370

TI 370
Media Integrated Local Area Networking (Milan™)
1.2 en

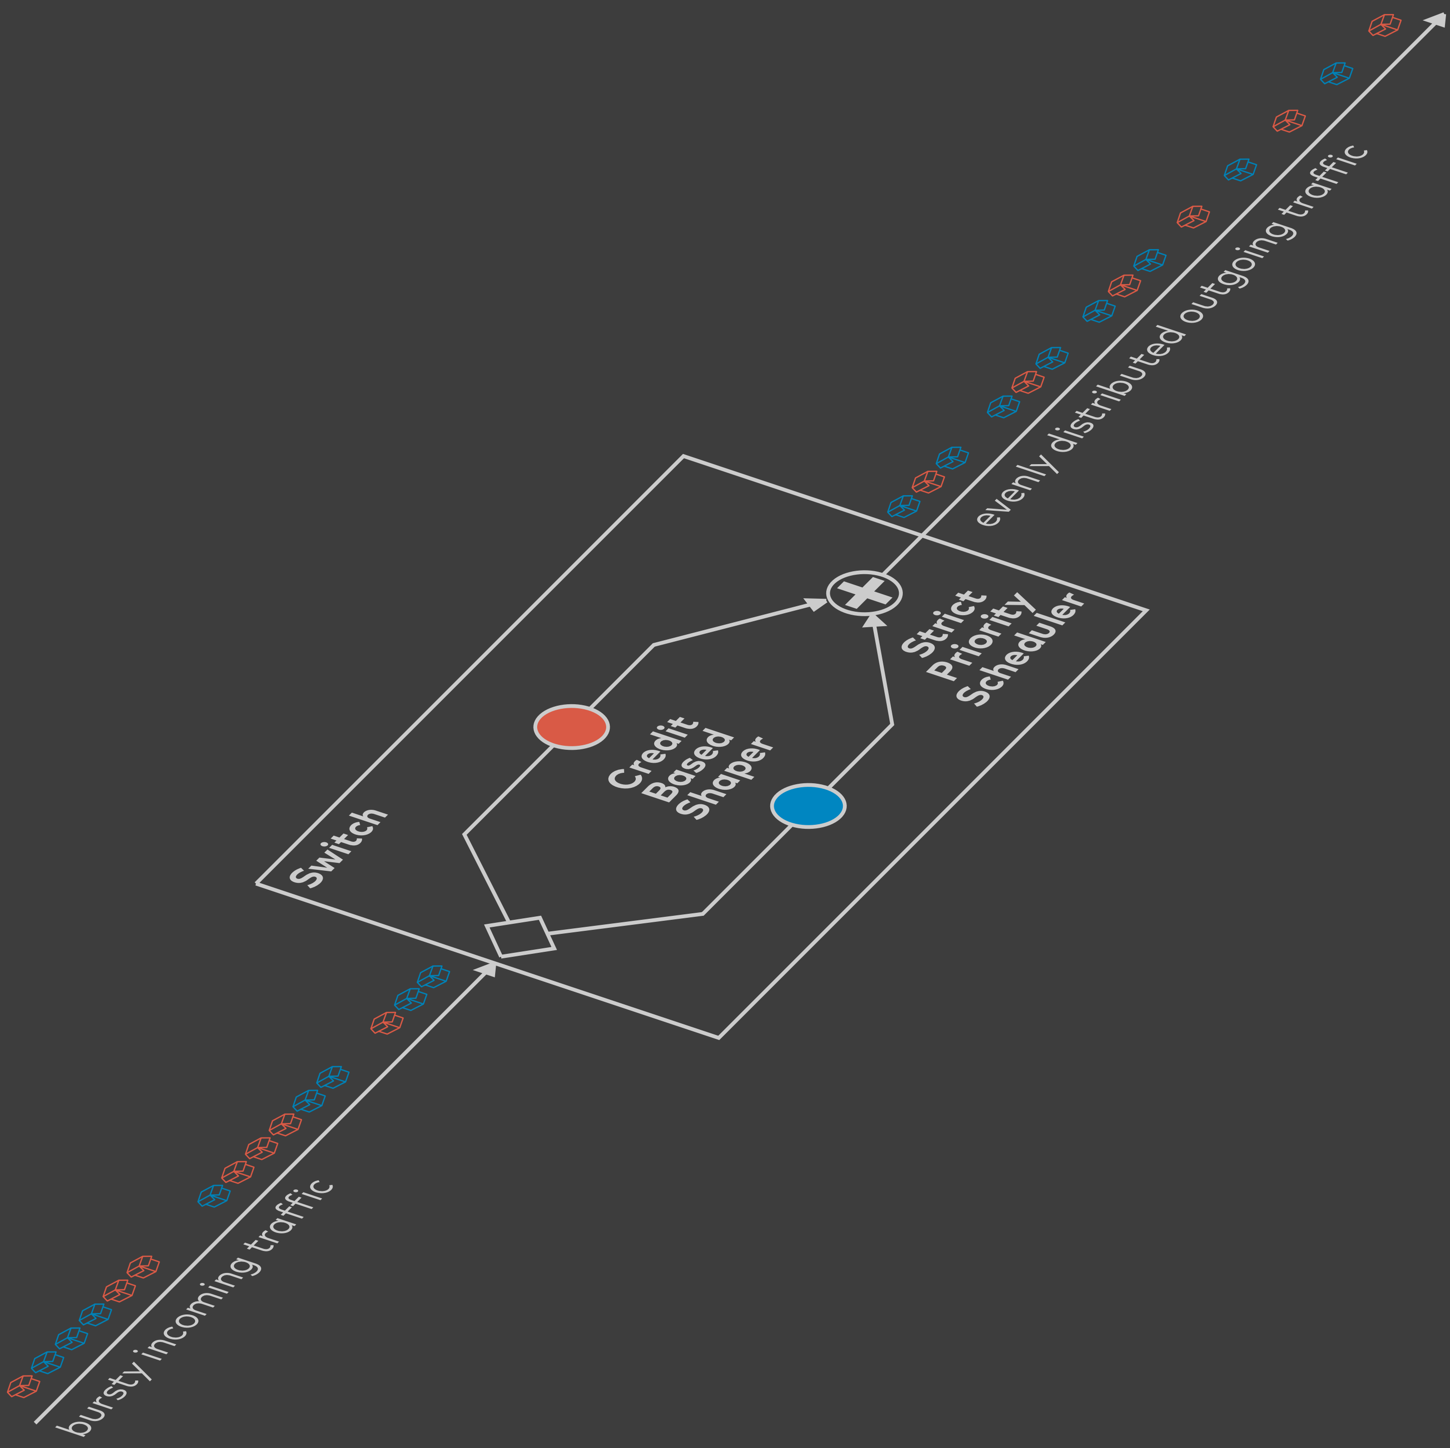

Notes on document version
All previous versions of this document are hereby no longer valid.

