

A multiyear effort by the IEEE 802.1 Audio/Video Bridging (AVB) Task Group is nearing completion on a series of enhancements to the legacy Ethernet standards. Here's what it means to embedded systems designers.

# Understanding IEEE's new audio video bridging standards

BY ROBERT BOATRIGHT

A multiyear effort by the IEEE 802.1 Audio/Video Bridging (AVB) Task Group is nearing completion on a series of enhancements to the legacy Ethernet standards that enable the delivery of time-synchronized, low-latency audio and video over Ethernet networks—with perfect Quality of Service (QoS)—while retaining 100% compatibility with legacy Ethernet networks.

IEEE 1722, the AVB Transport Protocol (AVBTP), has further built on the development of AVB by adapting IEEE 1394's comprehensive suite of media formats, encapsulations, and synchronization mechanisms for use in Ethernet AVB networks. Products using IEEE 802.1 AVB and AVBTP standards are now being introduced that enable the construction of highly interoperable Ethernet networks capable of streaming audio and video with perfect QoS.

This article provides an overview of the low-level IEEE 802.1 AVB Task Group standards (aka “the plumbing”)

and how AVBTP further leverages the benefits of reliable plumbing by defining a standard for media exchange to enable vendor-independent interoperability.

Understanding the purpose, interaction, and usage of these new protocols will give the embedded systems designer the background necessary to begin architecting a new generation of streaming audio and video embedded devices with price and performance heretofore unachievable.

**EVOLUTION OF LEGACY ETHERNET**  
Connectivity is an inescapable, recur-

ring theme common to modern embedded systems devices, and Ethernet has long since been declared the gold standard of networking. The ubiquity of Ethernet can be witnessed by how rare it is to find any microprocessor-powered product without an Ethernet jack.

The goal of those developing AVB is striking in its simplicity—extend Ethernet's unmatched data-networking capabilities to the realm of reliable real-time audio/video networking.

By design, the IEEE standards committees have focused on retargeting existing technology rather than inventing yet more new—and incompatible—standards. This focus on leveraging existing technology has accelerated AVB development by taking advantage of existing infrastructure. A key criterion during the AVBTP definition stage was ease of implementation in hardware.

### IEEE 802.1 AVB STANDARDS

Here's a summary of the core IEEE 802.1 AVB standards:

**IEEE 802.1AS Precision Time Protocol (PTP):**<sup>1</sup> Based on IEEE 1588:2002,<sup>2</sup> PTP devices exchange standard Ethernet messages that synchronize network nodes to a common time reference by defining clock master selection and negotiation algorithms, link delay measurement and compensation, and clock rate matching and adjustment mechanisms.

Designed as a simplified profile of IEEE 1588, a primary difference between 1588 and IEEE 802.1AS is that PTP is a layer 2—in other words, a non-IP routable protocol. Like IEEE 1588, PTP defines an automatic method for negotiating the network clock master, the Best Master Clock Algorithm (BMCA). PTP nodes can be assigned one of eight priority levels, presumably based on clock quality. BMCA defines the underlying

negotiation and signaling mechanism whose purpose is to identify the AVB LAN Grandmaster. Once a Grandmaster has been selected, synchronization automatically begins.

At the core of 802.1AS synchronization is timestamping. In short, during PTP message ingress/egress from the 802.1AS-capable MAC, the PTP Ether-type triggers the sampling of the value of a local real-time counter (RTC). Slave nodes compare the value of their RTC against the PTP Grandmaster and, by use of link delay measurement and compensation techniques, match their RTC value to the time of the AVB LAN PTP domain. After network time throughout the AVB LAN has converged, periodic SYNC and FOLLOW\_UP messages provide the information that enables the PTP rate matching adjustment algorithms. The result is all PTP nodes are then synchronized to the same "Wall Clock" time. PTP guarantees 1- $\mu$ s accuracy over seven network hops.

