

ICT TODAY

THE OFFICIAL TRADE JOURNAL OF BICSI

July/August/September 2020

Volume 41, Number 3

The background of the cover is a blurred photograph of a modern conference room. In the foreground, a black gooseneck microphone with a perforated grille is positioned diagonally. Behind it, a silver laptop is open on a table. The background shows other tables and chairs, with soft, out-of-focus lights from the ceiling.

AV OVER IP AND COMPLEX ENTERPRISE NETWORKS

PLUS:

- + Simple AV Control Concepts
- + Overcoming Challenges at Carnegie Mellon University
- + Wireless Communication: The First Line of Defense for Schools and First Responders

AV OVER IP AND COMPLEX ENTERPRISE NETWORKS



Distribution of audio and video signals over data networks, also known as AV over IP (AVoIP) or networked AV, potentially offers significant benefits for pro AV systems integrators and end users. Compared to traditional systems for AV distribution, networked AV systems can be designed with virtually unrestricted scalability and flexibility, in addition to the convenience and cost efficiency of standard data networks. The benefits of AVoIP can especially be realized in large installations for enterprises and other organizations.

The commercial AV industry is in the midst of a challenging transition with an emphasis on technologies that move AV operations to the network. Enterprise businesses, schools and universities, and entertainment venues that have long relied on purpose-built matrixed systems to route and distribute AV signals are now looking to converge these same essential functions within new or existing IT infrastructure.

However, the migration of AV systems to the IP network requires a great deal of planning. Detailed knowledge of network topologies, deployment issues, best practices, and quality-related considerations ensure a smooth transition to AVoIP and a reliable and consistent AV system moving forward. To achieve this, it is important to understand what is required to establish an AVoIP network from infrastructure through network architecture and protocols.

INFRASTRUCTURE

1 GbE Versus 10 GbE—The Debate

There is an ongoing debate in the industry about whether a 1 Gb/s or 10 Gb/s architecture provides the ideal network topology for AV systems. The discussion often gets mixed between infrastructure and AVoIP applications. For the cabling infrastructure, cables that can carry 10 Gb/s signals are advised, especially for new structures and buildings to future proof the networks either inside the data center or from remote customer premises equipment (CPE).

When the structure is wired with 10 Gb/s Cat 6A-capable cabling, an AVoIP application can be addressed. What type of AVoIP solutions should be used to move video and audio on the network? To answer this question, it is vital to understand the benefits of using 1 GbE versus 10 GbE AVoIP solutions.

1 GbE—The Incumbent

Installing and operating AVoIP encoders and decoders that are 1 GbE capable means that network devices (e.g., L2/L3 switchers) do not need to support 10 GbE; they can be limited to 1 GbE. At this point, the network becomes 1 GbE, and this solution delivers many features and benefits that will keep it viable for years to come. One benefit

is the cost of the network; 1 GbE network switches are far more cost-friendly and offer a capital expense advantage for companies. There are many 1 GbE network switches available. They provide the buyer with a good selection and the ability to choose one that fits both the budget and network specification requirements.

Another strong advantage of AVoIP network devices that run video traffic at 1 Gb/s over a 1 GbE network is the ability to run the signal over a longer distance. Data traffic that runs at a rate of 10 Gb/s over a 10 GbE network is more susceptible to electromagnetic interference (EMI) and radio frequency (RF) noise and runs at shorter distances. This can be a strong limiting factor when trying to run high-resolution video, such as ultra-high-definition (UHD) 4K/UHD at 60 frames-per-second.

When copper cabling is used as the cabling infrastructure, 1GbE network infrastructure devices benefit from the availability, cost effectiveness and maturity of Power over Ethernet (PoE) technology. Power over Ethernet can save costs and labor and expedite project completion. If PoE is not supported, every AVoIP device requires its own local power source that increases the complexity of the installation, thereby requiring specialized labor and potentially impacting installation times. To some degree, PoE can be used as a device management tool for AVoIP by enabling remote power cycles and power monitoring for edge devices.

