http://aimsalliance.org

ESSENTIAL CONSIDERATIONS FOR LIVE CONTENT PRODUCTION AND BROADCAST

Alliance for IP Media Solutions Key Considerations for Design and Operations – Updated MARCH 2018

JOIN THE ALLIANCE

The Alliance for IP Media Solutions (AIMS) is a non-profit trade organization founded by leading companies to foster the adoption of industry standards for the broadcast and media industry as it transitions from SDI to IP.

AIMS Alliance for IP Media Solutions

CONTENTS

INTRODUCTION

SMPTE has published the SMPTE ST 2110 Suite of standards for distributing audio, video and ancillary data over IP in response to the television industry's call for interoperability. Vendors worldwide have developed equipment and software in compliance with these standards, and broadcasters can now confidently pursue their transition to IP, using these guidelines as a starting point for their own operations.

The IP Transition is underway throughout the industry having leveraged the AIMS roadmap to identify and deploy equipment providing confidence in interoperability. These successes provide real-world confirmation that production infrastructures built on IP technology operate as intended, delivering benefits over traditional SDI-based facilities. As with other industries, the transition to IP technology will be swift.

This document provides help to broadcasters, smoothing their transition to IP by highlighting fundamental design considerations for the smallest studios, the largest, multi-location facilities, and everything in between. The good news for producers, directors, engineers and their staff is that traditional production workflows can remain unchanged, minimizing the need to re-train operations staff while starting a graceful migration to IP. In so doing, IP technology can be adopted while SDI-based equipment remains in operation.

AIMS
Alliance for IP Media Solutions

AIMS-Endorsed Specifications and Standards

The Joint Taskforce on Networked Media (JT- NM), comprised of four organizations; Society of Motion Picture and Television Engineers (SMPTE), Advanced Media Workflow Association (AMWA), Video Services Forum (VSF) and the European Broadcast Union (EBU) have developed a roadmap for IP technology. This roadmap is a living document, providing guidance on technology evolution. The AIMS roadmap reflects the work of the JT-NM providing clear guidance on those standards and specifications which provide the foundation for system level interoperability. Products developed in compliance with the AIMS roadmap meet the necessary requirements of the identified standards and specifications.

While constantly evolving, the main elements of the road map now make up the SMPTE ST 2110 Suite of standards for Audio, Video, Metadata and other media essence types, SMPTE ST 2022-6 for High Bit Rate rasterized video and SMPTE ST 2059-1 and 2059-2 for adopting IEEE 1588-2008 Precision Time Protocol (PTPv2) for broadcast operations. SMPTE ST 2022-7 for RTP Transport redundancy and AES67 for RTP audio essence are also incorporated. There are additional system level requirements for IP deployment. The Advanced Media Workflow Association (AMWA) has developed two important Industry Specifications; IS-04 for Registration and Discovery and IS-05 for Device Connection Management.

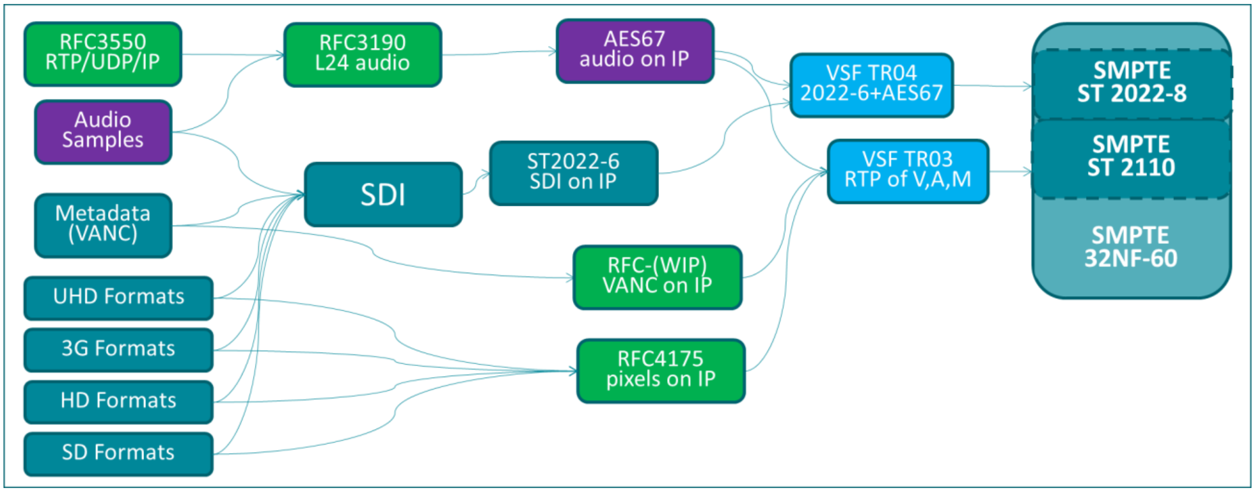

FIGURE 1: STANDARDS LINEAGE FOR SMTPE ST 2110 SUITE

As illustrated in Figure 1, the SMPTE ST 2110 standards were built from the foundation of the Video Services Forum (VSF) TR-03 recommendation as the documented practice for an interoperable IP essence transport. SMPTE ST 2022-8 (proposed) provides the bridge for rasterized video in RTP (SMPTE ST 2022-6) to co-exist in an ST 2110 system fulfilling the function of VSF TR-04.

These specifications apply to transporting audio, video and metadata essence for capture, processing, monitoring and storage. Every effort was made to leverage existing standards from

SMPTE, AES, IEEE and the IETF to ensure the best fit as well as to enable the broader ecosystem of content creation, production and distribution. For example, RFC-4175 provides the basis for how pixels are mapped into RTP payloads. AES67 is used to map 24-bit linear PCM audio and SMPTE ST 2038 is the basis for RFC8331 which is in turn the basis for SMPTE ST 2110-40 providing meta-data transport over RTP.

This standards advancement is especially timely and provides a standardized yet flexible infrastructure for production of next-generation formats such as Ultra HD (4K), UHD-HDR and 1080p. SMPTE ST 2110-20 is video format independent and also ideally suited for current HD and SD formats. Ultra HD provides a rich feature set with various standardized curves for High Dynamic Range, a richer, Wide Color Gamut palette, and different resolutions. All these different features enable artistic breadth for content creation. The resulting programs can then be easily distributed over various channels to multiple different viewing screens. In the end, the ability to monetize content is enhanced due to format agility.

Visual quality, latency and processing parameters of SMPTE ST 2110-20 all meet the high standards set for any video format, including UHD, with broad support from vendors offering SMPTE ST 2110 compatible products. It's also important to restate that SMPTE ST 2110 is format independent, and as such, may accommodate higher frame rates, image resolutions and processing accuracy as technology advances occur.