### IEEE 802.1Qat Stream Reservation

**Protocol (SRP):**<sup>3</sup> Legacy IEEE 802 Ethernet standards are characterized—and limited—by their inability to deterministically prioritize time-sensitive streaming media ahead of legacy asynchronous messages. To provide guaranteed QoS, the Stream Reservation Protocol ensures end-to-end bandwidth availability before an A/V stream starts. If bandwidth is available, it is "locked down" along the entire path until implicitly or explicitly released. SRP utilizes IEEE 802.1ak Multiple Registration Protocol as the messaging protocol to pass stream descriptors and resource reservation requests/results.

Switches implementing SRP can reserve—and defend—up to 75% of the available capacity on a given LAN link within an AVB "Defended Cloud."<sup>4</sup>

### IEEE 802.1Qav Queuing and Forwarding Protocol (Qav):

<sup>5</sup> The job of the Queuing and Forwarding protocol is to ensure that legacy asynchronous Ethernet traffic does not and can not inter-

# Your solution is here.

**Save time – and money**—with embedded software solutions built to run right out of the box. Get development started quickly, with no integration required and full support for popular tools. With Micro Digital you have low-cost, no-royalty licensing, full source code, and direct programmer support. So get your project off to a great start. Visit us at [www.smxrtos.com](http://www.smxrtos.com) today.

Free Evaluation Kits: [www.smxrtos.com/eval](http://www.smxrtos.com/eval)  
Free Demos: [www.smxrtos.com/demo](http://www.smxrtos.com/demo)

**ud** Micro Digital

RTOS INNOVATORS  
800.366.2491 [sales@smxrtos.com](mailto:sales@smxrtos.com)

[www.smxrtos.com](http://www.smxrtos.com)

ARM • ColdFire • Cortex • PowerPC • x86 • CodeWarrior • CrossWorks • GCC • IAR EWARM



fere with streaming AVB traffic. Fortunately for product designers, most of the burden of implementing 802.1Qav falls squarely on the shoulders of AVB-enabled Ethernet switches.

There are, however, defined pacing requirements that AVB media sources are constrained to follow. To further illustrate this point, imagine an AVB “talker” node in a live audio application. As analog audio is captured, digitized, and presented to the network, it is inherently isochronously paced. Such applications are natively suited for AVB transmission and will not be affected by 802.1Qav’s pacing requirements.

On the other hand, problems could arise if ill-behaved stored media applications (such as PCs) were allowed to send large bursts of irregularly paced audio/video samples. Think of it this way: you can’t send more than 1 Gbps down a gigabit Ethernet link. Because AVB guarantees the delivery of (up to) 750 Mbps of streaming data on that link, there is still capacity for 250 Mbps of asynchronous data. 802.1Qav defines how the link must be shared and how isochronous and asynchronous data are prioritized inside Ethernet AVB switches.

Previous attempts at solving what is inherently a shaping problem have relied on using memory to smooth burst-

ed traffic but at what cost? Memory buffers are the enemy of low latency and low BOM cost. AVB networks lower costs by eliminating the need for costly buffers while guaranteeing low end-to-end latency of less than 2 milliseconds over seven network hops.

A more precise definition of 802.1Qav can be found on the IEEE website: “This standard allows bridges to provide guarantees for time-sensitive (i.e., bounded latency and delivery variation), loss-sensitive real-time audio video (AV) data transmission (AV traffic). It specifies per priority ingress metering, priority regeneration, and timing-aware queue draining algorithms. This standard uses the timing derived from IEEE 802.1AS. Virtual Local Area Network (VLAN) tag encoded priority values are allocated, in aggregate, to segregate frames among controlled and non-controlled queues, allowing simultaneous support of both AV traffic and other bridged traffic over and between wired and wireless Local Area