The main disadvantage of 1 GbE networks is that they are bit rate constrained. A 1 GbE network actually runs a net bit rate (removing IP headers and other layers of signaling and protocols, and using regular frames) of around 900 Mb/s. This represents the amount of bit rate available to work with versus the approximate 9000 Mb/s in 10 GbE networks. This is not always a limitation but is something to be aware of when weighing the pros and cons of one versus the other.

10 GbE – The Need for Speed

High-resolution video consumes large amounts of data. For example, high-definition (HD) 1920 x 1080 resolution images running at 60 frames-per-second can consume up to 4.46 Gb/s while UHD video at 3840 x 2160 resolution at 60 frames-per-second can consume up to 17.82 Gb/s. A 10 GbE network can pass HD video without video

compression. Unlike 1 GbE-based AVoIP systems where video is always compressed in order to fit into the narrow bandwidth of 900 Mb/s, a 10 GbE AVoIP-based system often does not require compression in the case of HD video. When compression is required for UHD, the amount of compression required is very low, which is the most significant advantage of 10 GbE AVoIP systems; the amount of compression applied on video is minimal, if required at all.

Simply put, when comparing the video quality of the streamed video with uncompressed video, it is bit-to-bit identical to the source. What the source produces is exactly what the display will show. Yes, some compression is needed for UHD signals that can climb to 17.82 Gb/s, but the compression is very light versus compressing the same signal onto a 1 GbE network.

This 10 GbE advantage may steer designers and installers away from 1 GbE networks. After all, who wants low video quality images running in meeting spaces or lecture/study halls? However, advances in compression algorithms and increases in computing power have resulted in impressive video quality to a degree of visually lossless even when compressed to work in a 1 GbE pipeline. In other words, today's AVoIP encoders and decoders use advanced compression technology that makes it difficult to distinguish between the source and the compressed image when placing them side-by-side. For AV purposes, this means that a 1 GbE network can be sufficient.

A 10 GbE AV network requires little to no compression but uses the entire 10 Gb/s bandwidth provided by the network. The same 10 GbE network can accommodate 1 GbE AVoIP alongside existing enterprise network traffic with ease, thereby eliminating the need for a dedicated 10 GbE AV network.

NETWORK DESIGN

Topologies

A network topology defines how the various components of a network are arranged, including the various devices that the AVoIP system is connected to and how those devices are connected to each other.

Although there are a variety of different network topologies, AVoIP encoder and decoder devices will nearly

always be connected to one of three groups: single logical switch, star or fat tree. Each of these topologies has its own distinct advantages and disadvantages:

Single Logical Switch

By far, the simplest network topology is one that uses a single switch (Figure 1). In this type of system, each device is connected to a single logical switch that is not connected to any other switches or routers. Note that this type of topology could refer to a single physical switch or to multiple physical switches that are stacked and act as one logical switch.

Advantages: Simple, easy to manage, limited network impact of video traffic, inexpensive.

Disadvantages: Limited scalability options, no network redundancy, single point of failure.




FIGURE 1: Illustration of the single logical switch.

Star

For systems needing more scalability than that offered by a single logical switch, multiple logical switches can be each connected to a single switch for intercommunication. In this star topology (Figure 2), the central switch is often referred to as a “core” switch; the other switches are referred to as “access” switches. This type of system will typically have AVoIP encoders and decoders connected to the access switches.

Advantages: Scalability, multiple access switches can create network path redundancy.

Disadvantages: Physical network layout can make cabling very expensive, since all access switches must be directly connected to the core; single core switch creates a single point of failure.

