

AES67 PRACTICAL GUIDE



Content

1 Basic Principles of AES67	2
1.1 Synchronization.....	2
1.2 Multicast packet transport	2
1.3 Quality of Service.....	3
1.4 Session information.....	4
2 Guidelines for Configuration of an AES67 System.....	5
2.1 System planning	5
2.2 Device configuration	11
2.3 Check for proper synchronization (PTP)	12
2.4 Stream configuration	14
3 Outlook on emerging Technologies and Industry Standards.....	26
3.1 ANEMAN.....	26
3.2 NMOS	26
3.3 SMPTE ST 2110.....	27
4 Conclusion	28

Abstract

AES67 is a standard for high-performance audio-over-IP interoperability. By implementing the standards definitions and requirements, devices otherwise adhering to specific AoIP protocols that don't interoperate with each other (i.e. Dante, Livewire, QLAN or RAVENNA) can establish direct connectivity and become interoperable.

However, while the standard defines what protocols and functions need to be supported, it still leaves various choices open to the implementer and – by nature – doesn't serve very well as a user's guide; thus, some background knowledge on networking in general and on AoIP-related topics in particular is certainly helpful when planning for an AES67 setup.

This guide explains some of the choices and ambiguities left open by the standard and describes how to circumvent the most commonly observed obstacles when setting up an AES67 network.



1 Basic Principles of AES67

AES67 is intended to run on standard packet switching networks build from COTS¹ infrastructure. If configured properly, other traffic can share the same network without degrading the audio streaming experience.

AES67 builds on these fundamental principles:

- Synchronization
- Multicast packet transport
- Quality of Service
- Session information

While other AoIP solutions offer enhanced functionality (i.e. stream & device discovery, GPIO transport etc.), AES67 has deliberately not defined any requirements in this respect, because they are not required to establish interoperability on the most basic level. Furthermore, various industry standards covering these functions already exist or are emerging² and can be implemented if applicable.

1.1 Synchronization

Synchronization is based on distribution of a common wall clock time to all participating nodes with sufficient precision. AES67 specifies the IEEE1588-2008 standard (also known as PTPv2 - Precision Time Protocol version 2)³ to be used for time distribution. Note that PTPv2 is not backward-compatible with PTPv1 (IEEE1588-2002). PTPv2 includes a Best Master Clock Algorithm (BMCA) which ensures that the best available master clock is elected to serve as the Grandmaster for all participating AES67 nodes. In the guidelines section, hints are provided on how the Grandmaster selection can be modified, if required. Once a node is synchronized to the wall clock time served by the Grandmaster, any desired media clock can be generated locally. If the synchronization precision is accurate enough, all locally generated media clocks will have the same frequency (i.e. 48 kHz) and they may even be accurately phase-locked to each other.

With PTP it is possible to achieve accuracy in the sub-microseconds range (deviation of local clock with respect to the Grandmaster). However, in most cases this requires the deployment of PTP-aware switches. Fortunately, for most audio applications single-digit microsecond accuracy is still good enough, which usually is achievable with standard, non-PTP-aware switches.

1.2 Multicast packet transport

While PTP is based on multicast packet transport, AES67 also mandates for multicast support of audio stream packets. While basically any COTS switch supports multicast traffic, only managed switches provide multicast management to effectively avoid network flooding. Unmanaged

¹ Conventional-Off-The-Shelf

² DNS-SD / mDNS, NMOS IS-o4 etc.

³ https://en.wikipedia.org/wiki/Precision_Time_Protocol

switches (or improperly configured managed switches) will treat multicast traffic like broadcast traffic, forwarding any incoming multicast packet to all switch ports. With high levels of audio stream traffic this will result in network flooding and can result in a total network lock-up.

Managed switches provide multicast traffic management through IGMP⁴ support. With IGMP, registered multicast packets are forwarded to their designation ports only. Consequently, all AES67 nodes are required to support IGMPv2 which is used to tell the network which streams are to be received. The IGMP (join) requests need to be updated periodically in order to maintain the multicast forwarding. This is ensured by enabling the IGMP querier function in one of the participating switches, which sends out periodic IGMP queries triggering the nodes to renew their IGMP requests. Once a stream connection is terminated, an IGMP leave request will be sent out to signal termination of a particular multicast flow to that node.

One of the benefits of managed multicast traffic is its scalability: any multicast flow is only sent once by a particular sender into the network. If more than one receiver requests the same flow, the network switches will clone packets as required. With IGMP the network inherently optimizes the traffic, so that a particular multicast flow will be present on any involved link just once.

1.3 Quality of Service

Quality of Service (QoS) is another fundamental principle which needs to be supported by the network – again, only available with managed switches. Proper QoS configuration ensures that the most critical packets – PTP and audio stream traffic – receive prioritized forwarding on their way through the network. AES67 mandates for support of *Differentiated Services* (DiffServ)⁵, a QoS scheme where different types of traffic can be categorized into service classes. DiffServ works with 64 different priority tags – DSCP⁶ values – which can be applied to individual IP packets. End nodes can apply different tags to different traffic classes; switches can then examine the individual priority tags and forward packets on a preferred basis. Put simply, with DiffServ a network works in a very similar way to the boarding procedure at airports - priority passengers (first and business class) board the plane first (and at any time), while economy class passengers have to wait in line as long as priority passengers are still queuing up.

However, while recommendations exist in the standards on how to assign DSCP values to certain types of traffic, a network administrator is free to configure these values according to the individual application requirements. Specifically, in larger network environments with a variety of shared traffic classes, QoS configuration requires special attention.

AES67 requires the use of 3 traffic classes and recommends certain DSCP values:

- PTP traffic should be tagged with DSCP EF (46), translating into *expedited forwarding*, receiving the highest forwarding priority (first class passengers)

⁴ Internet Group Management Protocol

https://en.wikipedia.org/wiki/Internet_Group_Management_Protocol

⁵ https://en.wikipedia.org/wiki/Differentiated_services

⁶ Differentiated Services Code Point – a number in the range of 0..63



- RTP traffic (audio packets) should be tagged with AF₄₁ (34), translating into *advanced forwarding* with the second-highest forwarding priority (business class passengers)
- All remaining traffic should have no priority tagging, which is BE (0) for *best effort* (economy class)

Since the network may be used to transport other types of traffic which needs certain prioritization, (i.e. voice data or video flows), the network administrator may have to change these values or adopt the switch configuration accordingly. Since not all AES67 devices support DSCP reconfiguration (and / or use other values by default), other strategies may have to be applied (see guidelines section).

1.4 Session information

In order to connect to an available stream and process its audio data, a node needs technical information about the stream. This is called session description data (SDP⁷) and contains information about the multicast address of the stream, the encoding format and packet setup (i.e. bits per sample, sampling frequency, number of channels in stream, number of samples in packet) and its relationship to the reference time. Without this information, a receiver would not know how to connect to the stream and how to decode the packet content.

While AES67 clearly defines any required SDP attribute and their allowed parameter ranges, it is silent on the required method to convey this information. Session discovery (which would allow for system-inherent detection of available streams) has also been deliberately excluded from the standard requirements. While a number of protocols exist to announce available streams and transport the related SDP data, the creators of the AES67 standard felt that it would have been too stringent to actually mandate for a specific method; instead, they decided to just mention some widely used protocols and to leave it to the device manufacturer to decide on and implement those methods which would best suit their typical application environments⁸.

While most devices support mDNS/RTSP⁹ (the default RAVENNA method) and SAP¹⁰ (Dante devices in AES67 mode), not all devices support both methods, and some don't even offer manual read-out / entry of SDP data. If there is no common method of sharing the required SDP information between two devices, stream connection setup may be impossible or at least very difficult. The guidelines section provides hints on how to potentially circumvent this problem.

⁷ Session Description Protocol - https://en.wikipedia.org/wiki/Session_Description_Protocol

⁸ The AES67 standardization Task Group simply assumed that with no mandatory protocol in place a device would provide a method to read / enter the SDP data by manual or other device-specific means (i.e. GUI or other configuration interface).

⁹ mDNS = multicast DNS (also known as Bonjour) for service discovery, RTSP = real-time streaming protocol for SDP transport

¹⁰ Session Announcement Protocol, an experimental multicast protocol for periodic session announcement and SDP transport



2 Guidelines for Configuration of an AES67 System

2.1 System planning

Before wiring up a system, careful planning is advised. Configuration, system monitoring and debugging will be much more efficient if the general system layout and other vital aspects have been given ample thoughts.

