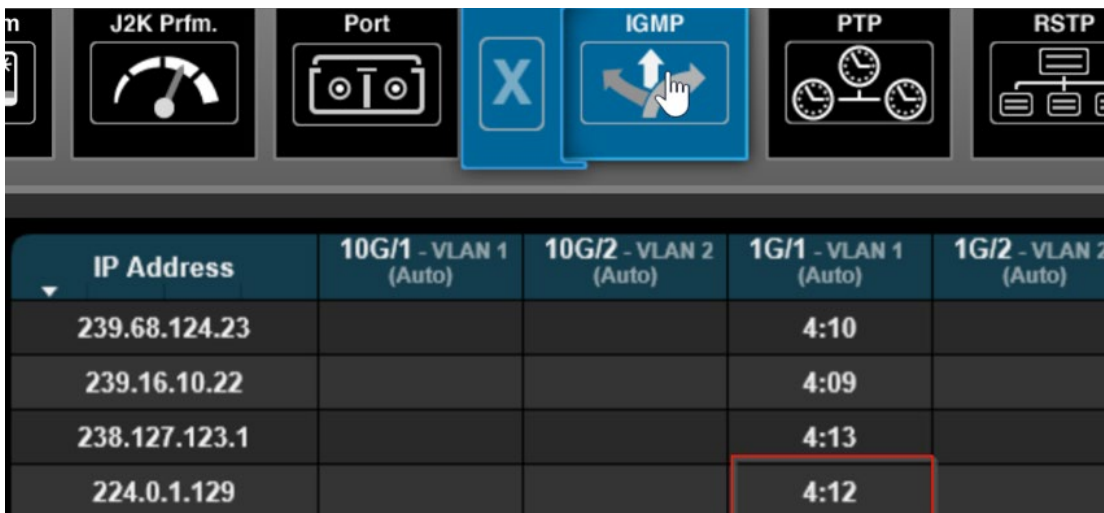
General issues with multicast and issues with the IGMP querier can be distinguished from other network issues by switching the stream distribution from multicast to unicast (remember to copy the SDP again – it changes when setting the stream to unicast). If the unicast stream stays stable, look for the issue in the realm of multicast.

In the “Status” > “Switch” > “IGMP” menu of the V__line you can also check whether you receive multicast packets by looking at the multicast address vs. interface table. if you see no entry for the multicast IP address or you don't see a timer running on the interface where you expect to receive the data, the V__line does not receive the respective multicast data. The following screenshot shows missing multicast data for PTP (224.0.1.129; no timer running on any interface):



IP Address	10G/1 - VLAN 1 (Auto)	10G/2 - VLAN 2 (Auto)	1G/1 - VLAN 1 (Auto)	1G/2 - VLAN 2 (Auto)	1G/3 - VLAN 1 (Auto)	1G/4 - Mgmt (Auto)	CPU/1
238.127.123.1							4:04
224.0.1.129							4:07

If the data is correctly received on interface 1G/1, you will see a timer running, that never reaches 0:00:



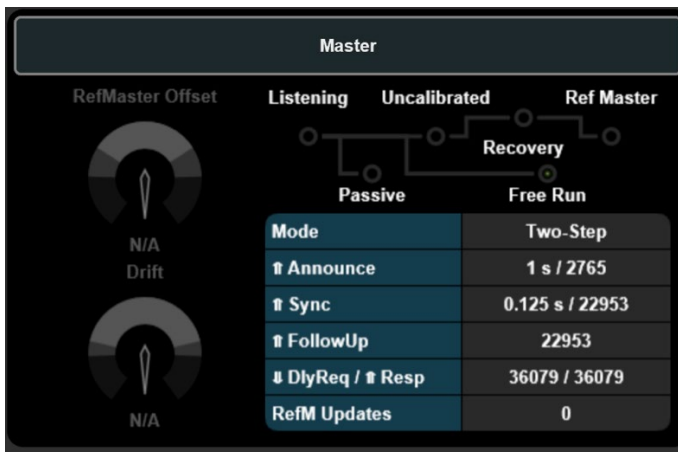
The image shows a network management interface with several tabs: J2K Prfm., Port, IGMP, PTP, and RSTP. The IGMP tab is selected and highlighted in blue. Below the tabs is a table with the following data:

IP Address	10G/1 - VLAN 1 (Auto)	10G/2 - VLAN 2 (Auto)	1G/1 - VLAN 1 (Auto)	1G/2 - VLAN 2 (Auto)
239.68.124.23			4:10	
239.16.10.22			4:09	
238.127.123.1			4:13	
224.0.1.129			4:12	

12.2 PTP

PTP relies on multicast, so if you are having trouble establishing streams, check for issues with multicast first, then come back to PTP.