Moving the Roadmap Forward

The AIMS agenda does not end with SMPTE ST 2110. The latest ratified objective to be reached on the roadmap provides a common means of identifying and registering devices across all workflows and locations based on the Network Media Open Specifications (NMOS) IS-04 developed by the Advanced Media Workflow Association (AMWA). The processes embodied in AMWA-NMOS greatly simplify building and expanding IP production facilities by automating configuration of device connectivity in all environments, from the simplest to the most complex.

For our purposes here it's important to recognize that, no matter where you are in your decision to move to IP, the industry has achieved consensus on a foundation to interoperability that allows broadcasters to move forward with implementation of IP-based production at whatever pace suits their needs.

Figure 2 below shows the relationship between the AIMS Roadmap and the Joint Task Force on Networked Media (JT-NM) roadmap. The nature of IP technology requires these roadmaps to be dynamic and changeable, while still providing stable operating points.

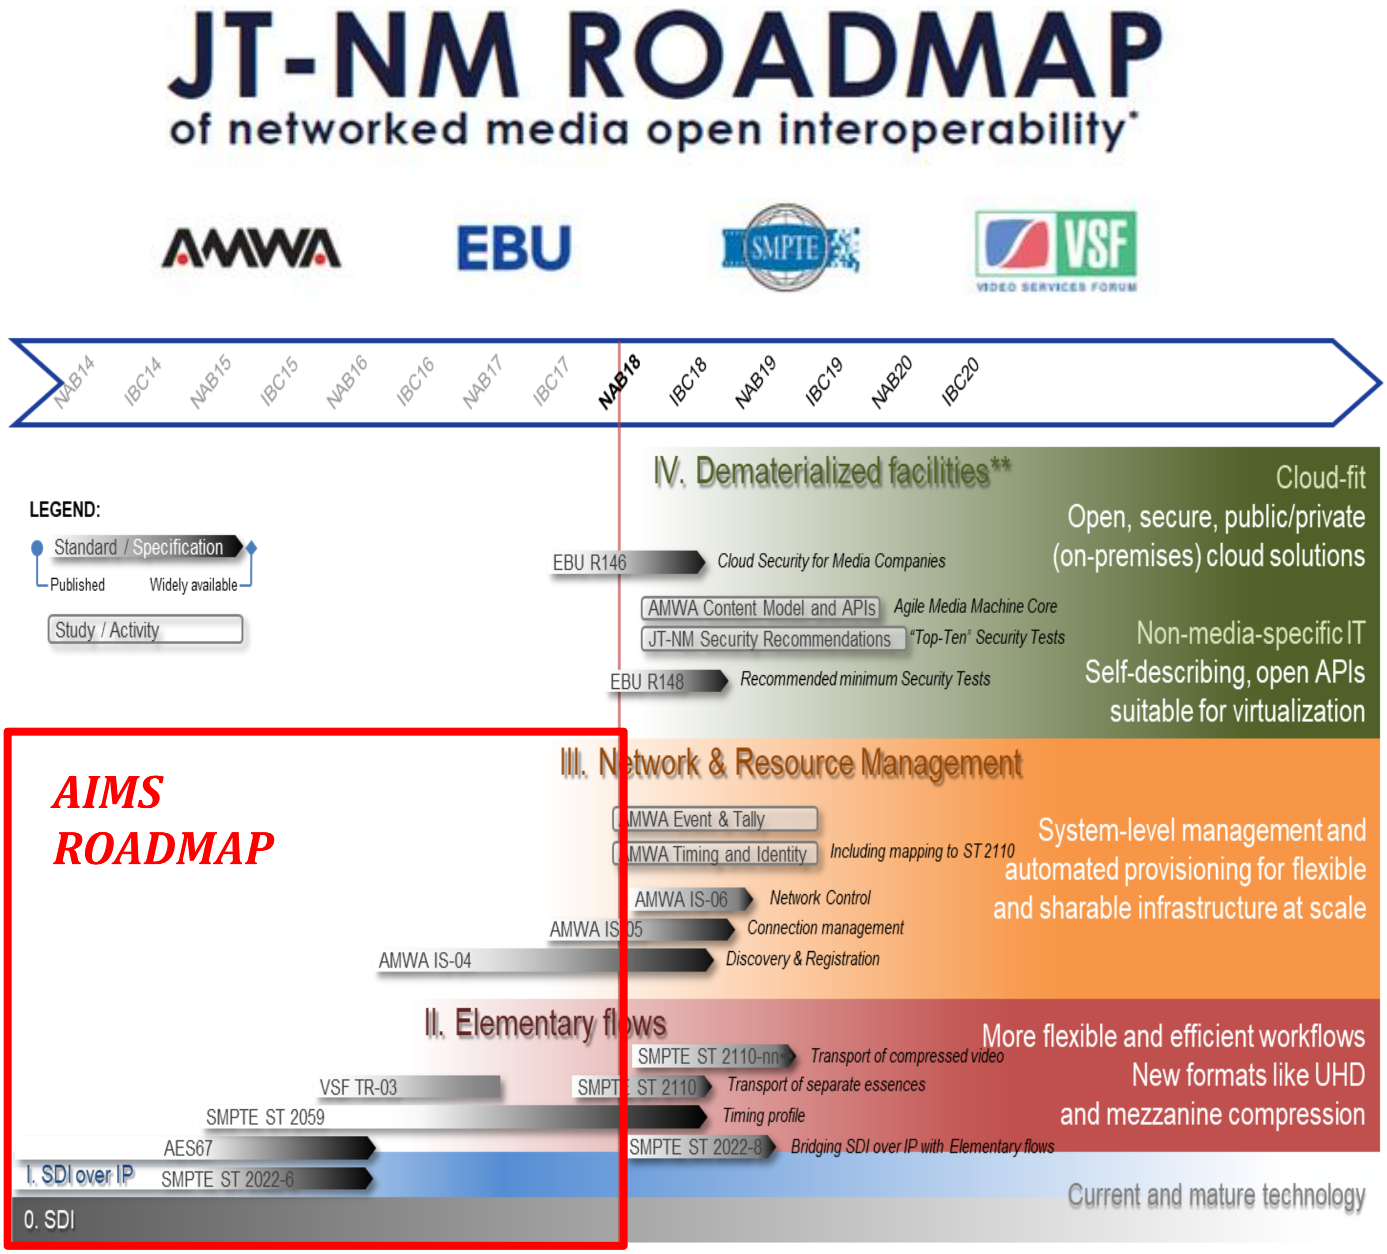

FIGURE 2: AIMS ROADMAP AND JT-NM ROADMAP ALIGNMENT

No Dumb Questions, Many Great Answers

Because SMPTE ST 2110 was developed with consideration of traditional formats and workflows, broadcasters may choose their angle of approach, starting incrementally with new islands of IP operation within their existing SDI facility, or jumping with both feet into a new green field project completely based on IP.

The following guidance is provided in part, to show fundamental considerations for a range of system designs covering small, larger, distributed and mobile systems. In a second aspect, optional design strategies may be presented for these scenarios. Again, these are guidelines which should help start the thought process. IP Technology is very agile, enabling rapid deployment of new, different workflows. Dictating a single approach would contradict the desired outcome of using IP. AIMS members are working together to provide a rich ecosystem of products as shown in Figure 3.

