October 2018 Digital Edition

The Next Generation of AV and IT Networks

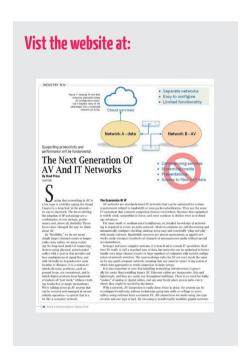By Brad Price, Senior Product Manager, Audinate

Saying that networking in AV is a hot topic, is a bit like saying the Grand Canyon is a deep hole in the ground – no one is surprised. The forces driving the adoption of IP technology are a combination of cost savings, performance and, above all, flexibility. These forces have changed the way we think about AV.

By "flexibility," we do not mean simply larger channel counts or longer cable runs. We mean replacing the long-used model of connecting devices with physical point-to-point cables with a system permitting endless combinations of signal flow and with virtually no dependencies upon location or distance. It is a system in which old noise problems, such as ground loops, are nonexistent and in which digital products from hundreds of makers all "just work" without clocking headaches or magic incantations. We're talking about an AV system that can be secured and managed to ensure reliable operation — a system that is a lot like a computer network.

The economics of IP — designing, installing, maintaining

AV networks are standards-based, IP networks that may be optimized for certain requirements related to bandwidth or even packet prioritization. They use the same IT equipment that connects computing devices everywhere. Because this equipment is widely used, competition is fierce and costs continue to decline as technology advances.

For most small- to medium-sized installations, no detailed knowledge of networking is required to create an audio network. Modern solutions, like Dante, are self-discovering and automatically configure clocking, making setup easy and essentially "plug and play" with simple software. Bandwidth concerns are almost

nonexistent as gigabit networks easily transport hundreds of channels of uncompressed audio without special accommodation.

In larger and more complex systems, it is beneficial to consult with IT specialists. Real time AV traffic is still a standard type of data, but networks can be optimized to better handle very large channel counts or large numbers of endpoints with careful configuration of network switches. The system design rules for AV are very much the same as for any good computer network, meaning that one must be aware of any points at which data aggregates to avoid congestion in large setups.

It is also important to note that installing networking infrastructure is generally far easier than legacy AV. Ethernet cables are inexpensive, thin and lightweight, and are easily run throughout buildings. There is no need for bulky "snakes" of analog or digital cables, and one may freely place access jacks everywhere they may be needed in the future.

With a network, system integration is easily made. Once in place, the system can reconfigure endlessly without getting into walls or ceilings to move cables, and it uses software from a common PC. All connections are made using one type of cable and one type of jack. By choosing to install readily available gigabit networking equipment, an installation is effectively future-proofed for years to come.

Lessons/ideas from IT and security IP networks
For many years, AV systems lacked any security beyond lock and key since the closed nature of legacy technologies required nothing more. IP networks, however, are a different story. Like any computer network, an AV network has many endpoints facing many people and are accessible by many devices, creating opportunities for theft, vandalism and more.

Fortunately, the lessons that IT have learned over the years are available to help the AV community. Practical design choices in networks, such as isolation techniques, can help mitigate unwanted access, as can products that impose a user and privileges hierarchy upon access to system control.

A notable new product in this area is Dante Domain Manager. When used with Dante networks, access is limited to only authorized users on a zone-by-zone basis and complete audits are maintained, ensuring only authorized changes are made to devices and channel subscriptions. This approach closely mirrors traditional IT solutions for security and will expand to include more device types and granularity over time.

How AV will remain unique
To the network equipment, traffic is traffic, and audio and video are simply data types. What makes AV unique is the real-time nature of the data as experienced by human listeners and viewers; interruptions that are perfectly tolerable in other scenarios are simply unacceptable here.

IP networks already comply with all the standards required for accurate, reliable transport of data that is real-time on a human scale. The Precision Time Protocol (IEEE1588) standard has existed since 2002 and permits networks to be time-aligned with great accuracy: one microsecond or less. This, coupled with practical design choices in network topology and switch configuration, are unique attributes of AV networking at scale.

Audio AND Video
Audio networking is very much a force in the AV industry, and has become the leading installation type. With video, things are a little more complicated.

The fundamental difference between audio and video networking is bandwidth consumption. Compared with the capacity of modern gigabit switch, audio presents a very mild load in most cases - 64 channels of uncompressed 48kHz audio requires only about 100Mb/s, or 10% of a single gigabit port. Video, especially when uncompressed, can easily consume far greater amounts of data - an uncompressed 4K stream using 10-bit color requires 9.2Gb/s (!).

