SMPTE

MOTION IMAGING JOURNAL

Covering Emerging Technologies for the Global Media Community

Engineering Design Challenges in Realization of a Very Large IP-Based Television Facility

By Steve Sneddon, Chris Swisher, and Jeff Mayzurk

Introdução

Em 2018, Telemundo Center foi o lar do maior ambiente SMPTE ST 2110 do mundo, com mais de 12.000 fontes em HD, e 150.000 fluxos multicast em áudio e vídeo, produzidos para a Copa do Mundo da FIFA na Rússia! O presente artigo da uma verdadeira aula para explorar as principais considerações e desafios na construção em grande escala de uma instalação de produção em IP seguindo as diretrizes da SMPTE 2110. Para os gestores de projetos, o tópico "Lições Aprendidas" será uma aula aparte. Boa leitura!

Por: Tom Jones Moreira

Abstract

In 2016, NBCUniversal (NBCU) broke ground on Telemundo Center, the new Miami-based global headquarters for Telemundo Enterprises. The facility, which features 13 production studios and 7 control rooms supporting scripted episodic content, daily live news, and sports programming, went live in Spring 2018 with the coverage of the Fédération Internationale de Football Association (FIFA) World Cup. At the time of launch, Telemundo Center was home to the largest SMPTE ST 2110 environment in the world, consisting of over 12,000 unique high-definition (HD) sources and 150,000 multicast streams across audio and video. The technical decision to use a software-defined video network infrastructure was essential to supporting the scale and flexibility of a facility of this magnitude. This article explores the major considerations and challenges in building such a large-scale, all-Internet Protocol (IP) broadcast production facility, including design factors around switching video flows, redundancy, control and orchestration, Precision Time Protocol (PTP) master clock systems, and handoffs to multimanufacturer SMPTE ST 2110 devices as well as non-IP-enabled devices. We also review our experiences and lessons learned with utilizing software-defined network (SDN) control plane and routing commands that abstract the underlying physical and link-level connectivity.

Keywords

Control room, Internet Protocol (IP), network, production, router, studio, uncompressed, video over IP

Introduction

In 2016, NBCUniversal (NBCU) broke ground on Telemundo Center, the new global headquarters for Telemundo Enterprises, located in Miami, FL. This new facility is comprised of offices and studios for Telemundo Network, Telemundo Studios, Telemundo International, and Universo Network, in addition to, being the home of NBCU Internationals Latin America offices. One of the goals of this project, in addition to bringing business units together, was to make the facility as technically future-proof and flexible as possible, in order to best serve Telemundos needs in an evolving media landscape.

Telemundo Center is the manifestation of our commitment to the Hispanic market and a representation of our core values of innovation, collaboration and transparency, said Cesar Conde, Chairman, NBCUniversal Telemundo Enterprises and NBCUniversal International Group. Latinos are a growing cultural, political, and economic force influencing every aspect of our country. Telemundo Center is the only facility that can fuel the preferences and demands of this dynamic audience, while driving unlimited growth and opportunity for our company, our employees and our community for years to come. [1]

Telemundo Center is the mani - festation of our commitment to the Hispanic market and a representation of our core values of innovation, collaboration and transparency, said Cesar Conde, Chairman, NBCU Telemundo Enterprises and NBCU International Group. Latinos are a growing cultural, political, and economic force influencing every aspect of our country. Telemundo Center is the only facility that can fuel the preferences and demands of this dynamic audience, while driving unlimited growth and opportunity for our company, our employees and our community for years to come. [1]

Telemundo Center opened its doors in mid-2018, premiering with 31 days of coverage of the Fédération Internationale de Football Association (FIFA) World Cup 2018. The facility is now a hub of content creation, delivering daily live news, entertainment shows, sports programs, and scripted episodic content across multiple media, including broadcast, cable, and digital.

The use of video over Internet Protocol (IP) using SMPTE ST 2110 is vital to the future-proof design of

Single-tier topology consists of large core switches to which endpoints are directly connected. These core switches would typically have a large number of low-speed ports (10 or 25 Gb/s). They may also include a few higher bandwidth ports (40 or 100 Gb/s) for various specialty or bulk transport purposes, including high-bandwidth endpoints, facility interlink, or connection between core switches.

Single-tier topology offers two major advantages. First, no design consideration is required for the sizing of inter switch bandwidth and uplink aggregation. Second, hardware deployment and the resulting infrastructure are generally simple compared to leaf-spine or any other distributed topology. This leads to reduced cost, complexity, maintenance, and wiring when compared to a leaf-spine topology. **Figure 1** compares both leaf-spine and single-tier topologies.

In the design of Telemundo Center, we considered leaf-spine topology due to the benefits listed above. It is also more in line with the current trends in network design for nonbroadcast applications and can therefore be considered the more common or standard model. However, we ultimately pivoted to a single-tier topology because it is simple and economically cheaper. In fact, we deployed a pair of discrete single-tier networks, named production and acquisition. Each of the production and acquisition networks is serviced by a large core switch with over 2,000 10-Gb/s ports to connect endpoint devices.