Check the master's state:

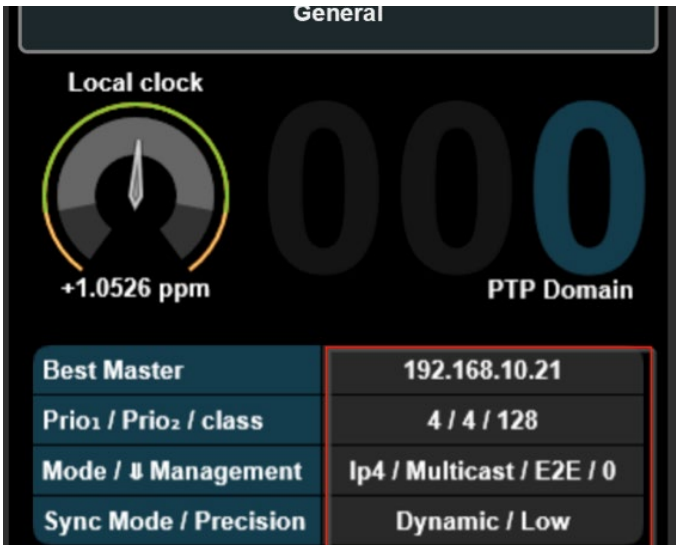


The master should send out messages (“Announce”, “Sync”, etc. increasing). If the master is not synchronized to an analog reference, it should be in state “Free Run”, else in state “Ref Master”.


If it is in “Passive” mode, another device has become master, probably due to higher priority.

If it is in state “Recovery” the V_line might have been rebooted and is waiting for manual confirmation that it should take over the role of Grandmaster again (only in “Redundant” mode; confirm in “Settings” > “Switch” > “PTP” > “Resume”).

The client should see the correct Grandmaster and its parameters:



You might see the switches IP address here, if the switch runs in Boundary Clock mode. Check the client's state:



The client should be sending out messages (“Ann/DlyResp”, etc. increasing). Once the client has successfully synchronized to the master’s time, it will be in state “Calibrated”.

If the client cannot reach the “Calibrated” state, check the offset from master and its jitter window:



If the offset is too big or the jitter range is too big, the client cannot synchronize. You can increase the tolerance of the V__line in “Settings” > “Switch” > “PTP” > “Timing Precision” (set to “custom” and enter a new acceptable value). Please keep in mind that this only increases the tolerance of the V__Line; it doesn’t make the PTP better. The unit may still run into trouble with other devices with a lower tolerance level.

If the clients do not recognize the master or if you still cannot get a reliable synchronization, check that all the configuration values (Domain, Sync Interval, Delay Request Interval, etc.) are the same for all devices involved (masters and slaves).

If you cannot get PTP to work reliably (due to multicast issues or jitter which is too high), you can still use the the V__Line unit’s “Virtual Black Burst” feature to synchronize. This, however requires you to use one of your transmission streams.

13 FAQ

13.1 General

13.1.1 How much bandwidth is used? What connectivity do I need?

For some examples of bandwidth used refer to the “Bandwidth examples” section in this guide. In general, 1GbE between switch and device should be sufficient for many audio applications, 10GbE for SD; HD and 3G video and higher bandwidths (25, 40, 50, 100GbE) for many video streams and 4K resolutions.

13.1.2 What cabling shall I use?

For short distances you can use copper cables (CAT6 or better) for 1GbE. For other connections we recommend fiber; multimode vs. singlemode depends on the distance you need to cover.

For short distances and higher bandwidths (10GbE+) we recommend the use of Active Optical Cables (AOCs) to reduce cabling cost.

