



# AVnu™ Alliance White Paper No-excuses Audio/Video Networking: the Technology Behind AVnu

---

8/24/09

## Author

Michael Johas Teener

## Contributors

Allen Huotari, Yongbum Kim, Rick Kreifeldt, and Kevin Stanton

## Abstract

Recent work by IEEE 802 working groups will allow vendors to build a standards-based “time sensitive” network with the appropriate quality of service for professional A/V use, in a cost effective fashion. This new set of standards, currently being developed by the IEEE 802.1 Audio Video Bridging (AVB) Task Group, provides three major enhancements for 802-based networks:

1. Precise timing to support low-jitter media clocks and accurate synchronization of multiple streams,
2. A simple reservation protocol that allows an application on an endpoint device to notify the various network elements in a path so that they can reserve the resources necessary to support a particular stream, and
3. Queuing and forwarding rules that ensure that such a stream will pass through the network within the delay specified by the reservation.

These enhancements require higher layer applications to initiate those features and may require some changes to the link technology (almost none for Ethernet), and are compatible with all the other functions of a standard 802.1 bridge. As a result, the proprietary professional AV solutions that have historically been built by painstakingly adapting and fine tuning IT grade networking equipment can be abandoned in favor of a standards based approach.

This paper outlines the advantages of new-generation 802.1 AVB protocols and gives an introduction to those new protocols and capabilities.





## About AVnu Alliance

The AVnu Alliance is an industry forum dedicated to the advancement of professional-quality audio video transport by promoting the adoption of the IEEE 802.1 Audio Video Bridging (AVB), and the related IEEE 1722 and IEEE 1733, standards over various networking link-layers. The organization will create compliance test procedures and processes that ensure AVB interoperability of networked A/V devices, helping to provide the highest quality streaming A/V experience. The Alliance will promote awareness of the benefits of AVB technologies and intends to collaborate with other organizations and entities to make use of this work in their respective efforts to provide a better end-user A/V experience.

The Alliance is focused on applications of these technologies in the Automotive, Professional, and Consumer Electronics markets.

© 2009 by AVnu Alliance. All rights reserved. AVnu™, AVnu Alliance™, and AVnu design and logos are trademarks and of the AVnu Alliance. All other names and logos are trademarks and/or service marks of their respective owners. Specifications and content subject to change without notice.

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS, IMPLIED, OR STATUTORY. AVNU ALLIANCE MAKES NO GUARANTEES, CONDITIONS OR REPRESENTATIONS AS TO THE ACCURACY OR COMPLETENESS CONTAINED HEREIN. AVnu Alliance disclaims all liability, including liability for infringement, of any proprietary or intellectual property rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any proprietary or intellectual property rights is granted herein.



# Introduction

## Some historical background

Connections between A/V equipment have traditionally been analog single-purpose point-to-point one-way links. Even when the data in the A/V streams transitioned to digital, the point-to-point one-way link architecture was often reused. For example, audio connections moved from analog to I<sup>2</sup>S or SPDIF/AES, and the professional video world moved from RS-170A in analog days to SDI as digital production started to HD-SDI now, and consumer video has moved from analog composite to analog component to HDMI. Some of the need to preserve point-to-point one-way links was indeed necessary, since appropriate local area network technology was very expensive at the time the particular connection was first deployed (e.g., even a single I<sup>2</sup>S stream could not be reliably carried by 10Mbit/sec CSMA/CD Ethernet when I<sup>2</sup>S/SPDIF connections were first used in audio systems).

This dedicated connection model unfortunately resulted in a rat's nest of cabling in both professional and high-end consumer applications as shown in Figure 1.

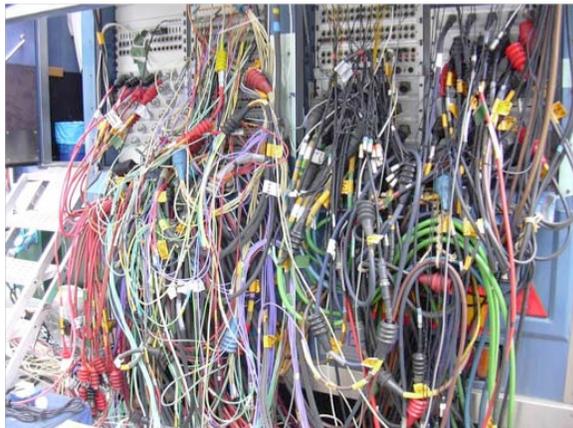


Figure 1 - Existing HD studio wiring

There have been several attempts to get around these problems:

1. specialized pro A/V technologies such as IEEE 1394/FireWire,
2. various non-standard wireless digital audio distribution systems for home theaters,
3. expensive and inflexible entertainment networks for automotive applications, and
4. adaptations of standard IT-type networks such as CobraNet.