Version 1.2:
Notes on "Milan Manager" added.

Refer to:
⇒ Chapter 5.1 “Milan Manager” on page 13

General information

TI 370 Media Integrated
Local Area Networking (Milan™)

Version: 1.2 en, 08/2024, D5370.EN .01

d&b audiotechnik GmbH & Co. KG
Eugen-Adolff-Str. 134, D-71522 Backnang, Germany
T +49-7191-9669-0, F +49-7191-95 00 00
docadmin@dbaudio.com, www.dbaudio.com

Contents

1 Milan

Media Integrated Local Area Network (Milan) is a network protocol for real-time media that has been defined by the Avnu, a consortium of manufacturers mainly from the pro audio, network technology and electronics industries. It is based on Time-Sensitive Networking (TSN) which is a collection of Institute of Electrical and Electronics Engineers (IEEE) standards that make Ethernet deterministic by default. This targets Milan at applications that must deliver precise synchronization and low latency in large, converged networks, even when operated at high load, which is why d&b audiotechnik sees it as the network backbone for our systems.

Milan is not a proprietary product. Rather, manufacturers are free to implement it on a variety of hardware platforms as required. At the same time, a stringent standardization and certification process ensures interoperability of Milan certified devices.

1.1 History

Due to being based on open standards that developed over time, the terminology used is not always consistent. Therefore, it makes sense to look at a timeline that relates the overall evolution of Ethernet, the specific work of the respective IEEE working groups regarding the subject of this document and the work of the Avnu to each other.

1983	The IEEE standardizes a set of wired computer networking technologies under the name "Ethernet". Initially, data could be transferred over a coax cable at 10 Mbit/s.
1991	The IEEE standardizes 10BASE-T Ethernet, which first uses twisted pair cables with a data rate of 10 Mbit/s.
1995	The IEEE standardizes 100BASE-T or Fast Ethernet with a data rate of 100 Mbit/s over twisted pair cables.
1999	The IEEE standardizes 1000BASE-T Gigabit Ethernet with a data rate of 1 Gbit/s over twisted-pair cables.
2004	Within the IEEE, the "Residential Ethernet" study group is formed, later to become the Audio-Video Bridging (Milan) Task Group.
2009	The Avnu Alliance is founded as an industry consortium to develop open interoperability standards and certifications and commercialize the technology.
2011	The IEEE's AVB Task group publishes a set of technical standards that serve to provide time-synchronization, low latency, and reliability for real-time media transfer over switched networks under the common name "AVB".
2012	The Milan Task Group is renamed to TSN Task Group to reflect the much wider scope of not just ProAV, but many other industry's applications that time-synchronized, low-latency data transmission addresses. Based on TSN, several profiles are defined that address the special needs of these various industries. AVB is the profile for the ProAV industry.
2018	The Avnu publishes the Milan protocol, which is based on the AVB profile of TSN. Milan declares a specific implementation of several IEEE standards to ensure interoperability of devices from different manufacturers. Milan conformity is verified by a stringent certification process that devices must pass to carry the Milan logo.

2 Relevant technical background

2.1 Preliminary remark

This chapter employs simplifications where deemed meaningful. The goal is to explain principles and their differences at a level of detail that is appropriate for the scope of this document and a normal, interested user who is not a network expert.

For further reading, an up-to-date list of applicable IEEE standards can be found in the latest revision of the Milan specification which can be downloaded from the Avnu website: https://avnu.org/resource/milan-specification/

Also, due to changes in terminology to make technical terms more inclusive, multiple variations of the same term may exist, depending on when documentation was released and/or standards were last updated in this regard. This should be noted because it may cause confusion when comparing different documents about the same subject matter. As a prominent example, the long-established hierarchical and universal terminology of "master/slave" is increasingly replaced by other terms, which can differ depending on the context ("leader/follower", "transmitter/receiver", and others).

2.2 How real-time media transfer is achieved in Ethernet networks

As initially developed, Ethernet fulfilled the goal of networking computers together in establishing a common standard for the transfer of packet-oriented data between devices. Together with higher-layer protocols such as TCP, the design focus was on establishing a connection and ensuring that data packets reach their destination. Attributes such as latency or synchronization were not part of these design goals.

With the widespread proliferation of network technology in all fields private and commercial and with Ethernet evolving towards much higher bandwidths toward the end of the 1990s came the desire to transmit uncompressed, real-time audio and video content. Here, deterministic, low latency and precise synchronization are crucial.

To this end, additional standards have been developed around the turn of the millennium to leverage, when implemented, hardware that is already widely in use. They have enabled several (semi-) proprietary Audio over Ethernet and Audio over IP solutions such as CobraNet, EtherSound, Dante, AES67 or Ravenna that were or still are commercially highly successful for this reason.

Latency management by forced preferential forwarding

To achieve low latency, preferential forwarding of specific packets is forced by implementing Quality of Service (QoS) techniques that introduce classifiers at the packet level (DiffServ, "Differentiated Services"). These can be read and used by so-equipped network switches to apply a class-specific forwarding rule, buffering all lower-priority traffic (e.g., a print job or non-critical control data) until concurrent higher-priority traffic (e.g., synchronization data, media data) has passed. This method is referred to as Strict Priority scheduling.

The illustration below shows how a Strict Priority scheduler in a switch handles incoming traffic of different classes. Packets ingressing the switch through different ports and which are forwarded from one outgoing port are processed according to their time of ingress and priority. In the example, traffic of lower priority is held back significantly to let packets of higher priority pass first, even though both ingress the switch at the same time. The result is a continuous outbound stream of packets, sorted by class.

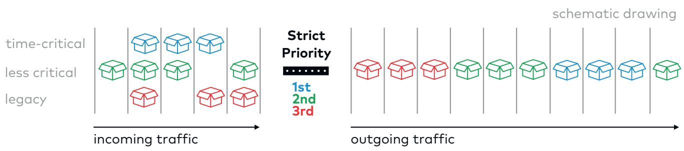

Synchronization via PTPv1

Synchronization is achieved by implementing the Precision Time Protocol (PTP). It has been published in different versions, the first of which (PTPv1) has been released in 2002. A prominent protocol to use PTPv1 is Dante.

The overarching principle is always identical: One network device, in case of PTPv1 an endstation (a network device with audio I/O), must provide a common time reference for all others to synchronize to. This endstation is usually elected by an automated process and is called the Leader Clock.