2.1.1 Network infrastructure

2.1.1.1 Managed switches

In most cases, an AES67 system requires an administrable network (due to QoS and multicast requirements), which mandates the deployment of managed switches. Managed switches provide means for accessing the switch configuration, which, in most cases, is achieved by an internal web browser providing user-friendly access through any web browser. Other switches (mostly enterprise-grade switches) may offer a command line interface (“CLI”) for more complex configuration tasks. While most switches have a useful out-of-the box default configuration, it is always advisable to check and verify the required settings.

2.1.1.2 Topology

While AES67 is strictly based on IP and can thus run on any “standard” network topology, it is always a good rule to minimize the number of switches (“hops”) any particular stream will navigate in the final network. A small network may consist of only one switch, which of course makes configuration relatively easy. As the network becomes larger, star or tree topologies¹¹ come in to play. In larger corporate networks spanning multiple subnets, it can be essential to have a deterministic route for any given connection – in this case a leaf-spine architecture¹² would be the most preferred topology.

2.1.1.3 Bandwidth

In any case, it needs to be assured that ample bandwidth on any given path is available. While individual devices may have more than enough bandwidth available on a 100 Mbit/s Fast Ethernet (FE) port, the total required bandwidth to accommodate all streams on a backbone link may easily require Gigabit speed (GbE). It is a good idea to use GbE switches exclusively for your infrastructure. If you need to accommodate several hundred channels of audio, particularly if you plan to share your network with other IT services, you may even consider upgrading your backbone infrastructure to higher speeds (i.e. 10 GbE or above).

Despite the nominal link rate of the switch ports it may also be advisable to check for the maximum switching capacity. Some switches (specifically at the lower cost end) may offer a large number of ports, but won't be able to cope with the total traffic when all ports are heavily loaded. Check for

¹¹ https://en.wikipedia.org/wiki/Network_topology

¹² A Beginner's Guide to Understanding the Leaf-Spine Network Topology
(<http://blog.westmonroepartners.com/a-beginners-guide-to-understanding-the-leaf-spine-network-topology/>)



terms like “backplane speed” or “non-blocking switch fabric” etc. if you expect a high load on your switch.

2.1.1.4 “Green” is evil

While preserving energy is usually a good idea, it impedes proper operation of any low latency real-time audio over IP technology. With the energy-saving function switched on, most switches will not forward single incoming packets immediately, but will wait for a few more packets to be sent down a specific link. This will result in an increased packet delay variation (PDV) which directly affects the PTP operation (end nodes fail to settle into a stable sync condition or exhibit a large time jitter). Simple disable any energy-saving functions on all switches.

2.1.1.5 Cabling

This may sound like odd advice, but ensure that you are always using quality patch cables. The required grade for GbE is Cat5e, but it doesn’t hurt to use Cat6 or Cat7 cabling, especially if you need longer runs close to the maximum allowed Ethernet cable length (~ 125 m). Special care needs to be taken with mobile installations where cables often come on a drum for multiple uses: cable quality will degrade over time as twisted pairs tend to slacken inside the cable. This may lead occasionally to dropped packets despite signaling an otherwise proper link status.

2.1.2 IP addressing

Even if you are planning a small or medium-sized installation running on a single LAN, IP addressing is required. In general, there are three methods to assign IP addresses (and every device, including the switches, requires an IP address):

- DHCP: an automatic IP address assignment which requires the presence of a DHCP server; in most cases this can be one of the switches, if a dedicated DHCP server is not present. While this method is very convenient and you don’t have to fiddle with address administration, subnet and gateway configuration, the disadvantage is that the assigned IP addresses are not immediately known (however, a device GUI will reveal its current IP address in most cases) and that devices may not receive the same IP address again once repowered or reconnected to the network.
- Zeroconf¹³: an automatic IP address assignment which doesn’t require a DHCP server. Devices entering the network assign themselves an available IP address in the pre-defined zeroconf IP address range 169.254.0.0/16. While this is also a convenient method for device network configuration in small LAN setups, it exhibits the same disadvantages as DHCP (you will get different IP addresses each time), plus one can’t even select the IP subnet range.
- Manual / static IP configuration: This method requires devices to be configured individually, and IP addresses are assigned on an administrative basis. While this is quite a lot of work, especially in larger environment, it provides full control on how subnets and devices are configured. Since IP addresses remain unchanged after repowering or reconnecting to the

¹³ https://en.wikipedia.org/wiki/Zero-configuration_networking



network, a device can be safely preconfigured offline. A spreadsheet or a device database is essential to manage the network configuration.

2.1.3 Multicast

In order to avoid multicast packet flooding, your switches need to be configured for proper multicast traffic registration and forwarding by activating IGMP. Three versions of the IGMP protocol exist¹⁴; AES67 requires IGMPv2 to be supported by the network. You can also configure your switches to support IGMPv3; they will, by definition, automatically revert to version 2 once any device is issuing IGMPv2 messages.

Next, the IGMP snooping function¹⁵ needs to be activated, and forwarding of unregistered multicast traffic needs to be disabled.

In order for IGMP snooping to work properly, an IGMP querier needs to be present on the network. This function can usually be invoked on any managed switch. Although a network can accommodate multiple IGMP queriers (and will automatically select one), it is safer to have only one IGMP querier enabled, preferably on a switch sitting close to the root of your network topology.

On larger networks or when employing enterprise-class switches, further multicast traffic management configuration may be required: some switches can be configured to forward any incoming multicast to a so-called multicast router port; this may or may not be desirable, depending on your network situation¹⁶.

2.1.4 QoS

Since clock and audio traffic require high forwarding priorities, AES67 end nodes support DiffServ QoS and assign certain DSCP tags to those IP packets. The switches need to be configured to support DiffServ QoS and prioritized forwarding. Most switches have layer 2 CoS QoS¹⁷ enabled by default; this needs to be changed to layer 3 DiffServ QoS. Once enabled, check the priority assignments – a managed switch usually has at least 4 priority queues per egress port and AES67 operating with the recommended / default parameters requires this configuration:

- DSCP EF (46) (clock traffic) → highest priority queue (4)
- DSCP AF41 (34) (audio packets) → second-highest priority queue (3)
- All other DSCP values (remaining traffic) → lowest priority queue (0)

Note: On some networks running other important / prioritized traffic other priority configuration may be required; however, it is advised, that PTP traffic always receives highest priority treatment. RAVENNA and Dante use other DSCP defaults (CS6 (48) for PTP, EF (46) for audio), but unlike Dante, most RAVENNA implementations allow DSCP reconfiguration at the end nodes to match the AES67 defaults (or any other desired configuration). For guidelines on how to interoperate AES67 with Dante devices in AES67 mode, refer to the respective chapter later in this guide.

¹⁴ https://en.wikipedia.org/wiki/Internet_Group_Management_Protocol

¹⁵ https://en.wikipedia.org/wiki/IGMP_snooping

¹⁶ See also https://en.wikipedia.org/wiki/Protocol_Independent_Multicast

¹⁷ https://en.wikipedia.org/wiki/Class_of_service



Finally, check that the forwarding policy for the egress scheduler is set to “strict priority forwarding” (at least for the PTP traffic class, but also recommended for the audio traffic class).

Note that on larger / corporate networks, specifically if stretching across WAN connections, DSCP tags may not be respected by edge routers (they may be configured to not trust the DSCP markings originating from the local subnets and may even delete them). This will break the tight priority forwarding requirements and may lead to increased packet delay variations, resulting in longer latencies and degraded clock accuracy. Furthermore, after traversing any WAN link employing this “DSCP no-trust” policy, the DiffServ priority mechanism may be irreparably broken for any subsequent local network segments, eventually resulting in AES67 ceasing to work at all after traversing a WAN link. You may have to consult with your network administrator to discuss options to remove or bypass this constraint, if it exists.

2.1.5 PTP

Planning for PTP deployment is a topic on its own which may exhibit many complex facets, especially if your network is larger and stretches several subnets. Larger networks in most cases require PTP-aware switches (Boundary or Transparent Clocks) in strategic positions in the network. Due to the complexity which may be involved in configuration of such networks, we limit the discussion of PTP planning to a single LAN segment without PTP-aware switches.

2.1.5.1 PTP parameters

In most cases, PTP-aware switches are not required in LAN segments up to a medium size (several tenths of end nodes). With standard COTS switches, proper QoS configuration should result in a decent PTP performance. However, there are a few parameters of choice:

- Domain number: unless required for certain reasons, leave the domain number to the default value (0).
- SYNC message interval: all AES67 devices are required to operate with the PTP Default profile which has a default sync message interval of 1 second (2^0). Other choices under the Default profile are 2^1 and 2^{-1} - we recommend setting the SYNC message rate to 2^{-1} for faster settlement and better stability.