## AVBTP network stack.

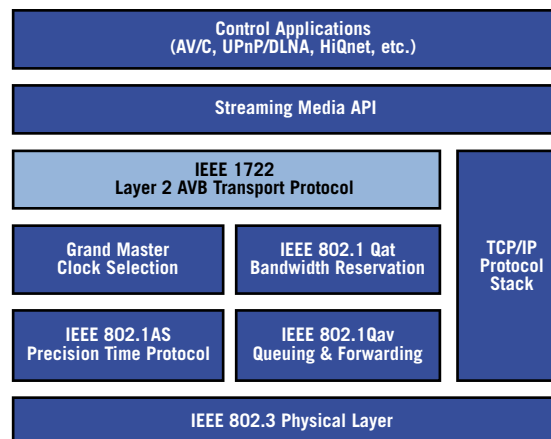


Figure 1

Networks (LANs).”<sup>6</sup>

What a great description! Even better, it means that all the hard work and heavy lifting is accomplished inside the Ethernet switch. These same changes are already shipping in standards-based Ethernet AVB-capable switches and differentiate AVB from the single-sourced, expensive, and proprietary niche media networking technologies that have come before.

## WHERE DOES AVBTP FIT?

In short, the purpose of AVBTP is to act as a wire—with delay—to allow the logi-

## IEEE 1722 packet construction.

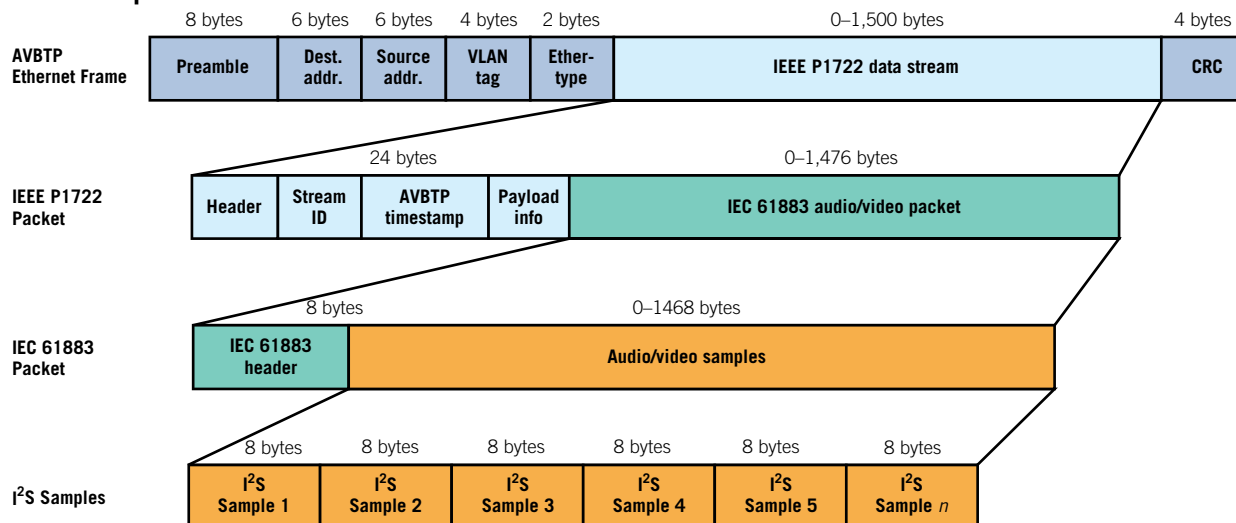


Figure 2