The Leader Clock periodically sends out Sync messages that contain its current time of day. All other endstations, the Followers, compare the time stamp of incoming Sync messages to their own clock to determine a conversion formula that translates Leader Clock time to Follower Clock time:

> *"When I received the Leader Clock's Sync message, it read 10:00:00 h. At this moment, my own clock reads 09:59:00 h. So, I need to subtract 00:01:00 h from incoming time stamps in Leader Clock time to get the corresponding time of day on my clock."*

Additionally, every Follower needs to determine the propagation time or path delay of Sync messages through the network to offset its conversion formula accordingly. To this end, all Followers periodically send Delay Request messages to the Leader Clock which in turn responds with a Delay Response message. The measured round-trip time between sending the Delay Request and receiving the Delay Response is divided by two to determine the path delay.

> *"I know that a Sync message from the Leader Clock takes 00:00:02 h to reach me. This means that I must add this time to the value of every Sync message before applying the conversion formula."*

Being able to convert the time stamp of Sync messages to their internal time and knowing the path delay allows multiple endstations to play out a given audio sample at the desired time coherently.

This process is called "End-to-End Synchronization" because the Delay messages traverse the whole network, from Follower to Leader Clock, and back.

PTPv1: End-to-End synchronization

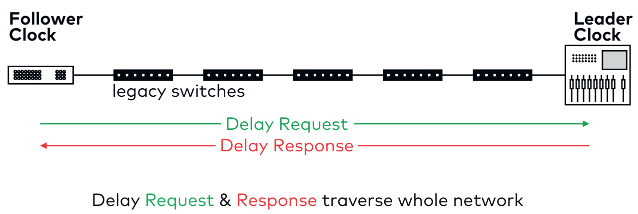

Synchronization via PTPv2

First published in 2008, PTPv2 is an improvement over PTPv1 that makes it suitable for large networks at high load and with a high endstation count. Specifically, PTPv2 introduces the option to utilize specific network switches which are "time-aware".

In contrast to legacy switches, instead of just forwarding packets, time-aware switches know the residence time of packets within the switch and add this information to the packets. This makes the determination of path delay much more precise. However, PTPv2 is not downward-compatible to PTPv1. Common protocols that utilize PTPv2 are AES67 and Ravenna.

2.3 Challenges in converged networks

The schemes described in the previous section perform well in networks that

1. transport only one dominant kind of traffic.
2. have an abundance of free bandwidth.

In many systems, these conditions are easily met. For historic reasons alone, different departments (audio, video, lighting, etc.) tend to run separate physical networks. Also, and in many typical applications, only a fraction of the bandwidth of a Gigabit Ethernet link is used, and higher bandwidths for inter-switch links are increasingly commonplace.

Example:
128 channels of 96 kHz/24-bit audio require less than 400 Mbit/s of bandwidth or less than 40% of a Gigabit Ethernet link.

The situation changes when network links are operated closer to capacity, which is more likely to happen in a converged network that combines all traffic of multiple departments on the same physical network, and especially across multiple switch hops. In this regard, it is irrelevant whether the network is subdivided into V-LANs, since traffic of multiple V-LANs competes for the same bandwidth and scheduling queues when traversing inter-switch links.

Strict Priority scheduling under high load can lead to bursts in traffic as accumulated lower-priority data that has been buffered by switches while forwarding higher-priority packets is transmitted all at once. This can overload ingress buffers on switches downstream when converging with other traffic, and results in significant variations in throughput latency and late, or worst case, dropped packets.

In the same way, the path delay of Delay *Request/Response messages* can vary significantly when traversing the network to the Leader Clock and back. As the calculation of the synchronization offset assumes a path delay that is identical for both directions, this introduces timing errors ("jitter"), which also vary with network load. Audio systems with many endstations that are expected to play out audio at the same time suffer from degraded coherence when this happens.

PTPv1: End-to-End synchronization at high network load

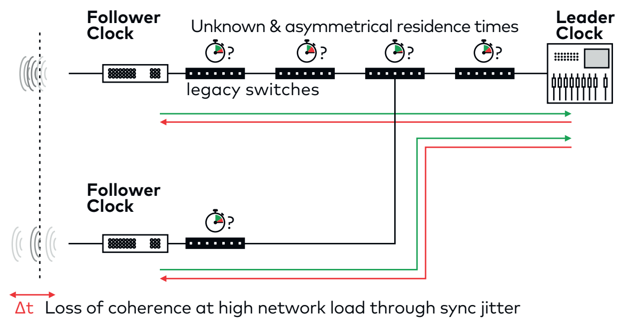

2.4 Technical improvements of TSN

TSN is still based on Ethernet, but it adds a collection of standards geared at maintaining Quality of Service regarding synchronization and latency in converged networks at high load. Compliant devices (endstations **and** switches) implement these standards at the hard- and firmware level, which means they do not need additional user configuration. This results in the following relevant functionalities.

Synchronization of all network devices including the network switches

TSN implements a special profile of PTPv2, called Generalized Precision Time Protocol (gPTP), which was first published in 2011 and formally added to PTPv2 in 2019. Some of the main differences are that it supports Layer 2 transport exclusively and that it **mandates** the use of time-aware network hardware instead of just making it an option.
This also means that TSN does not deal with IP addresses at all, which would be OSI Layer 3.

gPTP implements Peer-to-Peer (P2P) instead of End-to-End (E2E) synchronization, which means that *Sync* and *Delay* messages do not traverse the whole network, but only travel one hop to an adjacent switch. This is possible because all switches on the network must be time aware.

As a benefit, timing stability is almost unaffected by network size and even by extreme load conditions, staying within fractions of a microsecond at worst. This is in stark opposition to PTPv1 environments, where variations can reach several orders of magnitude at high network load. In addition, P2P synchronization considerably reduces the required network traffic and shortens sync times.

The Leader Clock, called a Grandmaster in gPTP terminology, is still elected automatically. Switches are favoured over endstations by the election algorithm.

gPTP: Peer-to-Peer synchronization

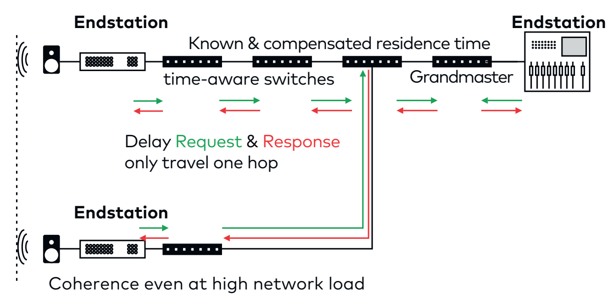

Additional traffic shaping to avoid bursts of packets

Traffic shaping describes methods by which switches manage their packet throughput. In the context of real-time media transport, this relates to ensuring that important data (such as synchronization and media data) is forwarded without significant variations in latency. Strict Priority forwarding as described before is a basic form of traffic shaping, as it prioritizes certain packets over others without any further constraints.