AES67 also defines its own PTP profile, the Media profile. If all AES67 devices on the network support this profile (this is not a requirement), you can reduce the SYNC message interval down to 2^{-4} - we recommend that you keep the SYNC message interval at the Media profile default value of 2^{-3} .

- ANNOUNCE message interval: ANNOUNCE messages are required to establish the best master clock currently available on the network. We suggest that you keep the ANNOUNCE message interval at the default value of 2^1 (applies both for the PTP Default and Media profiles) and the ANNOUNCE message timeout interval at 3. Note: it is very important, that ALL devices have the same setting, otherwise the BMCA may not work as expected and devices may not synchronize properly.



- DELAY REQUEST intervals: no need to deviate from the default values (2^0) either (unless you know what you are doing). Keep the delay measurement mode configured to end-to-end (E2E) delay measurement.

2.1.5.2 BMCA parameters

For best synchronization results, you want to make sure that the best available master clock on the network is actually taking this role. If you have a dedicated Grandmaster device, all settings are usually in place by default to let this device become Grandmaster.

However, if you experience that this is not the case or if no dedicated Grandmaster device is present, you may have to dig a bit further into the BMCA parameter configuration in order to resolve any problems or make sure that only those devices qualify for BMCA competition which exhibit a decent PTP Grandmaster functionality by design (usually a device with a very precise and stable internal clock circuitry or which can be connected to an external reference signal, i.e. a word clock or a black-burst input).

The BMCA is an exactly specified algorithm that each devices has to follow to come to the same conclusion on the best available master clock on the network; any failure to fully and correctly implement the BMCA (even in end nodes which never can become Grandmaster at all) may result in improper synchronization results (yes, we have seen this). The BMCA relies on the ANNOUNCE messages being distributed in the network. The ANNOUNCE messages contain certain parameters about the clock quality which are compared in certain precedence:

1. *Priority 1 Field*: This is a user settable value. The lowest number wins. Normally this is set at 128 for master-capable devices and 255 for slave-only devices. However, if you want to overrule the normal selection criteria you can change Priority 1 and create any pecking order you wish.
2. *Clock Class*: This is an enumerated list of clock states. For example, a clock with a GPS receiver locked to Universal Coordinated Time (UTC) has more class than one which is free running and set by hand to your wrist watch. There are also states for various levels of holdover when a clock which had a GPS receiver lost the connection.
3. *Clock Accuracy*: This is an enumerated list of ranges of accuracy to UTC, for example 25-100 ns.
4. *Clock Variance*: This is a complicated log-scaled statistic which represents the jitter and wander of the clocks oscillator over a SYNC message interval¹⁸.
5. *Priority 2 Field*: You guessed it, another user-settable field. The main purpose at this low end of the decision tree is to allow system integrators to identify primary and backup clocks among identical redundant Grandmasters.

¹⁸ Quote from D. Arnold, Meinberg: "In fact it is so complicated that if you can accurately determine it for a clock then you get three credits toward a degree in mathematics." – also see <http://blog.meinbergglobal.com/2013/11/14/makes-master-best/> for more in-depth information on BMCA



6. Source Port ID: This is a number which is required to be unique, and is usually set to the Ethernet MAC address. Essentially this is a coin toss which is guaranteed to break a tie.

For practical purposes, the Priority 1 field is the most important. Start with keeping the value at the device default setting (should be 128 for devices which can become GM and 255 for devices which are slave-only). If you don't have a dedicated GPS-referenced GM device in the network you may either decrease the *Prio1* for certain devices you want to become preferred GMs, or decrease the *Prio1* field for those devices, which should become GM only in case there is absolutely no better GM available on the net¹⁹.

In any case, and regardless of the intended size of your network, always make sure that the PTP distribution results in the desired accuracy in any particular network segment before proceeding with setting up any stream traffic. A good indicator is the clock offset (calculated time offset from PTP master) indication offered by most end nodes. Indicators may vary between devices, most feature at least a status indicator or a numerical offset display. If you see a "green" light or see offset numbers in the single-digit microseconds or sub-microseconds range, you are usually good. Remember to check those indicators from time to time during regular operation.

2.1.6 Discovery

As described in the introduction, session description data is required to connect to an available stream and decode its content. While the parameters required and their proper line-up are defined by the session description protocol (SDP), AES67 does not define a mandatory method to transport the data; hence, manual read-out and entry is assumed as the minimal common ground.

Most AES67 systems or devices provide means of discovering available streams on the network and support protocol-based communication of these SDP parameters. The methods and protocols supported usually relate to the native networking solution those devices adhere to; RAVENNA, Livewire and Dante all offer discovery and connection management functionality, which of course includes the transfer of SDP data. Unfortunately, they all use different methods and protocols, rendering them incompatible with each other:

- RAVENNA uses DNS-SD²⁰ for discovery and RTSP²¹ for SDP transfer
- Livewire uses a proprietary protocol, but also supports the RAVENNA method
- Dante uses different methods – a proprietary method based on mDNS²² for native stream operation and SAP²³ for AES67 formatted streams.

Since Dante devices don't have means for manual read-out or entry of SDP data, there is no practical way to establish connections between Dante devices with activated AES67 mode and any

¹⁹ Some AES67 devices have default Prio1 values of 248 which usually will result in those devices not becoming GM, unless there is absolutely no better GM available on the network. In those cases, check the achievable accuracy of all participating nodes as – depending on the clock quality of the then selected GM – they may not be able to successfully settle into stable synchronization.

²⁰ DNS-based service discovery - https://en.wikipedia.org/wiki/Zero-configuration_networking#DNS-SD

²¹ Real-time Streaming Protocol - https://en.wikipedia.org/wiki/Real_Time_Streaming_Protocol

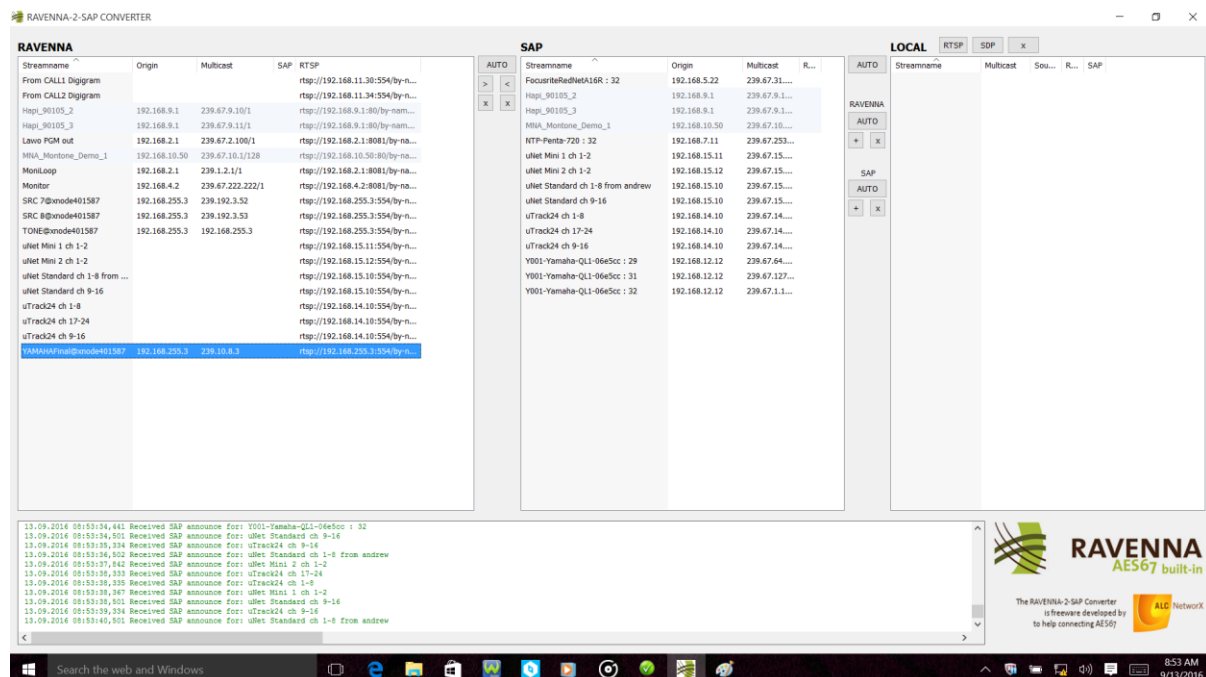
²² Multicast DNS - https://en.wikipedia.org/wiki/Multicast_DNS

²³ Session Announcement Protocol - https://en.wikipedia.org/wiki/Session_Announcement_Protocol

other AES67 device. For this reason, some device manufacturers have decided to include SAP support.

