BROADCAST AUDIO
THE BRIDGE
Connecting IT to Broadcast

Your Practical Guide To AES67- Part 1

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.

While practically every audio-product vendor claims full compatibility with familiar audio standards, the truth is sometimes less so. This tutorial will examine the AES67 standard and how it enables interoperability between different audio solutions. This article will also review basic network design and operation. The goal is to enable engineers to build and operate AoIP networks with confidence.

Getting audio encoder A to properly connect to console B and then route the combination feed to some remote site C should be easy. After all, if each of the devices above meet the AES67 standards, what could go wrong? Turns out, there are a few things that might prevent the above scenario from working properly. Let’s pull the covers back on AES67 and audio networks to learn how to leverage the AES67 technology in audio networks.

Behind the standard

AES67 is a standard for high-performance audio-over-IP *interoperability*. When a manufacturer follows the standard's definitions and requirements, devices that adhere to other other AoIP protocols that do not interoperate with each other (i.e. Dante, Livewire, QLAN or RAVENNA) can become interoperable.

It is important to understand that while the standard defines what protocols and functions need to be supported, it does not tell vendors precisely how to implement them. Thus, solutions differ. Simply understanding the terms used in AES67 is not sufficient for a working knowledge of audio networks. Rather, building audio networks requires some background knowledge on networking in general and about AoIP-related topics in specific. Let's get started.

Basic principles of AES67

AES67 is designed to run on standard packet switching networks over a COTS (Commercial Off-The-Shelf) infrastructure. If configured properly, the same network can be shared with other traffic without degrading the audio streaming performance.

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 at the most basic level. Furthermore, various industry standards covering these functions already exist or are emerging and can be implemented if applicable. Let's examine the four AES67 technical principles.

Synchronization

Network synchronization is based on distribution of a common wall clock time to all participating nodes with sufficient precision. AES67 specifies that the IEEE1588-2008 standard (also known as PTPv2 - *Precision Time Protocol* version 2) be used for time distribution. 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.

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. Part 2 of this series will examine system planning with tips on how the Grandmaster selection process can be modified, if required.

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 good enough, which usually is achievable with standard, non-PTP-aware switches.

Multicast packet transport

PTP is based on multicast packet transport. AES67 also requires multicast support for audio stream packets. While basically any COTS switch supports multicast traffic, only managed switches provide multicast management to effectively avoid network flooding. Unmanaged 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 causing a total network lock-up.

Managed switches provide multicast management through IGMP (*Internet Group Management Protocol*) snooping support. With IGMP snooping, only those multicast packets that have been registered through IGMP are forwarded to their designated ports. Consequently, all AES67 nodes are required to support IGMPv2, which is used to tell the network which streams are to be forwarded.

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. The periodic IGMP queries trigger the nodes to renew their IGMP requests. Once a stream connection is ended, an IGMP leave request is sent to terminate 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 just once on any involved link.

Quality of Service

Quality of Service (QoS) is another fundamental principle the AES67 network must support. Again, this is 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 support of *Differentiated Services* (DiffServ), a QoS scheme where different types of traffic can be categorized into service classes. DiffServ works with 64 different priority tags – DSCP values – that can be applied to individual IP packets.

End nodes can apply different tags to packets belonging 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 resembles 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. In larger networks that may carry a variety of shared traffic classes, QoS configuration requires special attention.

AES67 supports QoS parameters. Even so, proper utilization requires the use of DiffServ and priority tags to make this happen.

Recommended DSCP values

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

- PTP traffic should be tagged with DSCP EF (48), translating into *expedited forwarding*, receiving the highest forwarding priority (first class passengers)
- 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 lower or no specific priority tagging, which by default which is BE (0) for *best effort* (economy class)

Because the network may be used to transport other types of traffic that needs certain prioritization, (i.e. voice data or video), the network administrator may have to change these values or adjust the switch configuration accordingly. Because not all AES67 devices support DSCP reconfiguration, or even don't use the recommended default values, other strategies may need to be applied.

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* (SDP). It contains 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 cannot connect to the stream and decode the packet.

One of the benefits of managed multicast traffic is its scalability.

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 a specific method. Instead, the standard writers decided to just mention some of the widely used protocols and to leave it to the device manufacturer to select a method that works for them.

While most devices support mDNS/RTSP * (the default RAVENNA method) and SAP (Dante devices in AES67 mode), not all devices support both methods, and some don't even offer manual read-out/ SDP data entry. In cases where 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. Part 2 of this series will offer some suggestions to overcome this issue.

*In computer networking, the *multicast Domain Name System* (mDNS), also known as Bonjour, resolves host names to IP addresses within small networks that do not include a local name server. *Real Time Streaming Protocol* is used to establish and control media sessions between end points.

Making Connections with ANEMAN

ANEMAN is a free software from Merging Technologies that 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 already boarded the project started by Digigram and Merging and many more are coming.

ANEMAN is free and available at www.merging.com/aneman.

This concludes Part 1 of the two part series, *Your Practical Guide To AES67.* We have reviewed some basic principles of AES67: synchronization, multicast packet transport, QoS and system information. By now you should have a general understanding of how AoIP networks work, the key components needed, timing and addresses and various adjustments and settings required.

Part 2 of this series, Your Practical Guide To AES67 will appear on The Broadcast Bridge website on August 31 and can be found at the above link. This second installment will focus on system planning, network infrastructure, IP addressing, QoS, PTP, discovery and device configuration. Don't miss this important tutorial.

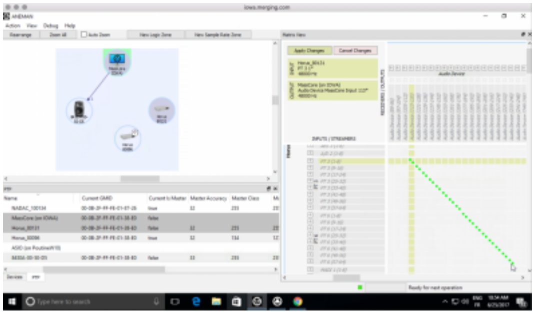

Aneman is a useful utility to help configure AoIP equipment that may be designed for alternative network topologies. Click to enlarge.

Finally, an expanded version of this guide on AES67 networking is available from its author, RAVENNA.