The specialized A/V technologies for professional, home, and automotive use were *too* specialized, in that they did not have any easy kind of interoperability with regular IT networks like Ethernet. This limited their market to those applications that needed those particular services, like video cameras or professional audio equipment.

The adaptations of standard IT networks had the opposite problem: they were built from commodity equipment, but getting a professional level of service<sup>1</sup> required very tight control over how the equipment was used and managed. The introduction of even one “unmanaged” device could cause the whole network to fail.

## Requirements for A/V streaming

So, what's wrong with standard IT networks? Before that can be answered, the requirements for A/V streaming must be understood:

1. It must be possible to synchronize multiple streams so that they can be rendered correctly in time with respect to each other. At its simplest case, this might be guaranteeing “lip synch” so that the audio and video aspects of a

---

<sup>1</sup> “Professional level of service” means very low and deterministic delays around 2ms, no dropping of packets, and the ability to tightly synchronize endpoint playback of a stream to within a microsecond or so.



movie or television show are not out of synchronization. A much more stringent requirement comes from the need to keep multiple digital speakers properly in phase: for the professional environment, this means keeping streams synchronized within approximately one microsecond.

2. The worst case delay for a stream in the network, including buffering delays at the source and destination, must be low and deterministic. For almost all consumer applications, this means that the network must not significantly contribute to the user-interface delay: the time from when the consumer requests an action (e.g., presses a button on a controller) until that action is readily perceived by the consumer (e.g., “play”, “pause”, “forward”, “reverse”, etc.). This is something on the order of 50ms. For live performance, studio, or gaming applications the requirements are much more rigorous, on the order of 2ms.
3. Finally, applications must be able to get a high level of confidence that the network resources needed are available and will remain available as long as the application needs it. This is sometimes called a “reservation”, and sometimes this is called “admission control”. The intent is for an application to notify the network of the requirements for a stream ahead of time, and have the network lock down the resources needed for that stream and, if they are not available, to notify the application. Typical resources needed by A/V streams are throughput and specific bounds on delay.

#### The problems with existing approaches

Almost all current network equipment is based on IT requirements: move the data through the network as quickly as possible with minimum cost and minimal management. This is an

excellent approach as long as there are no hard limits on delay or synchronization requirements. IT-oriented networks do not always, however, meet the requirements of the previous section:

1. There is no concept of “time” in an IT network – There is nothing in the network infrastructure itself that can aid in synchronization or provide any kind of precision timing mechanism.
2. Delays can be too high – Although delay through a network may, on the average, be very low, there is little effort made to limit that delay. In an IT network, delivering data reliably is viewed as much more important than for the data to be delivered within a specific time.
3. The network itself does not prevent network congestion, so data can be lost if buffers are inadequate or link bandwidth is insufficient for offered traffic – IT networks count on higher level protocols to handle congestion (e.g., TCP) by throttling transmission and retransmitting dropped packets. This is adequate when long delays are acceptable, but will not work where low deterministic delays are the requirement.

The typical way these last two problems are handled today is with buffering ... but excessive buffering can cause delays that are annoying in the consumer environment, and completely unacceptable in a professional one.

Another way to allow existing IT-oriented networks to be used for A/V streams is to “manage” the network at a higher layer or to impose very strictly defined, inflexible configurations. For example, in the professional market, there are a few systems in place that can provide adequate delays and guaranteed bandwidth, but they require a single proprietary solution, are initially configured using a special “system generation” program, and need to be reconfigured every time a new device is added.



CobraNet is an example of this kind of architecture.

The vision: “no compromise” streaming quality

Several years ago, an effort was started within the IEEE 802.3 (Ethernet) working group to define a “residential Ethernet” which would directly address the challenges of A/V streaming. It quickly moved over to the IEEE 802.1 working group, since that is where all the major work needed to be done, and because that group is responsible for all the “cross network” bridging specifications. In particular, the group wanted to ensure that the technology was scalable from consumer applications (home/auto) to very high professional standards.