2.1.6.1 RAV2SAP

ALC NetworX has released the RAVENNA-2-SAP converter (RAV2SAP), a freeware tool²⁴ to convert between RAVENNA and Dante discovery method. It translates selected stream announcements from one side to the other and makes the SDP data available accordingly. It also features manual SDP data entry and read-out and can thus help to diagnose any connection problems or integrate any devices which do not support RAVENNA or SAP.



RAV2SAP SCREEN SHOT

RAV2SAP is a Windows application which needs to run on a PC which is connected to the audio network. RAV2SAP only monitors and transmits discovery and SDP-related data traffic, no audio is passed through the PC (unless your PC also hosts an AES67-capable virtual sound card).

2.2 Device configuration

2.2.1 IP configuration

Select the method of IP assignment: DHCP, Zeroconf, manual / static. In case of static IP assignment, make sure you don't assign any IP address twice and that the subnet mask matches your intended network configuration. In some cases, a gateway needs to be configured (even if it won't be used). If no gateway is present, just enter the IP address of one of your switches. An Excel spreadsheet helps tracking IP configuration. If you prefer automatic IP configuration, check if IP parameters have been properly assigned through DSCP or Zeroconf.

²⁴ RAV2SAP is available on the RAVENNA web site - <https://www.ravenna-network.com/using-ravenna/support/downloads/>



2.2.2 PTP configuration

Check / configure all relevant PTP parameters the device offers; follow the guidelines given in section [2.1.5 PTP](#).

2.2.3 Device-specific configuration

Some devices need further configuration to interoperate properly with other AES67 gear. Here are a few commonly observed settings which may have to be configured individually:

2.2.3.1 AES67 mode

Some devices require you to activate the AES67 mode (i.e. Dante devices), other devices support AES67 natively (RAVENNA, Livewire).

2.2.3.2 Multicast address range

Despite not being fully AES67-compliant, some devices only support a limited range of multicast addresses for AES67 interoperation (i.e. Dante devices). The range needs to be configured properly with all devices; note that this may even affect devices which don't exhibit this limitation, as AES67 streams would only be identified / accepted when their multicast address is within that configured range. As some devices don't have a general device-level configuration for multicast address range (they can work with any valid multicast address in the range $239.x.y.z$), this may have to be respected when configuring individual streams.

2.2.3.3 Discovery

While most devices also use their native discovery method for announcement of AES67 streams, some devices offer to enable other discovery options on demand (i.e. enable SAP support).

2.2.3.4 Audio-related configuration

Some devices support different sampling rates, but only one may be selected at any given time (usually because a device only has one clock circuitry). AES67 calls for support of 48 kHz, but other sample rates may be used; make sure you select the desired sample rate. Further device-specific parametrization may be required, check with the operator's manual.

2.3 Check for proper synchronization (PTP)

Once all devices have been configured, check for proper synchronization. This is important because all devices on the network derive their locally generated media clocks from the network clock distributed with PTP.

2.3.1 Grandmaster selection

It is advised that you select a device as the preferred master beforehand and set every other device to slave-only mode. Follow the steps under [2.1.5.2 BMCA parameters](#). Once properly configured, all devices should indicate that they are listening to the same Grandmaster (IP address and / or GM-ID should be identical). If you see different GM-IDs, the BMCA did not work as intended and at least one device is assuming a false GM role. Here's a quick checklist:



- Check if (PTP) multicast traffic is forwarded to all nodes (nodes need to receive the ANNOUNCE messages from all other devices for proper BMCA execution). Although the PTP multicast address (224.0.1.129) is a well-known multicast address which should be forwarded by a switch by default, an IGMP request may need to be issued to activate forwarding in certain switches.
- Check *priority 1* values of devices which assume a GM role unexpectedly and compare with the settings of the designated GM. You may have to lower the priority (increase the *priority 1* value) of that particular device or assign "slave-only" operation. Alternatively, increase the priority (lower the *priority 1* value) of the designated GM.
- It may also help (even just for analysis) to select a different device to become GM by adjusting the *priority 1* fields accordingly, or by temporarily removing suspected devices from the network.
- If you have PTP-aware switches in the network, it may help to switch PTP support off to diagnose the situation. If the situation corrects after switching off PTP support, you need to carefully check all PTP-related settings in the PTP-aware switches.

2.3.2 PTP accuracy

Check PTP accuracy on all nodes – slave devices generally inform about proper sync status. They either have a sync indicator (traffic light or any other graphical means) or they indicate the current offset from master numerically; in most cases single-digit microseconds are usually sufficient, sub-microseconds are perfect.

If you don't have proper sync on all end nodes, you have to resolve this situation before proceeding any further (i.e. configuring streams). You may check on these potential issues:

- SYNC message rate too low: some devices require a certain sync message rate in order to reach a stable locking situation. Try to decrease the SYNC message interval at the chosen Grandmaster (i.e. try a SYNC message interval of 2^{-2} or 2^{-3}).
- QoS not properly configured: PTP traffic needs to receive the highest forwarding prioritization. Check if PTP packets are marked with a proper DSCP value²⁵ and if all participating switches in the network are configured to store packets with this DSCP value in their highest priority queue.
- Removing traffic load: If you are unsure about properly configured QoS you may also try to remove any foreign traffic on the network to reduce the bandwidth utilization (i.e. to remove potential network overload). The simplest approach would be to unlink devices or network segments which are not relevant to AES67. You may also start building your network from scratch by incrementally plugging in devices and check each time for proper synchronization.

²⁵ Some end devices offer individual configuration for DSCP markings; otherwise you may have to use tools like Wireshark (<https://www.wireshark.org/>) to examine packets on the network.



- Mixed switch configuration (FE / GbE)²⁶: In some cases a mixed use of switches with different network link speeds may cause synchronization issues. In most cases, this will result in a permanent offset from master only without necessarily affecting the synchronization stability. The node may settle into a synchronized condition, but most likely larger latency settings will be required for streams coming from / going to this node due to a permanent displacement between local and network time. It is a good advice to only use GbE switches in the network and connect end nodes with FE interfaces directly to the GbE switch ports.
- PTP-aware switches: as mentioned earlier, PTP-aware switches are certainly valuable (or even required) to improve synchronization (especially in larger networks), but they may make things more complicated and require deeper knowledge for proper configuration. Check if PTP-aware switches are part of your network and try switching PTP support (temporarily) off. Make sure that all configuration requirements for COTS switches are in place (i.e. QoS, IGMP etc.).

2.4 Stream configuration

Once your network is prepared as described above, you are ready to configure streams. While AES67 calls for support of multicast and unicast transport, we will focus on multicast streaming only as this is the method commonly available on all AES67 devices. Configuring and connecting to multicast streams generally always follows these two basic steps:

1. Configure and start a multicast stream on the sender node
2. Make the related SDP data available to the desired receiving node
3. Connect to the selected stream

Execution of these steps usually varies between individual devices; in this guide we use screenshots from the RAVENNA Virtual Sound Card (RVSC) which is based on the RAVENNA framework developed by ALC NetworX. Consult the respective Operating Manuals of other devices to execute these steps accordingly.

2.4.1 AES67 stream format

Since the main focus of AES67 is on interoperability, the stream format variations to be supported by all devices are pretty narrow:

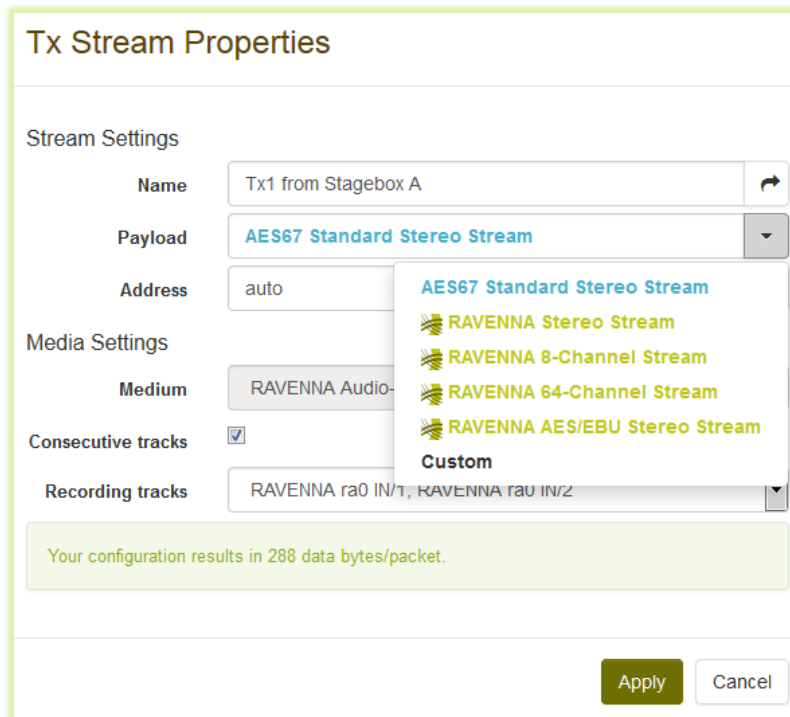
- Sample rate: 48 kHz
- Data encoding: linear PCM with 16- and 24-bit (L16 / L24)
- Number of channels per stream: 1..8
- Packet time (number of samples per packet): 1 ms (48 samples per channel per packet)

²⁶ FE – Fast Ethernet (100 Mbit/s), GbE – Gigabit Ethernet (1000 Mbit/s)

Other variations are recommended, but not required to be supported. Therefore, we focus on a typical AES67 stream setup: a stereo stream with L24 encoding running at 48 kHz with 1ms packet time.

