THE BROADCAST BRIDGE

Connecting IT to Broadcast

Audio Over IP - Making It Work

Essential Guide

EG

ESSENTIAL GUIDES

Introduction from Thomas Riedel

Thomas Riedel, CEO Riedel Communications.

Hello, and thank you for downloading this Essential Guide, covering the world of Audio Over IP.

Riedel and The Broadcast Bridge strongly feel that there is a lot of misinformation out there concerning IP and how it is affecting broadcast workflows. Therefore, we decided to take the initiative to enlist one of the industry's top technical writers, Tony Orme, to help lead us down the path to greater understanding.

Over the coming months, we will take you on a 12-part journey that will not only define IP terms and spell out what they mean for you, but also empower you to begin your own transition towards realizing the promises of broadcast IP infrastructures.

Riedel has been very active in IP for several years, including membership in the standards organizations, participation in interoperability events and plugfests, and innovation with several of our products. We hope to share our experiences with you so that your transition is as smooth as possible.

I hope you enjoy this Essential Guide to Broadcast IP Infrastructures and that it becomes the kind of document that you can refer back to again and again. Please share it with your friends and please let us know how we're doing by sending an email to marketing@riedel.net.

Best regards,

Thomas Riedel

Audio Over IP - Making It Work

By Tony Orme, Technology Editor at The Broadcast Bridge

Sponsored by Riedel Communications

Introduction

Building and operating IP networks is much more than just about saving money on infrastructure costs. Its success is deeply rooted in the ease of flexibility, scalability, and inter-connectivity that it can provide. And one of the greatest benefits of IP is that the protocol and underlying hardware is independent of the data being carried, therefore, distributing and interfacing between different formats is easier than ever.

Traditional broadcast systems consisting of SDI, MADI, or AES work perfectly in isolation but become incredibly complex when signals need to be routed between them. De-embedding audio from the SDI feed to multiplex it into the MADI stream, or even distribute down an AES cable, increases cost and specialization. And differing sampling formats provide even greater challenges as we must configure each piece of equipment to the studio standard before transferring across the node.

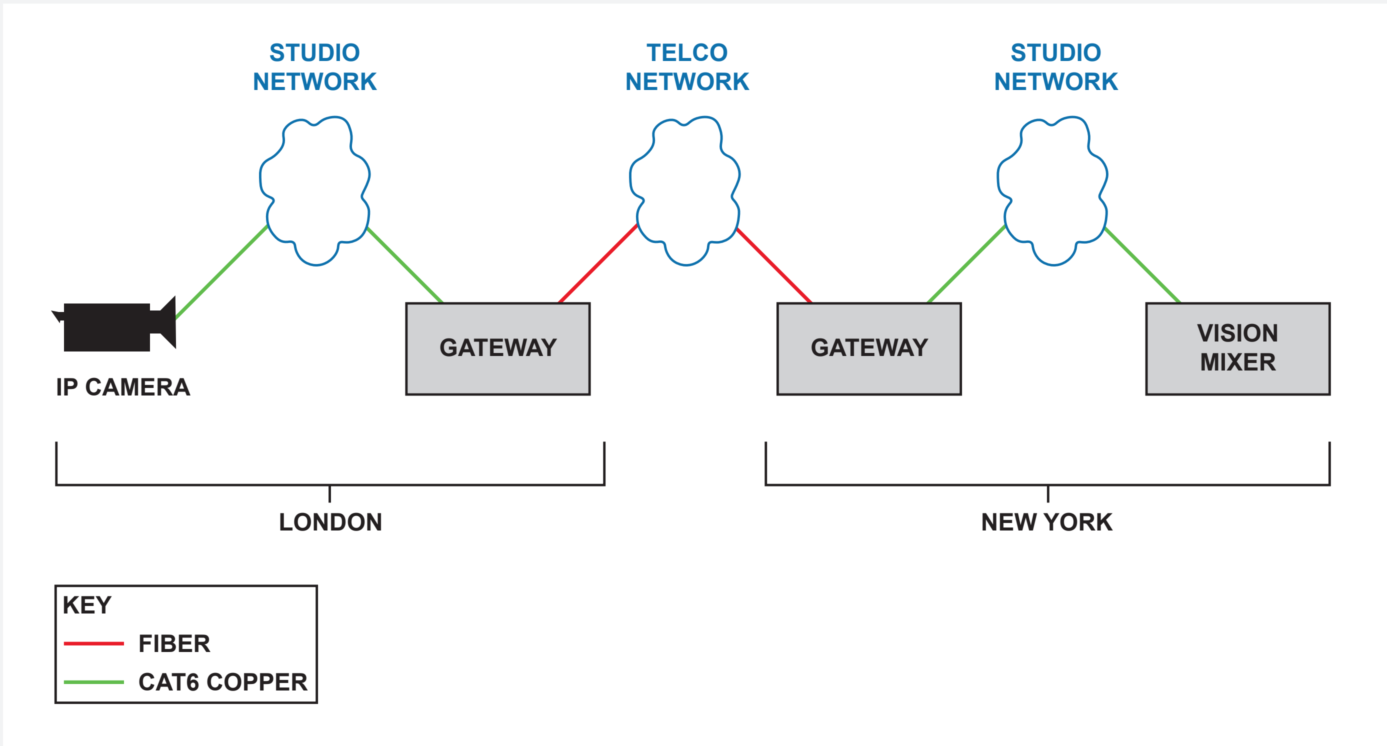

Diagram 1 – IP moves independently of the underlying hardware and communication channels.

Monitoring presents its own challenges since we would need systems with many different interfaces to access the whole of the signal chain. Sometimes, not all interfaces are available on one monitoring solution so two or more would need to be installed resulting in increased costs, complexity, and valuable real-estate in the studio or outside broadcast vehicle.