Each endpoint in the facility is connected to one or the other of the production and acquisition networks, divided according to operational affinities. These affinities represent functional requirements for each group of endpoints to primarily consume video within its group. All endpoints associated with studios and control rooms are connected to the production core, including cameras, graphics devices, production switchers, and playback servers. The acquisition core supports incoming remote feeds, outgoing distribution, disk recorders, and post-production.

SDN and Hardware-Controlled Network

The video network at Telemundo Center is a software-defined network (SDN), meaning that there is a software controller that instructs the control plane of the core switches on how to route video flows. This software controller understands the physical network topology, with ingress ports and stream information for each source flow, as well as egress ports and host information for each endpoint consuming video flows. It also provides a user interface to issue route requests and then instructs the core switches to direct flows from ingress to egress ports.

An alternative to an SDN would be a more traditional hardware-controlled network, where packet forwarding decisions are individually made in the hardware control plane of each node of the network. In the absence of a central controller, each network switch operates autonomously, forwarding packets based on a set of predefined rules. Route requests may be issued directly to network nodes by endpoints via Internet Group Management Protocol (IGMP).

Video Network Redundancy

As mentioned above, the Telemundo Center video network is an SDN-controlled single-tier architecture. In fact, the video network is comprised of two separate such networks providing redundant architecture based on the SMPTE ST 2022-7 standard for seamless protection switching.[2]

In an ST 2022-7 environment, two video transport networks always actively transport video. These two active networks can be thought of as X/Y networks where every endpoint device is simultaneously connected and is active on both the X and Y networks using double the amount of physical network interfaces that would be required in a nonredundant network. Half of the network interfaces on the endpoint device are connected to the X network and half the network interfaces are connected to the Y network.

Endpoints transmitting flows send packets to both X and Y networks simultaneously. Endpoints receiving flows receive and process packets from both X and Y networks and compare Realtime Transport Protocol (RTP) timestamp and sequence numbers to assess packet timeliness and perform a hitless merge of a single stream for decode or downstream processing based on packets from either network.

Although ST 2022-7 is an important feature of a large IP video implementation, the protection model offered does have limits. One important limitation is that ST 2022-7 cannot protect against the full failure of an endpoint. If a power issue, for example, causes the complete failure of a transmitting endpoint, no video flows will be transmitted to either redundant network. Because of this limitation, highly important video should also be backed up using additional redundancy models, for example, models that might be utilized in an SDI plant. A completely redundant model for the most critical feeds would employ discrete source and destination feeds across diverse hardware, each transported redundantly with ST 2022-7.

An additional limitation of the ST 2022-7 seamless protection switching model is when there are multiple network path failures on the X and Y networks. For example, consider a transmitting endpoint that is dual fiber connected to the X and Y cores and a receiving network device that is dual fiber connected from the X and Y cores. In the event of a fiber failure on the X path from the transmitting device to the X core, along with a simultaneous failure on the Y path between the Y core and the receiving endpoint, video is now lost between the transmitting endpoint and the receiving endpoint, even though each endpoint still has a one good link up. Even more interesting, other endpoints in the system will still be able to send and receive flows with these endpoints because they still have both their links up. This situation can lead to confusing troubleshooting that defies traditional broadcasting source/destination testing logicwhere a source is available to all but one destination and a destination is available to all but one source.

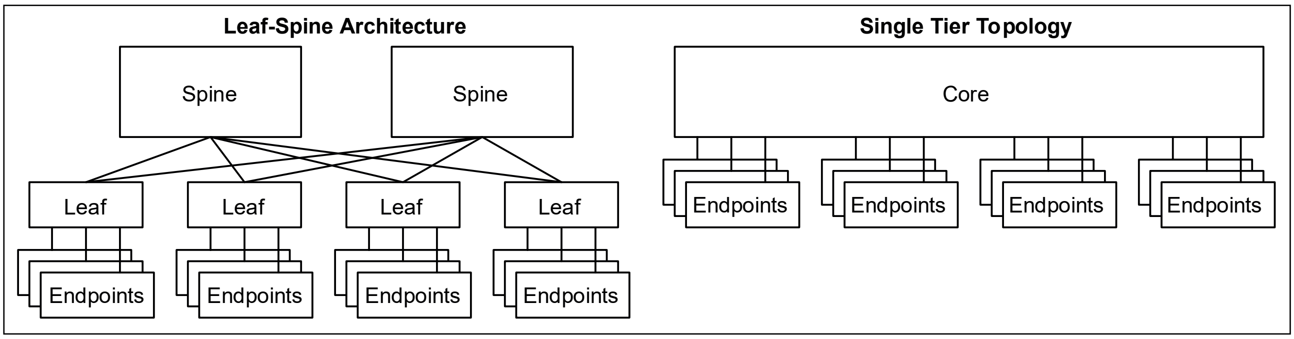

FIGURE 1. Simplified block diagram illustrating leaf-spine and single-tier topologies.

Telemundo Center. At the time of launch, Telemundo Center was home to the largest ST 2110 environment in the world, with more than 12,000 unique high-definition (HD) sources and 150,000 multicast streams across audio and video. This study explores the major considerations and challenges in designing and building such a large-scale, all-IP broadcast production facility.