Audio Mixers and Processors	OTT Encoders
Branding Devices	PCI-Form Factor Capture/Playback Cards
Broadcast Encoders	Production Switchers
Cameras	PTP/Sync Generators
Encode/Decode Contribution Links	Replay Servers
FPGA IP	Routers
FPGA IP supporting ST 2110	SDI-IP Gateways
Graphics Compositors	Standards Conversion Products
Ingest Systems	Video Processing Products
Integrated Playout/Channel-in-a-Box	Video Servers
Master Control Systems	WAN/LAN Bridges
Multi-viewers	Displays

FIGURE 3: PRODUCTS ANTICIPATED TO SUPPORT ST 2110 IN 2018

ACCOMMODATING HYBRID OPERATIONS WITH SDI-BASED FACILITIES

Gateways, or devices that convert to or from SDI to IP, perform a host of functions. They provide conversion between SDI and IP transport. They provide aggregation of one or more essence streams into a 10 GbE, 25 GbE or even higher bandwidth network segment. They provide signal buffering to ensure proper time alignment. They can provide clean transitions between IP streams. They can also include CODECs for saving bandwidth. Gateways are the new glue for your facility providing a way to integrate IP islands into SDI, or SDI islands into IP. They provide the necessary capability to connect distributed studios, or distributed facilities. They provide the ability to span different PTP time domains or manage addressing across different control domains. They enable nearly any facility design you might consider and provide the ability to tailor costs based on your exact needs. There are many different options offered by the current marketplace, so be sure to investigate what the market has to offer as part of designing your IP project.

IP-Based Production for UHD

There are SMPTE SDI standards defining the transport of UHD signals, most frequently over one or more coaxial cables. A 12 Gbps data rate is most common for 4Kp60 UHD video today and it is possible to implement systems using 12 Gbps SDI. But higher frames rates, such as 120 fps, and wider color gamut using 12-bit depth are being proposed. Higher resolution 8K formats have just been introduced. Deploying IP technology offers more agility if you are unsure about how these format changes will impact your future. To this point, there is more than one way to manage unpredictable bandwidth growth. Buying the highest capacity IP network gear currently available is one option. Using low latency, visually lossless codecs is another option. There are a number of codecs to choose from, and SMPTE is being asked to consider a generic codec transport format as a part of the ST 2110 suite. And, there is a third option based on using groups of lower bandwidth streams which can become essential with high frame rate video. Because SMPTE ST 2110 is based on bandwidth agnostic RTP streams and uses RTP time stamps to indicate Top Of Frame (TOF), it is possible to use any, or all of these techniques at the same time, in the same network. The best answer depends on your needs. You may be bandwidth limited and working with an installed infrastructure. You may use high frame rate video, perhaps for replay, and need to accommodate its higher bit rates for a small number of signals in your plant. Every facility has budget constraints.

UHD production usually requires both high dynamic range (HDR) and standard dynamic range (SDR) program output formats and may rely on both HDR and SDR source material. Processing for up and down mapping between these formats is an additional requirement. Monitoring signal flow is also critical. Consideration for Wide Color Gamut (WCG) is another factor. It will be necessary to balance the need for HDR/WCG monitors and their high band width signals with SDR monitors using lower bandwidth. And in a third aspect, there will be a need to flexibly arrange workflow based on the visual monitoring requirements for a given production. Typically, monitors capable of WCG and HDR intensity are used only for the most critical viewing locations. Multi-viewers are used uniformly to provide high quality viewing where either a larger number of images must be seen together, or where budgets dictate fewer monitors. Multi-viewers enjoy a big boost with IP distribution. The fact that an IP link carries multiple videos, the number of which can be changed based on their bandwidth, means that monitoring systems are more flexible and easier to configure than ever.

As part of remote productions, contribution feeds may be compressed and carried over IP, using a Specification such as VSF TR-01:2013. While not a part of the studio per se, this studio contribution link can be managed with a gateway function as well. This will be discussed in more detail starting on page 19 but, is pointed out here for consideration in the overall system design and budget process.

There are a number of options for integrating and managing remote feeds, signal monitoring, SDI to IP conversion for various bit rates including UHD, into the system design as bridges or islands. Gateways can provide the conversion from SDI to IP as well as aggregation of multiple streams into high capacity links: 10/40 GbE, 25/50 GbE and 100 GbE. This helps provide a bridge between UHD IP islands, or IP UHD facilities, back to legacy SDI equipment. Here again, gateways provide conversion between transports and any necessary encode and decode functions. Figure 4 shows a basic idea of this.

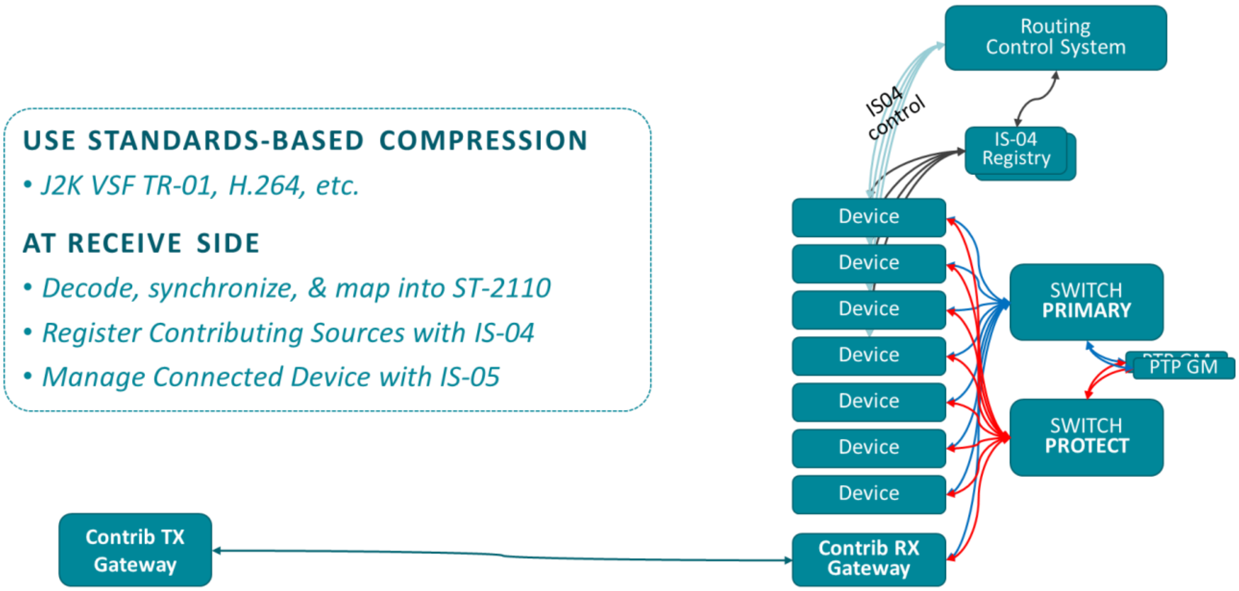

FIGURE 4: GATEWAYS CAN ENABLE REMOTE CONTRIBUTION

Because there are a number of options, you may consider using a spread sheet to calculate the costs, and the differences. Be sure to consider the cost per essence stream as one figure of merit to evaluate design options based on different signal flow topologies.