This implies that for workflows involving uncompressed video, 10Gb/s or higher, networks are required. While this technology can be purchased today, it is far more exotic and expensive than the ubiquitous 1Gb/s networking equipment we see everywhere. This does not mean, however, that video cannot run over 1GB/s connections. With compression techniques, such as JPEG 2000, high quality post-production video can easily be sent to networked endpoint displays and monitors. Over the next few years, we can expect to see this approach become commonplace, alongside further adoption of very high speed (>10Gb/s) networks for production locations.

Converged vs. segregated — where each architecture is best deployed
In the early days of media networking, only very-low-level protocols were used that prevented the audio system from being able to coexist with other data. That, plus concerns about bandwidth, led to common design choices that segregated audio from all other networking. This practice remains in place today.

Segregated systems are a good idea in some instances. For example, a very large system that carries a huge channel and device count may tax a Gigabit network close to its limits, and thus be a good candidate for "stand-alone" operation. Segregation also aids in designing secure networks, as the audio network has no connection to other

sources of data. Some installers also prefer segregated systems, as it means that they do not have to worry about clashing with an IT department.

But in many cases, segregation is not the most useful option. In settings such as conference spaces, corporate boardrooms, hotels and hospitality centers, data access and audio are frequently required on the same devices and locations. In those instances, only a converged network allows computers and AV devices to access necessary services and resources, and makes the system more coherent and easy for end users to operate.

There is no universal advantage to either approach. Instead, by using the common platform of IP networking, one may tailor solutions to best fit a client or situation.

10G vs. 1G — how each has its application
As mentioned, audio today is well served by commonplace 1GB/s networks. To achieve this same level of continuity with video, far greater bandwidth - and devices that can consume that bandwidth - is required.

10Gb/s networking and higher is still considered somewhat exotic, but exists and is common in areas of very high bandwidth consumption. The backbones of large networks use this technology (and much more) today. Over time, it is likely that the costs of 10Gb/s equipment will come down just as other technologies have, making its way into more everyday products.

In audio, 10Gb/s switch ports are useful for "trunking" data from one switch to another in very large systems (typically >1000 channels). In video, 10Gb/s and higher will be required to transport virtually any uncompressed video and will be useful for streaming multiple sources of compressed video simultaneously to many endpoints. At present, a 1GB/s network can only handle a small number of high quality compressed video streams.

 Moving up in the stack —security, permissions and more
Security is a constant concern for IT administrators. A robust industry and toolset has grown around these concerns. For AV, the good news is that much of the legwork has already been accomplished.

Much practical AV network security is still done via obfuscation and limited physical access – if networks are kept segregated, this is a logical approach. But in larger systems that are part of broader networks - like in conference rooms and corporate spaces in office buildings, or even a large stadium - the need for a coherent, user-based method becomes clear, since this is precisely how IT departments manage access.

Dante Domain Manager is a first-of-its-kind product and is a user-driven management system for Dante audio networks. With Dante Domain Manager, users log in using credentials created for the Dante network, or with common methods

such as LDAP and Active Directory. Once logged in, users may only interact with the audio network according to the permissions set by the administrator. Different people may have different levels of access to different areas of the network, and devices are kept locked within defined groups, or "domains," so that new devices cannot be added without permission.

Because Dante is a multi-vendor solution, Dante Domain Manager works with any brand of audio product using Dante audio networking, enabling levels of security that would otherwise be very difficult to achieve in a mixed-manufacturer system.

Scaling the network

Real-time media relies on technologies originally designed to work in smaller networks. The key time standard, IEEE1588, is specified to work only within what is called a subnet, or broadcast domain, in networking. This is due to the required use of multicast messages (i.e., those sent to simultaneously to specified members of the network) for time alignment between devices. Until recently, this requirement has limited most practical systems to small sizes from a networking perspective: around 250 devices, which is still large for AV.

To scale beyond this, there must be time alignment between subnets. Dante Domain Manager does exactly that, automatically configuring clocks within subnets and linking them to maintain the same levels of performance as before, but now at a much larger potential size. With Dante Domain Manager, AV networks can scale to thousands of devices and continue to be managed from a single point of control.

For systems without Dante Domain Manager, single networks are limited to a few hundred devices each. If these systems do not need to communicate with one another, multiple "stand alone" AV networks can be deployed to achieve larger scale using VLANs to contain each one. While this does limit some functionality, it is often a practical solution.

IP networking is here to stay as a primary means of connecting the myriad devices we use for audio, video, and nearly everything else. Advancements in real-time media are making it easier to deploy and manage IP-based AV securely and at very large scale, adopting methods and techniques from the world of IT management.

The benefits of networked AV remain as promising as ever: easier deployment, near infinite flexibility and flawless performance. As the industry grows, these promises are being delivered on a larger scale and are changing the way audio and video are getting done.