2.4.2 Creating an AES67 stream

Invoke the stream creation function ("create Tx stream", "create session source" or alike) and fill in the parameters as required:

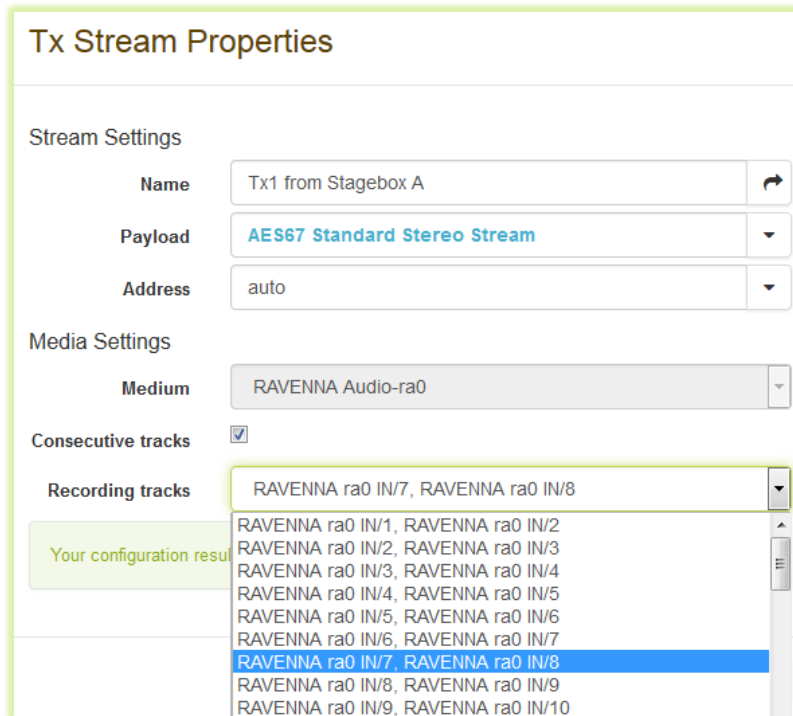


RVSC SCREEN SHOT: TX STREAM PROPERTIES (PAYLOAD FORMAT SELECTION)

- *Name*: assign a meaningful name for this stream (not required by AES67, but helps to identify this stream when discovery is used)
- *Payload*: select from pre-defined stream formats (here: AES67 standard stereo)
- *Address*: enter desired multicast address²⁷ or leave at "auto" for automatic assignment

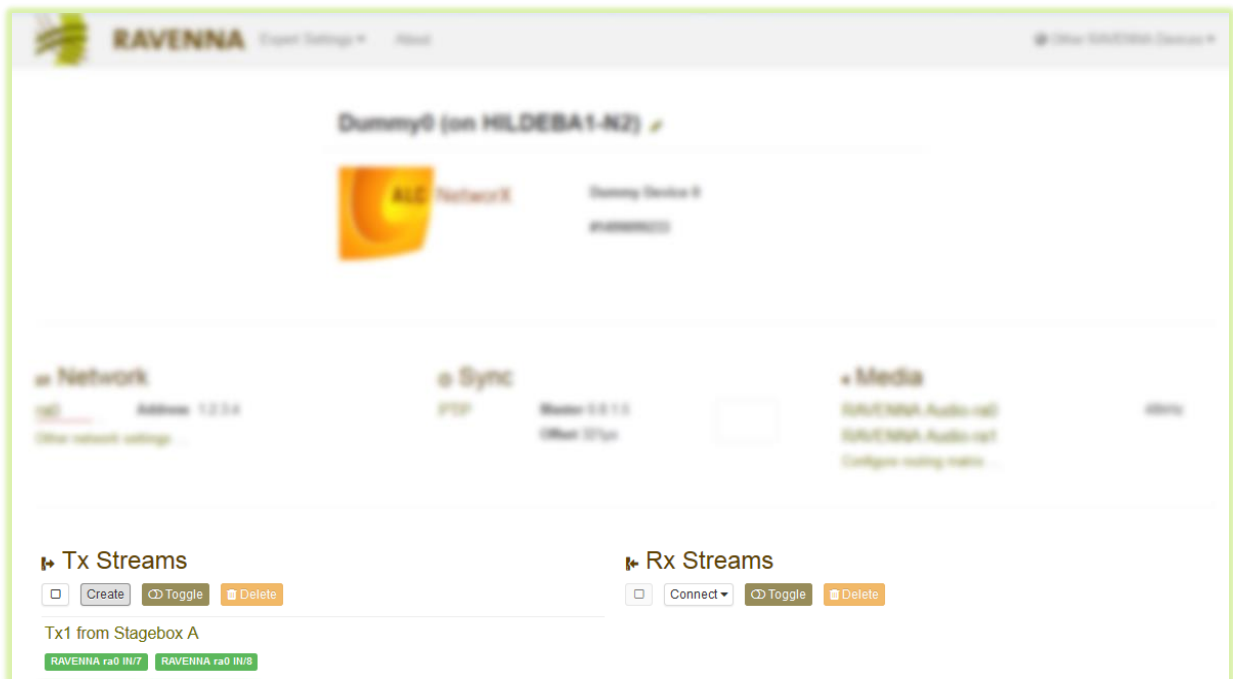
The most essential choice is of course on the number of channels in the stream and which channels to incorporate from the individual device. Select the desired audio channel pair from the drop-down menu:

²⁷ See 2.2.3.2 *Multicast address range* for further hints



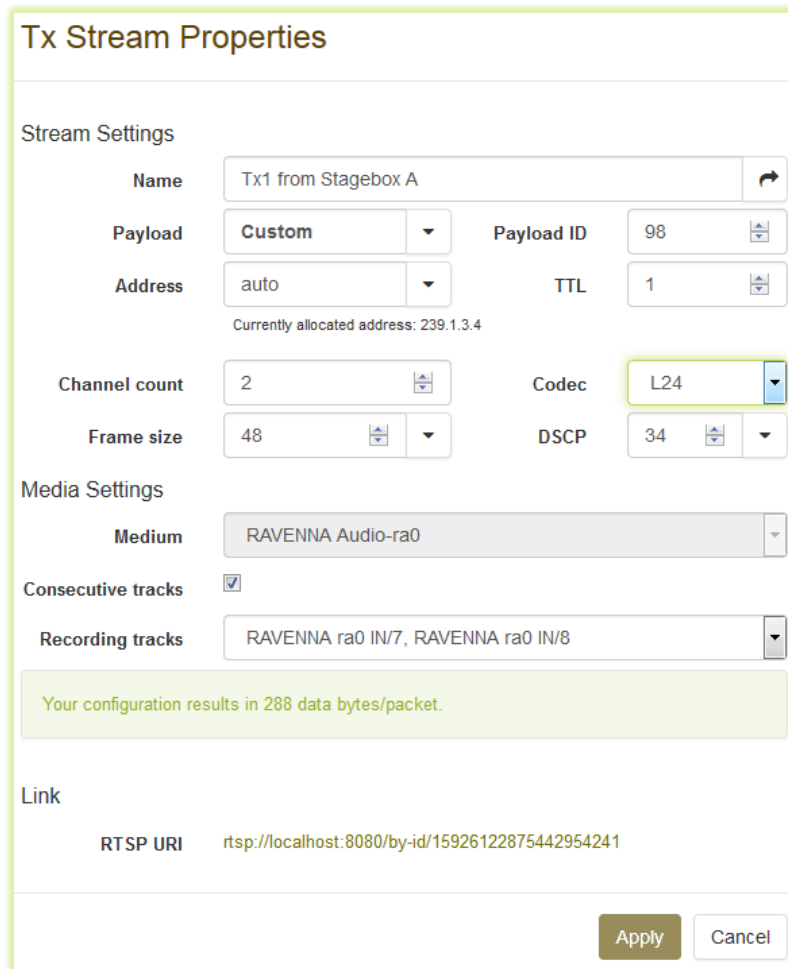
RVSC SCREEN SHOT: TX STREAM PROPERTIES (AUDIO CHANNEL ASSIGNMENT)

Once selected, hit "Apply". This will create an AES67 stereo stream in L24 / 48 kHz data format from audio channels 7 + 8 with an automatically assigned multicast address. The stream is immediately started and available on the network. The stream packets will reach the first switch where they are dropped unless another device has registered to this stream by IGMP.



