SMPTE Meeting Presentation

Practical Transition Strategies of SDI Facilities Utilizing Newer IP Baseband A/V signals

Scott Barella, Chief Technical Officer

Utah Scientific, Inc.

Written for presentation at the

SMPTE 2016 Annual Technical Conference & Exhibition

Abstract. *Operators who have significant investments in large-scale SDI plant are going to face technical and budgetary challenges as they transition to newer IP-based topologies while still maintaining current SDI workflows. This paper has been written to provide some practical guidance for operators, system designers and manufacturers, with a focus on constructing 'islands' within current infrastructures. Given that SDI is the predominant topology, the center of these plants lies within SDI routing cores. Therefore, any transition will have to address key SDI architectures and augment an existing plant. The paper addresses a number of key parameters such as control systems, SDI-to-IP conversion, and 12G SDI. The conclusion of the paper points to an 'island' approach as a practical means of providing a low-risk transition to the world of uncompressed IP video, audio, data and timing.*

Keywords.

(The SMPTE disclaimer is on a footer on this page, and will show in Print Preview or Page Layout view.)

Introduction

The push to transition from SDI to uncompressed IP audio and video has triggered development of a number of approaches over the past few years, as implementation catches up with conceptual thinking. Standards-based methodologies are beginning to emerge, such as SMPTE 2022-6, TR-03 (SMPTE 2110 draft) and TR-04. There are also a few other offerings, such as ASPEN and NDI.

As the details of these IP solutions emerge, operators must now consider how to manage their plant operations and structures in a completely new environment that will include new cable, new connectors, and new management of the operational plant. However, operator workflows will most likely remain the same, at least in the short to medium term. So the question is how will operators navigate this new IP technology and implement IP in their current SDI plants that have served them well? This paper will address the most practical means to accommodate this transition and usher in a new way of thinking about the manner in which broadcast professionals deal with uncompressed audio and video over IP.

Background

Current SDI Topologies

It's a given that many operators using large SDI core routers will not get rid of them overnight, so it would be logical that these core SDI devices will remain as part of most transition plans. Many operators have investment in very large routing cores that are 1056 inputs by 2112 outputs. These systems are not only extremely large but also present an enormous challenge in terms of dissecting individual signals to be selected for IP conversion.

The mainstays of today's video operations are SD-SDI (270Mb/s), HD-SDI (1.5Gb/s) and 3G-SDI (progressive formats at 3Gb/s). When converted to uncompressed IP, these signals fit nicely into the 10Gb/s routing lanes of current IT infrastructure switch ports. However higher resolutions of UHD-1 (12 Gb/sec) and UHD-2 (24 Gb/sec) are becoming more prevalent, among outside broadcast (OB) operators in particular, and right now it remains unclear as to how to approach these larger signal payloads, given that the payloads exceed the increasingly common 10G/s backbone standard. While the topic of emerging UHD-1 and UHD-2 carriage leads to the discussion of mezzanine compression and newer 25Gb/s switch ports, this paper is going to focus on transition plans for the current uncompressed HD/3G formats.

While most baseband equipment is supplied with SDI I/O's, few vendors supply equipment with native uncompressed IP I/O's. For example, cameras are just now starting to show IP I/O's, but multiple and potentially incompatible IP implementations have been created because the broadcast market has been unclear about the type of IP to incorporate in broadcast equipment. The proliferation of various uncompressed IP formats such as 2022-6, TR-03, AVB, and others has led to indecision about IP and therefore manufacturers have avoided placing IP I/O's in their equipment. Given that some production and broadcast equipment has IP while other equipment does not, this inconsistency is something system engineers should consider as part of the transition plan from SDI to IP. Clearly, as time goes on, more broadcast equipment IP I/O's will appear more frequently, especially given more market clarity as to the specificity of the type of uncompressed IP formats such as 2022-6, TR-03 or both for example.

IP Topology Choices

SMPTE 2022-5/6/7

The SMPTE 2022 family of standards started with compressed MPEG-2 video and audio in the form of 2022-1 which was primarily aimed at creating an IP version of the Asynchronous Serial Interface (ASI) signal. Because of the 'bursty' nature of IP Ethernet traffic, it became clear to early designers of this standard that even though the IP packets were sequenced using Real Time Protocol (RTP) packets, there was a real danger of disrupting the video streams with out of order packets. In an IT environment, data can appear in bursts, but that is not a good approach in the deterministic and time-critical world of video. As an answer to this problem, a Forward Error Correction (FEC) method was used to send redundant packets in rows and columns, effectively creating the basis of 2022-2 FEC. When 2022-6 was included in the family, it used a similar method of encapsulation but used uncompressed audio and video from SDI rather than compressed video from ASI.