Rigid Systems Lack Scalability

Broadcast systems have stood the test of time and continue to deliver reliable high-quality audio. However, the price we pay is lack of flexibility. Rigid systems make scalability difficult and adopting new formats extremely challenging.

Moving to IP solves many of these challenges. Audio signals of different formats and sample rates can be distributed over the same network, as the physical interfaces needed for interconnectivity consist of standard "off the shelf" types often found in the IT industry.

IT Continues to Innovate

Fiber infrastructures deliver unparalleled bandwidths with modern routers increasing from 10Gbps to 100Gbps in recent times, and 400Gbps and more in the near future. Riding on the back of IT research and development, broadcasters can take advantage of the new speeds and innovation, with the certainty that there is always downward pressure on price.

Interfacing to telecommunication companies (telco's) in broadcasting has long been a painful and expensive task, requiring specialized interfaces due to the limited adoption of SDI, MADI and AES in wider industry, thus resulting in greater complexity and increase in costs.

IP is Ubiquitous

IP is now the go-to interface for telco providers and is ubiquitous in the IT industry. An abundance of data circuits championed by IT has reduced prices and increased bandwidth and availability. And moving to IP systems allows broadcasters to hook into telco's with ease and reduced costs, taking advantage of the downward pressure on prices demanded by IT customers.

One of IP's greatest strengths is that the datagrams it sends are independent of the underlying hardware. IP is equally happy to send its data over ethernet, as it is over frame relay or Wi-Fi. And as a user, we do not know or care about the physical layer of connectivity between different nodes.

To make IP systems easier to understand and interface to, the Open Systems Interconnection (OSI) model breaks network operations into seven layers to give an abstracted view and demarcation of the structure of networks. A User Datagram Protocol (UDP) will work on any IP system regardless of the source - assuming the manufacturer has complied with the specification for IP.

Ethernet is Pervasive

Ethernet networks are the most generic form of distribution for data streams within a broadcast facility. Speeds continue to increase as switch manufacturers plough more money into research and development. Data rates of 10Gbps are easily affordable with 25Gbps, 40Gbps, and 100Gbps in mainstream use. Simplistically, one 10Gbps ethernet or fiber channel can carry three, 3G-SDI services.

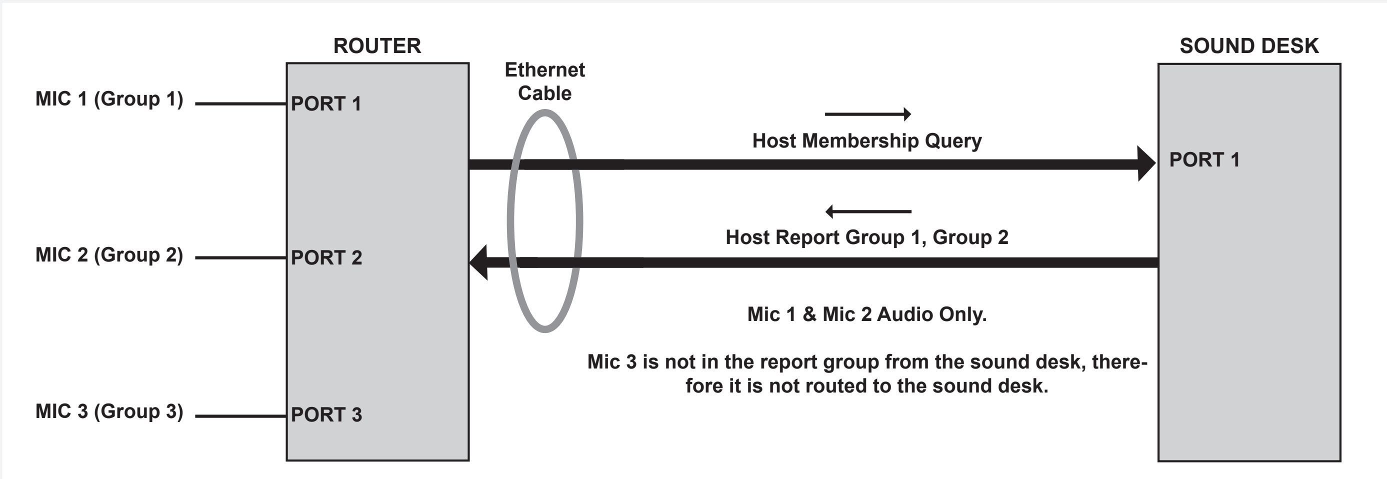

Diagram 2 - IGMP Host Membership Query Messaging.

Routing takes place at both layer-2 and layer-3 levels. Layer-2 refers to ethernet switching and layer-3 to IP routing. Layer-2 switching is faster than layer-3 as only one level of datagram needs to be decoded. If layer-3 routing is required, then both the ethernet and IP headers must be decoded, thus increasing the processing time in the switch and reducing the speed of routing.

Domain Routing

Although we must set the source and destination addresses within each device, such as the camera, microphone, or monitor stack, the internal network will only switch at level-2. It's only when an IP datagram is routed outside of the network, or domain, that IP routing occurs.

IP datagrams distributed within a studio tend to travel over a common medium such as ethernet's copper CAT6, but when the signal leaves the station it may well be converted to fiber. The underlying physical medium and its associated addressing changes, but the IP source and destination addresses stay the same.

To make multiple copies of microphone outputs in a traditional workflow, we would use a distribution amplifier with each output going to a unique destination. In effect, we've provided "one to many" mapping.

In IP the same concept is possible but, instead of using distribution amplifiers, we use multicasting. Multicasting is a protocol used in switches, routers, and end devices such as cameras, microphones, and monitors.