13.1.3 Can I combine multimode and singlemode SFPs?

For every connection only one type of SFPs can be used. A multimode SFP on a sending device needs multimode fiber and a multimode SFP on the receiving device. However, if a device offers multiple SFP connections, you can mix multimode and singlemode as long as the individual connection uses only one.

13.1.4 What multicast scheme shall I use?

We recommend using the administratively-scoped multicast range as described in IETF RFC2365 (239.0.0.0/8). That should avoid collisions with other multicast addresses used in e.g. routing protocols.

In order to structure your multicast range, apply a logical pattern, e.g. separating building areas or essence types (Audio, Video, Data).

Also take a look at the section “Multicast Address Considerations” in this guide.

13.1.5 Do I need to worry about oversubscription?

Oversubscription describes the situation that a link, either between a device and the switch or between switches, has more data to transport than it has capacity. In the case of time-critical data such as audio and video an oversubscription usually means loss of signal and should thus be avoided. At the moment this requires careful network design. In the future more advanced network control mechanisms such as Arista MCS, Cisco DCNM, et.al. will handle oversubscription scenarios in a better way.

13.2 Switches

13.2.1 Which switch types does Lawo support?

The networking guide lists a number of switches we have worked with in various projects. The list also details whether the switch has been used for audio or video and whether was used with PTP E2E or Boundary clock. This list is updated with the guide.

13.2.2 Which switch types does Lawo recommend?

The switch requirements vary a lot with the application, size of the system and the potential integration into an existing infrastructure. Considering these variables, we cannot recommend a particular switch model, but we are happy to discuss the requirements after understanding the concrete needs.

13.2.3 Can you provide switch configs?

The switch configuration equally varies with the application, size of the system and the potential integration into an existing infrastructure. We thus cannot provide more than the configuration examples listed in this guide. For more details, please contact your local Lawo representative to discuss a solution that meets your specific needs.

13.3 Redundancy

13.3.1 Does Lawo support SMPTE ST2022-7 redundancy?

Most of Lawo's equipment supports essence stream redundancy according to SMPTE ST2022-7.

13.4 PTP

13.4.1 What is the difference between the One-Step mode and the Two-Step mode in PTP?

For almost every PTP message that is exchanged between the master and the slave the timestamp can either be placed directly into the message (one step) or a separate message containing the time when the first message was sent can be transmitted (two step). Since the PTPv2 standard IEEE1588-2008 defines that slaves must be able to cope with both modes there is no need to run all PTP devices with the same mode.

A more detailed explanation can be found in the blog of the PTP equipment manufacturer Meinberg: <http://blog.meinbergglobal.com/2013/10/28/one-step-two-step/>

13.4.2 Do I need PTP or can I work without it?

In IT networks PTP replaces the known synchronization methods such as Black Burst, Tri-Level or Word Clock and the new SMPTE standards such as ST2110 mandate the use of PTP for essence synchronization. It is perfectly suited to provide phase-accurate audio and video.

Many Lawo products are still capable of using the traditional synchronization methods, but we recommend to complement existing infrastructures by PTP – which, with the correct synchronization equipment can run alongside e.g. the established black burst.

13.4.3 Which PTP Grandmasters do we work with?

The networking guide lists a number of PTP Grandmasters we have worked with in various projects. This list is updated with the guide.

13.4.4 How many Grandmasters do I need in a system?

As with traditional synchronization methods you need multiple generators, if you need redundancy. We recommend using at least two Grandmasters.

If you use PTP in a redundant network setup (SMPTE ST2022-7), each Grandmaster should be connected to both networks in order to avoid different times in the two networks (e.g. in case only one GM loses its GPS connection).

13.4.5 My Grandmaster does not support enough clients / how do I scale PTP?

Scaling PTP requires careful system design and adjustment to the application. However, speaking generally, you can use PTP boundary clock on the switch (if supported). Then, the only client to the Grandmaster is the switch, which in turn handles all other clients. Please bear in mind that there are also limits on the number of PTP clients a switch can support. We also recommend the usage of PTP “hybrid mode” in which the DelayRequest/DelayResponse messages between the client and the Grandmaster are exchanged using unicast, thus lessening the load on the network and devices.