RVSC SCREEN SHOT: OVERVIEW WITH 1 TX STREAM CREATED

If you need to access other stream configuration options, you can select “*Custom*” from the payload format drop-down menu which opens all available parameter fields:



RVSC SCREEN SHOT: TX STREAM EXTENDED PROPERTIES

You can now change the number of channels in the stream, use a different encoding format (i.e. L16 or AM824²⁸), change the packet time (frame size) or assign another DSCP value to this particular stream.

2.4.3 Accessing the SDP data

In order to connect to this stream, the desired receiver device needs access to the respective SDP data. While AES67 specifies the required SDP data, it does not mandate for a specific method to convey this data. Most AoIP solutions offer means for advertising / discovering available streams and transporting the SDP data automatically. If no common method is available between sender and receiver device, manual SDP data transfer is assumed. Alternatively, the RAV2SAP converter (see **2.1.6.1 RAV2SAP**) may be used to translate between different discovery methods and / or aid manual SDP data transfer.

²⁸ AM824 is a special payload format available in RAVENNA devices to transport AES3 signals in a fully bit-transparent manner.



2.4.3.1 SDP data transfer by a common discovery method


The streams available on the network will be directly visible in the desired receiver and the SDP data will be transferred when setting up the connection; no further action has to be taken at this stage.

2.4.3.2 Manual transfer of SDP data

If manual transfer is required, the SDP data can be copied (and later pasted) by opening the related SDP data record. The RVSC provides a link to the SDP data set when creating the Tx stream:

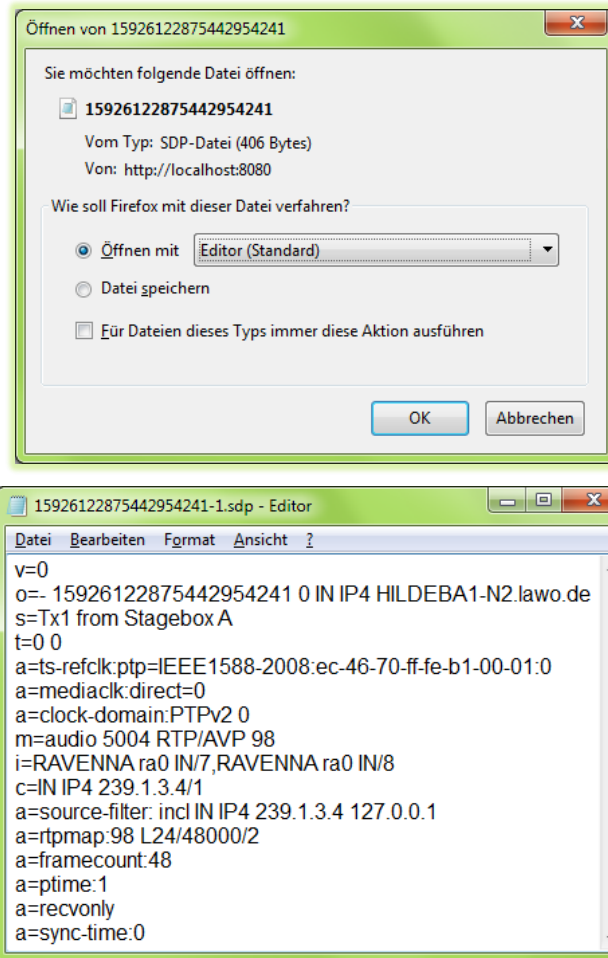
RVSC SCREEN SHOT: TX STREAM SDP LINK

The link allows direct access to the SDP data by any device supporting RTSP. Since RTSP is very similar to HTTP and the SDP data is formatted in ASCII text, the link can also be used to access the SDP data with any browser; simply copy and paste the link into your browser address field and replace "rtsp://..." with "http://...".



ACCESSING SDP DATA WITH BROWSER (BY HTTP)

This will open the Windows text editor (or any application linked to text files):



OPENING SDP DATA WITH EDITOR

The SDP data can now be copied (or saved) and used for setting up a connection to this stream at the desired receiver.

2.4.4 Receiving an AES67 stream

Connecting to an existing AES67 stream requires inputting the SDP data of the respective stream to the desired receiver. This can either be done manually or with support of a discovery & connection management method.

2.4.4.1 Connecting to an AES67 stream with discovery support

If a common discovery method is supported by both the sender and receiver, connecting to a stream should be as easy as identifying the desired stream in the receiver's user interface (usually by name) and executing the connection function.

In the RVSC, hit open the Rx creation dialog ("Connect Rx Stream"):



Rx Stream Properties

Stream Source

Name Codec Sample rate Hz Channels

Tx1 from Stagebox A

Channels

Receiver Settings

Label (not labeled)

Delay (samples) 512

Syntonized Mode

Request unicast

Channel count 0

Media Settings

Medium RAVENNA Audio-ra0

Consecutive tracks

Play tracks

Apply Cancel

RVSC SCREEN SHOT: RX STREAM PROPERTIES (STREAM SOURCE SELECTION)

In the "Stream Source" drop-down box, select the desired stream from the list. The related SDP file will automatically be accessed and all relevant parameters are filled-in:



Rx Stream Properties

Stream Source

ravenna_session:Tx1 from Stagebox A ▼

Name Tx1 from Stagebox A **Codec** L24 **Sample rate** 48000Hz Show raw SDP

Channels RAVENNA ra0 IN/7 RAVENNA ra0 IN/8

This is an AES67 compliant stream.

Receiver Settings

Label (not labeled)

Delay (samples) 512 ▲▼

Syntonized Mode

Request unicast

Channel count 2 ▲▼

Media Settings

Medium RAVENNA Audio-ra0 ▼

Consecutive tracks

Play tracks RAVENNA ra0/1, RAVENNA ra0/2 ▼

RVSC SCREEN SHOT: RX STREAM PROPERTIES (STREAM PARAMETERS FILLED-IN)

Next, you can select the desired latency by adjusting the *delay* value accordingly. The RVSC provides the ability to enter the latency individually per stream to accommodate for any relevant packet delay variation, individual stream setup and correlation between any streams on the network, if desired. The number indicates the desired playout delay of an individual audio sample with respect to its sampling time at the sender. The configured number must be large enough to cover the original packet time plus any jitter the packet may experience while being transported across the network. Since the packet time in AES67 is 1 ms, the delay needs to be larger than 48 samples plus sufficient delay to cope with the packet jitter²⁹.

Other devices or AoIP systems may offer predefined (sometimes even system-wide) latency classes like low / medium / high or the equivalent.

The final step is to assign the channels being transported in the stream to the desired output channels of the device:

²⁹ A suggested minimum value to start with is 144 samples (3 ms). A higher value may be required in larger networks or with higher traffic load, smaller values may work with lower traffic or just one or two switches in place. In case the number is too small, the receiver will indicate an error (packets coming in too late).



Rx Stream Properties

Stream Source
ravenna_session:Tx1 from Stagebox A

Name Tx1 from Stagebox A Codec L24 Sample rate 48000Hz Show raw SDP


Channels **RAVENNA ra0 IN/7** **RAVENNA ra0 IN/8**

This is an AES67 compliant stream.