Multicasting uses the principle of destination opt-in or opt-out. Using the Internet Group Management Protocol (IGMP), a destination device communicates with a router to request a copy of a multicast data stream. If the multicast stream is available on that router, the stream will be sent to the requesting port. If it is not, then the IGMP request is sent further upstream until the source device is found.

Switch Duplicates Datagrams

If a microphones output is sent to the sound console and monitor station, both these devices will send an IGMP message to the switch to request copies of the stream. The switch will duplicate the multicast stream and output it to the requesting port.

The alternative to multicasting is for the microphone to be manually configured to send its datagrams to each requesting device. In this case there will need to be two distinct streams of data leaving the microphone, one for the sound console and one for the monitor station. In doing so, we've already doubled the data rate, and adding another destination such as an audio recorder will triple the output. With multichannel sound recording, this method soon becomes unsustainable.

SDN's Overcome Limitations

Multicasting is an efficient method of distribution as the switches duplicate the datagrams, resulting in the load being taken off the microphone, optimizing its IP output. However, IGMP relies on configuration information being typed into the routers resulting in complex network administration. To overcome these limitations, broadcasting is looking to Software Defined Network (SDN) techniques.

The flexibility and scalability offered by IP is unprecedented. We're not just saving money by using commercial-off-the-shelf products, but also by providing a hardware interface that is independent of the data travelling on it. The result is future-proof broadcast operations.

To Switch or Route?

To fully leverage the benefits of IP networks we need to think in IT terms. Just replacing the acronym MADI or AES with IP is insufficient as all we end up with is a very complex, poorly utilized, static network.

Point to point connectivity has provided broadcasters with guaranteed bandwidths and exceptionally low latencies. With their telephony background, Telco's have provided analogue and digital connectivity to route audio and video all over the world. The price is increased complexity and cost.

Ironically, Telco's have been distributing broadcast signals using IP networks for many years without us being fully aware. Gateways, at the interface to broadcasters, provided the necessary format translation from SDI, AES, or analogue, to IP and hid the underlying network from us.

Understanding how Layer-2 switching and Layer-3 routing operates, and how it relates to the Open Systems Interconnect (OSI) model is key to understanding computer networks and their implications for latency and timing.

Confused Terminology

It doesn't help that broadcast engineers refer to signal switching devices as routers, such as an SDI router or audio router. The terminology becomes even more challenged as Layer-2 switches sometimes incorporate Layer-3 routing functionality. These switches are often referred to as multilayer switches.

Networks consist of two geographical zones – LAN's (Local Area Networks) and WAN's (Wide Area Networks). A LAN is generally associated with a single building, office or studio, and the WAN is a collection of many LAN's to build a much bigger network. The largest of all networks is the World Wide Web.

Solve Congestion

One of the simplest LAN's consists of a single Ethernet hub. Data received on each port on the hub will be duplicated to all other ports with no consideration for security or frame filtering. This works well for a small home or office network with a few connected computers and printers but is completely unworkable for broadcast infrastructures requiring high speed, low latency, and secure links. Congestion and collisions would soon occur resulting in highly distorted audio and video.

The Ethernet switch was introduced to solve the congestion issue. Switches automatically learn the Ethernet source and destination Media Access Control (MAC) addresses of the devices connected to each port. Using this information, the switch sends frames to ports only destined for the connected device, thus greatly reducing the potential for network congestion and collisions.

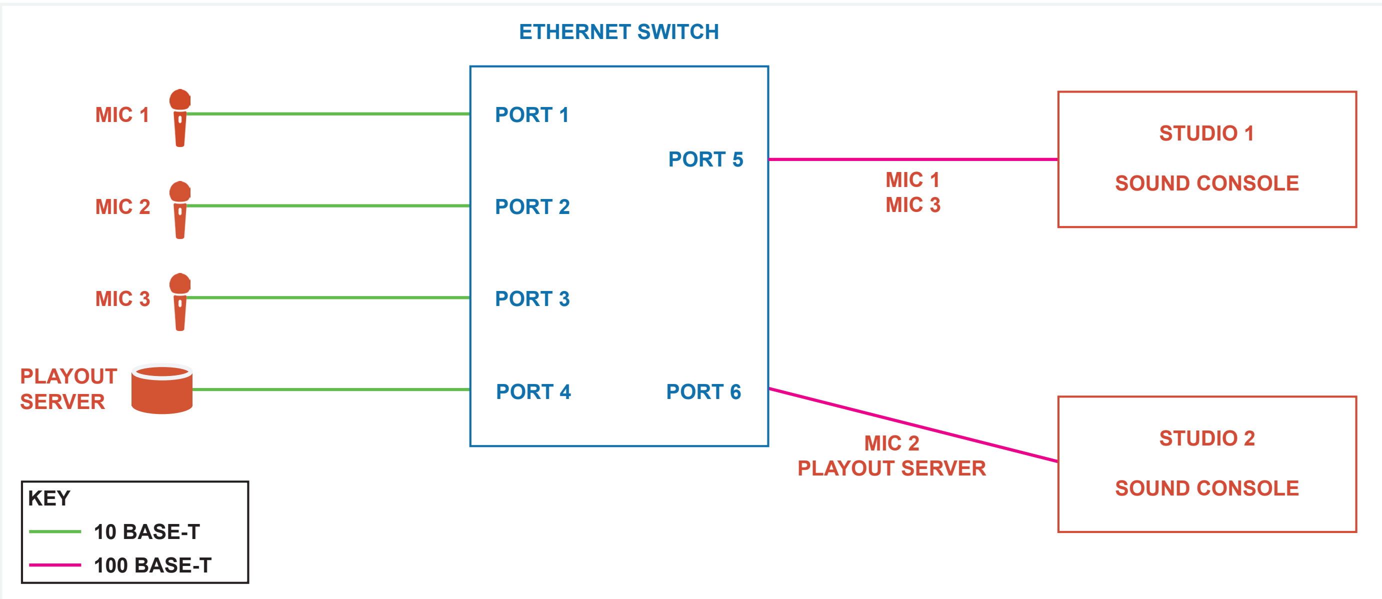

Diagram 3 - An Ethernet switch transfers frames to ports that are destined for a specific device.