13.5 Ravenna

13.5.1 What is the Payload?

When you create a TX stream, the resultant network packet size is determined by three parameters which are collectively known as the payload:

- Frame Size = the number of samples per channel per network packet. The smaller the frame size, the more often the sender transmits packets. This results in a lower sending latency, but also a higher demand on the network’s bandwidth. In Lawo devices, the frame size limits the number of TX streams which can be created by each device
- Codec = the encoding method used for the digital audio. For example: L16 = 16-bit linear PCM; L24 = 24-bit linear PCM; AM824 = 24-bit linear PCM + 8-bit metadata, a non-standard format commonly used in AES/EBU
- Channel count = the number of channels to be encoded: mono, stereo, 8-channel, etc.

It is the payload which forms the bulk of the RAVENNA network packet size. In short, the more channels per stream, the bigger the payload.

To calculate the payload:

$$\text{payload (bits)} = \text{channel count} * \text{bits per sample} * \text{number of samples per packet}$$

e.g. the payload of a stereo stream using the L24 codec and frame size of 48 samples per packet = $2 * 24 * 48 = 2304$ Bits.

13.5.2 Do I need to stick to the Payload Presets?

No, but if you choose a custom setup, you MUST remain within the maximum network packet size, and understand how the packet size affects the sending latency and network bandwidth requirements.

13.5.3 How does the Payload affect Latency?

The bigger the payload, the bigger the network packet size, and the longer it takes to assemble the data before each packet is sent. This delay is known as the sending latency.

To calculate the sending latency:

$$\text{Sending Latency (s)} = \frac{\text{Number of samples per packet (frame size)}}{\text{sample rate}}$$

e.g. the sending latency of a stream using 48 samples per packet at 48kHz = $48 / 48,000 = 0.001\text{s}$ (1 millisecond)

Note that this figure is a nominal amount, as the actual delay will be subject to other factors such as jitter within the sending device.

13.5.4 What is the Network Packet Size?

In the Ethernet standard, every network packet consists of two parts: header + payload

- the header contains key IP information such as the address and is a fixed size (in the Ethernet standard) = 82 Bytes
- the payload varies in size (see above)

The Ethernet standard defines a maximum network packet size of 1460 bytes, otherwise known as the MTU (Maximum Transmission Unit). To comply with this, RAVENNA packets must not exceed the MTU. Therefore, it is important to know the network packet size and which parameters affect it. When you create a TX stream in the RAVENNA Web UI, the resultant network packet size is displayed. If the size exceeds the MTU, a warning appears, and you must take steps to reduce it (see Creating a TX Stream).

Note: RAVENNA does not support jumbo frames.

The network packet size of a RAVENNA stream is calculated as follows:

$$\text{packet size (bits)} = \text{header size (82 Bytes} \times 8) + \text{payload (number of channels} \times \text{bits per sample} \times \text{number of samples per packet)}$$

e.g. the packet size of a stereo stream using the L24 codec and 48 samples per packet = $(656) + (2 \times 24 \times 48) = 2960$ Bits (or 370 Bytes).

13.5.5 What if the Network Packet Size is too big?

When you create a TX stream in the RAVENNA Web UI, the resultant network packet size is displayed in the “Source Properties” window. If the size exceeds the MTU, a warning appears.

To deal with this, you will need to reduce one of the payload variables - for example, change the codec from L24 to L16, or reduce the frame size if it is possible to do so (the frame size limits the number of TX streams).

Usually the easiest solution is to reduce the channel count and split the audio into multiple streams. This works well providing phase coherence is not required.

13.5.6 Calculating the Bandwidth of a Stream

The math for the bandwidth calculation is as follows:

$$\text{Nominal bandwidth per stream (Mbit/s)} = \text{packet rate (packets per second)} \times \text{packet size (bits)} \times 10^{-6}$$

The packet rate is the nominal number of packets per second sent out by a device = $\text{sample rate} / \text{number of samples per packet}$.