It might seem obvious, but the decision to migrate to IP depends on the ratio of IP to SDI equipment. The more IP connected equipment included in the facility, the more likely it is that an IP router be used. Additional requirements, such as the need to develop a campus environment, or receive remote contribution feeds may also be more easily facilitated by with an IP routing infrastructure.

ROUTING SIGNALS IN COMBINED SDI AND IP FACILITIES

Traditional Router Control system functionality will be inherited by new, Logical Broadcast Controllers. These new products will typically provide the interface between traditional router control surfaces, SDI routers, IP Routers and managed IP endpoint devices, such as gateways. Incorporating traditional router control surfaces into new systems reduces operator training. Controlling IP routers and Gateways provides the ability to track bandwidth and utilization to properly program streams and may apply QoS priorities for the combined network of IP and SDI devices. The logical broadcast controller also provides the ability to manage every router port with a uniform policy. If desired, it can oversee every connection. Working together, IP and SDI routers can ensure that the only connections made, are the ones requested by and needed for workflow operations. Endpoint device connections are authorized by the logical controller, rather than the IP router itself. This prevents over-subscription and also provides an additional layer of access control.

The logical broadcast controller can support either, or sometimes both, Internet Group Messaging Protocol (IGMP) or Software Defined Network (SDN) operations. They manage connections between devices, both physical and virtual, and the bandwidth required. The bandwidth of these connections is provisioned to prevent oversubscription. They also provide the ability to control every connection in the system. No connection is enabled, or executed, without approval by the logical broadcast controller. Company operational policy can be distilled into automation rules, which in turn, can then be carried out by the controller. With IP transport, devices, including the router itself, are free to share monitoring data. This data is then used by the logical broadcast controller to manage the network, connected devices and provide status and alarms based on correct system functionality

Some equipment may provide different physical ports for essence media and control data. Video data constantly consumes large amounts of bandwidth. PTP traffic consists of short messages. Combining these traffic types in one network segment can lead to packet jitter called Packet Delay Variation (PDV). Using two separate networks may help reduce this undesired effect.

IGMP AND SDN CONTROL STRATEGIES

The Role of IGMP in IP Routing

Internet Group Management Protocol (IGMP) is a proven, standardized method for managing multi-cast streams. The general approach is that every sender in a broadcast IP system emits a multi-cast stream. A receiver can join a particular stream using IGMP requests. In this manner, a stream is forwarded, and may effectively be broadcast to any receiver requesting it, emulating traditional broadcast router behavior. Signal routing is achieved through the action of the IP router processing the IGMP command to ensure the correct stream is joined, and the correct packets are forwarded. This join does not occur at a specific point in time, such as the SMPTE RP-168 switch point. Therefore, the receiving equipment manages possible gaps in the streams. Additionally, it is the receiving device which needs to join the stream, but it needs something like the logical broadcast controller to initiate the join. This ensures there is adequate port bandwidth available for the stream and regulates the validity of device port connections in the facility.

In the event that the IP routing network is composed of more than one IP routed sub-net, it may become necessary to use Protocol Independent Multicast (PIM). PIM provides multi-cast routing information across the network to provide IGMP functionality to every possible client, or device. IP routers may provide additional bandwidth and path management resources internally.

Using IGMP relieves the logical broadcast controller of the need to directly manage the IP router. However, it is still possible for the IP router to communicate monitoring and status data regarding its operation and the performance of the streams flowing through it. In this way, the logical controller can access the state, or configuration of the connected network using a known API

The Role of SDN in IP Routing

Software Defined Networking (SDN) refers to a number of network technologies providing more flexible and agile networking configuration and provisioning for enhanced performance and monitoring. In this case, the logical broadcast controller could directly control the router to establish the packet forwarding paths between the ports of the router, or even, the ports of a multi-router fabric.

The logical controller would still manage bandwidth per port to prevent over-subscription, as well as manage port access by policing connection requests, among other system level tasks. By using this direct control approach, the processing latency of IGMP requests can be mitigated.

Endpoints can be controlled, but they can also be free running. For example, a camera could always provide a multi-cast output stream. With the SDN approach, it is only necessary to change the router configuration so that packets are forwarded to the necessary ports of the router. The receiver can join whenever. Alternatively with SDN, the router paths can be established, thereby reserving the bandwidth. A second command then instructs the camera to start sending, and the receiver to join. In this case, a number of paths may be established in advance of any action taken for senders and receivers. The logical broadcast controller must ensure that the necessary bandwidth is reserved when the paths are created in advance of the new signal flows being sent and received.

Receiving devices, such as gateways, will need to manage the impact of these different switching options. A receiver will need to drop and then join incoming streams. In the case of leave before join, there is some gap in the stream, which can be masked. Alternatively, in the case of join before leave, extra bandwidth is required to support both streams, at least for some period of time. And finally, a third approach would be a hard cut, much like a non-RP168 switch in today's SDI routers. All three of these cases are solved by using adequately sized memory buffers in combination with intelligent stream management.

ADDRESSING THE AUDIO QUESTION

IP audio may be used with, or without IP video. Both SMPTE ST 2110-30 and AES67 provide a standardized basis for audio essence transport over RTP. If fact, IP audio may be used together with SDI video, when the value of supporting both infrastructures is clear.

IP audio allows flexible grouping of audio channels into logical stream bundles, i.e. 8 channels of linear 24-bit PCM audio for a surround sound essence. But it also supports single-channel streams for maximum routing flexibility, or trunks of 64 channels (or even more), similar to MADI. The logical broadcast controller can support all of these different transports based on audio workflow needs. Live audio productions, for example in an OB van, often use many mono channels. Before these were mono inputs and a MADI, or other multi-channel output. Hybrid video routers today are typically based on 8, or 16, channel audio groups which are broken out into flat, mono-audio signals for routing. But, often, workflows are based on logical groups of channels, such as a surround sound master.

The key considerations for design are based on channel count and signal management. A completely flat, mono channel system has the highest stream count, and therefore consumes a larger number of IP addresses and a higher total bandwidth. But, each individual stream has the lowest bandwidth. Approximately 1.5 Mbps is needed for one, mono audio channel. Multiplexing audio channels into one RTP stream with perhaps, 8 or 64 channels, reduces the number of IP address, and increases the bandwidth for that particular multi-channel stream, but reduces the total bandwidth requirement compared to a single-channel streaming equivalent. Logical broadcast controllers can manage the concepts of trunks with breakout. The IP audio router should be evaluated for its ability to manage the scenario you need to implement.