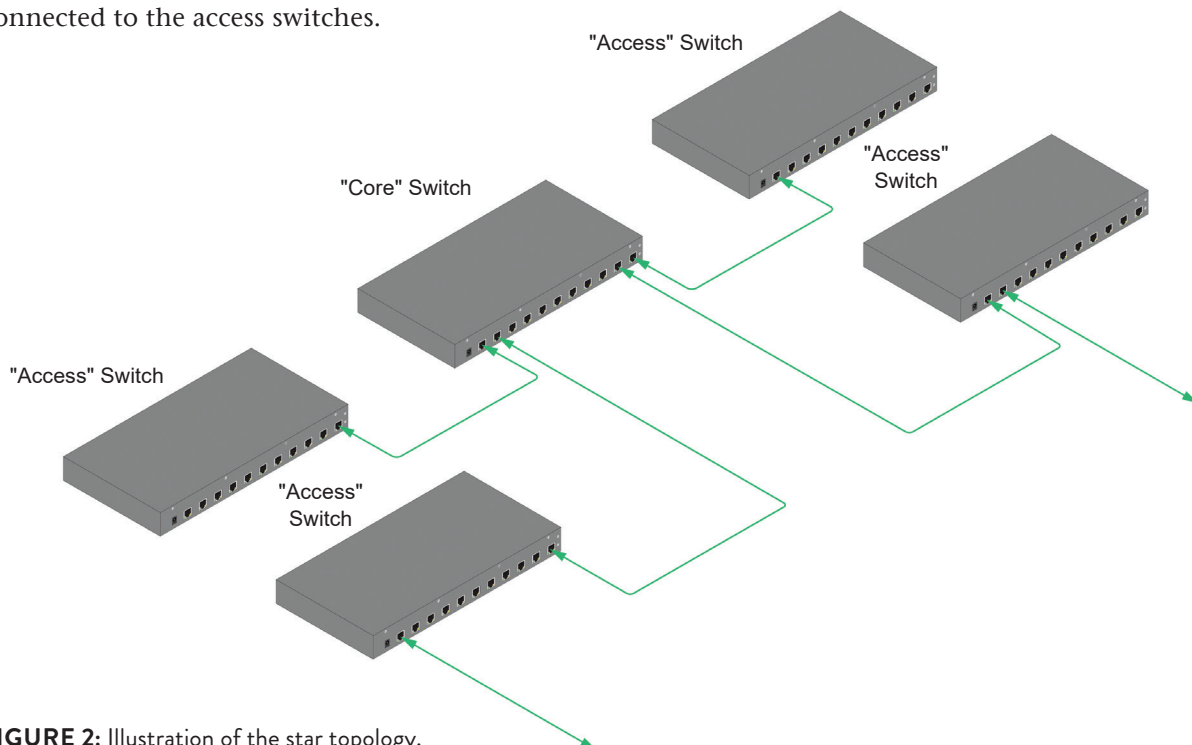


FIGURE 2: Illustration of the star topology.

Fat Tree

Very large network systems often fall into the category of a fat tree topology (Figure 3). Here, access switches are connected to a distribution switch in a way that resembles a star network. Multiple distribution switches are connected to a core switch. There can be multiple layers of distribution switches, and multiple core switches can be interconnected. Links further up the tree (closer to core switches) are “fatter” or higher bandwidth than links further down the tree (closer to access switches).

Advantages: Ultimate scalability, potential for high network path redundancy (protections when specific path is down due to switch failure).

Disadvantages: Management complexity, cost.

Network Segregation

There are essentially two ways to segregate IP networks. The first method is physical segregation. Referred to as an “air gap,” physical segregation physically isolates one network from the primary network infrastructure headend. That headend comprises multiple racks of network

switches, all of which belong to various forms of enterprise online and IP traffic.

Adjacent to those racks is the isolated network where the AV traffic is managed. The separation ensures that the AV and general traffic systems do not interfere with each other. There is no need to carefully prioritize one type of traffic over other IP traffic types. Since these AV network systems are separated from other networks in the organization, the complexity level of the network integration is reduced. On the other hand, network management can be trickier since there is a need to be able to access two different networks that are not necessarily talking to each other.

The second method uses a more logical approach. By utilizing virtual LAN (VLAN) or multi-protocol label switching (MPLS), AV/IT administrators can create virtual segregated managed networks that separate the high bandwidth AVoIP traffic from the rest of the IP traffic that flows in the networks. This method is widely used in cases when there is a preference to utilize already installed and operational network gear rather than buying a separate rack to segregate the IP traffic.

