ARE WE THERE YET?

Audio networking over Ethernet in 2012.

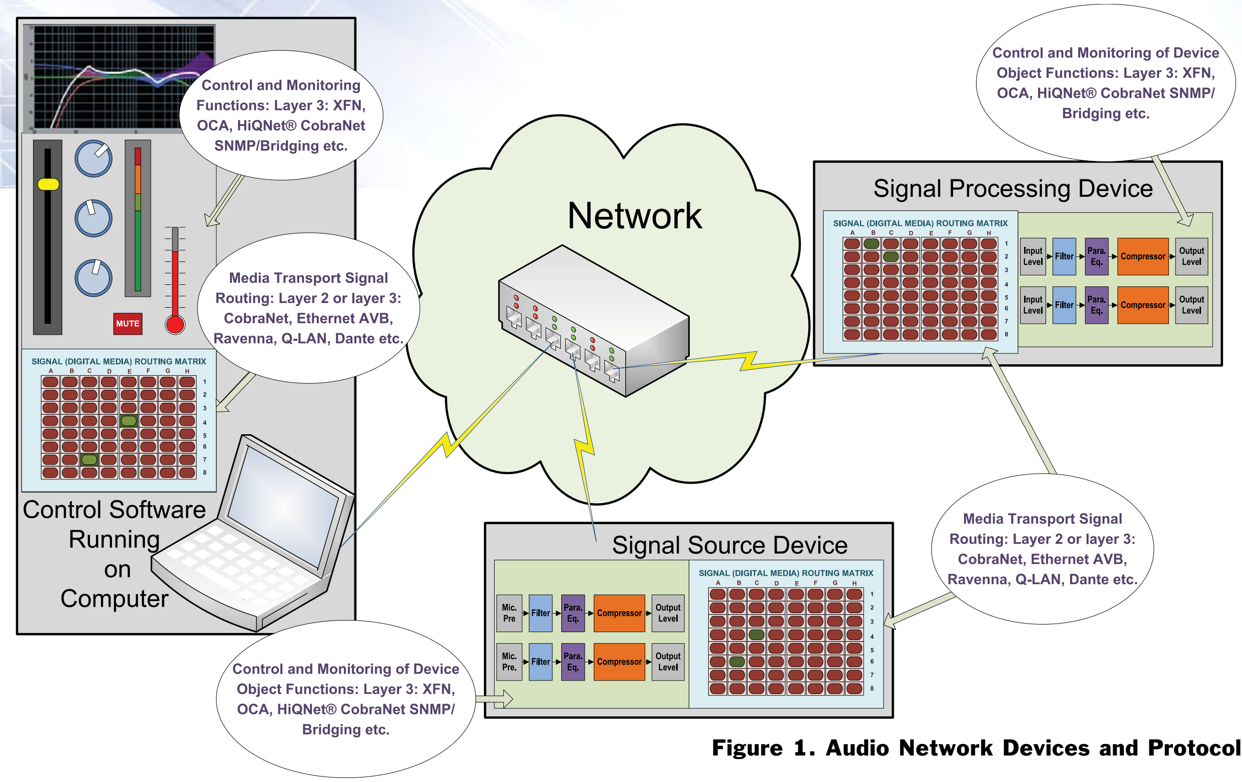

Figure 1. Audio Network Devices and Protocols

BY TIM SHUTTLEWORTH

The area of audio networking systems has seen some rapid development recently with new technologies being announced and discussed, promoted and critiqued. Systems integrators, design consultants and engineers are left struggling to comprehend all this technology and gain a perspective on what it all means for the short term and longer term future of audio systems.

Here, I will be addressing a couple of specific questions that I hear a lot from the systems integration community, and in doing so, will cover some basics of current network schemes for audio applications. Here are the two questions: Is Ethernet AVB the future? And, to plagiarize Bob Pease: What's all this layer 2 versus layer 3 stuff anyhow? These two questions are very much interrelated, and both must be addressed in order to provide a better understanding of where the various elements, protocols, methods and commercial solutions fit in a fully operational networked audio system.