Another consideration for audio is latency. AES67 calls for a 1 ms interval between packets; also referred to as packet time. While SMPTE ST 2110-30 also mandates support for this 1 ms packet time, it provides optional support for a 125 us packet time. This faster rate is useful in protecting lip sync in live production where there is a combined audio video workflow. Also, with the proposed SMPTE ST 2110-31, AES3 signals may be tunneled within an SMPTE ST 2110 compatible transport providing system workflow which includes encoded audio. Care should be taken to manage time stamp alignment and arrival if accurate latency compensation is required for subsequent processing. Different coding algorithms exhibit different performance.

AES67 was developed in advance of SMPTE ST 2110-30. The two standards are not identical, but SMPTE ST 2110-30 was based on AES67. Most AES67 vendors provide a SMPTE ST 2110-30 mode of operation as an option. Basic PTP settings and the use of SDP constraints must be accounted for to ensure interoperability, as an example, as documented in AES R16, '*PTP parameters for AES67 and SMPTE ST 2059-2 interoperability*'.

A SIMPLIFIED APPROACH TO REDUNDANCY

Redundancy and resiliency are foremost design considerations for systems of any size. The advance of IP technology into routing systems has led to the wide adoption of IT network redundancy models. SMPTE 2022-7 provides guidance for seamless protection switching between RTP essence streams. SMPTE 2022-7 compliant devices essentially include a 2 x 1, smart selection switch which examines the arriving RTP packets for accuracy and selects the correct, un-errored packets from each stream for downstream use. Including the secondary switch in the receivers, and the 1 x 2 fan-out in the senders makes it easy to implement fully redundant A and B networks. It is also possible to implement partially redundant topologies for cases where back up is not required for every signal in the facility. Figure 5 illustrates a simple example of this.

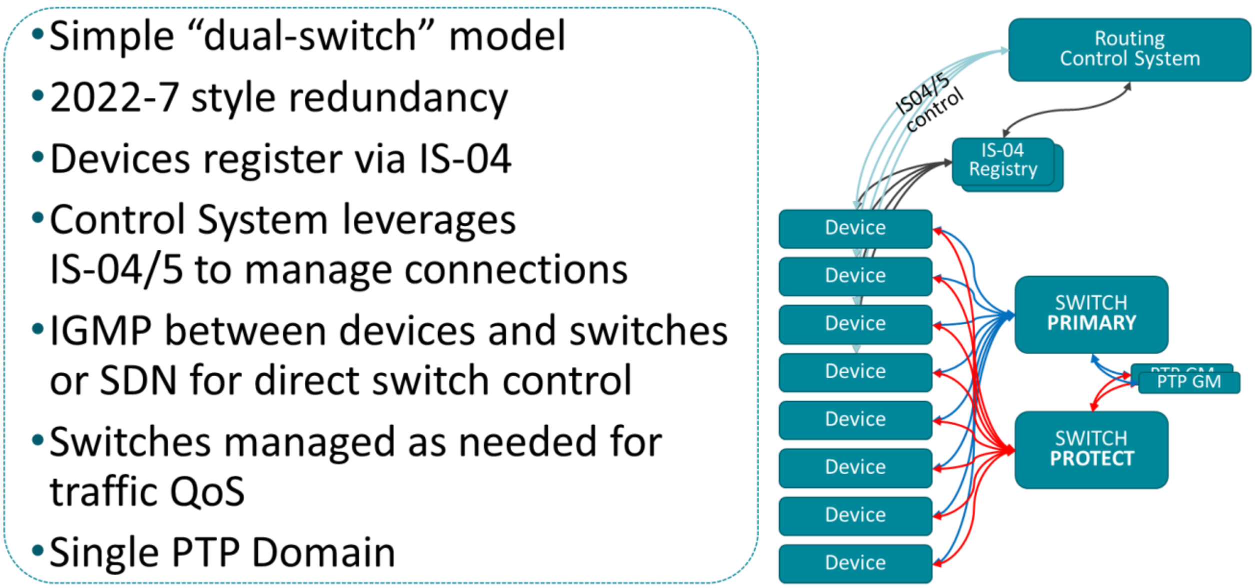

FIGURE 5: NON-BLOCKING, REDUNDANT CORE ROUTER SYSTEM

These smart packet selectors, or end nodes, do take time to make a decision and they must compensate for differential delay in the A and B network. SMPTE ST 2022-7 provides for 3 classes of differential path delay, or PD. Up to 450 ms is supported for WAN applications and 10 ms is supported for LANs. There is also a proposal in front of SMPTE supporting ultra-low latency PD for intra-facility use where the PD should be much less than 1 ms, SMPTE ST 2110-10 provides guidance for IP addressing conventions for these redundant streams.

Redundant Grandmaster clocks will be discussed later but, care must be taken with system design to insure timing packet delivery to downstream nodes and endpoints using dual physical connections for hitless switching, or redundancy. ST 2022-7 does not provide guidance on this.

Finally, ST 2022-7 provides a means to perform maintenance, install upgrades and if necessary, even replace failed equipment, while remaining fully operational.

FINDING AND USING STREAM ENDPOINTS: AMWA IS-04 AND IS-05

Device Discovery and Registration

In traditional SDI routing, devices which output SDI are wired to input ports (sources) on the matrix, and devices which expect SDI are connected to output ports (destinations) on the matrix. All signal switching happens inside the matrix, at the direction of the control system. The names (identity) of the sources and destinations are tied to the physical input and output connections of the matrix, through a routing database built when the system was wired.

In ST2110 IP systems, many different signals share the same physical connections; the identity of a source is NOT tied to its physical connection, but instead is associated with the sender and its associated stream details. The identity of a destination is likewise associated with a given stream receiver.

AMWA IS04 provides a dynamic linkage, a registry, between stream endpoints (senders and receivers) and the control systems which manage connections. When a streaming device comes up, it looks in the local Domain Name Server (DNS) to find an AMWA IS04 registry using the same technique your computer uses to find google.com. The endpoint then "registers" itself into the registry using a standard HTTP-based method, including a globally-unique identifier (GUID) for each sender and receiver. Routing control systems look in this dynamic registry for an inventory of devices, including their management details and supported protocols, and associate the GUIDs of the senders and receivers with named sources and destinations in the routing database. The IS04 method can be used in combination with current methods and protocols in a mixed environment as it gains widespread adoption.

Once the control system, or logical broadcast controller, knows about the stream sources and destinations (whether through IS04 or other methods), it needs to support signal routing transactions. In a traditional SDI system these switching operations were done inside the routing matrix using a protocol unique to the matrix manufacturer. In SMPTE ST2110 systems, streaming endpoints are made by many different manufacturers, and a common protocol is needed to tell receivers what stream they should take from the network in addition to forwarding packets correctly across the IP routing fabric.