An Ethernet frame is the smallest unit of bits on a layer-2 network. Frames are exchanged on the same LAN and provide a well-defined structure used for error detection and data link control. Frames contain source and destination MAC addresses and encapsulate IP packets when IP is used.

Broadcast Domain

Layer-2 networks, such as IEEE 802.3 Ethernet, use three types of delivery; unicast, multicast and broadcast. Unicast sends one single frame between devices. Multicast creates a "one to many" mapping from one device to many others. Broadcasting transmits frames to all devices in a network, also known as flooding the network.

A "broadcast domain" is a logical division of a network where all devices can be reached by a layer-2 broadcast message. This gives rise to the concept of a LAN being restricted to a building, office or studio.

WAN's join many LAN's together using layer-3 routers. Not all LAN's are Ethernet networks, so the function of a router is to join different networks and different network technologies together to form a secure, cohesive, and manageable system.

Layer-2 Detects Errors

In practice, a layer-2 Ethernet switch monitors the frames' CRC to determine if any errors have occurred in the source and destination MAC addresses, length and type field, and data payload. If the switch does detect a CRC error then it will simply drop the frame, thus resulting in a data corruption further up the IP stack.

IP is a data transfer protocol and no transmission medium is defined in the standard. The IP specification (RFC791) explicitly states that the IP protocol calls on local network protocols to carry the internet datagram to the next gateway or destination host. In other words, IP datagrams exist independently of an underlying medium on which to transport them.

Independence of an underlying network is one of IP's greatest strengths, but also provides some very interesting challenges. Throughout the history of broadcasting, the video and audio signals have been intrinsically connected to, and relied upon, the underlying transport medium. For example, AES-3 facilitates many data rates and encodes the data with the clock using bi-phase mark code (BMC) directly onto the wire to guarantee audio sample timing.

Timing is Lost

However, when using IP networks, the direct relationship between audio and video data, and sampled clock, is lost. We must adopt other strategies to reconstruct the video and audio signal such as RTP (Real Time Protocol).

Ethernet networks rely on VLAN's (Virtual Local Area Networks) to provide security. VLAN's split a network into logical units giving a unique number to each one. For example, VLAN-1 may consist of devices connected to ports 1, 3, 4, and 5, and VLAN-2 consist of devices connected to ports 2, 6 and 7. Any device on VLAN-1 will not be able to access devices connected to VLAN-2, even when sending a broadcast request.

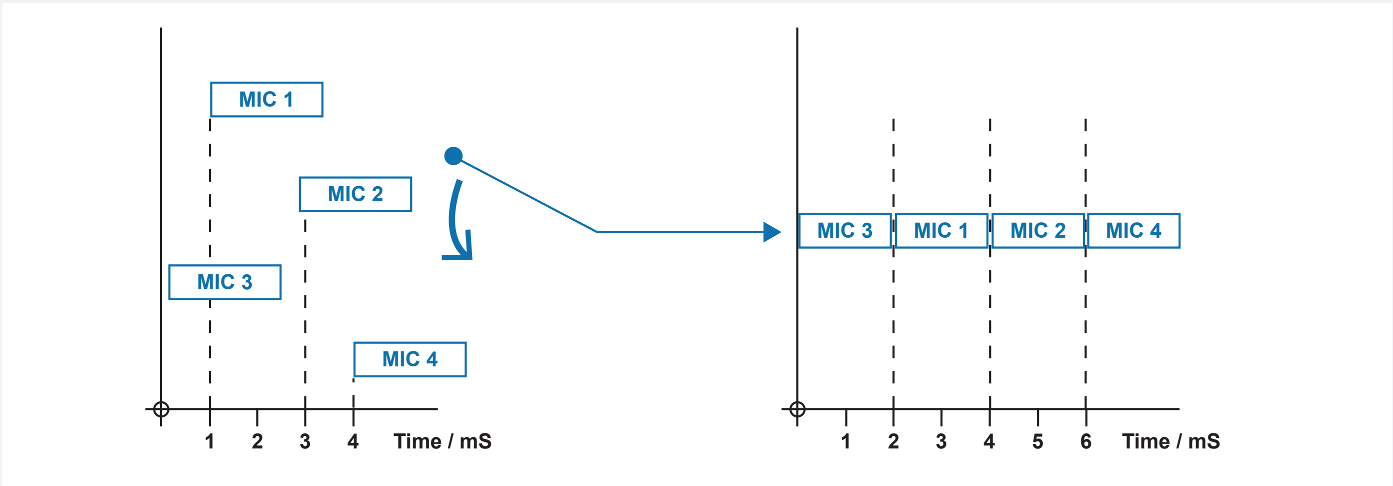

Diagram 4 - When designing a network, data-rates must be adequately calculated and provisioned to prevent unacceptable latency. Here, the diagram on the right simulates four microphones being multiplexed together on one port of a switch with insufficient bandwidth, each of the microphone packets are connected to a different port on the left diagram, and the frames queue and start to lag behind real-time.

Securely Route VLAN's

A LAN may consist of many connected layer-2 switches and if VLAN's were not used, then all devices connected to all ports within a network would be able to access each other – a clear security issue.