Ethernet Suite Push

The push toward the suite of IEEE standards generally referred to as Ethernet AVB has received a lot of attention and created a deal of momentum in discussion forums from online blogs to conference workshops. Despite all this dissemination, it remains poorly understood in relation to how it fits the overall landscape of audio devices, controllers, computer software and GUIs that make up a networked audio system. So, let me state right up front that Ethernet AVB will standardize the way media (that's audio signals and video signals) are routed over Ethernet. AVB concerns itself with packaging these digital media signals, and ensuring that a connection route delivers this media content such that it can be reconstructed accurately, reliably and quickly at the receiving device.

What it does not do is provide a means to control any aspect of the sending or receiving device beyond the routing of media content out of and into devices via the network. As with so many things, there is a caveat to that statement: In the future, under

---

Tim Shuttleworth is editor of the Emerging Trends *Report for the Audio Engineering Society's Technical Committee on Networked Audio Systems. He is engineering manager with Renkus Heinz Inc., and is a technical papers reviewer and papers session moderator for the recent AES 44th Conference on Audio Networking. He holds patents for techniques across the spectrum of audio electronics, and is a co-author of the Harman HiQNet Protocol. He can be reached at t.shuttleworth@sbcglobal.net.*

the auspices of IEEE1722.1, a means of limited control of device functioning, such as volume and mute *may* be defined.

First generation proprietary audio-over-Ethernet protocols such as CobraNet not only provided a means of routing digitized audio signals between devices on a network, but they also included some methods of controlling functions within those devices such as gain, equalization, limiter/compressor settings, and also means to return signal information from devices to a monitoring controller such as a software GUI—as level meters. CobraNet, to continue that example, also included a means to encode other command and control functionality into network data packets to allow manufacturer-specific (custom) non-audio command and control data such as remote power/sleep and temperature or fault reporting back and forth between control devices and operational devices.

Proprietary Methods

Many adopters of CobraNet utilized (and still do) only the digital audio transport aspect of CobraNet and used their own proprietary methods for control of signal processing objects and reporting of status information using separate Ethernet interfaces (NICs) on their products to handle control and command data. Some even use non-Ethernet networks for these command and control data, such as RS422 and RS485. This allowed manufacturers to develop control protocols that addressed their specific feature set efficiently, but required end users to use separate software to control products from different brands on the same network, even with all of them utilizing CobraNet as the audio signal distribution method.

This practice was railed against by serious industry systems users such as the late Albert Leccese of Audio Analysts, who pointed out that what the end user community wants is a common control platform allowing a single software program (GUI) to find, identify and fully configure and control every audio device on a network.

Ethernet AVB provides a license-free standard schema for media transport over Ethernet. What it does not do is move our industry any closer to the goal of a common control scheme across branded products. The connection and routing management software for CobraNet, CobraNet Discovery, provided the same utility as now offered by Ethernet AVB more than 10 years ago, albeit in a closed proprietary protocol requiring license fees and per-channel royalties, and with some Ethernet compatibility challenges.

Recently, efforts have focused on the need for an open standard protocol for device command and control allowing discovery and control of processing objects embedded within network node devices. It is these command and control protocol open standards, coupled with standards for media transport such as Ethernet AVB and AES X-192, that actually move us toward the ultimate objective of a cross-brand common control platform for an entire networked audio system.