Facility Overview

Telemundo Center, with a size of approximately 500,000 sq. ft, is located on 21 acres in Miami, FL. In addition to office space for 1,500 employees, the building features full production facilities to enable news, sports, and scripted entertainment for broadcast and digital outlets. To support these activities, the building has the following facilities.

- 13 production studios in various sizes up to 8,000 sq. ft.
- 5 live Production Control Rooms (PCRs)
- 72 edit seatsapproximately half of which are desktop edit and half are edit rooms
- 60 graphics creation seats
- a central video playback area
- a central graphics playback area
- a central camera shading area

A single central equipment room (CER) supports all the above activities. At its heart, is a redundant set of core IP video routers using SMPTE ST 2110. The CER also houses the fiber-optic patching core, with more than 10,000 strand presentations for plant-wide distribution. Fiber core frames feed out to intermediate distribution frame (IDF) rooms throughout the facility for secondary cross patching to local endpoints.

Even though Telemundo Center is the global headquarters for the Telemundo network, its technical infrastructure is limited to uncompressed studio production and post workflows. Master control and distribution services for Telemundo and other NBCU properties are handled mainly at other locations and then handed off to Telemundos network of affiliates for local programming and commercial integration and for broadcast transmission. Even the local Miami Telemundo affiliate is managed this way and is not located at Telemundo Center. Additionally, many incoming remote feeds are also received and normalized at a central location before handoff to Telemundo Center. Outbound program feeds from Telemundo Center are converted into serial digital interface (SDI) at the edge and handed off to service providers for interfacility transport. Therefore, from a technical perspective, Telemundo Center can be thought of as more of a pure production facility than one may have expected.

Network Architecture and Core Technologies

Many technological advances took place in the broadcasting industry while Telemundo Center was being built. The use of SMPTE ST 2110 and video over IP, in general, were just emerging as viable solutions for a facility the size of Telemundo Center and there were not many reference designs upon which to base the technology architecture. The following sections detail some of the architectures that were considered and explain what we ultimately selected in each area.

Leaf-Spine and Single-Tier Architectures

The primary requirements for the Telemundo Center video network were large-scale and nonblocking capacity, meaning that the architecture should not impose any limitations on the capability to route any source stream to any destination receiver. Two main classes of architecture were considered for these requirements: leaf-spine and single tier (**Fig. 1**).

Leaf-spine topology consists of a small number of large core switches (spines), along with many smaller switches (leaves). Generally, in leaf-spine topologies, endpoint devices are connected to the leaf switches and never to the spine switches.

For a nonblocking network, leaf switches have low-speed ports (10 or 25 Gb/s) for endpoint connections and higher bandwidth (40 or 100 Gb/s) uplink ports to spines. Total spine uplink bandwidth must equal the total endpoint bandwidth on each leaf. Total spine bandwidth is the overall network capacity.

Leaf-spine topology offers two major advantages. First, leaf switches can aggregate low-bandwidth-consuming endpoint devices, thus avoiding bandwidth wastage at the spine. Second, leaf switches can be distributed throughout a large facility and can also be located close to endpoints. This second advantage may simplify cabling distribution by keeping most links within a small area. The use of a leaf switch to service endpoints within a single rack is known as the *top of rack* switch model.

The problem may not move in a way that certain troubleshooting logic may suggest.

SMPTE ST 2022-7 seamless protection switching is one of the most powerful tools offered by IP video when compared with SDI. Redundancy is especially valuable in the case of very large IP networks with a failure block potentially equal to the entire video environment. But it is critical to understand the nature of the redundancy model and its limitations. The use of ST 2022-7 does not alone convey a bulletproof property to the video network and certainly not to the facility overall.

The use of a robust health monitoring and alerting toolset is recommended to keep support teams informed of actual or imminent failures. Seamless redundancy may have the effect of masking critical system faults, and in the event of link failure or other outage protected by ST 2022-7, care should be taken to repair the impacted leg and restore ST 2022-7 protection.

PTP and Reference Systems

In a traditional broadcast plant, a reference signal carried as a composite analog video signal, commonly referred to as black burst or genlock, is distributed to every piece of equipment that produces or processes video. Black burst itself is an analog video signal used as a common phase reference to synchronize the video generated throughout the plant, allowing for every source to be vertically in time with every other.

In a video over IP plant, a newer method of synchronization is used: Precision Time Protocol (PTP).[3] PTP is not a video-specific technology; it provides highly accurate clock information across all kinds of computer networks. However, SMPTE has standardized specific PTP profiles for use in professional media networks.[4,5] When designing an IP production plant, it is important to understand the mechanics of how a PTP clock system interacts with endpoint devices.

Traditional black burst distribution is a one-way clock signal. In contrast, PTP is two-way communication protocol in which endpoints both receive timing information and communicate back to the clock. Because of this clientserver relationship between clocks and endpoints, PTP clocks have a limit to the number of endpoints with which they can interact. For the purposes of the Telemundo Center network, there are two main types of PTP clocks.

PTP grandmaster: This is the ultimate PTP generator that sits at the top of a network. It can take time and phase data from an external source such as GPS and generate the master PTP signal on a network.