The rest of this paper shows how these emerging IEEE 802 specifications can be applied to meet these challenges. The general term for this technology is “Audio Video Bridging” or AVB.

## Summary of Audio Video Bridging

An “Audio Video Bridging” network is one that implements a set of protocols being developed by the IEEE 802.1 Audio/Video Bridging Task Group. There are four primary differences between the proposed Audio Video Bridging architecture and existing 802 architectures (from now on the term “AVB” will be used instead of “Audio Video Bridging”):

1. precise synchronization,
2. traffic shaping for media streams,
3. admission controls, and
4. identification of non-participating devices.

These are implemented using relatively small extensions to standard layer-2 MACs<sup>2</sup> and

<sup>2</sup> A “layer-2 MAC” is the part of the network interface that controls access to the network

bridges. This “minimal change” philosophy allows non-AVB and AVB devices to communicate using standard 802 frames. However, as shown in Figure , only AVB devices are able to:

- i) reserve a portion of network resources through the use of admission control and traffic shaping and
- ii) send and receive the new timing-based frames.

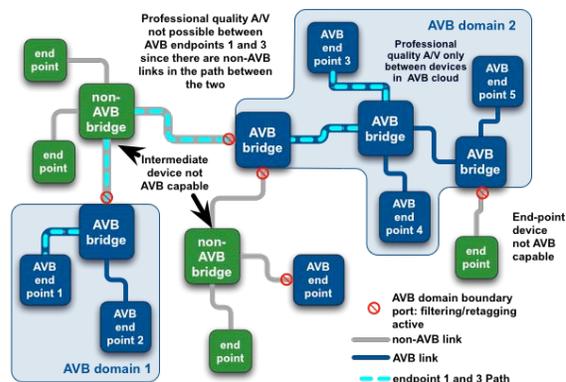


Figure 2 - AVB connections

### Precise synchronization

AVB devices periodically exchange timing information that allows both ends of the link to synchronize their time base reference clock very precisely. This precise synchronization has two purposes:

1. to allow synchronization of multiple streams and
2. provide a common time base for sampling/receiving data streams at a source device and presenting those streams at the destination device with the same relative timing. The protocol used for maintaining timing synchronization is specified in IEEE 802.1AS, which is a very tightly-constrained subset of another IEEE standard (IEEE Std 1588-2008), with

media (e.g., wire or frequency band). “MAC” is an acronym for “Media Access Controller”.



extensions to support 802.11 and also generic “coordinated shared networks” (CSNs – examples include some wireless, coax, and power line technologies). IEEE 1588 is currently used for industrial control and test and measurement applications.

An 802.1AS network timing domain is formed when all devices follow the requirements of the 802.1AS standard and communicate with each other using the IEEE 802.1AS protocol. Within the timing domain there is a single device that provides a master timing signal called the “Grand Master Clock”. All other devices synchronize their clocks with this master as shown in Figure 3.

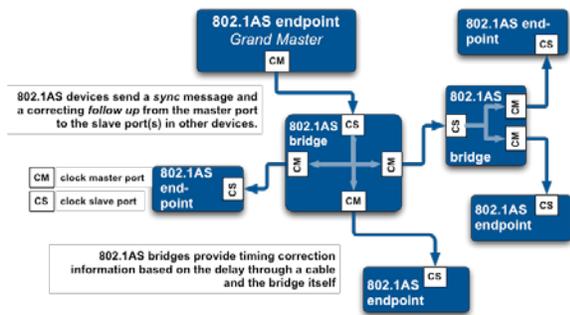


Figure 3 – Clocking hierarchy

The device acting as Grand Master can either be auto selected or can be specifically assigned (e.g., if the network is used in a professional environment that needs “house clock” (audio), or “genlock” (video), or if the timing hierarchy needs to be specified for other reasons).

AVB devices typically exchange capability information after physical link establishment. If peer devices on a link are network synchronization capable they will start to exchange clock synchronization frames. If not, then an AVB timing domain boundary is determined (as shown in Figure 2).

#### Traffic shaping for AV streams

In order to provide professional AV services, the AVB architecture implements traffic shaping using existing 802.1Q forwarding and priority

mechanisms but also defines a particular relationship between priority tags and frame forwarding behavior at endpoints and bridges.