## Presentation timestamps.

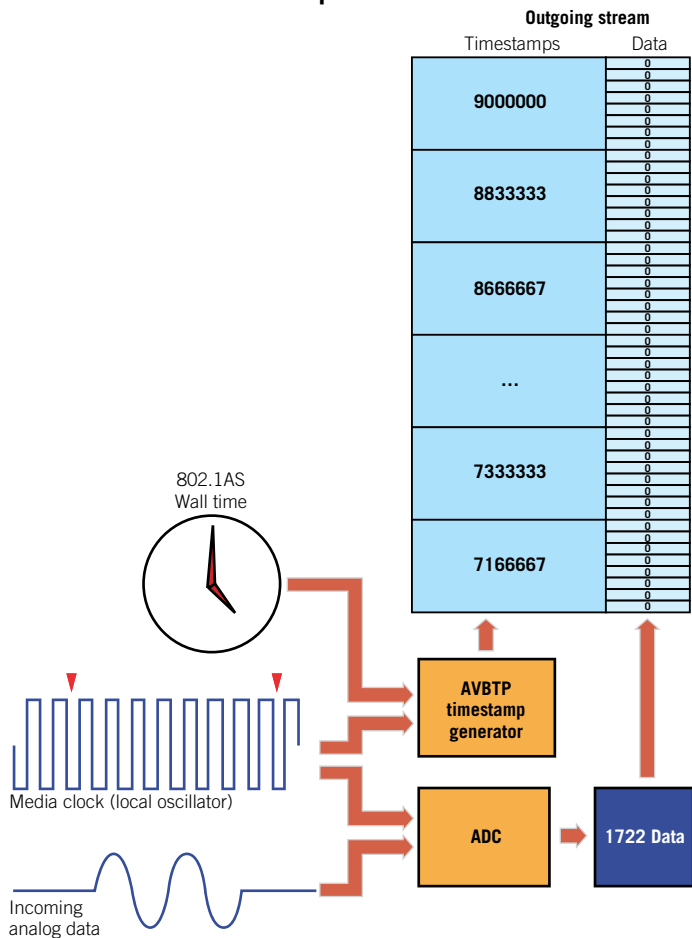


Figure 3

## Cross-time stamped clock recovery.

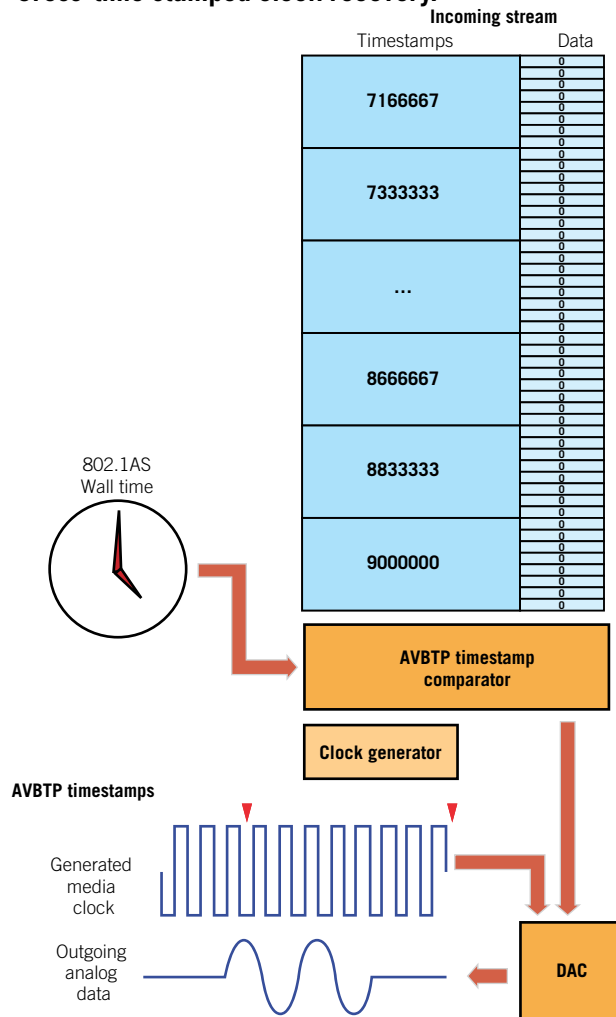


Figure 4

cal connection of physically distant CODECs. More long windedly, AVBTP abstracts the underlying network transmission channel to enable the virtual connection of distributed audio and video CODECs over reliable Ethernet networks.

No article on networking would be complete without the obligatory network stack diagram. As shown in **Figure 1**, AVBTP sits above the IEEE 802.1 AVB plumbing and below the application layer. It acts as the conduit between an Ethernet MAC and a streaming application.