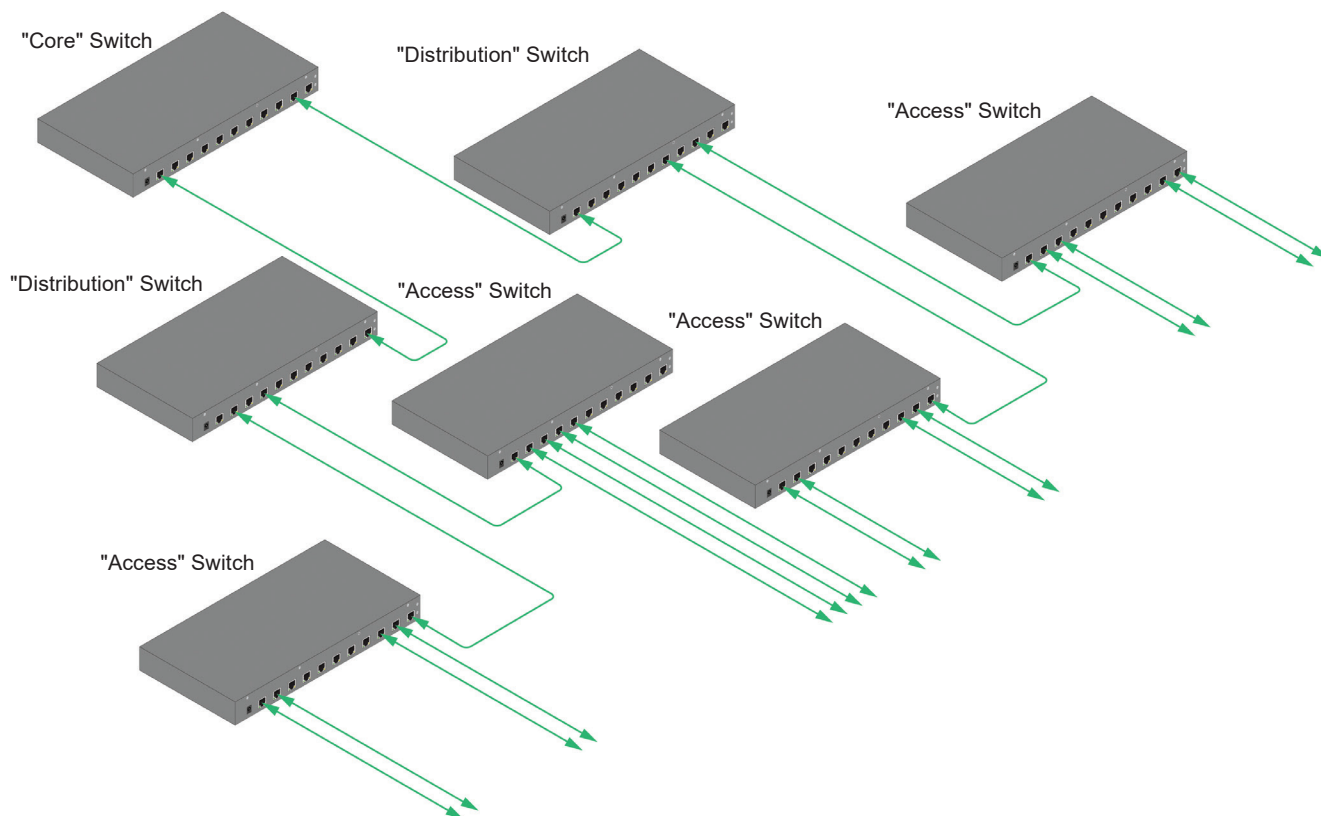


FIGURE 3: Illustration of fat tree topology.

Switch Fabric

No matter which of the network segregation methods is chosen, the switch fabric's non-blocking bandwidth specification is a key consideration. The network switches must have enough switch fabric bandwidth to support full non-blocking bidirectional gigabit bandwidth on all ports simultaneously. Information technology and AV administrators need to be mindful that when working with network switches that have downlink ports in the switch (i.e., ports that are connected to AV over IP encoders and decoders), physical line bandwidth may specify 1 GbE ports. However, the switch capacity to handle IP traffic (i.e., the actual bit and bytes of data) is lower than total downlink ports multiplied by connected active 1 GbE AVoIP devices. This creates a blocking architecture.