Traffic shaping is the process of smoothing out the traffic for a stream so that the packets making up the stream are evenly distributed in time. If traffic shaping is not done at sources and bridges, then the packets tend to “bunch” into bursts of traffic that can overwhelm the buffers in subsequent bridges, switches and other infrastructure devices (“bunching” is described in greater detail in the following sections).

#### Tagging requirements at the stream source and the bridge

AVB streams consist of 802 frames with priority tagging, with normal restrictions on format and length. The default 802.1 tagging for a particular market segment will be chosen to avoid potential conflict with existing uses of the 802.1 priority tags within that market segment

#### Traffic shaping at the stream source

Endpoint devices are required to very evenly transmit frames for a particular stream based on the AVB traffic class and the specific QoS parameters that were used when the stream was OK'd by the network (see “Admission controls” below). The specific rules for traffic shaping are described in the IEEE 802.1Qav specification, and are a simple form of what is known as “leaky bucket” credit-based shaping where the bandwidth reserved for a stream controls the time between the packets that make up the stream.,



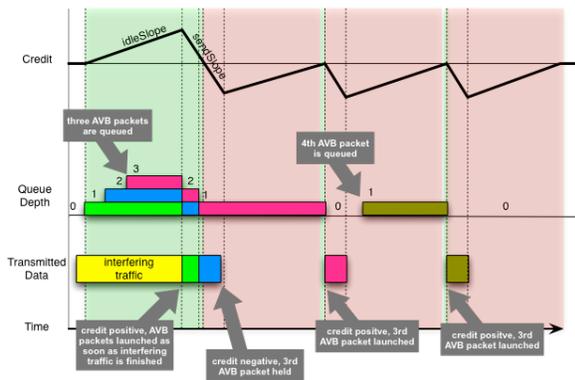


Figure 4 – Example Qav traffic shaping

### Traffic shaping at an AVB bridge

The traffic shaping mechanism used by stream sources is also employed by AVB bridges. AVB frames are forwarded with precedence over Best Effort traffic (i.e., reserved AVB stream traffic traversing an AVB bridge has forwarding precedence over non-reserved traffic) and will be subjected to traffic shaping rules (they may need to wait for sufficient credits).

Just like for stream sources, the traffic shaping rules for bridges require that frames be distributed very evenly in time, but only on an aggregate class basis rather than on a per-stream basis. This means that all the AVB traffic being transmitted out a particular port is distributed evenly in time measured using the QoS parameters of that class (this is the sum of the bandwidths of all the reservations for a particular AVB class for the particular port as made by the admission control process described below). This has the effect of smoothing out the delivery times (preventing “bunching” of frames<sup>3</sup>) as a stream propagates

<sup>3</sup> “Bunching” is a phenomenon where higher priority frames tend to group together as they are queued behind a lower priority frame currently being transmitted out an egress port. Without traffic shaping, these queued high priority frames are transmitted one after another in a “bunch”. Multiple “bunches” can arrive on different ingress ports on the next

through a network. The limited “bunching” has the very useful benefit of placing a relatively small upper limit on the size of the AVB output buffers needed at all egress ports on a bridge, independent of the number of hops in the path. This bounded buffer size is a key attribute that enables bounded delay and eliminates network congestion for admitted AV streams in AVB networks even if non-admitted traffic does experience congestion.

### Admission controls

Even though the preceding mechanism can reliably deliver data with a deterministic low latency and low jitter, it will only do so if the network resources (e.g., throughput on a port, buffer space in a bridge) are available along the entire path from the talker to the listener(s). In the AVB protocols, the term ‘talker’ is used to denote a stream source and ‘listener’ denotes a stream destination. In this architecture, it is both the talker’s and the listener’s responsibility to guarantee the path is available and to reserve the resources. The process to do this is specified by the 802.1Qat “Stream Reservation Protocol” (SRP). This protocol registers a stream and reserves the resources required through the entire path taken by the stream:

Talkers initiate by sending an SRP “talker advertise” message. This message includes a Stream ID composed of the MAC address of the stream source plus a talker-specific 16-bit unique ID and the MAC address of the stream destination. Additionally, the “talker advertise” message includes QoS requirements (e.g., AVB traffic class and data rate information), and accumulated worst case latency. Even though the address and QoS requirements are originated by the talker, the worst case latency is recalculated at every bridge so that the listener can communicate this information to higher layers to do media synchronization.

bridge at the same time and get queued behind another in-progress lower priority frame, resulting in an even larger “bunch”.