Streaming samples are connected to the application-side of the AVBTP interface. Supported media formats include a

variety of raw and compressed audio and video including I<sup>2</sup>S audio, IEC 60958 SPDIF, MPEG2/4 and H.264 Transport Streams, Bt.601/656 raw video, and so forth. A complete listing of the supported media types and an overview of the synchronization mechanisms appears later in this article.

On the bottom of the AVB Transport Protocol, connection is made to IEEE 802.1 AVB-capable Media Access Controllers (MAC). In essence, media streams are fed into the top of AVBTP and IEEE 1722 Ethernet packets pop out the bottom—and vice versa.

Vendor-independent interoperability is provided by defining supported

media types, their associated formats, and the location, organization, and orientation of media samples within the AVBTP Ethernet frame.

#### PACKETS WITHIN PACKETS (WITHIN PACKETS...)

A complete AVBTP Ethernet packet is shown in **Figure 2**. The AVBTP Ethernet frame is uniquely identified by the AVBTP Ethertype.

Within a P1722 packet, the content type is described by the Payload Info field.

The initial implementation of AVBTP supports IEC 61883 content however AVBTP has been designed to

allow future expansion to support other yet-to-be-defined media standards.

Figure 2 illustrates how IEC 61883-6 AM824 uncompressed audio samples are encapsulated in an Ethernet frame.

### SUPPORTED MEDIA TYPES

Continuing the theme of leveraging existing standards, AVBTP takes advantage of the media formats and synchronization mechanisms specified by IEC 61883 for use in IEEE 1394 FireWire devices.

Formats based on IEC 61883<sup>7</sup> parts 1–8 are:

- 61883-2 SD-DVCR
- 61883-4 MPEG2-TS Compressed Video
- 61883-6 Uncompressed Audio
- 61883-7 Satellite TV MPEG
- 61883-8 Bt.601/656 Video

- IIDC Uncompressed Industrial Cameras

A notable side benefit of basing AVBTP on IEC 61883 is the relative ease with which a FireWire gateway<sup>8</sup> to an AVB network could be implemented.

### SYNCHRONIZATION

Synchronization in an AVB network starts with the Precision Time Protocol but ends with synchronized media clocks. PTP is responsible for synchronizing all nodes in an AVB network to identical wall clock time; not for synchronizing media clocks. In other words, PTP does not actually transport synchronized media clocks but instead provides a low-level building block crucial for managing a distributed media synchronization system.

A crucial benefit of this approach is coexistence of multiple, independent

media clock domains on an AVB network. Unrelated audio and video streams can simultaneously exist in the same LAN.

### CROSS-TIMESTAMPING AND PRESENTATION TIMESTAMPS

AVBTP assumes that AVB node media clocks are clocked by free-running oscillators. It is also assumed that the node's internal concept of wall clock time has been synchronized to the PTP Grandmaster. AVBTP media clock sources embed "AVBTP Presentation Timestamps" in AVBTP streaming packets. **Figure 3** illustrates the relationship between PTP network time and AVBTP Presentation Timestamps.

Not all AVBTP streaming packets include an AVBTP Presentation Timestamps. The DBC (data block count) in the AVBTP header defines how often AVBTP Presentation Timestamps are

## Standards Make Sense

Standards improve quality and enable designers to share components across different projects. Today, ARM® Cortex™-M profile processors, combined with the Cortex Microcontroller Software Interface Standard (CMSIS) and optimized middleware from the industry's largest ecosystem, are setting the hardware and software standards for microcontrollers.

These standards enable leading vendors such as Luminary Micro, NXP, and STMicroelectronics to supply advanced microcontrollers, while maximizing code reuse across multiple platforms.

### Cortex-M3 Microcontrollers Make Sense

For more information visit  
[www.onARM.com](http://www.onARM.com)

*"STM32 microcontrollers revolutionize the market by combining high performance and low power with a scalable product range that fits every developer's needs."*



**Daniel Colonna**  
Microcontrollers Division  
Marketing Director



**Cortex**  
Intelligent Processors by ARM

**ARM** The Architecture for the Digital World®

© ARM Ltd. AD153 | 12.08



## cover feature

### AVBTP distributed PLL.




Figure 5

generated. A commonly used value for DBC is 8<sub>10</sub>. This means every eighth media sample, the value of the local PTP RTC is sampled, converted to a 32-bit value (in nanoseconds), and added to desired (if any) latency normalization value. More on latency normalization later.