TSN still uses a Strict Priority Scheduler, but precedes it, in case of Milan, with a Credit Based Shaper. This packet processing stage also distinguishes different classes of service but assigns "Credits" or "forwarding rights" according to class. High priority packets are still treated as such but are not fed into the Strict Priority Scheduler without regard to other classes of traffic, even when arriving at the switch in bursts. The outcome is traffic that is not exiting switches in monolithic bursts but evenly distributed over time and with different priority classes interleaved. For the price of a small additional overall latency, this keeps end-to-end latency deterministic and avoids dropped packets even at high network load.

High priority traffic
Lower priority traffic schematic drawing

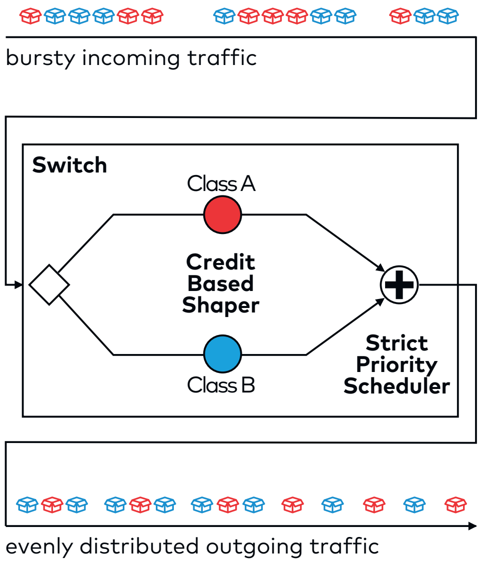

Bandwidth reservation

TSN employs bandwidth reservation, which works in several steps to ensure that the network is:

a. capable of transporting a media stream between the Talker (sending endstation) and a specific Listener (endstation that shall receive the media stream) before a transmission is started in terms of available vs. required bandwidth and
b. that an ongoing transmission cannot be interrupted by other traffic.
   Strictly speaking, not all TSN profiles implement bandwidth reservation in the same way, but for simplicity, the explanation given here relates to the ProAV profile.

Up to 75% of a link's bandwidth can be dynamically reserved for TSN traffic. All bandwidth not currently in use for TSN traffic, but at least 25% of a link's capacity, is therefore available for legacy traffic.

The network switches which must implement the necessary Stream Reservation Protocol (SRP), play a key role in this process, which works as follows:

1. The Talker advertises a stream to the network via multicast. Amongst other data, this Talker Advertise message contains information about the required bandwidth.
2. As it propagates through the network, every switch along the way reserves the needed bandwidth on the ingress port and declares it on all egress ports, if available.
3. If at any point, the bandwidth is not available, the switch returns Talker Failed to the Talker.
4. If a Talker Advertise reaches a Listener that intends to subscribe to this stream, it returns a Listener Ready message to the Talker along the same path.
5. This locks in the reservation in all switches along the path and initiates the actual stream on part of the Listener as soon as it receives the Listener Ready message.
6. Successfully reserved bandwidth cannot be used by other traffic and can only be actively released by the Talker or Listener or by the network being severed for long enough.

2.5 TSN Profiles

The general functionalities of TSN, combined with specific implementations in its various profiles, make TSN suitable not only for real-time media transport, but also for other highly time-sensitive applications that require tightly bounded latency. Examples are fly-by-wire controls in aviation, milling machines or beam-steering antennas for mobile networks. TSN allows for them to be remotely operated over a network without requiring local control units to achieve the desired performance and accuracy.
TSN therefore offers profiles for different industries, which implement specific functionalities as suitable.

- Automotive
- Industrial Automation
- Power Generation
- Mobile Communications
- ProAV

AVB is TSN's ProAV Profile.

3.1 Stream-based media transport, Dynamic Mapping, and patching

Milan streams

AVB, and with it, Milan, transports audio in streams. A stream is a container that encapsulates a defined number of audio channels of a given format (AM824 for legacy AVB, AAF for Milan). There can also be streams that only contain clock information, but no audio (CRF streams, see also section gPTP and Media Clock Domains).
In this context, all streams are transmitted as Multicast, which means that a Talker needs to transmit a stream only once, regardless of how many Listeners shall receive it.

Handling streams, whether as a Talker or a Listener, is a processor-intensive task, much more so than having more channels per stream. There is no general standard about how many Talker and/or Listener streams an endstation must provide or handle. The count varies based on the make and model of endstation, intended use and internal processing capabilities. Therefore, it is a relevant parameter to know when designing a system.

Dynamic Mapping - internal routing in Milan endstations

To give greater flexibility and to manage streams more efficiently, all Milan endstations with Listener capability provide an internal routing function, called Dynamic Mapping, between Listener streams and their physical outputs.
Milan endstations with Talker capability may also offer this to flexibly route physical inputs to Talker stream inputs, but it is not mandatory.

In the example shown here, a Listener shall output four audio channels (L, R, SUB, Fill) from a Listener stream to two physical outputs each. Without Dynamic Mapping, this would require two Listener streams. Also, it would still result in a lot of mouse clicks when using a channel-based patching approach.

Dynamic Mapping allows the duplication of audio signals from a Listener stream output to multiple physical outputs within a listener, as shown on the right, necessitating only one stream in total in such cases.

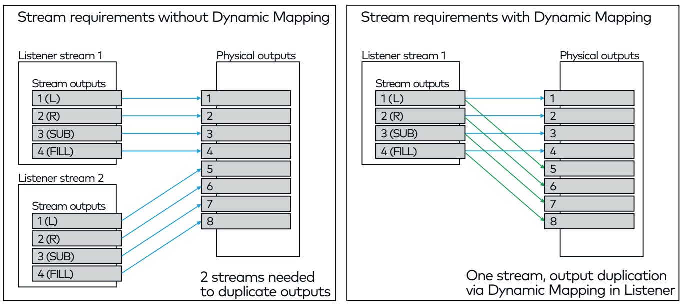

Functionally, Dynamic Mapping as shown in the previous example can be illustrated as shown below.

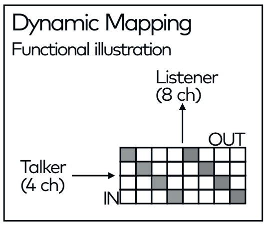

In total, Dynamic Mapping allows the user to tailor the stream output to physical output assignment of a Listener to requirements imposed by the application or the Talker and facilitates stream-based patching of audio connections which is operationally quicker than channel-based patches by saving mouse clicks.
A typical application would be sending the customary set of L/R/SUB/FILL audio signals to many endstations.

Depending on the controller software used, a procedure that looks like channel-based patching on the GUI is also possible, even though it still works with streams and all their requirements and constraints in the background. In this case, erasing all Dynamic Mappings in the Listener beforehand is recommended. The controller software will establish and remove them automatically based on the patching operations by the user.