Another consideration is the supported uplink bandwidth (i.e., from network switch to other network switch). Typical network switches are interconnected using dedicated or trunk ports with typical bandwidth between 10 Gb/s to 40 Gb/s. To ensure proper network operation, it is recommended to calculate how many simultaneous AVoIP streams will pass the trunk port. For example,

if the maximum bit rate of a given network trunk port is limited to 10 Gb/s, approximately 10 AVoIP streams of 1 Gb/s can pass the trunk port to the other switch. In the same scenario, if AVoIP that requires 10 Gb/s is used, the system will be able to move only one stream between the network switches, resulting in scalability and network bottleneck issues.

Network Protocols

IP networks are run and operated by protocols. Different protocols are applied to discover routing paths, move data, manage security, and many various applications and services. Therefore, it is important to understand the various types of data transmission:

Unicast—A unicast transmission means one host to a single host (Figure 4). If a packet has a specific designation address that is unique to one host, the transmission is unicast.

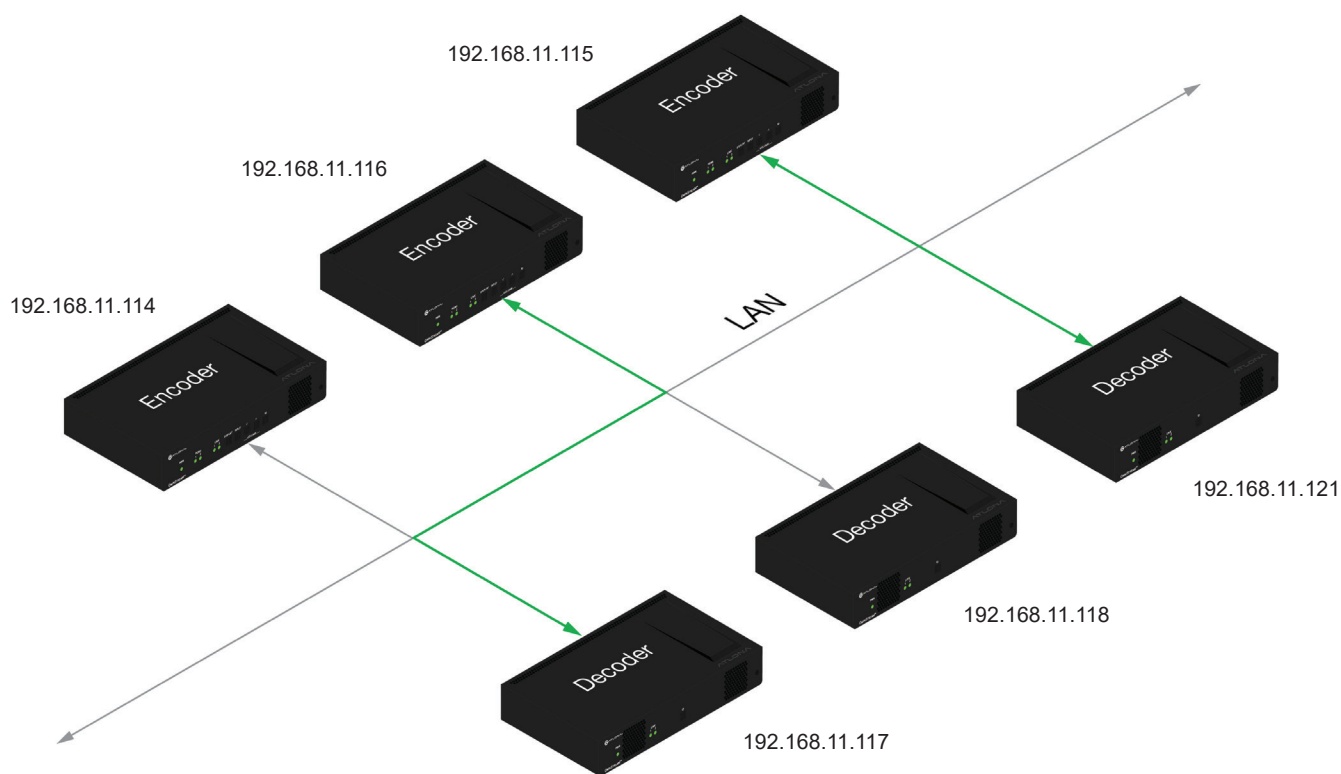


FIGURE 4: Diagram of unicast transmission.