### MEDIA CLOCK RECOVERY

Recreating synchronized media clocks follows a similar but reverse process.

The value of the DBC field in the P1722 header triggers how often the AVB media clock slave should recreate a media clock edge.

**Figure 4** illustrates the AVBTP media clock recovery mechanism. **Figure 5** illustrates one possible method of implementing what is essentially a distributed PLL. With only minimal additional filtering, sub-nanosecond jitter is easily obtainable.

### LATENCY NORMALIZATION AND PRESENTATION TIME

AVBTP Presentation Timestamps have another use besides media clock recovery. They also indicate to the AVBTP listener when, (in PTP network time,) to present received media samples to the output media interface. This simple but effective mechanism enables media synchronization across multiple media nodes.

AVBTP specifies that AVB end points have the ability to buffer 2 ms of media samples however, that value can be “dialed down” by an application to minimize latency. The purpose of this buffering is latency normalization. **Figure 6** (on page 35) illustrates this concept. It’s worth noting that this same mechanism inherently provides the ability to implement lip-to-ear audio/video synchronization.

### WHAT’S AN ENGINEER TO DO?

While the specifics of implementing an AVB-capable product are dictated by the ICs used, some rules-of-thumb are developing. For instance, AVB-capable

**Support for**

- STR7, STR9
- LPC2000, LPC3000
- AT91SAM
- STM32, Stellaris
- PIC32
- XE166
- ST10, S12X
- HCS08, XC800

and over 1000 other microcontrollers



*Reach the Top*

**Industrial**

**Debug Features**

- Support for all popular compilers
- Fast FLASH programming
- Software breakpoints in FLASH
- Test automation
- Long term support guaranteed

**CombiProbe Features**

- 128 MByte real-time trace
- Endless trace by streaming
- 4 bit ETM in continuous mode
- ITM via Serial Wire Viewer
- 4 bit program flow trace for PIC32
- Extensive profiling options



LAUTERBACH

DEVELOPMENT TOOLS

www.lauterbach.com

## Latency normalization.




Figure 6

switches will come from the manufacturer ready for use and require no further development; implementing AVB in an endpoint is somewhat more nuanced.

PTP hardware timestamping and rate matching algorithms are typically

wholly embedded in silicon leaving the BMCA and its priority assignments/interface to be implemented in host software.

SRP for an AVB endpoint is typically implemented entirely in host soft-

ware. The SRP and Qav firmware inside the AVB switch is responsible for defending the reservation.

AVBTP endpoint solutions seen to date have been implemented using a combination of dedicated hardware and

## DEFENDED CLOUD

Streaming QoS is guaranteed inside the AVB "defended cloud"—an interconnected LAN of AVB-capable endpoints and switches running the core AVB protocols: IEEE 802.1AS Precision Time Protocol (PTP), IEEE 802.1Qat Stream Reservation Protocol (SRP), and IEEE 802.1Qav Queuing and Forwarding Protocol.

To establish the defended cloud, a PTP Grandmaster is first elected. Next, the "wall clock" time of all AVB devices inside the defended cloud is synchronized using PTP. AVB "talkers" then advertise audio/video streams they can supply and finally, AVB "listeners" use SRP to reserve end-to-end bandwidth. At this point, guaranteed streaming of audio/video begins.

After the defended cloud is

### How the defended cloud works.




Figure 7

CONTINUED ON PAGE 37

embedded firmware with only minimal host application interaction required.

### THERE CAN BE ONLY ONE...

The economic realities of media networking dictate that ultimately there will be only one widely deployed media networking technology—and with the advent of AVB, that technology is clearly Ethernet. Worldwide, thousands of engineers are actively working on enhancing Ethernet infrastructure including wired and wireless physical layer devices, MACs, switches, redundancy, diagnostics, and so forth. For example, the IEEE Working Groups are already hard at work extending AVB to wireless 802.11 Ethernet.