Boundary clock: A boundary clock acts as a submaster clock on the PTP network. It connects to a PTP grandmaster to obtain PTP clock data and then acts as a master to downstream devices, including endpoint devices or other boundary clocks. Boundaries serve in this way to segment the PTP network into smaller zones. Boundary clocks are important because, as stated above, a particular PTP master can service only a finite number of endpoint devices. Adding boundary clocks in parallel to the existing boundary clocks is the proper way to scale a PTP distribution network as endpoint devices are added.

For Telemundo Center, there is a set of four redundant PTP grandmasters connected to a PTP distribution network. Also on the PTP network, there is a set of dedicated boundary clocks configured to synchronize with the grandmasters. Each of these boundaries, in turn, is connected to the 10 Gb/s ST 2110 video network to act as master clocks to endpoint devices. No switches or other network nodes are configured to act as PTP masters. The IP video network is configured to segment PTP traffic between these boundary clocks and endpoints to specify exactly which endpoints are locking to which boundaries, thereby ensuring that we do not oversubscribe the boundary clocks.

Finally, it is very important to properly configure the system to our needs. One of the foundational concepts of PTP is the best master clock algorithm, or the BMC algorithm. The BMC algorithm allows clock devices on the network to perform a voting procedure to elect one of the several available masters as the one to which they will synchronize. Though several factors play into the BMC algorithm voting procedure, tiered priority setting configured by system administrators on the clock itself is the major one. The BMC algorithm also considers more nuanced factors such as quality of signal and the source of the masters upstream synchronization. In an improperly configured PTP environment, the BMC algorithm may result in clocks assuming the role of a master, in contradiction to the system administrators intention.

SMPTE ST 2110

SMPTE ST 2110 professional media over managed IP networks[6] is a suite of standards for use in professional content production that describe the mechanism for using IP to transport video, audio, and metadata streams.

- SMPTE ST 2110-10/-20/-30 Addresses system concerns and uncompressed video and audio streams.[79]
- SMPTE ST 2110-21 Specifies traffic shaping and delivery timing of the uncompressed video.[10]
- SMPTE ST 2110-31 Specifies the realtime, RTP-based transport of AES3 signals over IP networks, referenced to a network reference clock.[11]
- SMPTE ST 2110-40 Maps ancillary data packets (as defined in SMPTE ST 291-1) into RTP packets that are transported via User Data Protocol/IP (UDP/IP) and enables those packets to be moved synchronously with associated video and audio essence streams.[12]

One of the key advancements of SMPTE ST 2110 is that all video, audio, and metadata are transmitted as separate IP multicast data flows. Having separate elementary essence streams over IP allows for a wide variety of content creation scenarios that would not be easily achievable if audio, video, and metadata were more tightly bundled together as in SDI.

Telemundo Center does not include any IP video in SMPTE ST 2022-6,[13] and compressed video streams are isolated to the facilitys edge. Incoming remote sources are normalized to ST 2110, and outbound feeds are converted into SDI before handoff to service providers for interfacility transport.

Design Considerations and Implementation

Video Standard HD, 3G, 4K

One of the often-advertised advantages of moving to an all-IP infrastructure is the promise of being format agnostic. SDI and baseband technologies had a strong bond between the bandwidth demands of their formats and their underlying transport medium. This model served SDI well, as the data rates required for SD (270 Mb/s), HD (1.5 Gb/s), and 1080p (3 Gb/s) exceeded those provided by commodity ethernet throughout the 1990s and early 2000s. In short, SDI baseband networks supported significantly higher data rates than were cost-effective for IP at the time. However, by the late 2010s, commodity ethernet had far exceeded the capabilities that could be developed economically for baseband video transport. IP transport technologies have a well-defined separation between their transport bandwidth abilities and their payloads. As commodity ethernet bandwidths increase, IP will be able to transport video formats with ever-increasing bandwidth demands.

While Telemundo Center is largely a 1080i 59.94 plant, we sized all infrastructure to allow for streams to run up to 1080p 59.94. For example, while six video streams at 1080i could be carried over a single 10 Gb/s link, we would instead provision bandwidth for those six streams as 2× 10 Gb/s. The net effect of this design is that the system has approximately 50% reserved bandwidth for future growth per link.

One of the main reasons we provisioned these links for 1080p was that we wanted to have a path to 4K UHD. Many devices in our production chain support 4K operation modes, wherein a set of four video ports typically run discretely at 1080i can also be run at 1080p in groups of four 3 Gb/s signals for quad-link 4K. By sizing IP bandwidth to 3 Gb/s reservations for each stream, we enabled future support for this kind of 4K operation mode.

Network Aggregation and Bandwidth Efficiency

One critical design consideration for IP video, especially in designs using a single-tier network topology, is the efficient port utilization on the core switches. Multiple IP streams flow across each network link, and each port can be subscribed up to its maximum available bandwidth.

A 10 Gb/s port would be considered significantly undersubscribed if it were transporting, for example, only one or two streams of 1080i 1.5 Gb/s. There is no technical problem with undersubscribing ports in this way, but it is an inefficient use of overall network capacity.