For networks that run AVoIP services in conjunction with other types of data, it is highly advised to prioritize AVoIP-related protocols (i.e., multicast and IGMP) over non-delay critical data.

Multicast—A multicast is described as a single transmission host to a group of receiving hosts. In this transport method, receiving hosts subscribe to a multicast group, and any receivers subscribed to that group receive the information forwarded to that multicast group address. The protocol that is used by receivers to subscribe to the multicast group is called internet group management protocol (IGMP).

Audiovisual over IP relies on multicast functionality to transmit and receive video and audio. Multicast provides efficiency, which is very important when moving high bit rate video over IP networks. Multicast eliminates the need to send the same high bit rate stream individu-

ally to each receiver, consequently reducing the load on the network and encoder devices (Figure 5).

Traffic Management

Most data flowing over the LAN and wide area network (WAN) connections, where the internet resides, use connection-oriented protocol layers. Packets sent have sequence numbers, and there is an acknowledgment mechanism in place that ensures packets are retransmitted if there is acknowledgement of no receipt on the receiver end. Thus, standard internet traffic, such as email or web browsing, will generally continue unabated even when quality issues (e.g., packet arrival order, packet loss) arise.

If a packet does not arrive, the user device—with the help of buffering algorithms, acknowledgments, and retransmissions—can still receive the entire data. These retransmissions and acknowledgments sound like mechanisms that are applicable for every type of data traffic. The catch is that latency and delay make these approaches to data transmission ineffective for real-time video and AVoIP.

Low-latency, high-fidelity AVoIP networks are “connectionless,” meaning that there is no feedback coming from the receiving end to indicate packet loss. This makes lower latency streaming over the network




FIGURE 5: Illustration of multicast transmission.

possible, but a genuine challenge arises collectively. What happens when packets get lost or arrive in the wrong order? What mechanism can be used to still benefit from the connectionless protocol layer and protect traffic, as much as possible, from arrival order issues and packet loss? This is where forward error correction (FEC) technology becomes critical.

Before diving into FEC and how it is used, it is first necessary to understand what causes the packet drop in the first place.

IP data packets are susceptible to packet drops, packet losses and packet arrival order misalignment as the facility's network scales, more network nodes are added, and traffic load increases. High-density, complex network topologies with multiple nodes and routes can yield three main packet drop scenarios:

- In scenario one, the network switch may drop packets due to buffering issues. Switches often build up a packet queue, which are waiting to be forwarded or delivered to the next hop. In some scenarios, these packets spend too much time in the queue and the buffers start to fill, especially when the bandwidth used is high and traffic rates peak. The switch will then start to drop packets, disregarding sensitivities such as latency.
- In the second scenario, each IP datagram packet carries a checksum value that is used to verify and detect errors that may have been introduced during data transmission. Checksum bit errors are translated into packet losses or packet drops by the receiver side.
- In the third scenario, packets may drop due to a divergent routing path. This results in rerouting and packets arriving in the wrong order at the destination. Highly complex networks will offer different routes to move signals to destinations, just as drivers can take the highway or back roads as they drive around their neighborhoods.

Error Correction

To fix errors and loss packets, the AV/IT administrator or system integrator needs to measure error rates. Software

tools that measure error rates are available, as are dedicated hardware appliances that can run traffic in and out and measure packet loss statistics over time.

Once the error rate and level are identified, the next step is to activate and configure the FEC mechanism.

Audiovisual, IT administrators and systems integrators are strongly advised to choose AVoIP encoder and decoder devices that support FEC mechanisms. An FEC-enabled encoder places video packets in an internal virtual table with L columns and D rows. The packets are aligned in columns and rows with a mathematical XOR calculation for each row and for each column.

The video packets, noted as 1-9 in Figure 6, together with the FEC packets, are sent over the network to the decoder.