One application of Layer-3 routing is to connect VLAN's together. Assigning different IP subnet masks to individual VLAN's makes routing easier to administer and more secure. If we use VLAN-1 for studio 1 and VLAN-2 for studio 2, then we can route the microphones from studio 1 to studio 2 by simply creating an entry in the layer-3 routing table. This is a greatly simplified example as there will be many VLAN's within a studio.

Sources of Delay

Moving frames from one port to another in a switch requires the use of look up tables so the frames destination MAC address can be associated with the correct destination port. Routers use a similar technique but rely on using the IP address within routing tables, and not MAC addresses, to determine where data packets are moved to. Although this method has become extremely efficient using content-accessible memory (CAM) techniques, an inherent delay is introduced in the process.

Static networks use manual routing tables to tell the router which port to send the IP datagram to. Routing tables tend to become bloated and difficult to administer, and in the event of a device failure, will require manual intervention to change the routing to a known good link. Dynamic routing fixes this.

If a WAN is designed to be resilient, with extra routing paths provided to compensate for link or device failure, then dynamic protocols such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) are used to facilitate dynamic systems. If a faulty link or router develops, these network protocols will detect and re-route data around them to effectively heal the network.

Adding Jitter

Unlike baseband MADI and AES, every single packet passing through a router or switch is processed to determine such values as the MAC source and destination addresses, CRC frame checks, and "Time to Live" counters in the IP headers. This adds variable processing time resulting in further jitter to frames and packets.

The combined effect of frame delays in switches and packet delays in routers, look up tables, and dynamic routing, leads to IP packets developing temporal jitter and latency. Data buffers in switches, microphones, sound consoles and all other IP host equipment are used to bring order back to the system. However, buffers add delay and too many concatenated buffers have a detrimental effect on the audio and video.

IP networks provide unprecedented flexibility for broadcasters and an incredible amount of research and development is being conducted by broadcast manufacturers to make IP systems work with the same reliability and quality of service broadcast engineers have come to demand and expect.

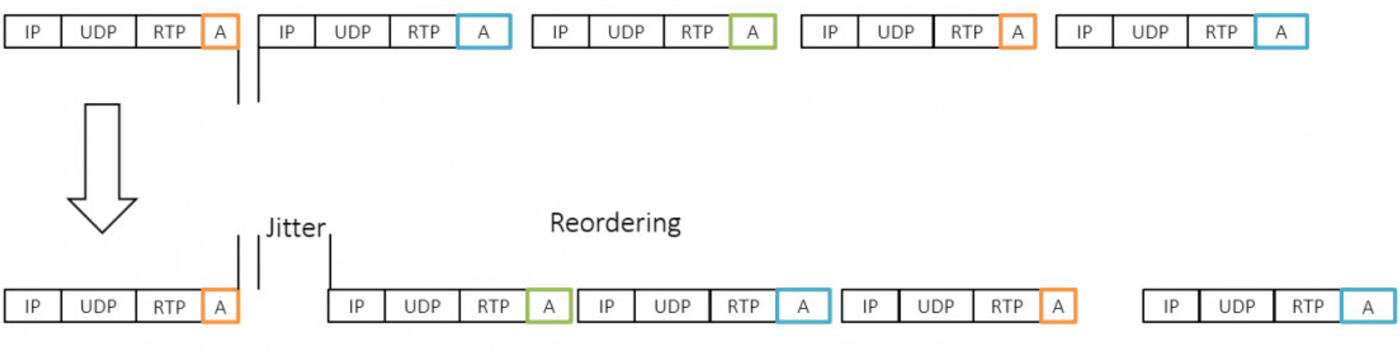

Diagram 5 - The top diagram shows a constant evenly spaced string of IP packets sent from a host device such as a microphone, the bottom diagram shows the packets with variable delay and re-ordering after traversing through switches and routers in a LAN or WAN.

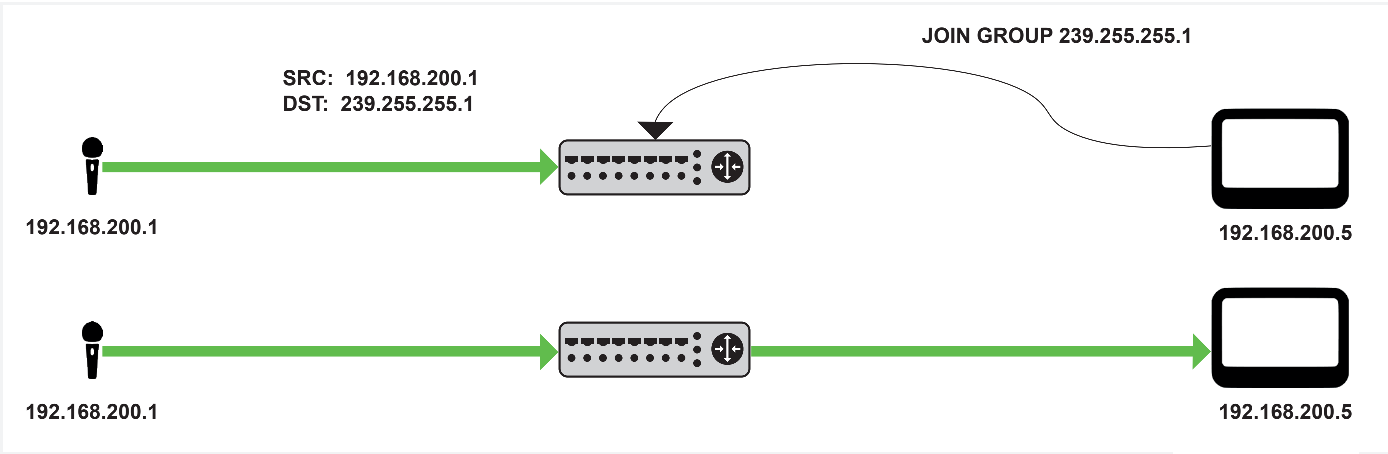

Diagram 6 - IGMP-snooping runs on the switch to determine which specific ports to move the layer-2 Ethernet packets to. This avoids network flooding.