One way to get around the issue of high port counts occupied with low-bandwidth devices would be to install smaller switches in areas of the plant to aggregate bandwidth more effectively, thereby freeing up higher bandwidth ports on the core switches. This method of aggregation is exactly the advantage of a leaf-spine network, but it can also be developed in a more targeted fashion in an otherwise single-tier network.

Aggregation of this sort can help unlock the full capacity of the network. Systems engineers need to strike a careful balance between avoidable port waste and excessive aggregation. With too much aggregation, some of the simple architectural advantages of a single-tier topology may be lost.

For the Telemundo Center build, we considered all levels of aggregation from a very aggressive model, where we would try to conserve as many core ports as possible to using no aggregation at all. We eventually ended up with a model of using limited or light aggregation. We accepted some portion of port waste on the core but provisioned some aggregation switches supporting banks of similar low stream-count devices. This model left us with more than enough ports in the core switches for future growth and allowed us to still have a simple core switch design.

Audio Transport Considerations

As we discussed earlier, an important feature of SMPTE ST 2110 is that video, audio, and ancillary data are transported as separate multicast streams that can easily be routed to different destination endpoints independently. Audio streams encapsulated as AES67[14] may be routed separately from video streams to an audio-only endpoint, such as an audio mixer.

For Telemundo Center, we had initially planned to use AES67 audio networking for everything, including streams that were part of audio-only sources, such as the output of studio microphone preamplifiers. The goal here was to have one media network for all video and audio.

As the project progressed, we found that large-scale AES67 audio-only deployments were not ready to fully interoperate with ST 2110. Due to this limitation, we designed to bridge the SMPTE ST 2110 video/audio transport environment with a more traditional audio-only router environment connecting production audio consoles to studio microphone preamps and interruptible foldback (IFB) in-ear monitors. The bridge between these two network worlds is a bank of bidirectional ST 2110 to multichannel audio digital interface (MADI) converters. These convertors serve as logical sources and destinations on their respective network to pass audio between environments.

Another key audio consideration is the limited standardization in the packaging of mono audio channels within ST 2110 AES67 multicasts. Some vendors have chosen a method of packaging one mono audio per multicast, whereas others have chosen 16 and some have chosen 4. This nonstandardization of packaging of multicasts has led to some difficulties interfacing audio streams between vendors.

Pooled Resource Management and Operational Presentation

This section mainly concentrates on exploring some of the ramifications of a large-scale IP build for production operations. First, we will consider traditional paradigms of video source presentation in an SDI video architecture for a large facility with multiple PCRs. For this purpose, we will use an SDI broadcast plant featuring several PCRs and a core router for shared resources. There are two main ways in which sources in this environment are presented to the PCR.

- *Local sources*These are sources directly wired to the production switcher and audio mixer in a PCR. These sources have a local-only identity and naming.
- *Core sources*These sources are wired directly to the core router but routed into the PCR and shared amongst all PCRs. The core router source may be called a *Remote* or *REM* and the local presentation would be Remote Source or RS.

For Telemundo Center, a necessary design strategy was the elimination of the split between core and local sources from a physical wiring point of view. Because of the nature of an IP video plant, where everything needs to be connected to an IP network, all of these sources are connected to a single nonblocking IP network.

Without care to simulate legacy presentation models, this new everything routable model may result in a scenario where previously local sources appear as remote sources.

The key problem with this model is that production operators will have only a globally unique nameset for all sources. Such global-only naming means that a director may be using CAM 1 and CAM 2 on one day, and CAM 3 and CAM 4 the next day. This variability will cause operational confusion.

Our goal for Telemundo Center was to present these global pooled resources to production operators in a way that appeared like local sources in the legacy SDI model. Our solution was the use of *virtual* sources and destinations to provide a localized presentation to operators in each control room. With the addition of this virtual layer, the director can always call CAM 1 in any control room, irrespective of which physical pooled resource is assigned to that localized virtual. **Figure 2** shows a simplified view of physical CAM sources and virtual CAM sources.

These virtual sources and destinations are configured within the router control system to act like physical sources

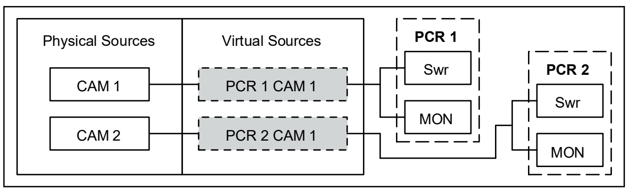

FIGURE 2. Simplified block diagram illustrating physical and virtual sources.

and destinations, but they do not consume any IP network bandwidth or physical ports. There is no practical limit to the number of virtuals that can be created.

The typical control room workflow is to route the localized virtual sources to local destinations, such as production switcher, mixer, and monitor inputs. These routes are defined in standard control room setup presets, generally locked and protected and not changed on a regular basis. Physical sources are then mapped to virtual sources on a per-show basis, as needed, and flow downstream to multiple receivers.