3.2 gPTP and Media Clock Domains

gPTP establishes a common general time reference, delivered by the gPTP Grandmaster, also termed a "wall clock" for this reason, for all network nodes, which then form a gPTP Domain. The gPTP Grandmaster is elected automatically out of all endstations and switches on the network, whereby switches have a higher priority than endstations in the election process.

gPTP Domains can only be formed and audio can only be passed between endstations connected by AVB switches. Inserting a legacy switch into a gPTP Domain splits it into two independent gPTP Domains with their own gPTP Grandmasters.

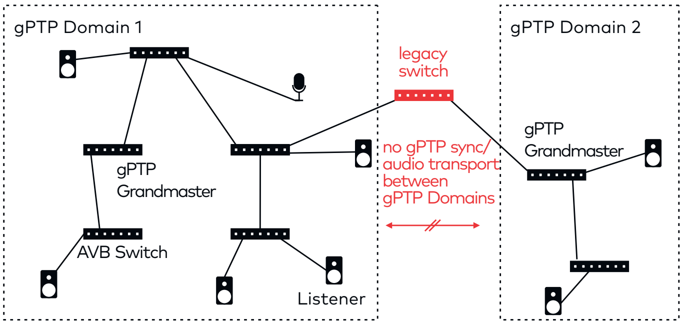

Media Clock via Clock Reference stream

In addition to a gPTP Grandmaster, endstations need a common rate at which samples are played out or recorded. This "Media Clock" is therefore different from the wall clock but relates to it, as sets of samples are delivered with a gPTP "Presentation time" and then processed at the specified sample rate.

Out of all endstations, a Media Clock Reference must be selected manually. This device then distributes its Media Clock as a separate Clock Reference Format stream (CRF stream). All endstations which subscribe to the same Media Clock Reference form a Media Clock Domain.

Multiple Media Clock Domains with arbitrary sampling rates and their own Media Clock References and associated CRF streams can coexist within the same gPTP Domain on the same LAN. Media (audio/video) can only be shared within the same Media Clock Domain.

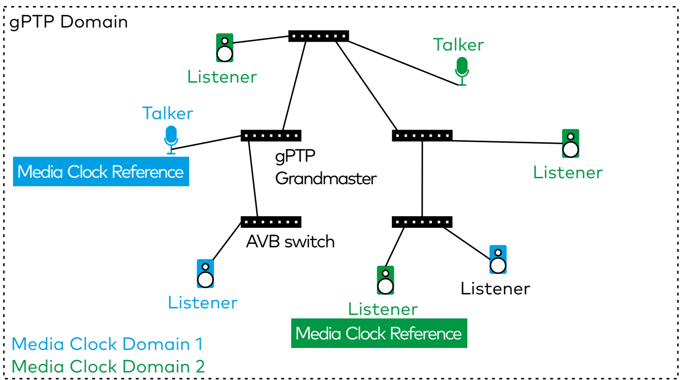

Specially equipped endstations can also be part of multiple Media Clock Domains. A typical application would be a matrix processor that routes audio signals between different rooms in a multi-venue building which are otherwise separate Media Clock Domains. However, this requires these endstations to perform internal sample rate conversion, which is why it is not a general feature.

Media Clock extraction directly from an audio stream

As an alternative to a CRF stream, Listeners can also extract the Media Clock from an audio stream they receive. This method is used in simple devices which can only receive one stream in total and therefore do not have the capacity for a separate CRF stream in addition to the AAF audio stream. In this case, the Listener joins the Media Clock Domain of the respective Talker.

3.3 Latency

Transporting a stream from Talker to Listener through the network needs time, mostly for traversing switches. The system must account for this latency by setting an appropriate user-defined value in the Talker. It can be set in steps between 125 µs and 2 ms.

The link speed and the associated worst-case latency per hop determine the permissible number of hops between Talker and Listener.

	100 Mbit/s network	**1 Gbit/s network**
Latency per hop (worst case)	280 µs	140 µs
Max hops within 2 ms	7	14
Max number of switches between Talker and Listener within 2 ms	6	13

A realistic scenario that uses a 1 Gbit/s network should never exceed the maximum possible hop count, and shorter latencies than 2 ms can usually be set. 1 ms (= 1000 µs) should account for all reasonable topologies, as illustrated by the graphic below. Using infrastructure-sized switches at FOH and in every amp city, a system rarely exceeds 5 hops.

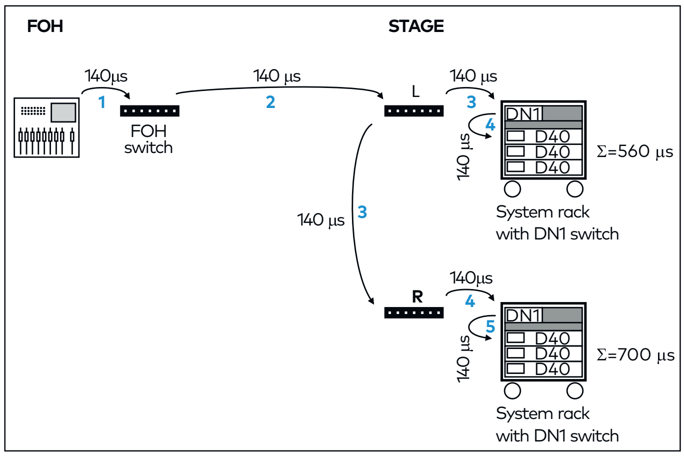

3.4 Endstation naming

Milan endstations offer several naming options to help with the functional structuring of a system (examples in brackets). Names can be freely assigned and do not impact patches.

- Group name (e.g. "FOH", "Stage")
- Endstation name (e.g. "Main console 1")
- Stream name (e.g. "Playback tracks" for Listener streams, "PA feed" for Talker streams)
- Channel name (e.g. "Stem 1 L", "Stem 1 R" for Listener channels, "L", "R", "SUB", "FILL" for Talker channels).

3.5 Interoperability definitions of Milan

While AVB as TSN’s ProAV profile defines general attributes, further standardization is necessary to assure interoperability between devices of different manufacturers, which makes the technology usable in typical, heterogenous applications.

These additional constraints, as defined by the Avnu, are the Milan protocol. While it contains more details at the protocol level, the user-relevant specifications are described below.

Stream formats

Endstations can only transmit to one another if they use a common stream format. Thus, every Milan endstation must support at least one Listener stream of the following properties.

Audio format	AAF
Channel count	1, 2, 4, 6, 8
Sampling rate	48 kHz or 96 kHz or 192 kHz
Sample bit depth	32 bits

The total number of Talker and Listener streams per endstation varies and depends on the device, its physical I/O count and resources and its intended application.

Redundancy

Every Milan endstation may but does not have to support redundancy. This is enacted using two designated network interfaces on the endstation (PRImary and SECondary), which are connected to separate primary and secondary physical networks.

Redundancy does not change the available stream count on endstations. I.e., an endstation with two Listener streams and a primary and secondary network interface supports two Listener streams per interface, but not four Listener streams on the primary interface if the secondary connection is not in use.

