Your Practical Guide to AES67—Part 2

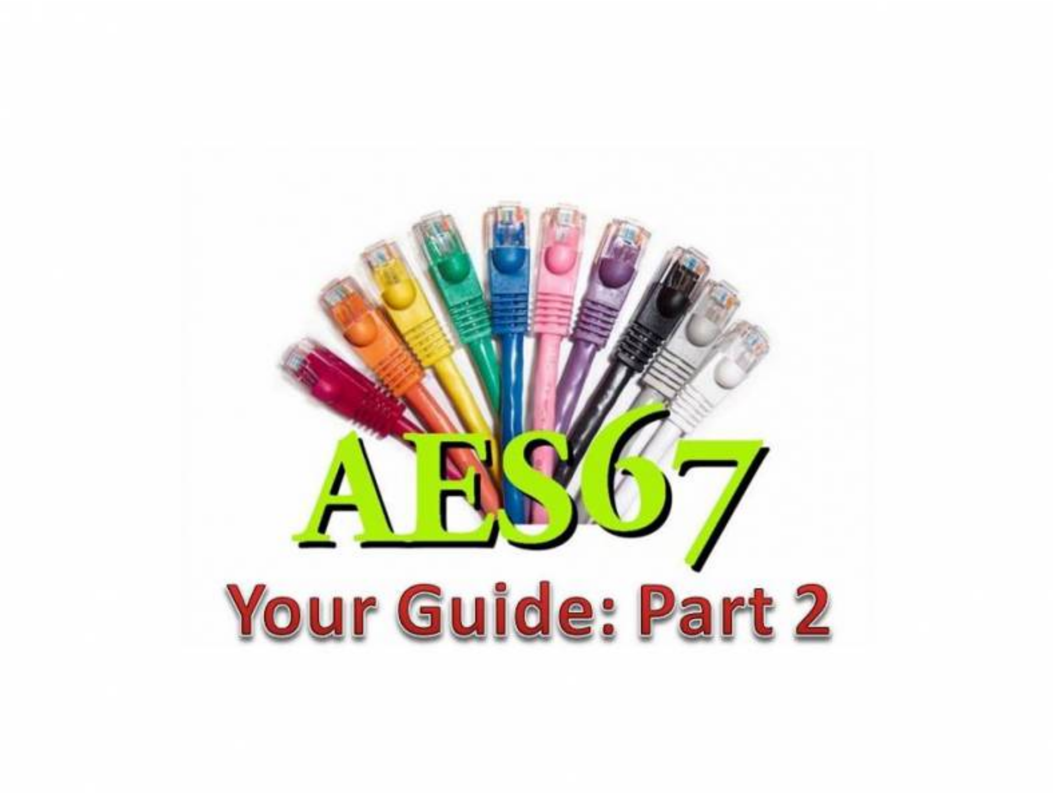

AES67 can be considered the 'glue' you may need to interface your audio gear and build an AoIP network.

The AES67 standard is sometimes misunderstood as the specifications on how all professional digital audio gear is supposed to work and interconnect. Not exactly. In fact, AES67 simply defines the requirements for high-performance AoIP (Audio-over-IP) interoperability. A manufacturer can implement AES67 anyway it wants, and there's the rub.

In Part 1 of this three-part series, we examined some background on the AES67 standards development. The tutorial examined the AES67 standard and how it enables interoperability between different audio solutions. The article reviewed basic network design and operation. A goal of this series of articles is to enable engineers to build and operate AoIP networks with confidence.

In Part 2, we will continue learning more of the advantages and capability of AES67 with a focus on key steps to configure an AES67 system. Let's get started.

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 thought.

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 have to 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 topologies 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 architecture 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 (especially 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 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). Therefore, for optimum network performance, all energy-saving functions on switches should be disabled.

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 as 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.

- Zeroconf : 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 you 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 environments, it provides full control on how subnets and devices are configured. Since IP addresses remain unchanged after repowering or reconnecting to the 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 exist ; 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 function 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 situation .

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 QoS 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.

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, especially 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 interval 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 (Best Master Clock Algorithm) 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 allow this device to become Grandmaster.

However, if 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 device 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 can never become Grandmaster) 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:

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.

RAVENNA

The IP-based Real-Time Media Network

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 a higher rating 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 has a GPS receiver loses the connection.

Clock Accuracy: This is an enumerated list of ranges of accuracy to UTC, for example 25-100ns.

Clock Variance:This is a complicated log-scaled statistic which represents the jitter and wander of the clocks oscillator over a SYNC message interval .

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.

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 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 field value for certain devices you want to become preferred GMs, or increase the Prio1 field value for the devices which should only become GM if there is absolutely no better GM available on the net .

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-SD/mDNS for discovery and RTSP for SDP transfer

- Livewire uses a proprietary protocol, but also supports the RAVENNA method

- Dante uses different methods – a proprietary method based on mDNS for native stream operation and SAP for AES67 formatted streams.

Because 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 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 tool 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.

ALC NetworX has released the RAVENNA-2-SAP converter, which is a freeware tool to convert between RAVENNA and Dante discovery methods.

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).

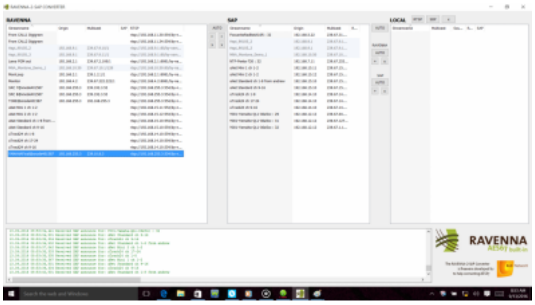

RAV2SAP is a freeware tool to convert between RAVENNA and Dante discovery methods. Click to enlarge.

Up next Part 3

Part 3 will conclude this multi part tutorial on the use of AES67 for AoIP networks. This concluding article will focus on some key steps to build a working AoIP network. Those steps include; device configuration, checking for proper synchronization (PTP) and stream configuration.

Lastly I will provide some thoughts on emerging technology and industry standards.

Editor notes:

Part 3, the conclusion in this series, can be found at this link.

Part 1 of this three-part series can be found at this link.