Multicasting

Multicasting is an incredibly powerful tool used in broadcast infrastructures to efficiently distribute streams of audio, video, and metadata. In this article, we look at the advantages of multicasting, how it works, and the alternatives that overcome some of its operational limitations.

Broadcasters use distribution amplifiers (DA) to deliver multiple baseband signals in "one to many" systems. Multicasting is operationally equivalent to the DA but works by duplicating packets so that we do not have to add extra cabling to a system.

To begin multicasting, network administrators define a set of IP-multicast addresses for the network they are operating with. Devices such as microphones are allocated one of the IP-multicast addresses and use them to stream packets to the network. The IPv4 multicast range spans from 224.0.0.0 to 239.255.255.255 providing over 248 million possible groups. Each group is divided into scopes defining different usage. For Audio over IP, the scope freely available to network administrators ranges from 239.0.0.0 to 239.255.255.255, just over 16 million addresses.

Multicasting is provided at layer-2 to keep latency and jitter low, and is referred to as IGMP-Snooping. A mapping between layer-3 IP multicast addresses and layer-2 Ethernet MAC addresses is defined allowing a unique MAC addresses to be created for each IP-multicast stream.

Without IGMP-Snooping, any multicast traffic received by the switch will be flooded to all ports on that switch, resulting in a network broadcast. This is a consequence of a switch not knowing which port to send a frame to, so it defaults by sending the frame to all ports. In the extreme, this causes excessive network congestion.

IGMP hosts, such as IP-Audio-Monitoring stations, will issue an "IGMP Membership report" encapsulated in an IP-datagram. The report, often referred to as a "join" message, requests one specific multicast group and an IGMP-snooping enabled switch will forward all frames to that receiver.

To stop receiving a multicast stream, the receiving host will issue an "IGMP leave" message.

During forwarding, the IGMP Querier (typically the switch) will issue IGMP Queries to all hosts, to check for frequent validation, if the hosts are still interested in the stream.

The switch often acts as the IGMP Querier and will issue IGMP Queries to verify the hosts, such as sound desks and monitoring stations, still want to receive the stream. If they report that they do not, or the query times out, the switch will stop sending the stream to the unverified port.

Hosts on different subnets can exchange multicast traffic, if routers between them support a multicast routing protocol, such as Protocol Independent Multicast (PIM). The routers will then forward the IGMP Messages, to the subnet, where the streams originate.

Using a recording studio as an example, each of the microphones would be configured to have its own multicast address. It is possible to use direct mapping from each microphone to each input on the desk, but this would restrict operation as no other device would be able to listen to the microphones and configuration would be extremely complex, thus placing restrictions on monitoring. Multicasting allows multiple destinations, such as the sound desk input and a monitoring station, to receive the same microphone output.

Multiple microphones could be configured by the network administrator to each have their own multicast stream. For example, if studio-1 is using twelve microphones, mic-1 would have IP-multicast address 239.255.0.1, mic-2 would have 239.255.0.2, all the way up to mic-12 which would have IP-multicast address 239.255.0.12.

Diagram 7 - Riedel Communications RSP-2318 and Extreme Networks X460-G2 with multicast L2 switching and IGMP-Snooping.

To keep jitter and latency low, each microphone, sound console channel, and monitoring station for studio-1 would connect to a port on the same layer-2 switch to form a LAN. However, to control network congestion and maintain security, studio-1 would use VLAN-1 and studio-2 would get its own VLAN-2. Any devices connected on VLAN-1 would not be able to access devices on VLAN-2 and vice versa. There may be many VLAN's in use depending on the granularity of security required.

As well as configuring the IP multicast address of each microphone, the source and destination MAC addresses of the frame must also be configured. The source MAC is set by the manufacturer of the microphone. But the destination MAC address is derived from a combination of the reserved MAC prefix 01.00.5E and the IP address of its multicast stream, this is usually set by the operating system or implemented by the vendor of the microphone. This forms the next three numbers of the MAC address derived from the IP address.

To receive a feed of the microphone outputs each audio channel on the sound console must subscribe to its associated IP-multicast stream.

The sound console would have a unique IP address and each channel would subscribe to one of the IP-multicast addresses. The sound console would send a "report" message for each channel to the switch to request a copy of the microphones multicast stream and hence its output. In this scenario, each microphone would issue a mono audio stream on a specific multicast address.

Source Specific Multicast is used to stop the possibility of multicast duplication. Supporting switches and routers will only forward multicast traffic if, the receiving host issues "IGMPv3" messages, specifying not only the multicast address it wants to receive, but also the source IP of the sending device.