Additionally, these virtual sources and destinations are given *namesets* localized to various control room contexts. For example, a single virtual source may be labeled CAM 1 in PCR 1, but labeled PCR 1 CAM 1 in PCR 2. Another virtual source would be labeled CAM 1 in PCR 2, but labeled PCR 2 CAM 1 in PCR 1. This allows not just localized presentation of resources, but also a naming convention for sharing virtual resources between control rooms.

As deployed at Telemundo Center, this solution provides a powerful and extensible mechanism for operational presentation of pooled resources. In our model, every virtual source and destination assignment can be described by a standard route event on the router control system. We have broad control over the mechanics of making device assignmentsoptions include individual or *ad hoc* source/destination routes, presets for common setups, or control from external orchestration systems using standard router control protocols. In fact, we deployed each of those mechanics for various device types and workflows, all running on top of the underlying virtual architecture.

IP to the EdgeHow Far to Go?

Some consideration should be given to the question of how far in a production plant an IP infrastructure should extend. From a network infrastructure simplicity perspective, an ideal situation would be for the ethernet to extend all the way to an endpoint. However, an endpoint in a live television studio may be an in front of the camera display technology such as on-set monitors and video walls, which are often exposed to movement, disconnection, and reconfiguration. While a fiber-based ethernet link can be ruggedized, it will probably never be comparable to a coax-based SDI connection from a reliability standpoint when plugged and unplugged often. Additionally, since an SDI connection needs very little configuration at the endpoint, it is more ideal for use in an environment used to that sort of reconfigurability.

At Telemundo Center, one possibility of how to make the transition from IP to SDI for feeding elements in the studios was to install SDI gateways in a central location to feed all the studios. However, because of the size of Telemundo Center and the size of the studios, we would have quickly exceeded the allowable cable length for SDI over coax. On the other hand, an ideal landing point for IP to SDI conversion would have been inside the studio broadcast service panels (BSPs). Unfortunately, at the time of the build, there was very little IP gateway vendor equipment that was both low profile and had sufficiently quiet fans to be usable inside a studio. We eventually placed the IP gateways for each studio in a network closet that was co-located with each studio. This solution was a good compromise, though since the time of the Telemundo Center build, vendors now have a larger variety of IP gateway equipment that would be acceptable to be located within a quiet studio.

Conclusion

Lessons Learned

Telemundo Center was a unique opportunity to build a large IP plant as a greenfield project. Over the course of design, installation, testing, and operational commissioning, we developed some key findings and recommendations that we hope will be helpful to other broadcasters seeking to implement IP solutions.

Try to Avoid Legacy SDI Coax Wiring Practices Where Possible

Even with the extensive use of SDI-IP gateway devices at Telemundo Center, we eliminated the use of jackfields and distribution amplifiers. While eliminating jackfields and distribution amplifiers may increase complexity in maintenance troubleshooting, it results in significantly streamlined physical builds in terms of time and wiring complexity. Infrastructural simplification also leads to a significant compression of required rack space and a reduction of passive gear in equipment rooms, which tends to result in orphaned power and cooling capacity in modern datacenters.

Develop a Strong, Well-Considered Fiber-Optic Management Plan

While fiber is not new or unique to IP video builds, IP will require a significantly increased use of fiber relative to SDI builds. Implement stringent standards around fiber cleaning and general fiber cabling management, including training programs for integration and maintenance teams. All fiber connections should be cleaned, inspected, and then cleaned again before being inserted into devices or bulkheads. Many IP builds suffer from improperly cleaned fiber connections that may result in data loss and force a recleaning process after the nominal conclusion of physical integration. Save time in the build by touching only each fiber connection once and cleaning it properly.

Complexity in Installation Phase Versus Configuration Phase of a Plant Build

In a typical SDI build, because each cable carries a single unidirectional video stream, the intent of a system can be understood from a wiring diagram and much of the configuration ends when the system is physically installed. In contrast, in an IP build, each cable carries multiple signals in each direction, and every cable connects an endpoint device to a network switch. After the physical installation of cables, a phase of programming the intent of the system into the network must be undertaken.

The wiring of an IP plant becomes the quicker part of the project, whereas the configuration phase becomes much longer and move involved. Installers and integrators must account for the extra time needed to configure software and network flows when planning an IP project.

The Perfect IP Network Is not a Minimum Requirement

Broadcasters should not feel that implementing IP has to wait until they can provide total interoperability in native IP on a single nonblocking redundant network for all media types. While total unification may be an industry goal in the long term, for the foreseeable future it is perfectly acceptable to make concessions to an idealized view of IP.

Vendor-Agnostic Fully Native IP Solutions Are Nearly Impossible to Achieve at This Time

Even where products share support for ST 2110 itself, there is no common control standard adopted, either for stream subscription or switching mechanics. Broadcasters will likely have three choices for integrating third-party devices into their IP plant for the foreseeable future. These options are SDI gateways, network address translation (NAT) solutions for static endpoint stream subscription, and device-specific or proprietary endpoint control application programming interfaces (APIs). Each option has its own advantages and disadvantages. For Telemundo Center, we selected an architecture in which most third-party devices are connected with SDI gateways. Gateways provided the most frictionless install and commissioning experience, without compromising any of the functional benefits of IP. As control standards and IP product support progress, we have the option to abandon gateways in favor of direct native connectivity.