In case of packet loss, the reverse XOR operation (XNOR) is used to calculate back the missing packet. As shown in Figure 7, packet 5 was dropped due to network issues. By using the XNOR operation, the decoder can restore the missing packet and recover the data without the need to retransmit the data again.



FIGURE 6: FEC diagram.




FIGURE 7: With XNOR, packet 5 that was dropped can be restored and the data recovered.

In cases where the network switch or switches are not loaded and are properly managed, the majority of the latency budget will be consumed by the AVoIP network devices.

The size of the FEC matrix can vary based on field installations and lab experiments. If burst errors are expected in the network, it is advised to have more columns. When possible, that column should be at least the size of the burst. Row FEC packets are good for protecting against random packet loss patterns.

It is important to note that activating and using FEC will result in generating additional data packets that are not part of the audio or video data stream. The FEC packets are additional traffic that will consume bit rate, and this needs to be considered when planning and allocating bit rate for video and audio (Table 1).

Quality of Service – Differentiates Services

Not all data types and packets are equal. There are data packets that need immediate forwarding and routing, and

there are other types of data packets that can observe a bit of delay before getting processed. A quality of service (QoS) policy must be implemented by the organization that runs the IP network.

The term “quality of service” can be thought of similarly to how an airline passenger is classified: first class, business class, or economy class. Through a process called packet marking, a value is assigned to the packet to determine how that packet should be treated as it travels over the network. When a passenger boards a plane, flight attendants are not concerned with how a passenger was classified. However, by looking at the boarding pass, the attendant can then provide the passenger with that assigned level of service.

The AVoIP traffic is delay-sensitive. As such, the traffic must be prioritized when arriving to a network switch. In the networking world, standardized protocols are used to implement policies. One of the protocols that is used to configure the network elements to classify traffic is differentiated service (DiffServ). In a nutshell, each data type will be assigned a differentiated service code point (DSCP) for classification. The DSCP value is assigned to each packet at the encoder, and it is mapped against five main precedence values: best effort, low, normal, high and highest.

For networks that run AVoIP services in conjunction with other types of data, it is highly advised to prioritize AVoIP-related protocols (i.e., multicast and IGMP) over non-delay critical data. For highly congested network

	Configured Bit Rate	Used for Video	Used for FEC
FEC disabled	900 Mb/s	900 Mb/s	0 Mb/s
FEC enabled, 4x4	900 Mb/s	600 Mb/s	300 Mb/s
FEC enabled, 10x10	900 Mb/s	750 Mb/s	150 Mb/s
FEC enabled, 20x20	900 Mb/s	818 Mb/s	82 Mb/s
FEC enabled, 4x4	450 Mb/s	300 Mb/s	150 Mb/s
FEC enabled, 10x10	450 Mb/s	375 Mb/s	75 Mb/s
FEC enabled, 20x20	450 Mb/s	409 Mb/s	41 Mb/s

TABLE 1: Bit rate chart for FEC, audio and video.

The image shows a configuration interface for an AV over IP encoder. It includes fields for Encoder (vc2_encoder1), Enable (a toggle switch), Destination IP address (225.0.0.7), Destination UDP port (1000), TTL (255), DSCP (Best effort), FEC enable (a toggle switch), FEC rows (15), and FEC columns (15). The DSCP field is highlighted with a red box.

FIGURE 8: Image taken from AV over IP encoder with QoS DSCP configuration set to Best effort.

systems, it is recommended to set the value to high. In cases when the network system is not congested, the DSCP value can be set to best effort (Figure 8).

Latency

Audiovisual over IP is used to deliver content and information to the audience that may be located in the same room as the content source or in other places within the enterprise. In cases when the audience is in the same room, such as a study hall, meeting space or conference room, the tolerance for delay or latency is very low. On the other hand, when the content source and audience

are in different locations, the tolerance for latency is higher and less strict.

Audiovisual over IP systems will introduce latency, mainly because of the video and audio processing that takes place within the devices. Some AVoIP devices will consume more time to process video and audio and others will require less. It is imperative to ensure that when AVoIP is installed, the minimum latency available is introduced.

Based on customer experiments and industry best practices, Table 2 can be used to better understand the tolerable latency map against use cases.