AMWA IS05 is a standard method for the control system to track all the stream details for each sender, and to inform the intended receiver about the details of the stream it should receive next. Using the Session Description Protocol (SDP) objects defined in ST 2110 to capture the stream details, IS05 provides a basic "universal driver" for streaming endpoints, avoiding the need to create custom drivers for proprietary protocols for every device in a system. See Figure 6.

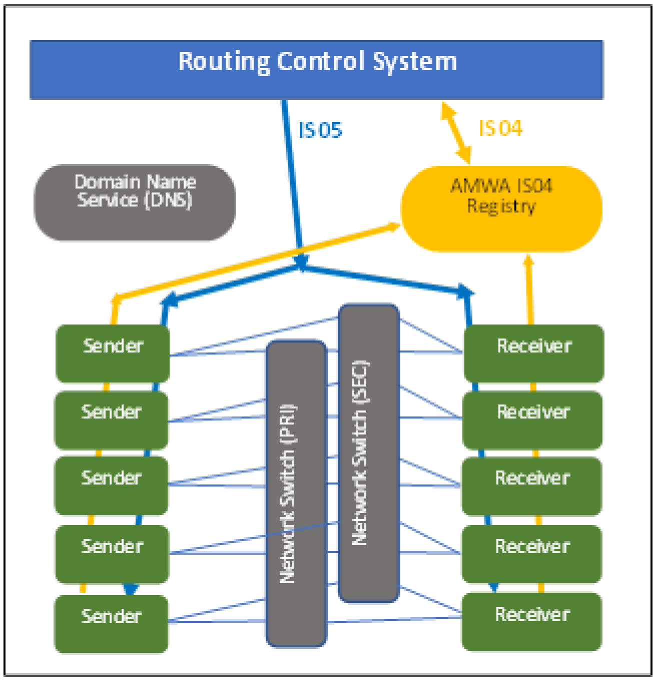

FIGURE 6: IS-04 AND IS-05 PROVIDE REGISTRATION, DISCOVERY AND CONNECTION MANAGEMENT

Taken together, AMWA IS04 and IS05 complement ST2110 to allow construction of complete, useful media signal routing systems.

SUPPORTING SYNCHRONIZATION

Time Stamps are the basis for plant synchronization. SMPTE ST 2059-1 and SMPTE ST 2059-2 provide a standardized approach for using the IEEE 1588-2008 Precision Time Protocol (PTPv2) for providing and distributing facility timing. SMPTE ST 2110 uses RTP time stamps in a specific way so that the traditional alignment points, such as Time Of Frame (TOF) are clearly and uniquely defined. These RTP time stamps are referenced to PTP so that the frequency and phase of any essence stream is known, so that any number of streams can be locked to each other and be kept in phase. A real time traceable reference maybe extracted from GPS providing an economical link to traceable real time via UTC. The result is a synchronized system referenced to an international standard.

Time stamps provide a durable, stream specific piece of data which is transported throughout the signal chain, and facilitates things like lip sync, and easy video processing alignment. Facilities may now be timed by using time stamps, rather than measuring lengths of coax, as was the case when using “Black Burst”. Distributed facilities can more easily re-align essence streams where processing is needed. Physics cannot be violated, and for longer distances, delay must be compensated. But in a studio, where streams are typically synchronized, and path lengths are relatively short, using time stamps can significantly reduce the number of SDI frame syncs in a given signal path, thereby reducing overall latency.

A PTP Grandmaster (GM) is like an SDI Sync Pulse Generator (SPG). It provides for plant synchronization and timing using digital time values sent in compliance with Precision Time Protocol. In addition to providing house sync, the Grandmaster can lock to GPS providing synchronization to a traceable source. The Grandmaster may include video sync pulse outputs providing cross lock to traditional SDI equipment. Grandmasters can also be redundant. Rather than use an SDI change over switch, there is a Best Master Clock Algorithm (BMCA) used by 2 or more Grandmasters, to determine and mutually share, their health and status. As part of this algorithm, one Grandmaster is promoted to be the true master, and the others become slaves or essentially backups to the master and are ready to take over if necessary. See figure 7.

- **Sync interval = 3 (8 Hz)**
- **Announce interval = 0 (1 Hz)**
- **Announce timeout = 3**
- **Communications Mode = Multicast or unicast without negotiation**
- **All end devices set to “Slave Only” mode**
- **For BC routers, all interfaces except those connected to the GM, are set to “Master Only” mode**

FIGURE 7: TYPICAL PTP PARAMETER RECOMMENDATIONS

A device with PTP capability can precisely lock to the network time being distributed as per the protocol and generate its own Ordinary Clock (OC). Systems can be built using PTP where only the edge devices are locked, and the streams flow through the routing fabric. This is actually quite similar to how SDI facilities work today. SDI routers may re-clock the digital signal, but they do not manage any latency or relative time alignment between signals.

IP Routers may implement Boundary Clock (BC) functionality. PTP messages are just that; digital messages whose jitter may increase as they traverse through packet routing networks. If the Packet Delay Variation (PDV) is too large, it will negatively impact jitter of the locally generated media clock. A Boundary Clock exists in a network to act as an intermediate point for the PTP protocol. BC is also utilized because of scalability and it can improve synchronization accuracy if all the switches in the network use BC. Here, the signal is “re-clocked”, and new messages whose PDV is greatly reduced are propagated throughout the network towards the nodes connected to the BC. A BC is a jitter, or PDV, scrubber. IP routers may include boundary clock capability to reduce PDV and ensure that very clean clocks may be recovered at the endpoints. Because a BC device is a protocol endpoint, PTP message traffic can be reduced for both the GM and other network connected devices. Lastly, BC provides some security features that can protect the network from a rogue PTP GM.

As an alternative to BC functionality, IP Routers, and other equipment which pass PTP messages through, may implement Transparent Clocks (TC). In this case, the timestamps in PTP messages are adjusted based on delays incurred traversing the equipment. The messages are modified to include offset values. The endpoint at the edge of the network will then use the offset adjusted messages, to precisely calculate the effective arrival time stamp, recover the clock and establish phase.

Based on how the router fabric is designed, there can be differing numbers of router hops. It might be necessary to include a router without either BC or TC capability. The number of devices connected to the router fabric is also a design factor for clock stability. A distributed facility has additional considerations due to latency and traffic volume between connected facilities. There may be additional considerations for the number and physical location of Grandmasters. These may include using GPS to lock Grandmasters together, different PTP domains, and BMCA operation. Attention to detail is needed but, robust, redundant facility timing is achievable.

ROUTER DESIGN CONSIDERATIONS

There are more options with IP routing, than with SDI. In SDI routing, one coax equals one SDI signal. In IP routing one port can potentially carry a number of IP streams. IP routing fabrics provide aggregation, or multiplexing, to create trunks which can be used to make even larger routers fabrics, if needed. Aggregation can also be used to multiplex a number of essence streams into a single, higher bandwidth network segment much like quad-link SDI interfaces that are used for UHD signals. For example, three 3 Gbps SDI signals along with their audio and metadata essence can be aggregated into a single 10 GbE segment. Perhaps up to seven 3 Gbps SDI signals, along with their audio and metadata, can be aggregated into a single 25 GbE network segment. This process is an important part of optimizing bandwidth in the overall router design. Gateways can provide aggregation along with conversion between SDI and IP.