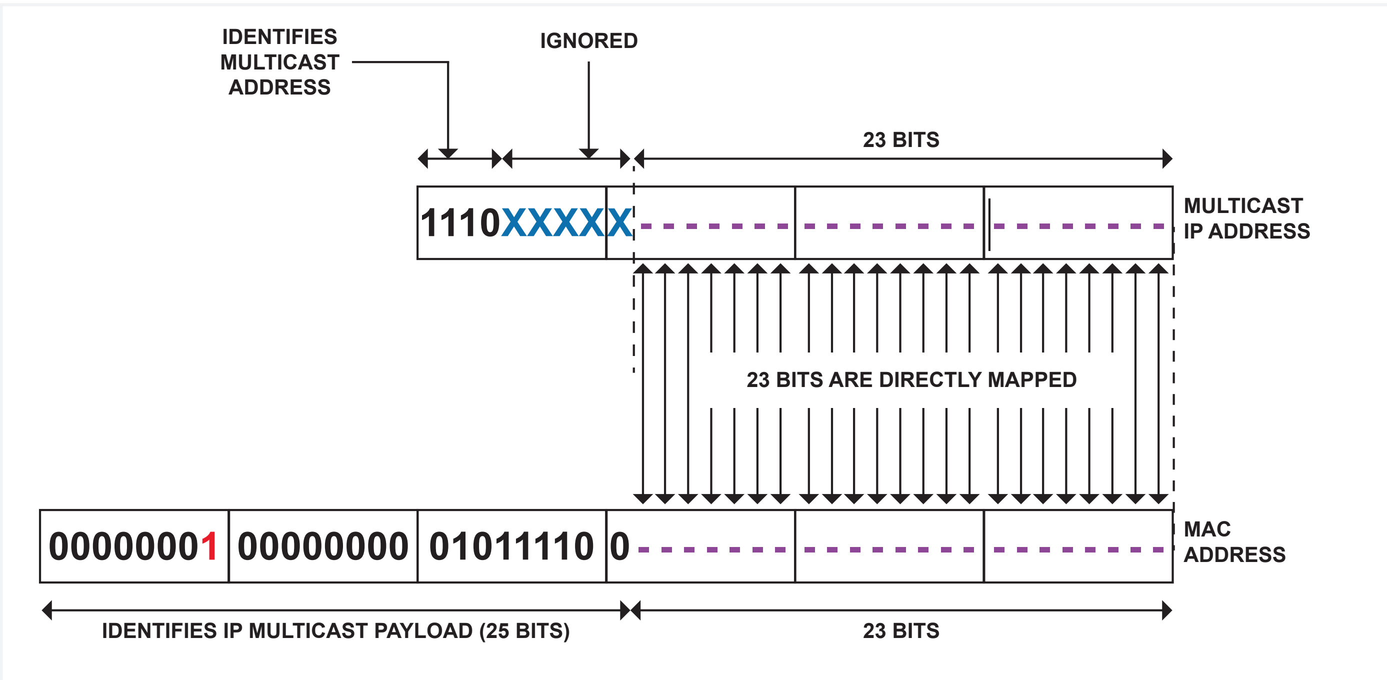

Diagram 8 - This diagram shows mapping of IPv4 multicast addresses to MAC multicast addresses.

Sponsored by Riedel Communications

RIEDEL

A device can only connect to one layer-3 router to exchange messages as the "Time to Live" of its IP header is set to one, a restriction of the specification, thus stopping the datagram from moving on to another router. IGMP facilitates communication between routers to allow streams to be accessed outside of a devices LAN.

Although IP-multicasting is incredibly powerful and extensively used, its limitations soon become apparent. IGMP is a program running on a CPU on the router and is constantly fighting for CPU time along with all the other protocols running on the router. As multicast streams increase and the number of devices wanting to connect to them also increases, so does the demand on the switches CPU. The streamed data reliability of the actual microphone output is not affected by the protocol but the speed and response with which devices can opt to switch to and from different streams.

Switches and routers state maximum numbers of multicast addresses to be used on the device for that reason. Typical values are between 1,024 and up to 10,000 addresses. But some devices exist that can only manage 128 addresses.

This is another reason software defined networking (SDN) is gathering momentum. SDN has a hierarchical view of the underlying hardware network and controls switchers and routers directly to overcome the speed limitations of protocols such as IGMP.

Vendors use different methods of remote control. Cisco use the command line interface (CLI) and REST (Representational State Transfer). Arista uses REST but it's different from Cisco's version. AMWAs IS-06 is looking to resolve these differences by providing a common interface for SDN's.

Multicasting works well when a connection needs to be occasionally made but can become slow to respond when there is heavy demand on the routers CPU. These limitations will be overcome as SDN develops.

Riedel - An IP Case Study

The Sponsors Perspective

Notre Dame's Campus Crossroads Project

An All-Encompassing IP-Based Production Infrastructure to Serve the Entire Campus

Notre Dame stadium. Photo credit Andy Fuller.

The preceding articles lay down the foundations of IP-based broadcast infrastructures and, more specifically, AoIP.

Here, we take a look at the landmark new installation at the University of Notre Dame that highlights one of the biggest advantages of IP-based systems - flexibility. In the past networks have required a lot of cables and interconnections, today, a single cable can deliver every signal to any endpoint in real time, and with outstanding quality, and robustness.

After years of standards development, the IP revolution is well underway now for media enterprises of all sizes and types. Spurred by ratification of SMTPE ST-2110 and AES67, many broadcast operations have started the hard work of turning the promise of IP into reality. Many, like the University of Notre Dame in South Bend, Indiana, are taking advantage of greenfield projects to "do it right from the first" - with all of the promises and pitfalls of being on the bleeding edge of a major technology movement.

Production Control Room During Football Game at Notre Dame's Rex and Alice A. Martin Media Center. Photo credit BeckTV.

A Universal Resource

The Campus Crossroads Project is the most ambitious building campaign in Notre Dame's storied, 175-year history. Although anchored to the university's iconic Notre Dame Stadium, the project is designed to create a year-round center of student activity and support a broad range of academic and faith-based programs in addition to athletics.

"We realized that, over time, the university has grown in such a way that Notre Dame Stadium has become the exact geographical center of the campus. It doesn't make sense for such a central facility to be used only seven times a year," said Jack Swarbrick, Vice President and James E. Rohr Director of Athletics at the University of Notre Dame. "The intent of Campus Crossroads is to bring year-round, everyday vitality to the stadium and make it a great resource for every aspect of the university."