4.1 How to set up a Milan-based system

Topology

- Use only AVB-certified switches to interconnect Milan endstations. Some switches might need to be set to "AVB mode".
- Deploy your network in a star topology to minimize the number of hops. Place switches with enough ports at FOH and in every amp city. Connect everything (consoles, matrix processors, stage boxes, system amplifier racks with AVB endpoint switches) to these central switches and avoid daisy-chaining of racks or other equipment.
- Consider fiber links for the FOH to stage connection. Typical cable lengths of 80 - 100 m are at the limit of what copper cables are capable of.
- If possible, use only etherCON connectors for copper links for a more robust mechanical connection to the individual devices.
- In case of redundant links, do not run the primary and secondary network cables next to one another from FOH to stage, if possible, to prevent one physical event from severing both at the same time.

  Note: Usually, AVB switches only support one gPTP Domain per switch. It is technically not possible (and it would not be a promising idea from a redundancy point of view) to enact a redundant AVB network using two V-LANs on the same switch.

Talker configuration

- Determine the worst-case number of hops in your network by counting the switches between your Talker(s) and the farthest Listener(s). On a 1 GBit/s network, the expected maximum network latency for your worst-case path is:

  (Number of switches between Talker and Listener + 1) • 140 µs

- Set your Talker to this value or the next higher one if the exact value is not available. Common controller software will also issue a warning if the actual value exceeds the configured latency.

Sampling rate

Set your sampling rate before patching any channels or streams. Changing it later will interrupt audio reproduction.

Media Clock Reference/Media Clock Domain

Associate all desired endstations to the same Media Clock Domain by subscribing them to the same CRF stream/assigning them to the same Media Clock Domain in the respective dialog (controller dependent).

Patching: Streams or channels?

- Consider patching streams instead of individual channels. For example, typical PA applications benefit from having L/R/SUB/FILL available at every system amplifier and using the input routing to select the desired channel or matrix them together as required.
- Before patching, check the Dynamic Mapping of all endstations to avoid unexpected results.

Redundancy check

When using redundant networks, it makes sense to check both the primary and secondary network individually. While d&b system amplifiers and audio network bridges indicate the presence of the respective Talker for both networks, it is still meaningful to check manually, especially when using heterogenous systems:

- When all streams are patched and audio is playing, first physically disconnect the primary network from the Talker. Audio should not be interrupted.
- Reconnect the primary network and wait for it to synchronise and settle. This process takes at least 10 s.
- Once the system indicates sync and has settled, disconnect the secondary network. Again, audio should not be interrupted at any endstation.

5.1 Milan Manager

Milan Manager is a cooperation between d&b audiotechnik and L-Acoustics to create a common Milan controller software that is user-focused and easy to operate.

It covers all basic operational requirements and will be constantly expanded with further features. Please refer to its online help section for details about individual functions. Milan Manager can be downloaded free of charge from the d&b audiotechnik and L-Acoustics websites.

> **Note:** Until this functionality is also implemented into Milan Manager, firmware updates for Milan endstations have to be performed using Hive. See the following section for guidance.

5.2 Hive controller (Quick start guide)

Even though it is increasingly a legacy Milan controller software and more of a developer tool, Hive is relatively widespread.

There is no official manual for Hive, so this section gives a quick overview of its basic operation.