For a stereo stream (24-bit, 48kHz, 48 samples per IP packet):

$$\text{Packet Rate} = 48,000 / 48 = 1000$$

$$\text{Packet Size (Bits)} = (656) + (2 \times 24 \times 48) = 2960$$

Therefore, the nominal bandwidth per stream (Mbit/s) = $1000 \times 2960 \times 10^{-6} = 2.96$ Mbit/s (or around 3 Mbit/s)

13.5.7 How does the Network Bandwidth affect RAVENNA Streaming?

If the amount of streaming traffic gets close to or exceeds the network bandwidth, then you may find that audio streams start to drop-out during playback.

13.5.8 How does jitter affect RAVENNA streaming?

Lawo's RAVENNA devices use a receiving buffer to deal with network jitter. For example, if packet 3 of an audio stream arrives before packets 1 and 2, the buffer holds packet 3 until packets 1 and 2 have arrived; the packets can then be played out in the correct order, resulting in a successful playout of the audio stream.

Too much jitter becomes a problem if the size of the receiving buffer cannot cope with the length of delay. For example, if packet 3 has fallen out of the receiving buffer before packets 1 and 2 arrive, then an audio drop-out may occur while playing out the audio stream.

Jitter can be introduced by any component in the network including the sending and receiving node. For example, in Lawo products, the amount of jitter introduced by the dedicated RAVENNA IO cards is negligible compared to that of a R3LAY PC (as RAVENNA streaming depends on other processes within the computer). Jitter is one of the main reasons why the number of RAVENNA streaming channels are much lower for R3LAY software applications, than for a console or router using dedicated RAVENNA hardware.

13.5.9 How long will it take for my audio stream to reach its destination?

It depends! Unlike a fixed point-to-point analogue or digital audio connection, the transport latency of an audio stream is variable, and is affected by the frame size, sample rate and amount of network jitter.

The total delay from sender to receiver (known as the total connection latency), is calculated as follows:

Total connection latency = sending latency + network latency + receiving buffer

- The sending latency is the nominal amount of time it takes the sender to assemble the data into network packets. This is affected by the number of samples per packet and sample rate (see sending latency). For example, the nominal sending latency of a stream using 48 samples per packet at 48kHz = $48 / 48,000 = 0.001\text{s}$ (1 millisecond)
- The network latency in data networks is usually very short in local area networks; in wide area networks it depends on the distance travelled.
- A receiving buffer is used in Lawo's RAVENNA devices to deal with network jitter. As a general „rule of thumb“ the receiving buffer should be at least two times the sending latency. For our example stream (using 48 samples per packet at 48kHz), the total connection latency would be just over 3 milliseconds.

13.5.10 How do I integrate Ravenna/AES67 and DANTE?

The integration of Ravenna/AES67 and DANTE requires some special attention in the network structure / design. Please get in contact with you LAW0 representative for more details.

14 Glossary

AoIP	Audio-over-IP
Avahi	A data network service (similar to Bonjour) that allows devices to publish and discover nodes running on a Local Area Network. Avahi is an example of a zeroconf networking implementation. Other zeroconf systems include Bonjour (licensed by Apple).
Buffer size	The buffer size sets the amount of data stored (in memory) before each data packet is transmitted or played out. In an audio system, the smaller the buffer size, the lower the latency, but the more susceptible to drop-outs.
COMi.MX	The name of Lawo's RAVENNA processing hardware device. The COMi.MX forms a sub-component of most of Lawo's RAVENNA IO cards.
DALLIS	Digital and Line Level Interface System; the name of Lawo's configurable IO device. Each DALLIS frame can be fitted with a combination of plug-in IO cards (analog, AES, MADi, RAVENNA, GPIO, etc.).
Ember+	A non-proprietary TCP/IP interface protocol. An Ember+ provider can "publish" parameters which may then be used by an Ember+ consumer. For example, to display information or enable control from a remote device. More information on GitHub: https://github.com/Lawo/ember-plus/wiki
Nova73	A stand-alone routing matrix with networking capabilities; this is a large matrix related to the mc ² series of Lawo consoles.
RAVENNA	A real-time, network-synchronized Audio over IP protocol. RAVENNA offers real-time distribution of audio and other media content within IP-based network environments. Parts of RAVENNA have been standardized by the AES as AES67.
Remote MNOPL	The remote control protocol RemoteMNOPL is a LAN-based client-server network byte order protocol to enable third party systems to control Lawo's digital mixing consoles or standalone routers.
RTP	Real-time Transport Protocol; a networking protocol that defines a standard packet format for delivering audio and video over data networks.
RTCP	Real-time Transport Control Protocol Works in conjunction with RTP. While RTP carries the media streams (audio and video), RTCP is used to monitor the transmission statistics and Quality of Service (QoS).
RTSP	Real-time Transport Streaming Protocol A networking protocol/URL address, commonly used in establishing point-to-point media sessions.
Sample Rate	The speed at which the Processing of the system takes samples respective to values from a continuous, analogue audio signal to make a discrete, digital one. For example, when running at 48kHz, incoming analogue audio is sampled at a rate of 48000 values per second.
SIP	Session Initiation Protocol; A networking protocol/URL address, commonly used within Voice-over-IP systems.
TDM	Time-Division Multiplexing; a common method of transporting signals via a point-to-point connection. In Lawo devices, TDM is used internally to transport audio along the backplane - e.g. from a I/O or DSP card to the routing matrix, and vice versa.