Additionally, the use of diverse IP media networks may be advisable for different use cases and media types. Such diversity may include audio-only networks for audio-only applications, as well as separate networks for compressed and uncompressed video. A well-designed and completely flexible IP plant solution may involve a number of different networks without any overall functional loss relative to a perfectly unified network. Consider the example of studio microphone sourcesthese will never be routed to production switcher inputs or multiviewer displays, so that capability should not be a requirement of any IP system build. Limitations of a diverse network can be overcome with media conversion gear at network interconnect points and by wiring devices to multiple networks where they need to send or consume a variety of media types.

Every IP De-Encapsulation Is a Unique Event, and not All Endpoints Are Guaranteed to Perform De-Encapsulation in an Identical Way

Areas where performance may vary are audio/video sync, route switching characteristics (seamless or not), and ST 2022-7 redundancy (hitless or not). An SDI video

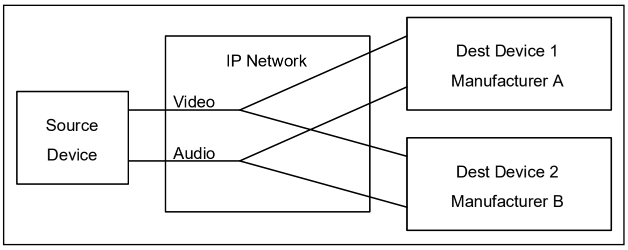

FIGURE 3. Simplified block diagram illustrating unique de-encapsulation events with differences in manufacturer implementations.

stream consists of three key components: audio, video, and ancillary data. In ST 2110, these three components are split into different multicast streams. Receiving endpoints typically use PTP time stamps to time-align the streams. However, one device, due to either intrinsic limitations or configuration error, may perform this synchronization differently than another. For a given source, we may see lip sync issues at one output but not another. **Figure 3** shows a pair of destination devices receiving the same set of multicasts from a single source device. They may perform the de-encapsulation differently.

Additionally, since different endpoints have their own internal mechanisms for genlock of de-encapsulated streams (such as in an IP to SDI gateway), reference timing for a given source may be measured differently between two different endpoints.

One takeaway from these observations is the need to be mindful about the value of quality control monitoring in an IP environment. Any encoded data streamSDI includedrequires decoding and interpretation by a receiving device. While IP video adds a new level of capabilities in a broadcasting plant, it remains an immature technology in many ways, as evidenced by the variation in de-encapsulation performance across endpoints. Understanding this variation is key to troubleshooting to maximize performance in an IP video environment.

Missing ComponentsWhat Do We Need?

As discussed earlier, the project to build Telemundo Center was relatively early in the technology lifespan of uncompressed IP, in general, and ST 2110, in particular. We faced numerous complications from incomplete or missing components during our design and build in 2017 and 2018, but many have since been addressed. The below list summarizes the remaining missing elements for a completely interoperable ST 2110 solution in early 2020.

Adoption of a Common Stream Connection Management Protocol Across Vendors

This need is directly addressed by Networked Media Open Specifications (NMOS) development ongoing with the Advanced Media Workflow Association (AMWA). A variety of other proprietary and niche options are available as well. But limitations on interoperability between endpoints and control systems remain a challenge for IP systems design. At the time of the Telemundo Center build, we often stayed with a single vendor family or opted to use the SDI versions of endpoint equipment with IP gateway devices, to avoid cross-vendor connection management integration. A fully compatible and broadly adopted stream connection management system is key to the success of larger installations that intend to use products from multiple manufacturers.

The Adoption of Advanced SMPTE ST 2110-40 Ancillary Data Stream Processing

It would be desirable if endpoint receiving devices could subscribe to multiple ST 2110-40 ancillary data (ANC) streams simultaneously and utilize them in a combined fashion. An example use case is the scenario in which some set of endpoints requires just closed captioning data, while another set requires both closed captioning and ANC triggers. In a fully realized IP solution, ST 2110-40 data senders would output ANC-only streams to the network. Receivers could then subscribe to any number of these atomic ANC streams and use them as needed. While the ST 2110-40 standard should allow for this scenario, it seems that no vendor has yet implemented such functionality. Indeed, there is limited existing product for IP-native ANC processing, even in a pass-through capacity. For Telemundo Center, all ANC encoding services, including closed captioning, are provided with SDI pass-through devices connected to the network via SDI gateways.

One potential interim step to a stackable/atomic ANC workflow directly to generic endpoints in ST 2110 would be a discrete subsystem for receiving multiple ANC streams for processing and explicit recombination. An ST 2110-40 ANC combiner could receive a stacked set of ANC streams and output a single stream representing the combined payload. This model of pregrooming ancillary data would alleviate the need for new stacked ANC processing features in generic endpoints while still allowing for a dynamic recombination workflow.

Audio Stream Packaging Standardization and Improved Flexibility

The ST 2110 standard specifies that devices support modes from 1 to 8 audio streams per multicast and it would presumably be left to the end user to determine the multicast channel packaging. However, in practice, endpoints produced by vendors currently support just a single channel count specification that does not necessarily match between vendors. With each vendor supporting different numbers of audio streams per multicast, handingoff of streams between vendors can become difficult, often requiring the use of an intermediate channel shuffler device to accomplish the handoff.