In cases where the network switch or switches are not loaded and are properly managed, the majority of the latency budget will be consumed by the AVoIP network devices. It is recommended that systems integrators who install AVoIP introduce less than 10 ms latency when the encoder and decoder are connected back-to-back.

Security

Network security is among the top priorities for an AVoIP installation or network device. Devices require security protocols around authentication and authorization. There are instances where rogue machines are plugged into a university network, for example, disguised as a legitimate element of the network. Proper authentication and authorization protocols prevent unlawful monitoring of traffic and theft of content and information.

Use Case	Description	Maximum Latency
Small conference rooms	Soft codec, with no sound amplification, ability to switch between several sources	< 18 ms
Large conference rooms	Presenter's voice is amplified (ceiling or on-wall speakers), audio DSP processing	< 18-25 ms
Auditorium and image magnifications	Video of presenter is projected to improve viewing ability, audio amplification is used	< 20 ms
Room overflow with interactive Q&A	Presenter in one room, overflow audience in other room with ability to interact (Q&A)	< 35 ms
Room overflow with no interactive Q&A	Presenter in one room, overflow audience in other room with no ability to interact (view only)	2-4 sec

TABLE 2: Tolerable latency map per use cases.

Security protocols, such as IEEE 802.1x, are ideal for implementation on network devices. Utilizing the IEEE 802.1x protocol on AVoIP networks enables the provision (e.g., assigning IP address, sending configuration) to take place only to authenticated devices and prevents rogue elements from disguising themselves as an encoder or decoder to gain access. Management protocols from control and network monitoring devices must also be secured for the same reasons.

All devices on the AVoIP network must also be immune to outside attacks and be able to fend off rogue attempts to access the network through device usernames and passwords. These attacks typically happen swiftly, and proper security protocols provide the appropriate level of access control and network protection.

Audiovisual over IP devices are tasked to stream and receive video and audio media over IP networks. Unlike other distribution technologies (e.g., circuit-based distribution), IP packets are susceptible to being hacked. Adding a layer of encryption to the content protects the service from being hacked and ensures that only devices able to decipher the content can decode and display the content. One leading technology that is used for this is the advanced encryption standard (AES), and it is advised to use AVoIP devices that support at least 128-bit AES.

SUMMARY

The migration to AVoIP in commercial AV is a path that several industries, including ICT, have already taken to great success. Internet protocol networks create enormous value given the highly flexible and scalable nature of the network architecture as conveyed in this article’s summary in Table 3. No longer are systems integrators and end customers limited by the amount of inputs and outputs per box to manage signal distribution. Audiovisual over IP also helps to overcome the traditional signal distribution limits of matrixed systems, and in medium-to-large AV systems, it will reduce capital outlay for the end customer.

AUTHOR BIOGRAPHY: David Shamir is a 20-year electronics and computer industry veteran focusing on streaming and video delivery in broadcast and audiovisual markets. He is responsible for many industry developments in encoder, decoder and transcoder technologies. Currently, he is director of product management at Atlona where he oversees the transition from point-to-point, circuit-based switching platforms to IP-based signal distribution for AV and control systems. David can be reached at david.shamir@atlona.com.

Recommendation		Reasoning
Cable Infrastructure	Category 6A cabling	Immunity to noise, ability to move high resolution to longer distance
Network Infrastructure	1 GbE network, including AVoIP devices	Lower Capex, switch selection, Power over Ethernet (PoE) maturity and availability, non-blocking switch fabric
Network Segregation	AVoIP network is separated using logical separation: VLAN, MPLS or air gap network (separate switch)	Improve management, lower risk of congestion
Protocols	Multicast, IGMP with QoS in case of congested networks	Reduce bit rate load, improve scalability
Latency	AVoIP encoder–decoder latency needs to be below 10 ms	Enable installation in conference rooms, auditoriums and room overflow applications
Security	Authentication and authorization using IEEE 802.1x protocol, content protection using AES-128 bit, and ability to change username and password of AVoIP devices	Improve immunity against rogue elements, and prevent unlawful monitoring of traffic and theft of content

TABLE 3: Summary and conclusions table.