Receiver Settings	
Label	RAVENNA ra0/1, RAVENNA ra0/2 RAVENNA ra0/2, RAVENNA ra0/3 RAVENNA ra0/3, RAVENNA ra0/4 RAVENNA ra0/4, RAVENNA ra0/5
Delay (samples)	RAVENNA ra0/5, RAVENNA ra0/6 RAVENNA ra0/6, RAVENNA ra0/7 RAVENNA ra0/7, RAVENNA ra0/8
Syntonized Mode	RAVENNA ra0/8, RAVENNA ra0/9 RAVENNA ra0/9, RAVENNA ra0/10
Request unicast	RAVENNA ra0/10, RAVENNA ra0/11 RAVENNA ra0/11, RAVENNA ra0/12
Channel count	RAVENNA ra0/12, RAVENNA ra0/13 RAVENNA ra0/13, RAVENNA ra0/14 RAVENNA ra0/14, RAVENNA ra0/15
Media Settings	RAVENNA ra0/15, RAVENNA ra0/16 RAVENNA ra0/16, RAVENNA ra0/17 RAVENNA ra0/17, RAVENNA ra0/18
Consecutive tracks	RAVENNA ra0/18, RAVENNA ra0/19 RAVENNA ra0/19, RAVENNA ra0/20 RAVENNA ra0/20, RAVENNA ra0/21
Play tracks	RAVENNA ra0/3, RAVENNA ra0/4

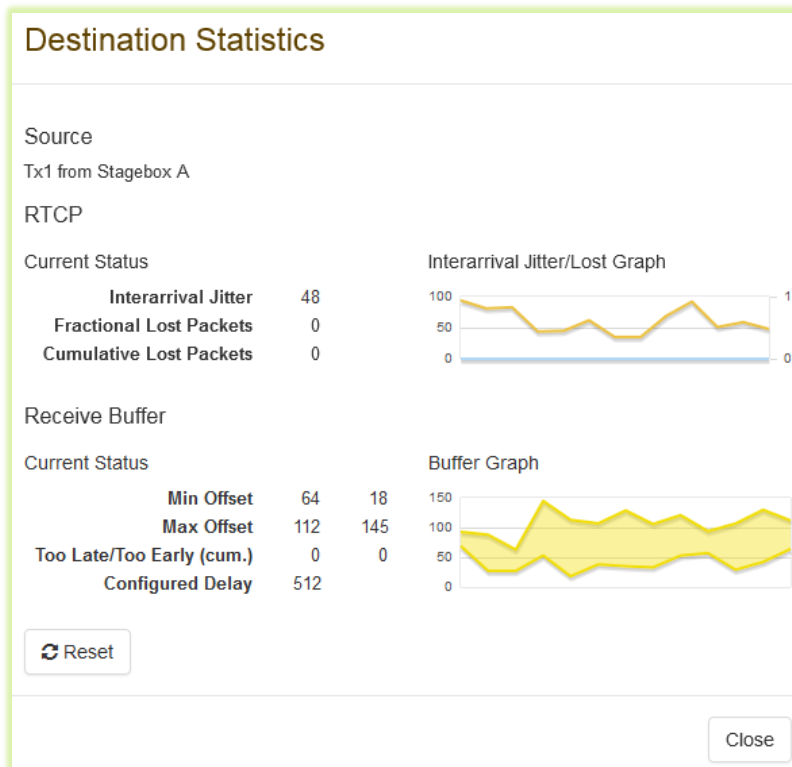
RVSC SCREEN SHOT: RX STREAM PROPERTIES (CHANNEL ASSIGNMENT)

Once assigned from the drop-down list, hit "Apply" and the receiver will connect to the selected stream.



RVSC SCREEN SHOT: OVERVIEW WITH 1 RX STREAM CONNECTED

The RVSC offers statistic displays where the current packet jitter can be visualized (other devices may offer numerical values to indicate current PDV):

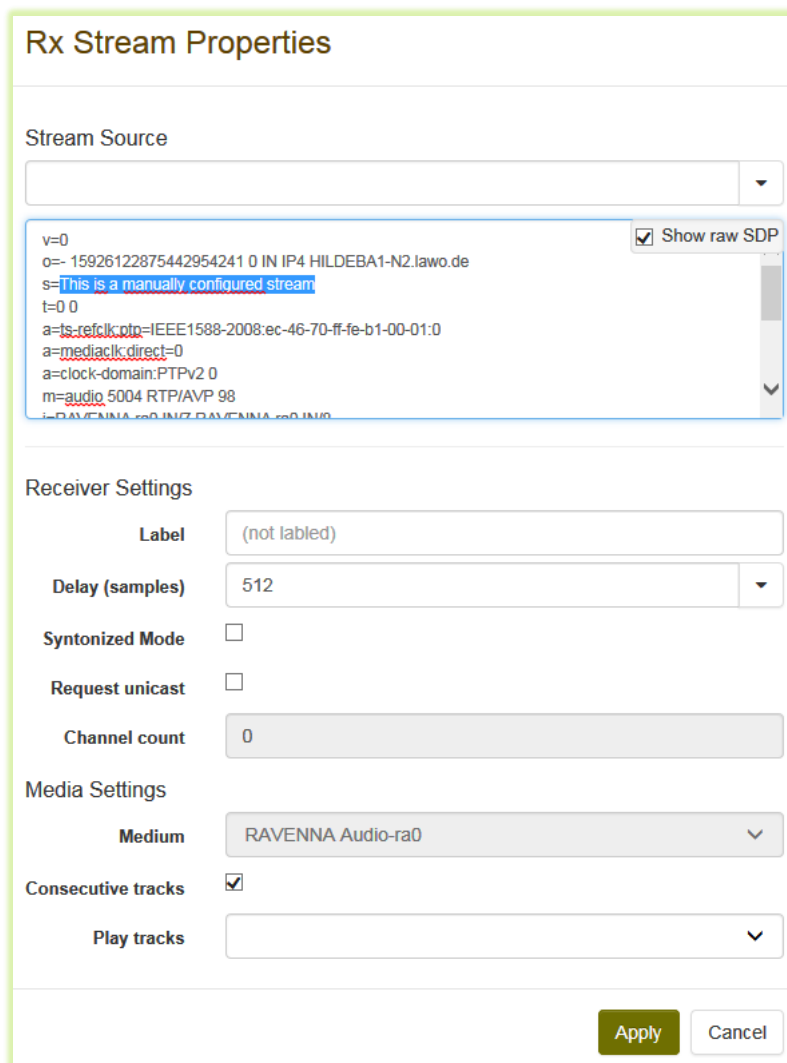


RVSC SCREEN SHOT: STATISTICS WITH PACKET JITTER AND RECEIVER BUFFER UTILIZATION

2.4.4.2 Connecting to an AES67 stream manually

If a common discovery method is not available, the SDP data has to be entered manually into the receiver. The means on how to enter the data varies among devices, a device may either accept the SDP data record as a whole (effectively using copy & paste) or it offers a form with individual parameter fields (which will then look similar to the Tx creation screen where you have to manually type in the respective values).

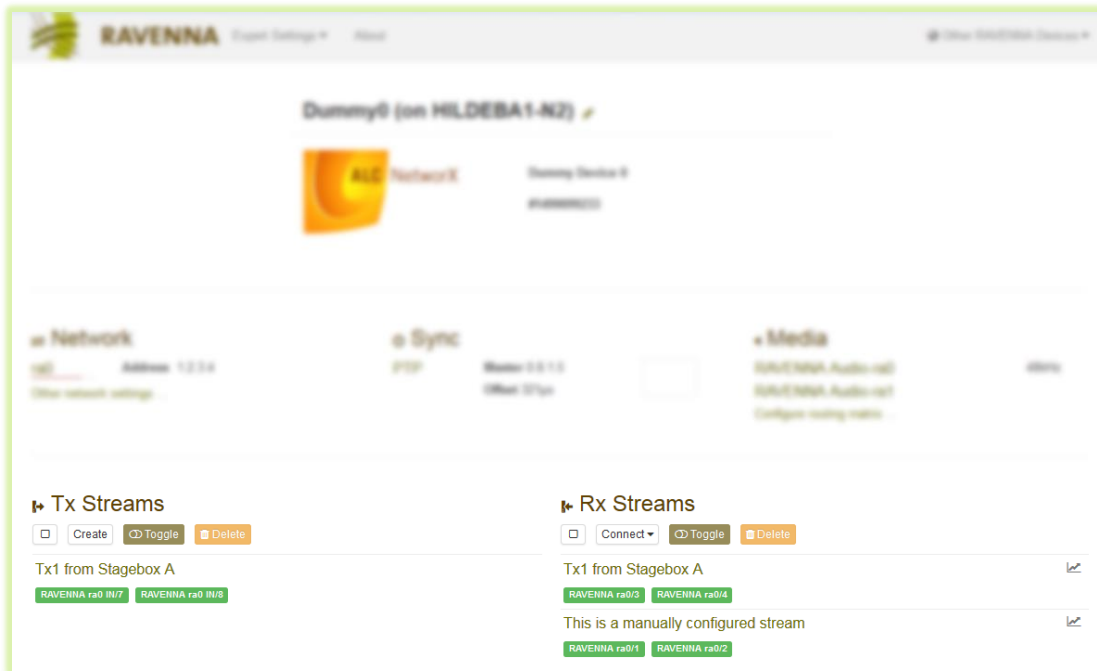
The RVSC offers the option to type in or paste a complete SDP data set. Open the Rx creation dialog ("Connect Rx Stream"), select "Show raw SDP" and double-click into the large empty field which just opened up. This field is now in edit mode and ready to accept the SDP data input³⁰. Simply paste the copied SDP data provided by the sender into this field and modify if necessary (i.e. you may assign a different name by changing the "s=..." line:



RVSC SCREEN SHOT: RX STREAM PROPERTIES WITH PASTED SDP DATA

³⁰ Opening this field in editing mode does not work with the Firefox browser (Mozilla); use Internet Explorer or Chrome instead.