In general, AIMS has supported the use of Commercial Off-The-Shelf (COTS) equipment especially with respect to IP routers and network switches. This was also heavily factored by the SMPTE ST 2110 standards committee. IP routers and network switches are being utilized specifically for the media industry especially as it relates to uncompressed video which, by its very nature, is made up of massive video bandwidths. This equipment is designed to provide very low latent operation under high traffic loads, extensive address space for multi-cast streams, and high bandwidth ports for 10/40 and 25/50/100 GbE capacities.

There are large routers, sometimes referred to as Core, or Modular routers. These are similar to traditional SDI routers in that they have modular Port I/O cards and Router Fabric cards in a common chassis. Routers can also be built using Leaf and Spine topologies. In this approach, Leaf routers act as aggregators with trunks connected to Spine routers. These Leaf and Spine routers are typically not modular. The Leaf and Spine topology is widely deployed by data centers.

It should be understood, that at least from a construction standpoint, a Core, or modular approach is analogous to a Leaf and Spine, where the leaves, port cards, and spines, routing cards, are all in one chassis. It is not uncommon to use a Modular router as the Spine element in a Leaf and Spine topology. It is instructive to note that port cards often contain routing capability.

There are many options to make larger systems, and to provide future scaling for port count and increases in bandwidth driven by advances in Ethernet technology. Clearly, any router purchased today, has some fixed inherent bandwidth capacity for its internal routing backplane. But, the beauty of Ethernet and IP routing is aggregation. A newer, higher bandwidth Core, Leaf or Spine, will accept aggregated signals from previously installed equipment. This provides the ability to scale up in the future without using a fork-lift to replace all the equipment.

The rules for designing system structured interconnect to scale now include the need to consider leaving ports, and the desired bandwidth, unused. In the future, these ports on Modular I/O cards, Leaf or Spine routers, are used as aggregation links to the additional routing capacity being installed.

In the case of a Core, or Modular router, the capacity maybe increased by adding cards, or by adding newer cards with higher capacity, up to the limit of the chassis backplane.

In the case of distributed systems when the aggregation trunk may become very large, flexible tie lines between routers may be installed. These may even span between diverse, geographic locations.

Gateway equipment and routing equipment may be rolled up into a common spreadsheet for analysis. Recall that the best figure of merit is based on optimizing the cost per stream supported by the overall routing, aggregation and conversion system design. Be sure to include the cost of interconnect in this system core. This includes the active optical cables (AOC). The other equipment in the system is now an endpoint attached to the fabric. The bandwidth of this equipment does impact port count so these AOCs should be included in the router fabric cost. The endpoint devices such as cameras and production switchers, may likely be considered separately.

Non-Blocking, Rearrange-able Non-Block and Blocking

Any router maybe designed for a given blocking requirement. There is no inherent reason a Leaf or Spine, or a Modular router, or a combination of them all, cannot be non-blocking. See Figure 8.

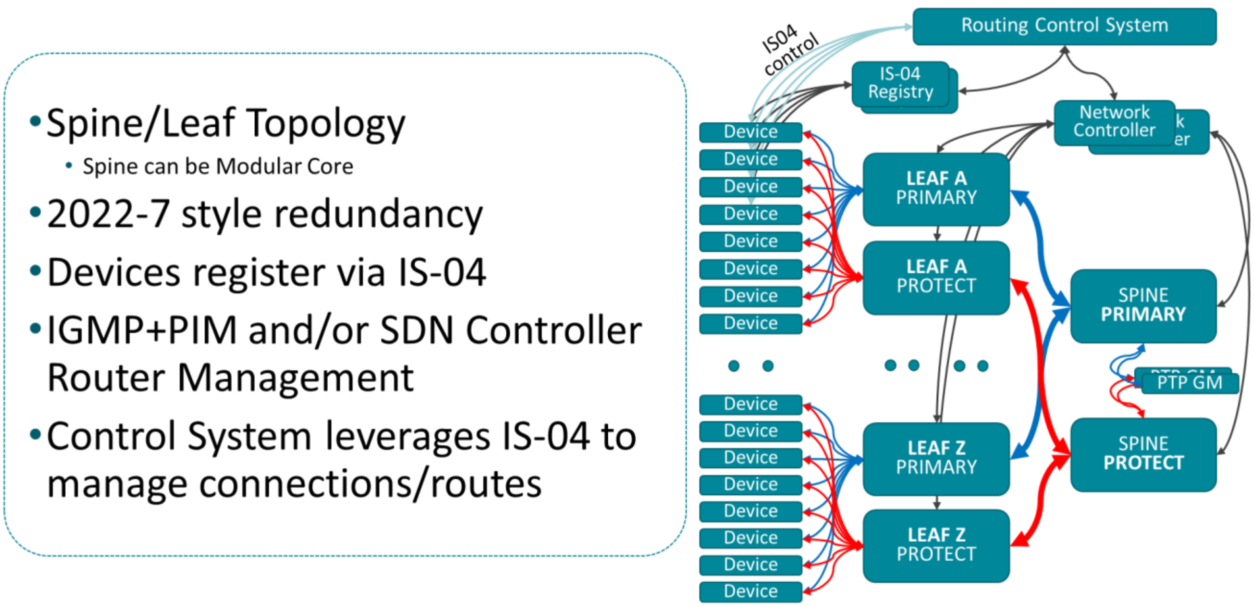

FIGURE 8: LEAF AND SPINE TOPOLOGY CAN BE NON-BLOCKING

There are three requirements for ensuring 100% non-blocking capability. First, the forwarding address tables of each router must be large enough to support the number of streams flowing through it. Second, the internal memory used route the data must be large enough to support the desired number of streams and their given SMPTE ST 2110-21 traffic profiles. Third, the bandwidth of an uplink trunk from a router, connected to a second router (analogous to Leaf and Spine) must be equal to, or greater than, the total bandwidth which could be routed to it by the network tributaries feeding the router generating the uplink.

Memory and bandwidth used together are equivalent to SDI crosspoints. The above rules ensure you have every SDI crosspoint equivalent and they can be controlled.

In many cases, it may be acceptable to have the flexibility to rearrange non-blocking routers. One example would be using aggregation trunks at tie lines. An IP router may have 3 Tbps of capacity, but you need only share 500 Gbps of capacity between two production areas. Unlike SDI tie-lines, which are very rigid, this trunk bandwidth may be flexibly provisioned to accommodate varying flow counts.

In a different example, there may not be enough memory, or bandwidth to guarantee 100% non-blocking in the IP router. This can occur after a large number of multi-cast routes have been created. However, there can still be enough extra memory and bandwidth so that if a block is detected, the stream may be routed around the block by rearranging other connections to traverse different paths. This must be done carefully so as not to interrupt a stream resulting in dropped packets, or unacceptable latency shift which might exceed a receiver's buffer.