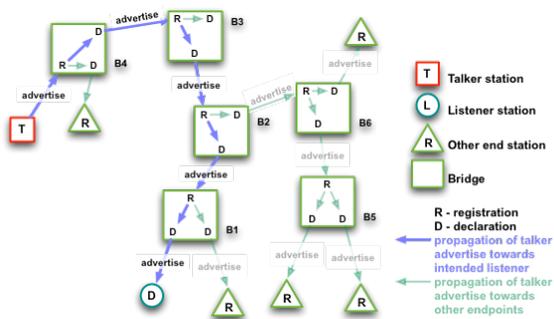


Figure 5 – Successful reservation (talker advertise)

All AVB intermediate bridges receiving a “talker advertise” message check for bandwidth availability on their output ports. If the bridge has sufficient resources available on that port, then the “talker advertise” is propagated to the next station.

If the resources are not available, rather than propagating the advertise message, the bridge sends a “talker failed” message. Included in this message is a failure code and bridge identification such that a higher-layer application can provide error checking or notification. An intermediate bridge receiving a “talker failed” will just pass on the message out towards the listener.

When a listener receives a “talker advertise” message, it will know whether the resources are available, and if so, the latency for the path. It can then respond with a “listener ready” message that is forwarded back towards the talker. Intermediate bridges use the “ready” message to lock down the resources needed by the stream and to make the appropriate entries in their forwarding database to allow the stream to be sent on the port that received the “ready” message. When the talker receives a “ready” message, it can start transmitting the stream.

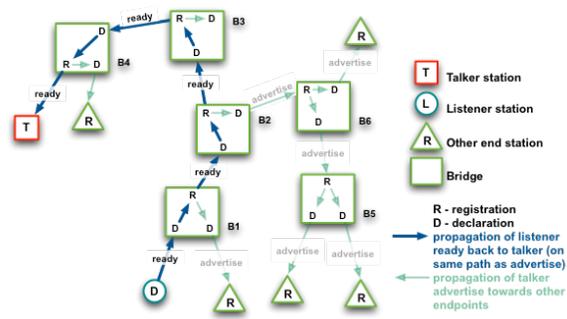


Figure 6 – Reservation acknowledge (listener ready)

The talker can explicitly tear down a stream by de-registering the “talker advertise”, and a listener can disconnect by de-registering the “listener ready”. A de-registration message propagates through the network in the same manner as the original registration.

There are also implicit methods used to tear down a connection and release the allocated resources. For example, the listener must periodically resend registrations and “ready” messages, and talkers must periodically resend “advertise” messages. That way any receiving device (including intermediate bridges) could automatically release assigned resources and notify higher layers if the appropriate registrations and reservations were not received due to a system that, for example, suddenly lost power.

#### LAN-specific considerations

Although the intent of the AVB Task Group is to provide a LAN-technology-independent method for requesting and providing streaming services, the characteristics and architectures of different LAN technologies require specific ways of implementing those services as outlined in the next few sections.



### *IEEE 802.3 / Ethernet Links*

Today Ethernet devices predominantly support full-duplex operation at 100Mbps or greater.

Thus, since the total available bandwidth available over such an Ethernet link is both known and constant, an AVB reservation over those Ethernet links combined with the appropriate traffic shaping assures both throughput and delivery latency parameters are met for packets of reserved streams. Since the bandwidth and delivery timing cannot be assured between two devices in an older shared CSMA/CD Ethernet using hubs, these older technologies are not supported by AVB.

AVB's Ethernet time synchronization standard, 802.1AS, leverages and simplifies deployed IEEE 1588-2008.

### *IEEE 802.11 / Wireless LAN*

To date, the AVB support planned for 802.11 is limited to time synchronization. 802.1AS provides for accurate time synchronization over 802.11 links, in part by invoking MAC-specific timestamp-reporting primitives defined in draft 802.11v. The time synchronization protocol defined by 802.1AS has been designed to be resilient to the transmission characteristics that are possible on wireless medium.

### *Coordinated Shared Network Links*

Several MAC/PHY specifications and standards are currently deployed or being developed which operate over existing wires within the home (e.g. AC power lines, coax cabling,). These wires are electrically "Shared" between multiple devices (not point-to-point like Ethernet), so to provide predictable performance, transmission of information onto the wire is "Coordinated" to avoid collisions by one of the devices on the network. Such networking technologies are typically called Coordinated Shared Networks (CSN).