Table 1 diagrams where the different media transport and device control schemes map against the full implementation reference example of CobraNet, complete with control and monitoring of audio signal processing objects and non-audio parameters. Fig-

Table 1: Media Transport and Command and Control Protocols

CobraNet	Ethernet AVB	AES X-170 Draft Standard	OCA Alliance Proposed Protocol	Harman Proprietary Protocol
Non-Audio Control Data (temperature, fault, remote power) encoded/decoded/managed by CobraNet Passbridge or Packet Bridging		X-170 (XFN) Command and Control Protocol	OCA (Open Control Architecture) Command and Control Protocol	HiQNet Command and Control Protocol
Audio DSP Object Control and Monitoring using "CobraNet SNMP Library" commands				
Packetized audio routing and connection management. CobraNet Discovery	Ethernet AVB IEEE802.1Qav, IEEE802.1AS, IEEE7122.1, IEEE1733, etc.			
Audio stream digitization/packetization and timing control. CobraNet packets, CobraNet Conductor and beat packet protocol				

ure 1 diagrams how these protocols interact with network devices; controllers, computer software GUIs and signal processors, and the processing objects and parameters within them.

How Do Ethernet Layers Relate?

Armed with a better understanding of how the available protocols relate to operations across the network (such as transport of media over the network), and within devices on the network (control of signal processing and reporting of faults and so forth), let's now look at how the concept of Ethernet Layers relates to all this. A common game played at children's parties in Britain is called "Pass the Parcel." In this game, the host wraps something of value to the participants (in this case, a toy or chocolate bar, for instance) in multiple layers of wrapping paper. The parcel is passed around the attendees until a signal to stop when the current possessor of the parcel removes one layer of wrapping.

The process continues until a lucky participant removes the final wrapper and gains possession of the prize. Ethernet Layering is a lot like this game, except that the outer layers have to be added according to strict rules governed by the authority in charge of the Ethernet game, the IEEE. The first layers of wrapping can be applied by the "host" according to rules they dictate as long as the intended recipient is in possession of an understanding of the same rules. We'll call these Layer 3. As long as the layers of wrapping can be interpreted by both the sending and receiving application, the prize—the data payload at the core of the packet—will reach its intended destination and be usable and valuable.

But the outer layers must be fully compliant with the standardized rules of Ethernet protocols so the entire network infrastructure understands the wrappers. We'll call this Layer 2. The wrapping layers for Layer 2 have to be understood by the entire network infrastructure because these are the layers of data that determine how and where the data packet is to go.

Started As OSI

This whole concept started out as the OSI (Open Systems Interconnection) model that defined layers of wrapping for network data packets. As developers of Ethernet systems became more familiar with implementing these layers, they became aggregated, at least in discussion vocabulary, into the three layer model most commonly referred to in the Ethernet community. It is this contemporary 3 Layer model being described in the wrapping of the parcel analogy. Layer 1 is where the actual connection to the physical network is made. Above Layer 3 sits the software application that's being used to do some work, such as browse the internet, transfer files from computer A to computer B, or communicate instructions to a remote device and, of course, for streaming media.

Table 2 diagrams the relationship between the OSI model's layers and the more commonly used 3 Layer TCP/IP

Table 2: The Network Layer Schemes

	OSI Model		3 Layer TCP/IP Ethernets Model	
This layer passes data back and forth from/to the application software and provides utilities	7	Application Layer	Layer 3	Network management, Utility Protocols, etc., and Applications
This layer converts between the application data format and standard network data	6	Presentation Layer		
This layer manages the application to application connections across the network	5	Session Layer		
This layer typically communicates up through ports and communicates down using TCP or UDP	4	Transport Layer	Layer 2	Standard TCP/IP Layer
This layer supports IP addresses and subnetting	3	Network Layer		
This layer is where the MAC address applies	2	Data Link Layer	Layer 1	Ethernet Layer
This layer is the physical connection	1	Physical Layer		

Ethernet model we're actually talking about. (TCP/IP refers to the Transmission Control Protocol/Internet Protocol suite of computer network protocols, often called the Internet Model or IP Model, which defines a standardized method to enable computers to communicate over a network and encompasses the functions described in the left column of Table 2. It's the basis on which both Ethernet local area networks (LANs) and the internet work.)

The important point about all this is that Layer 2 wrappings have to be handled and managed by the network infrastructure to ensure correct routing and prioritization, as well as appropriate bandwidth allocation. This means that the Ethernet switches have to be able to manage the wrappings at Layer 2. The Layer 3 wrappings are of no concern to the Ethernet switches, and conversely, the Layer 3 software need not concern itself with getting the data packet across the network.