To allow for a fully functional multivendor environment, vendors either need to support all channel packaging options as defined in ST 2110, or they need to agree on a single channel packaging scheme across all vendors.

Adoption Across Vendors of Common Methods for IP-Based Trigger and Tally Data

In the legacy SDI environment, tally and triggers were communicated via a variety of physical connections and protocols, including general purpose input/output (GPIO), serial connections, and IP-based solutions. While such device control considerations are out of scope of the ST 2110 standard, it would be desirable in IP builds to deprecate wherever possible the need for GPIO and serial connectivity and replace those functions with IP-based solutions. One major driver for this need is the increasing use of virtual computing solutions in the place of legacy hardware appliances. For such virtualized devices, non-IP (serial/GPIO) interfaces tend to be either impossible or impractical to implement.

A variety of possible solutions exist to address this problem, including specifications address in NMOS IS-07 and preexisting proprietary protocols. Broad support for IP-native triggering and tally protocols would dramatically simplify builds from the perspective of hardware, wiring installation, and dynamic functionality.

Final Thoughts

As demonstrated by the implementation at Telemundo Center, large-scale SMPTE ST 2110 deployments are not only possible, but also provide a level of flexibility and scale unattainable with a traditional SDI broadcast plant. However, as with the adoption of most new technologies, SMPTE ST 2110 raises several considerations, namely the fundamental shift from hard-wired connectivity to a system defined by software configuration. As more broadcast engineering teams move toward SMTPE ST 2110, we expect the industry to evolve, filling many of the gaps we identified in this inaugural installation.

Acknowledgment

The authors would like to offer heartfelt gratitude to the entire NBCU and Telemundo Operations and Technology team that spent countless days, weeks, and months on the Telemundo Center build. The massive scale and innovation of this project would not have been achievable without the extraordinary dedication and expertise of the engineering teams from New York, Los Angeles and Miami. They would also like to thank the solutions providers and vendor partners who brought their resources and innovation to bear on Telemundo Center.

References

1. NBCUniversal, NBCUniversal Telemundo Enterprises Celebrates New Global Headquarters, NBCUniversal, 9 Apr. 2018. [Online]. Available: http://www.nbcuniversal.com/article/nbcuniversal-telemundo-enterprises-celebrates-new-global-headquarters
2. SMPTE, ST 2022-7, Seamless Protection Switching of RTP Datagrams, 2019.
3. Institute of Electrical and Electronics Engineers (IEEE), IEEE 1588-2008, Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, 2008.
4. SMPTE, ST 2059-1, Generation and Alignment of Interface Signals to the SMPTE Epoch, 2015.
5. SMPTE, ST 2059-2, SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications, 2015.
6. SMPTE, SMPTE ST 2110 FAQ, Accessed: 1 Jul. 2019. [Online]. Available: https://www.smpte.org/smpte-st-2110-faq
7. SMPTE, ST 2110-10, Professional Media Over Managed IP Networks: System Timing and Definitions, 2017.
8. SMPTE, ST 2110-20, Professional Media Over Managed IP Networks: Uncompressed Active Video, 2017.
9. SMPTE, ST 2110-30, Professional Media Over Managed IP Networks: PCM Digital Audio, 2017.
10. SMPTE, ST 2110-21, Professional Media Over Managed IP Networks: Traffic Shaping and Delivery Timing for Video, 2017.
11. SMPTE, ST 2110-31, Professional Media Over Managed IP Networks: AES3 Transparent Transport, 2018.
12. SMPTE, ST 2110-40, Professional Media Over Managed IP Networks: SMPTE ST 291-1 Ancillary Data, 2018.
13. SMPTE, ST 2022-6, Transport of High Bit Rate Media Signals Over IP Networks (HBRMT), 2012.
14. AES, AES67-2018, AES Standard for Audio Applications of NetworksHigh-Performance Streaming Audio-Over-IP Interoperability, 2018.

About the Authors

Steve Sneddon is senior vice president of production engineering at NBCUniversal (NBCU), where he oversees the design, build, and support of the technical facilities and infrastructure across NBCUs live production locations in New York for MSNBC and the NBC Network, Telemundo Center in Miami, FL, and Universal City Studios in Los Angeles, CA. He has nearly 20 years of experience in large-scale broadcasting technology and has spent the past 15 years at NBCU.

Chris Swisher is director of production systems engineering at NBCUniversal (NBCU), where he focuses on the design and build of new television production facilities based around SMPTE ST 2110 and related technologies. He also serves as a liaison to manufacturing partners to inform the development of SMPTE ST 2110-based technologies and control systems for use in large-scale production environments.

Jeff Mayzurk is senior vice president of operations and technology at NBCUniversal (NBCU), where he oversees production and technical operations supporting Telemundo and the companys Spanish-language content businesses. He led the design and construction of Telemundo Center in Miami, FL, which opened in 2018. Prior to Telemundo, Mayzurk held senior positions in studio and broadcast operations for NBCU and Comcast in Los Angeles, CA.

SMPTE