15 Standards

The following tables lists standards applicable to technologies discussed above. The list is not exhaustive. Some standards like the IETF RFCs are freely available while others like the SMPTE are not.

AES67	"AES standard for audio applications of networks - High- performance streaming audio-over-IP interoperability"
IEEE 1588-2008	"IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems"
IETF RFC 1112	"Host Extensions for IP Multicasting"
IETF RFC 2236	"Internet Group Management Protocol, Version 2"
IETF RFC 3190	"RTP Payload Format for 12-bit DAT Audio and 20- and 24-bit Linear Sampled Audio"
IETF RFC 3497	"RTP Payload Format for Society of Motion Picture and Television Engineers (SMPTE) 292M Video"
IETF RFC 3550	"RTP: A Transport Protocol for Real-Time Applications"
IETF RFC 3551	"RTP Profile for Audio and Video Conferences with Minimal Control"
IETF RFC 4175	"RTP Payload Format for Uncompressed Video"
IETF RFC 4541	"Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches"
IETF RFC 4566	"SDP: Session Description Protocol"
IETF RFC 768	"User Datagram Protocol"
IETF RFC 791	"Internet Protocol"
IETF RFC 4601	"Protocol Independent Multicast - Sparse Mode (PIM-SM)"
IETF RFC 5771	"IANA Guidelines for IPv4 Multicast Address Assignments"
IETF RFC 7042	"IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters"

SMPTE ST 2022-6:2012	Transport of High Bit Rate Media Signals over IP Networks (HBRMT)
SMPTE ST 2022-7:2013	Seamless Protection Switching of SMPTE ST 2022 IP Datagrams
SMPTE ST 2059-2:2015	SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications
VSF TR-01	Transport of JPEG 2000 Broadcast Profile video in MPEG-2 TS over IP
VSF TR-03	Transport of Uncompressed Elementary Stream Media over IP
VSF TR-04	Utilization of ST-2022-6 Media Flows within a VSF TR-03 Environment
SMPTE ST 2110	Professional Media over Managed IP Networks

IP NETWORKING GUIDE FOR VIDEO AND AUDIO APPLICATIONS

© 2020 Lawo AG. All rights reserved. Windows is a registered trademark of Microsoft Corporation. Other company and product names mentioned herein may be trademarks of their respective owners. Product specifications are subject to change without notice. This material is provided for information purposes only; Lawo assumes no liability related to its use. As of March 2020.

HEADQUARTERS

Lawo AG
Rastatt
GERMANY
+ 49 7222 1002 0
sales@lawo.com

INTERNATIONAL OFFICES

BENELUX	+ 31 6 54 26 39 78
CANADA	+ 1 416 292 0078
CHINA	+ 86 10 6439 2518
NORWAY	+ 47 22 106110
SINGAPORE	+ 65 9818 3328
SWITZERLAND	+ 49 7222 1002 0
UK	+ 44 333 444 5296
USA	+ 1 888 810 4468

RENTAL SERVICE

+ 49 7222 1002 0
rental@lawo.com



www.lawo.com