Hive controller - Overview of the most important controls

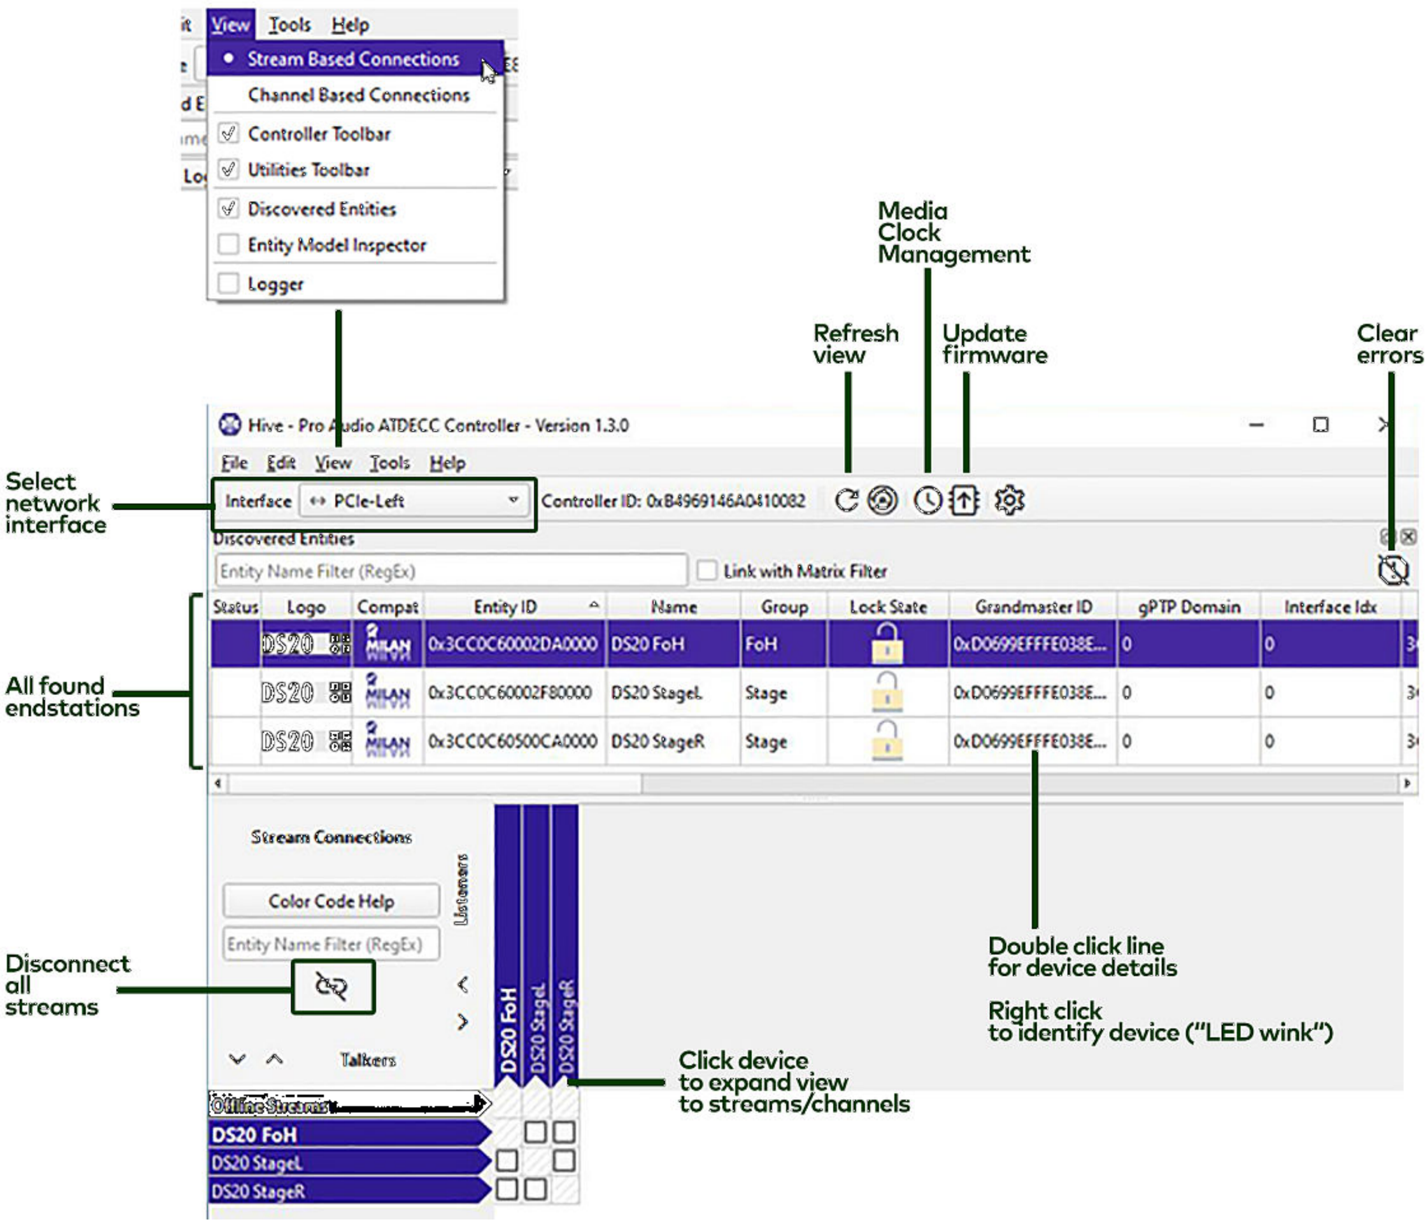

5.3 Basic configuration and use

Configure interface

⇒ From the «View» menu select:

- Controller Toolbar
- Utilities Toolbar
- Discovered Entities

Select network interface

⇒ Use the drop-down menu to select the correct network interface of your computer.

↳ No IP address configuration is required.

If in doubt, make sure your firewall does not block the application.

Discovered endstations ("Entities") should appear after a few seconds.

Configure endstations

Apart from names, endstations can only be reconfigured if no streams are connected to them.

1. Disconnect the respective endstation if necessary or use the "Disconnect all streams" button.
2. Double click an endstation in the list to open its Device View.
3. Configure:
   - Sampling rate
   - Latency (if this endstation is the Talker)
   - Names
   - Stream formats for Listener and Talker streams

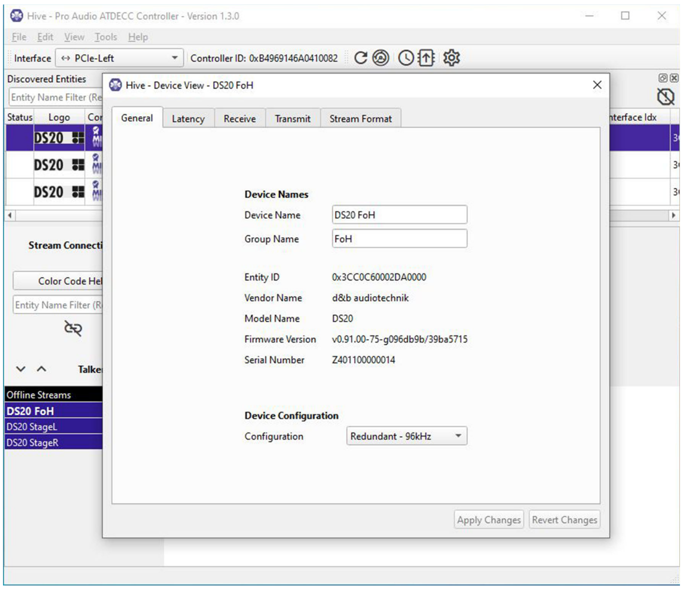

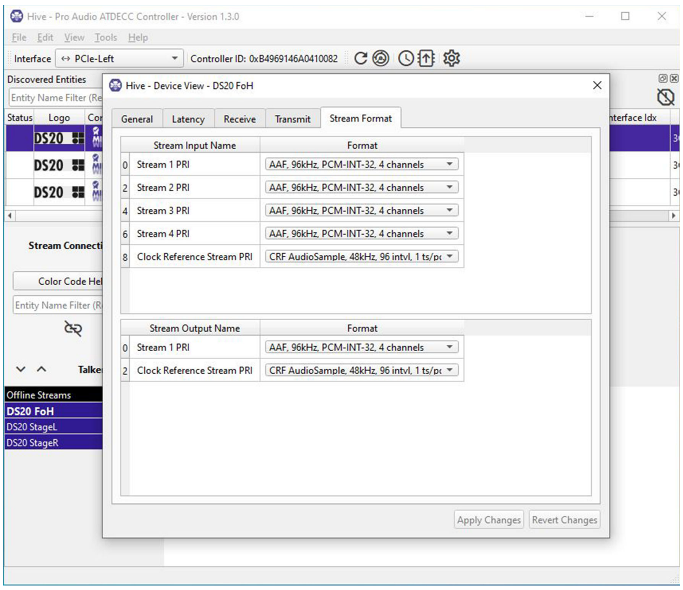

Configure a Media Clock Domain

1. Open Media Clock Management.
2. Add a Media Clock Domain by clicking «(+)».
3. Assign endstations to the Media Clock Domain.
4. Set the Media Clock Reference in the "Master" column.