Table 3 describes the relationship between these Ethernet layers and the media streaming and command and control protocols we've been concerned with for getting our networked audio system up and running. Media transport using Ethernet AVB is a Layer 2 activity. This means that the infrastructure of the network, *i.e.*, the Ethernet switches, must understand the protocols because they are going to take responsibility for routing the streaming media data from source to destination together with all the necessary timing data, and ensuring reserved bandwidth along this route guarantees delivery within the agreed delay time (latency).

→

Table 3: Audio Related Protocols and their Ethernet Layers

TCP/IP Ethernets Model	
Layer 3	Network management, Utility Protocols, etc., and Applications
Layer 2	Standard TCP/IP Layer
Layer 1	Ethernet Layer

Audio Related Protocols		
Control Protocols: OCA, X-170, HiQNet	Layer 2 Media Transport: AVB, CobraNet	Layer 3 Media Transport: Ravenna, Dante, X-192

Proprietary Schemes

Note that, in Table 3, there are some proprietary media transport schemes that operate at Layer 3 (Dante, Ravenna, etc.). These schemes use the preexisting Layer 2 protocols to manage routing and timing without reliance on the new and imminent IEEE standards for Ethernet AVB. This provides the advantage that AVB-compliant Ethernet switches are not required to build a network using these schemes. As at this writing, AVB-compliant Ethernet switches are not shipping extensively; you can't go to your local computer store and buy one.

There's some advantage for the Layer 3-based methods in the short term. The trade-off, however, is that some quite serious care and attention is required to correctly select, set up and configure the network switches and infrastructure. I'm not going to attempt to explain how these schemes achieve reliable connections and timely transfer of media content across the network. That is something for their promoters to explain. However, being equipped with an understanding of how Ethernet layers work and how they relate to media transport schemes and command and control protocols will enable a better understanding of those explanations.

Implementation Details

Another important point in relation to the proprietary schemes being offered to device manufacturers is that a lot of the implementation details have been figured out by the scheme's vendor. Manufacturers must pay for the privilege of using the protocol and its prepackaged implementation; usually a license fee accompanied by a per-unit royalty, and in exchange, they receive a lot of support incorporating the scheme into their products: software and firmware that's ready to go, reference designs for hardware and even hardware modules for purchase.

To adopt a standard such as Ethernet AVB or OCA (should this become a standard as its promoters intend), a manufacturer must invest in its own implementation of the standard and commit to maintaining compliance with the standard as it evolves in the future. Of course, there are non-brand third-party developers working on implementation of the standards approach, and manufacturers can partner with these design houses rather than invest in in-house design resources.

Hybrid Business Model

An interesting hybrid business model is being provided in relation to the X-170 command and control protocol. In this case, a product developer can choose to implement the standard itself or can buy a prepackaged implementation being offered as the commercialized version known as XFN. A similar business model may, I suppose, emerge from the OCA Alliance group. Software and firmware for these protocols can be created in a portable non-platform-specific form such as the C programming language and compiled by a manufacturer to suit its particular hardware processors.

So are we there yet? If "there" is at a point where all manufacturers can provide networked devices that are fully interoperable and can be configured and controlled by a common protocol and hence be managed by a single user interface software, the answer is no. What's even more concerning is that we are seeing the emergence of competing standards to add to the abundance of existing competing proprietary protocols.

It's hard to be optimistic that the convergence required for this much desired interoperability is even in prospect. Ethernet AVB and wide availability of AVB-compliant Ethernet switches certainly will make configuring and setting up networks much simpler and more reliable. That's no small achievement, and definitely something to look forward to.

But it looks like end users will be stuck with competing protocols and their user interface control and monitoring software applications for the foreseeable future. A thorough understanding of what part each of the various commercialized audio networking offerings plays in creating a complete audio system and how to select the best combinations for each system configuration will continue to be necessary. ■