Hit the "Show raw SDP" field again, apply desired latency setting and assign channels as in the previous example and hit "Apply". The stream is now being connected to and it shows up with the edited name in the Rx section of the overview screen:

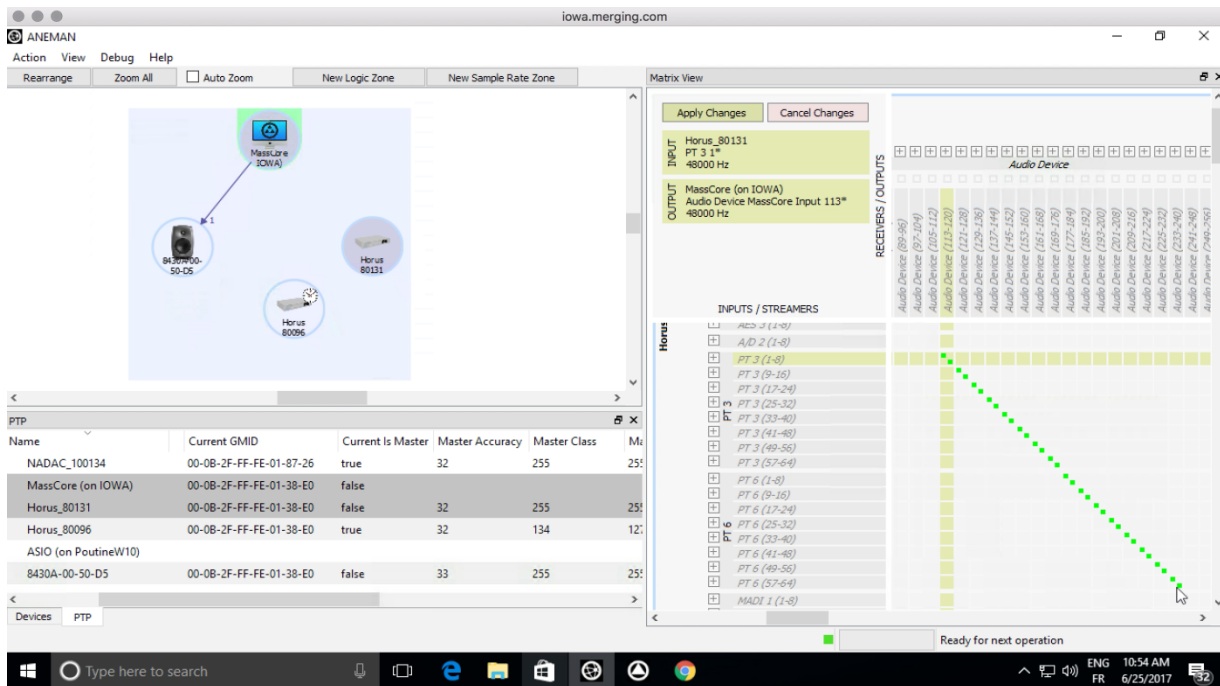


RVSC SCREEN SHOT: OVERVIEW WITH 2 RX STREAMS CONNECTED

3 Outlook on emerging Technologies and Industry Standards

3.1 ANEMAN

ANEMAN (Audio Network Manager) is free software from Merging Technology which allows connection between AES67 devices much like Dante controller does for Dante. However, because of the wide range of AES67 devices, compatibility with ANEMAN is achieved on a per manufacturer basis. Several manufacturers have already boarded the project started by Digigram and Merging and many more are coming.



The screenshot shows the ANEMAN software interface. At the top, there's a menu bar with 'Action', 'View', 'Debug', and 'Help'. Below it are buttons for 'Rearrange', 'Zoom All', 'Auto Zoom', 'New Logic Zone', and 'New Sample Rate Zone'. The main area is divided into three sections:

- Network Diagram:** A central hub labeled 'MassCore (on IOWA)' is connected to three peripheral devices: '8430A-00-50-D5', 'Horus 80131', and 'Horus 80096'.
- Device Table:** A table listing devices with columns for Name, Current GMID, Current Is Master, Master Accuracy, Master Class, and M.

Name	Current GMID	Current Is Master	Master Accuracy	Master Class	M.	
NADAC_100134	00-0B-2F-FF-FE-01-87-26	true	32	255	25	
MassCore (on IOWA)	00-0B-2F-FF-FE-01-38-E0	false	32	255	25	
Horus_80131	00-0B-2F-FF-FE-01-38-E0	false	32	255	25	
Horus_80096	00-0B-2F-FF-FE-01-38-E0	true	32	134	12	
ASIO (on PoutineW10)	8430A-00-50-D5	00-0B-2F-FF-FE-01-38-E0	false	33	255	25
- Matrix View:** A grid showing connections between inputs and outputs. The 'INPUTS / STREAMERS' list includes 'Horus 80131 PT 3 1*', 'MassCore (on IOWA) Audio Device MassCore Input 113*', and 'ASIO (on PoutineW10)'. The 'RECEIVERS / OUTPUTS' list includes various 'Audio Device' entries. A green diagonal line indicates connections between the input and output streams.

ANEMAN SCREEN SHOT: OVERVIEW WITH CONNECTION MATRIX

ANEMAN will be released end of June 2017 (www.merging.com/aneman).

3.2 NMOS

NMOS (Networked Media Open Specifications)³¹ – hosted by AMWA (Advanced Media Workflow Association)³² – are a growing family of specifications which are available to both suppliers and end users, at no cost, to support the development of products and services which work within an open industry framework. The goal for this initiative and the Open Specifications is to deliver interoperability and increase the choice of products across a wide range of suppliers, allowing flexible, cost-effective system designs.

Following the JT-NM Reference Architecture, NMOS addresses functionalities beyond the transport and synchronization of media over networks to further foster interoperability between devices from different manufacturers. Current work includes:

³¹ <http://www.nmos.tv>

³² <http://www.amwa.tv>



- IS-04: discovery & registration - mechanisms, protocols and APIs to enable discovery of available resources, devices and streams and their current status on a network of any size
- IS-05: connection management – mechanisms, protocols and APIs to enable connection management among devices and management systems
- IS-06: network control - mechanisms, protocols and APIs to gather topology information, available device and network resources to enable management of flows on the network

Further work is conducted on timing and identity models and applicability of dematerialized (virtualized, cloud-based) environments. Wherever possible, the specifications are being developed using Internet standards or Internet-friendly techniques. They are complementary to and co-exist with industry specifications and standards; for example, VSF TR-03, SMPTE ST-2110 and AES67.

3.3 SMPTE ST 2110

SMPTE ST 2110 is an emerging suite of standards defining synchronized elementary essence transport (video, audio, metadata) on IP networks. The basic principles are very similar or even identical to AES67; in fact, AES67 has been chosen as the standard for audio essence transport.³³ Since this is work-in-progress, no specific source can be cited here; stand-by for public announcements or check the SMPTE web site³⁴ for further information.

³³ Transport of linear PCM audio data will be defined in ST 2110-30. While ST 2110-30 defines a few minor constraints with respect to AES67, it is safe to say as of today that any AES67 device will be ST 2110-30-compliant.

³⁴ <https://www.smpte.org/>



4 Conclusion

As we said at the beginning, while the AES67 standard defines the protocols and functions to be supported, it still leaves various choices open to implementers, which necessitates some background knowledge on networking in general and on AoIP-related topics in particular in order to get past first base.

This guide provided tips on configuration and hints to circumvent the most commonly observed obstacles when setting up an AES67 network. We have seen how to prepare for network setup, with particular emphasis on IP addressing, network topology and switch configuration, followed by device configuration to be checked before connecting to the network. Some tools have also been presented that will help establishing interoperability or even manage devices. More tools and industry standards are emerging, enabling an easier and faster system design for efficient operation of audio networks.

As a final reminder, always follow these four steps: prepare the network, prepare the devices, check for PTP synchronization, and then connect audio.



RAVENNA
AES67 built-in