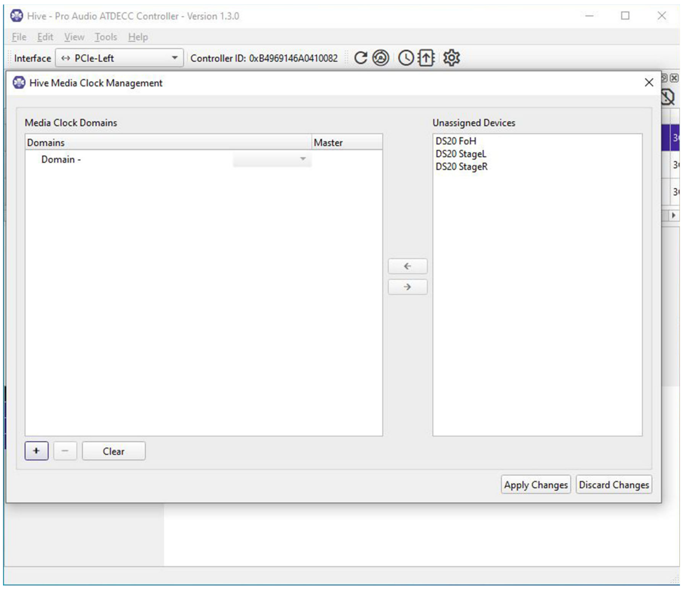

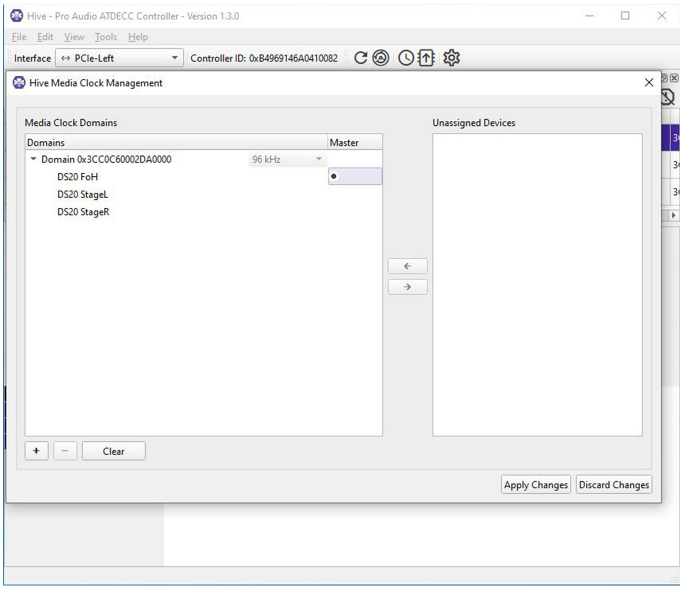

⇒ In the connection matrix, this will show up as a connected CRF stream.

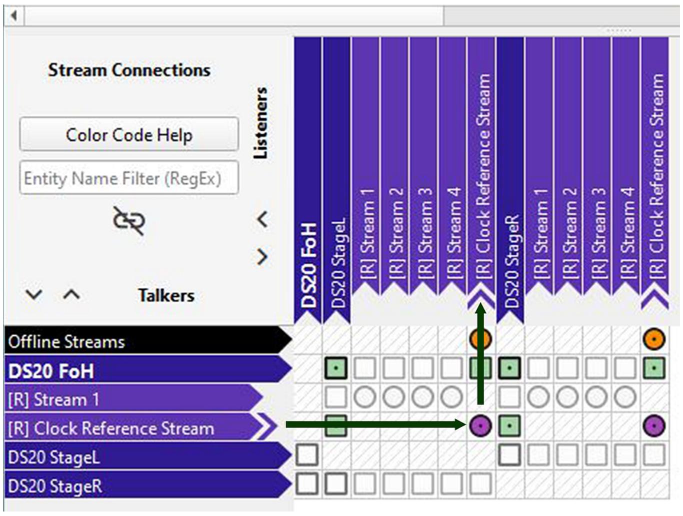

Check/configure Dynamic Mapping

1. Right click an endstation in the Listener row to configure its Dynamic Mapping.
   ↳ To prepare for channel-based patching, delete all the links of Listener streams to physical outputs. Drag from the physical output into a white area.

   To prepare for stream-based patching, configure the links of Listener streams to physical outputs as required by dragging from a stream channel to the physical output. Several physical outputs can be connected to one stream channel.
2. Double check the Talker stream Dynamic Mapping (insofar the endstation supports it) by right clicking the Talker in the Talker column.

Patching

⇒ Set marks in the connection matrix to patch or unpatch as required.

Error handling

Unplugging a patched endstation will raise an error in the "Status" column.

⇒ Clear by clicking "Clear errors."
   ↳ Refresh view.

6 Glossary

Bounded latency
Bounded latency for packets means that the time they require to travel through a network is guaranteed not to exceed a pregiven value. It is one of the key factors in a TSN network.

Bridge
The technical term for "switch" used by the IEEE.

Converged network
This term describes a network that combines all network-based communication, e.g., media transport and control data of all kinds. This contrasts with using multiple logically or physically separate networks. Converged networks mandate standards and hardware that guarantee precise synchronization and bounded latency to function.

Endstation
An endstation is any network device that provides audio I/O and that is not exclusively a switch.

gPTP Grandmaster
The Grandmaster is the central clock within a TSN network, also termed the "wall clock" (visible to all devices on the network) that provides a common time reference with nanosecond precision. Out of all endstations and switches on the network, the Grandmaster is automatically elected using the Best time Transmitter Clock Algorithm (BTCA). The algorithm favours switches over endstations.

IEEE
The Institute of Electrical and Electronics Engineers is one of the leading associations for electronics and electrical engineering and related disciplines. It has published a significant amount of the world's professional literature in the related fields and hosts many standardization committees that define the design and interaction of devices and applications in many technical fields.

Listener
A Listener is an endstation that can receive streams.

Media Clock Domain
All endstations in an AVB network that shall exchange audio must be part of the same Media Clock Domain. This means that they all follow the same Media Clock Reference regarding the rate at which they play out or record audio samples. There can be multiple Media Clock Domains in the same physical AVB network, each with their own Media Clock Reference.

Media Clock Reference
The Media Clock Reference (previous known as Media Clock Master) provides the rate at which samples are played out within a Media Clock Domain. In contrast to the gPTP Grandmaster, this is not a "time of day" index. Blocks of samples in a Milan AAF stream contain a presentation time that relates to the gPTP Grandmaster time at which the first sample shall be played out, with the subsequent samples following at the rate given by the Media Clock Reference.

Presentation time
The gPTP Grandmaster time index, set by the Talker, at which a sample is to be played out by Listener.

Stream
A stream is a container that contains a defined number of audio channels of a given format. There can also be streams that only contain clock information, but no audio. In a TSN network, all streams are transmitted as Multicast. Multicast means that even if there are multiple Listeners for a specific stream, the Talker needs to transmit it only once. This is opposed to Unicast, where every Listener receives its own stream, even if it contains the same audio channels as an existing one, multiplying the network load for the Talker.

Talker
A Talker is an endstation that can send streams.

Time aware switch
In legacy switches, the residence time of a packet (the time required for a packet to ingress the switch, get processed, and egress again) is unknown to the switch. This adds unknowns to the determination of path delay which in term causes synchronization jitter. Time aware switches do have a concept of time. They add information about the residence time to packets, so that the synchronization process can take it into account. The Avnu Alliance Certified Product Registry lists all certified endstations and switches:
https://avnu.org/certified-product-registry/

Traffic shaping
Traffic shaping is a technique to manage bandwidth utilization in a data networks to optimize or guarantee performance by preventing congestion. A traffic shaper handles different kinds of packets by predefined rules which govern the rate at which they can be sent.

370

www.dbaudio.com

d&b audiotechnik