The logical broadcast controller will play an integral role in configuration, control and monitoring of dynamic bandwidth allocation in the router fabric, tie-line trunks and edge device control as part of meeting non-blocking goals for the overall router fabric.

It is important to discuss capacity and potential blocking issues with your preferred switch provider.

DESIGNING PRODUCTION FACILITIES EXTENDING OVER MULTIPLE LOCATIONS

IP-based production can be integrated across multiple facilities with all the benefits that accrue in single-location environments. Producers can collaborate in real-time throughout dispersed locations, across campus, towns or continents, over linked, high-speed terrestrial backbones. In many cases, these network connections are fast enough to provide real time "feel" for human operators. The range of facilities which can be integrated into distributed networks includes mobile operations, fly packs and fixed studios of any size.

The same technology is used for distributed systems but, some parameters require extra attention. Synchronization, latency, encoding and addressing are areas of interest. Gateway functionality, either as a standalone equipment, or integrated into another device can be provided to manage these features.

Synchronization between locations is facilitated using PTP, and the RTP time stamps in the signals. In this way, every device is locked to a common clock, assuming physically separate Grandmasters are locked to GPS. Even with frequency lock, phase lock, or alignment, may be hard to achieve due to the absolute propagation delay across the network. Offset values can be measured and used to manage differential latency. Actual offset between the essence streams is only eliminated if an identical amount of delay is present. This means buffering the early signals to line up with the late ones. Asynchronous signals, those not locked to a common time base during their creation, will still require the use of frame syncs. See Figure 9.

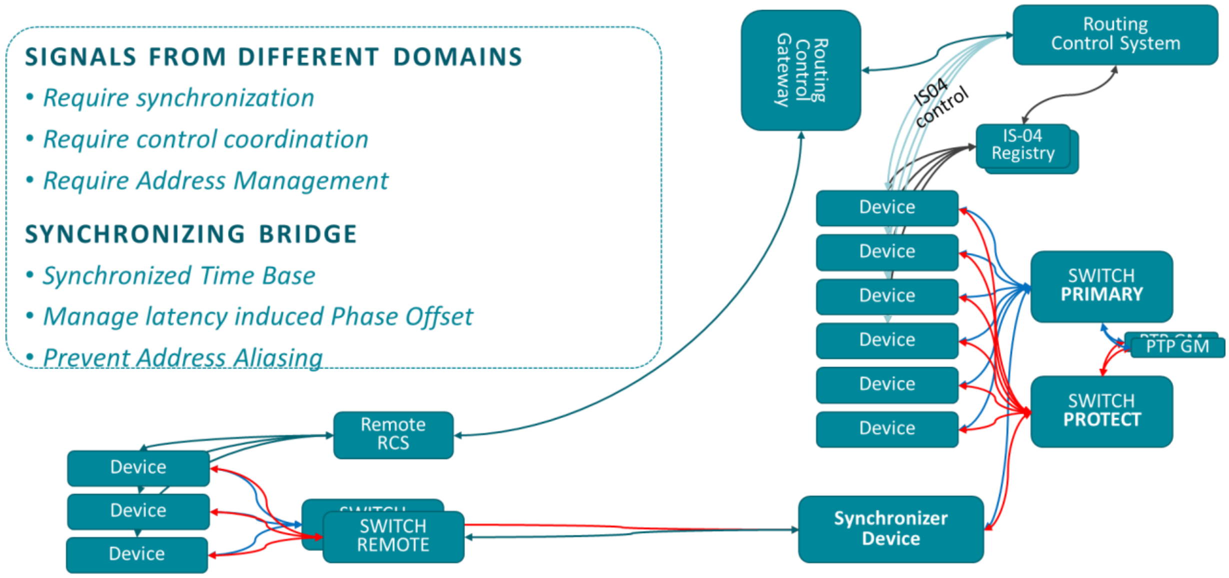

FIGURE 9: MANAGING REMOTE DOMAINS

Long distance signal contribution may include encoding to save bandwidth. This function results in delay, which must also be compensated for.

There may be the need to translate IP addresses of arriving signals to ensure unique identity in the local domain. Network Address Translation (NAT) can be used for this. The broadcast controller and a gateway card, or router, maybe used to carry out this function to ensure unique addressing is provided.

Finally, if the goal is to create a completely unified facility, distributed across various locations, they should all share a common registration and discovery database. The AMWA IS-04 specification provides guidance for this. And, there are numerous ways to synchronize distributed databases such as might prove useful for a distributed system. And again, the logical broadcast controller is an invaluable tool in making this information available to all the devices attached to the network or spun up on a node within the network.

Disaster Recovery is another design aspect for distributed facilities. In this case, decisions should be made about which equipment is located where, so that in the event on facility is not fully operational, another facility, or facilities, can cover, or take over for it. Such a resilient mesh replaces the concept of 1 + 1 local redundancy. However, very careful network design is required to ensure uninterrupted connection between all locations in the mesh. The alternative is to consider including localized grandmaster clock sources, storage capacity and compute capacity so that each facility could run as if it were completely independent.

Modern software techniques would enable something like the logical broadcast controller to operate as a distributed process, caching critical data in distributed network attached storage. Using BMCA, a grandmaster capable source could be located at each facility, and locked to a GPS signal.

CONCLUSION

IP Technology provides a richer tool set for designing facilities which provide every capability of today's SDI designs. SMPTE ST 2110, SMPTE ST 2059, AES67 and AMWA IS-04 and IS-05 provide key milestones on the AIMS roadmap where standards existing and extensive, successful, interoperability testing has occurred, and more will come.

There are many more system design options than before, and as such, it is important that any design start based with fundamental analysis. Perhaps the first decision is whether or not the facility is distributed and whether disaster recovery should to be considered a part of the distributed operational model. The number of streams, type and bandwidth will help size the capacity for routers and trunks between facilities. Next might be the consideration for endpoint devices. Will they include an IP interface natively, or need a gateway? If virtualized processes will be used, general compute resource requirements should also be included.

This overall view should help focus the requirements for synchronization. Port counts, workflows and distribution needs provide clarity for decisions on the use of OC, BC or TC, and how to provide for solid BMCA operation. This can then funnel back into which router topology to use, and which type of routers you need to implement your fabric. Finally, the Logical Broadcast controller presides over the system. Its feature set of capabilities should be clear. And the scale for each of these, or its capacity should also be clear.

While not a solution, this paper should have provided guidance on the technology you need to leverage for the design of your new IP system, and hopefully, it presented options for you to consider. Undoubtedly, your design team will have more questions but, there are many AIMS member companies ready to help you compare alternatives and finalize the best plan for your needs.

CONTACT US

Tina Lipscomb - ADMINISTRATOR
23117 39th AVE SE
Bothell, WA 98021 US
+1 4258706574
tina@aimsalliance.org

AIMS
Alliance for IP Media Solutions