If the CSN provides an access method with bounded latency (as most do), and if accurate link-specific time stamping or clock distribution

is available for the CSN, then extensions can be defined to take advantage of them

### *Identification of participating devices*

Since the whole AVB scheme depends on the participation of all devices between the talker and listener, any network element that does not support AVB (including so-called "unmanaged bridges") must be identified and flagged. The process to do this is described in the developing IEEE 802.1BA "Audio Video Bridging Systems" standard, which specifies the default configuration for AVB devices in a network.

For Ethernet, the method specified by 802.1BA to determine if its peer is AVB capable is a combination of 802.3 link capabilities (determined during Ethernet link establishment) and the link delay measurements done by IEEE 802.1AS. An AVB capable Ethernet port uses AVB operation if:

1. the link is full duplex 100Mbps or greater, and
2. the 802.1AS protocol discovers exactly one peer, and
3. the round-trip delay to the responding AVB device is no more than a worst case wire delay (computed from the IEEE 802.1AS "PDelay" exchange) Note: the worst case wire delay is less than that of a non AVB switch, and
4. an SRP reservation request or acknowledge is received on the port.

Other layer 2 connections will have their own specific methods to identify cooperating peers.

Even though a port may be enabled for AVB operation, there is a possibility that a complete end-to-end AVB connection cannot be made to another endpoint device that is AVB enabled. For example, in Figure above, devices in AV domain 1 cannot establish an AVB connection to devices in AV domain 2. An AVB connection can only be assured if a successful reservation is made using SRP and SRP "talker advertise"



messages will be marked as “failed” when they are sent out a non-AVB port.

### Higher layer protocols

For applications to take advantages of the features of AVB, there needs to be some coordination with portions of the higher layer communication protocols in between. In addition, some transport protocols have been adapted to provide information for applications to use AVB.

An application can implement synchronized distributed rendering using 802.1AS and higher layers. Specific audio samples and/or video frames carried by higher-layer protocols are given an associated presentation time (in terms of the shared 802.1AS clock) by the media source that is also an AVB talker. Each media renderer, that is also an AVB listener, renders the referenced audio sample or video frame at the 802.1AS presentation time. For example

1. If an application uses the IETF “Real Time Protocol” (RTP), it can use a new RTCP payload format defined in IEEE 1733 that correlates the RTP timestamp with the 802.1AS presentation time. The applications at the renderer(s) then use that correlation to translate the RTP timestamp to the presentation time stamp allowing the renderer(s) to start playing at the same time and keep playing at the same rate.
2. Applications using HTTP can also take advantage of AVB’s time synchronization by carrying a presentation time. E.g., MPEG transport streams that require clock synchronization between the server and client can include Transport Time Stamps (TTS) as defined by ARIB TTS [ARIB STD-B24] which are derived from the 802.1AS clock. Similarly, an application could utilize clock synchronization through methods described in ISO 13818-1 Annex J which includes a discussion of various clock recovery schemes proposed for MPEG2 Transport Streams over jitter inducing networks, and figure J.2 illustrates a

simple way to use the 802.1AS clock for this purpose.

3. Applications using IEC 61883 formats can use procedures defined in IEEE 1722 to sample the 802.1AS clock at the start of an A/V data block and then add the worst case transport delay to the sample time to get a presentation time which is inserted into the 1722 packet.

If the media source is not a real-time source (e.g. a media file on a mass storage device), the presentation times can be generated based on the nominal media rate. If the media source is a real-time source (e.g. a microphone), the presentation time can be constructed by the talker based on its observation of the 802.1AS time in relation to the microphone’s sample clock.

Other higher layer services can use AVB in a similar way. Existing connection management schemes, for example, can use the AVB SRP reservation services by mapping their internal stream identifiers with the SRP stream ID. Once a connection is established, streaming can start. E.g., applications using RTP transmit RTCP packets defined by IEEE 1733 that correlate the SSRC to the SRP stream ID. Furthermore, listener applications using 1722 use the SRP stream ID to discriminate between different streams.

## Conclusion

We have illustrated how the open standard AVB protocols work, and how their synchronization and quality of service fit into the architecture of a modern audio/video network.

For more information about AVnu please see <http://www.AVnu.org>.