The technology centerpiece of the project is the brand-new, 18,000-square-foot Rex and Alice A. Martin Media Center, home of Notre Dame Studios. In addition to supporting Fighting Irish Media, the video production and storytelling arm of the Notre Dame athletics department, Notre Dame Studios serves a wide-ranging group of campus customers including public affairs and communications, university relations, student affairs, and many more. Output ranges from digital content in support of academics and faith-based initiatives to production of videoboard shows in Notre Dame Stadium, Purcell Pavilion, and Compton Family Ice Arena.

Taking the IP Plunge

"We set out to create a state-of-the-art facility to meet our needs not just for today, but for years to come. It was important to pick best-of-breed technology that would allow us to change direction as needed, and also to keep up with the speed of technology advancement," noted Dan Skendzel, Executive Director, Notre Dame Studios. "IP was the obvious approach for meeting these requirements."

Notre Dame Studios turned to BeckTV for systems integration and design consulting on the IP infrastructure. Scott Rinehart, the university's Director of Broadcast Technology, commented, "We had plenty of open and honest conversations with BeckTV about whether to play it safe with baseband or take a chance on whether IP video has reached enough of a tipping point to invest in it for the future. A big deciding factor was the leading-edge technology vendors like Riedel. IP is where they're spending their R&D dollars now, so it makes sense to take advantage of their next-generation products."

He added, "The more we got into it, the more we realized that IP would be the right fit for the long-term goals we want this project to meet. Its expandability and flexibility will allow us to get maximum use out of our facilities throughout the year. For instance, we can use a studio for an academic event in the morning and then quickly turn it around for a sports event in the evening. IP also offers a foundation for us to bring more advanced technologies into the classroom, such as virtual and augmented reality."

An IP-Based Comms Foundation

At the core of Notre Dame's new IP-based communications infrastructure is Riedel's Artist digital matrix intercom system and SmartPanel app-driven user interfaces. Intercom systems offer a great use case for demonstrating professional broadcast IP implementations, since modern solutions like Artist support AoIP through industry-standard interfaces including AES67, Dante, AVB, and Ravenna. Building a future-proof intercom foundation is as simple as installing an appropriate interface card into the intercom matrix frame and then connecting to end devices via CAT5 or fiber using an IP switch.

"Riedel is all about embracing standards, and our Artist Digital matrix Intercom ecosystem has supported AVB and AES67 for years. Combined with AES67 client cards in the Artist matrix, our SmartPanels are the only SMPTE 2110-30-compliant intercom panel solutions on the market," said Joyce Bente, President at Riedel North America. "The panels are full-bandwidth, compared to competing IP panels with just 7KHz of audio bandwidth and higher latencies over LANs. That means you could conceivably layer your intercom atop your existing IP infrastructure to save cabling and lower costs."

Lessons from the IP Trenches

As with any major technology shift, Notre Dame Studios' IP adoption hasn't been a completely smooth ride. However, with IP-optimized solutions from Evertz, Telos Alliance, and Riedel in use, Notre Dame Studios is well on its way towards a stable and turnkey IP operation.

"There have been plenty of pain points along the way, mostly about having to 'unlearn' what we know about baseband signal distribution. The learning curve on this is steep, and we have a long way to go before we fly solo," Rinehart noted. "At this stage, nobody fully understands all of the ins and outs of a ground-up IP facility, but the BeckTV and Riedel teams have been great resources as we've talked things out and come to decisions."

One persistent issue, in which the routers were not passing timing information correctly, offered a particularly valuable lesson for the engineering and IT teams and emphasized the necessity of close cooperation between both. Over the course of several months, the IT department agreed to a solution in which PTP data would be delivered first, media data second, and all other data third, but this still didn't eliminate excessive jitter/offset on the PTP clock across all switches. It was only after Notre Dame Studios upgraded their Cisco switches that PTP came within spec and these timing problems disappeared.

"The big lesson here is that the industry is a long way from plug-and-play gear for IP networks. We also learned the importance of getting engineering staff up to speed on the knowledge sets they'll need for the new IP paradigm, as well as making sure the IT team truly understands our objectives," Rinehart said. "I will say that the support we've gotten from Riedel has been fantastic. We struggled with the router timing issue for a while, but we were finally able to get it resolved through the combined efforts of Riedel, Cisco, and our networking team."

Storied Past, Exciting Future

The challenge of managing and maintaining IP networks will require a whole new set of skills for technicians and new levels of cooperation between technicians and their IT departments. IT networks are collaborative in nature, so it will be up to all parts of the broadcast world – manufacturers, operators, engineers, and IT professionals – to work together on delivering the best results for technical infrastructures. But, as the Notre Dame Studios can attest, the rewards are worth the effort.

Rinehart observed, "We're not the tip of the spear, but we are awfully close. At the moment, building an IP facility from the ground up is by its nature highly complex, so we've gone through plenty of challenges – but things are smoothing out and we're confident this facility will carry us for many years without a major 'forklift' upgrade. It's an exciting future."

In a recent article in a sports tech journal, he took a more philosophical view. "There's a quote sitting right outside my door that I look at every day, all day, and it says one of Notre Dame's greatest assets has been the boldness of its vision, the ability to see possibilities and connections where others saw only obstacles and fragmentation. I look at that, and it hits me every day that this place is willing to accept calculated risk. It's what we do; it's in the DNA of the university system to do that. And, luckily, we had essentially a greenfield possibility here to do that."

THE BROADCAST BRIDGE

Connecting IT to Broadcast

WP
WHITE PAPERS

EG
ESSENTIAL GUIDES

MEDIA

CS
CASE STUDIES

Find Out More

For more information and access to white papers, case studies and essential guides please visit:

thebroadcastbridge.com

6/2018

Sponsored by Riedel Communications

RIEDEL