Work is already afoot to use Ethernet AVB technology for consumer, professional, and automotive streaming audio/video networks.

In the consumer space, AVB will allow the sharing of streaming devices throughout the home. Digital video recorders (DVR), media servers, network-attached storage (NAS) boxes, AM/FM/Satellite tuners, and so forth will reside in a centralized equipment closet where A/V streams will be served throughout the home. UPnP/DLNA and AV/C-based distributed discovery and control Ethernet applications are already well understood and, with the advantage of AVB's reliable transport, will become even more widely deployed. Ultimately, the confusing bar-

rage of incompatible plugs found on the back of consumer devices will be replaced by the ubiquitous RJ45 jack.

Professional audio and video installations will likewise benefit from a single network. For example, modern houses of worship typically have distinct distribution solutions for audio (CobraNet), video (BNC/coax), and data (Ethernet). AVB will offer substantial costs savings by only requiring

**Work is already afoot to use Ethernet AVB technology for consumer, professional, and automotive streaming audio/video networks.**

the installation of a single network that can be shared by audio, video, and data nodes while offering additional benefits such as centralized control, redundancy and enhanced diagnostics.

Automotive infotainment networks have been common for over a decade but have run into the brick wall of limited bandwidth and inflexible topologies allowed by niche networks. Conversely, Ethernet AVB networks offer highly scalable bandwidth and flexible topologies increasingly demanded by

today's media-rich vehicles. Infotainment transport and driver's assistance solutions using AVB in automotive networks are already under development.

### THE MEDIUM IS THE MESSAGE

AVB-capable ICs with AVBTP implemented entirely in silicon have already been introduced by multiple silicon vendors—and more are slated to come online soon. Highly flexible networking of professional, consumer, and automotive A/V products using Ethernet is now possible.

The time for widespread networking of audio and video systems has arrived and a new wave of elegant, intelligent, and easily connected media-centric products will be the result. ■

Robert Boatright is the director of research for Harman International's Corporate Technology Group and serves as chair of the IEEE P1722 AVB Transport Protocol Working Group, is a voting member of the IEEE 802.1 Working Group, and is a member of the Audio Engineering Society. His background includes both the hardware design and engineering management for a wide variety of professional consumer and professional consumer audio, video, and graphic applications. He can be reached at [rboatright@ieee.org](mailto:rboatright@ieee.org).

### ENDNOTES:

- 1 P802.1AS: IEEE Standard for Local and Metropolitan Area Networks—Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks
- 2 IEEE 1588:2002 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems
- 3 P802.1Qat: IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks—Amendment 9: Stream Reservation Protocol (SRP)
- 4 See "Defended Cloud" sidebar
- 5 P802.1Qav: IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks—Amendment 11: Forwarding and Queuing for Time-Sensitive Streams
- 6 Taken from [www.ieee802.org/1/pages/802.1av.html](http://www.ieee802.org/1/pages/802.1av.html)
- 7 IEC 61883-[1:8]:2008
- 8 Annex C, IEEE P1722 draft v1.3

### CONTINUED FROM PAGE 35

established, synchronized (within 1µs) streaming traffic is guaranteed to be delivered with less than 2mS latency—over seven network hops. Up to 75% of any given link inside the defended cloud can be reserved for streaming traffic. This results in at least 25% of network bandwidth inside the AVB cloud remaining available for legacy "best effort" traffic.

Legacy devices can communicate through a defended cloud using any standard Ethernet protocol but this traffic is relegated to "best effort" status, as the figure in this sidebar illustrates. Delivery of guaranteed protocols such as TCP/IP remains guaranteed—only the timeliness is potentially affected. A key takeaway message is that because AVB has been developed by the IEEE 802 Working Group, it retains 100% backwards-compatibility with legacy Ethernet.