SMPTE 2022-6 is currently the most predominant IP format for uncompressed audio and video. When introduced, it was also one of the first formats to address carriage of uncompressed video over IP. The 2022-6 implementation as an IP format addressing uncompressed video and audio began a number of years ago when development engineers continued from 2022-1 using compressed video. 2022-1 addressed the need to convert compressed ASI MPEG-2 transport streams to IP. As time proceeded, amendments to the SMPTE 2022 specification were adopted to include uncompressed SDI formats.

At the recent Video Services Forum (VSF) interoperability testing in August 2016 nearly all of the 35 companies that participated in the 2022-6 testing were able to successfully send and receive these IP signals from 2022-6 IP transmitters to 2022-6 IP receivers. Given that 2022-5 allows for using FEC, 2022-5/6 will make the decision an easy choice for operators who are carrying signals across long haul networks that require FEC as well as containing all the essential elements of Audio, Video, and Ancillary Data, that are all bound together in one unicast or multicast IP address. WAN providers that haul signals from one city to another are particularly fond of this IP choice because it fits their current business model of providing point-to-point video delivery.

The 2022-5/6 signal construct also takes into account all of the synchronization components of the SDI signal because of the fact that it preserves the SDI blanking interval along with the ancillary VANC data. Because 2022-5/6 is very close to our current SDI formats, it seems logical to many developers and operators that it should be utilized for intra-facility use. In other words, when program signals need to be carried long distance much like ASI whose audio and video are unlikely to be manipulated then the 2022-5/6 is a good IP choice for this kind of application where one composite signal contains all the video, audio and data components without the need to de-embed and re-embed the audio.

SMPTE 2022-7 allows for redundant/backup seamless switching, so for those operators who feel they need extra signal protection and redundancy, 2022-7 allows for switching between two identical IP streams with no interruption of video or audio when the switch occurs. This is a similar to the concept of switching two identical ASI streams, except that the switch it is done at the IP level.

TR-03 (aka SMPTE 2110 Draft)

In November of 2015, The VSF's Studio Video over IP committee published Technical Recommendation 3 (TR-03). The main thrust of this document was the culmination of the broadcast industry's input for creating a standard that was based on solely on IP. Furthermore the committee found that 2022-5/6/7 didn't address the true need to readily separate video, audio, and data streams. In a live studio environment using SDI, there is implicit binding of audio and video. Multi-level breakaway creates the need to constantly de-embed and embed audio, which is not ideal. Therefore the VSF committee set itself the task of defining an IP format that could manage breakaway. This is done by treating video, audio and data as separate IP streams using separate multicast addresses.

TR-03 video will utilize the internet standard RFC 4175. This standard eliminates the need for the vertical blanking interval, as well as other historical encapsulations of additional signals. It provides a means to incorporate just the video essence that creates lines of video no matter the desired resolution and frame rate.

This recommendation also incorporates the Audio Engineering Society's AES 67 for the carriage of uncompressed audio signals and allows this IP audio signal stream to be transported separately and thereby eliminating the need to 'unbind' the audio from the video as it is in the 2022 IP format. This standard took years for audio engineers to come to agreement regarding a single uncompressed standard for audio not just in paired channels but rather multiples of channels of audio.

Timing and synchronization may be one of the more important aspects to the TR-03 IP topology and the TR-03 document points to the use of SMPTE 2059-1/2. This standard is not explained easily. Its roots come from the IEEE 1588 standard using the Precision Time Protocol that emits very precise time packets in the IP stream much the way that Network Time Protocol (NTP) works - only with much greater precision. SMPTE 2059 allows for video and audio signals to be 'stamped' with this very precise time. This time stamping is not possible in SDI, so a transition to IP brings an unexpected benefit beyond genlock video timing: Timestamping will help broadcast engineers and users to address lip synchronization errors. This feature alone has many touting IP as a complete signal package; even in SDI a separate coaxial cable carrying analog composite black signal or tri-level digital sync signal must be used to synchronize devices in a broadcast plant.

Ancillary data will also be a separate essence that is not bound to either audio or video in the way that current VANC data is. In SDI, the data is captured within the video and some metadata is captured in the Dolby audio for example. Separate from audio and video, however, this data stream will soon be a new RFC reference that will be rolled in to 2110-40 where the data is delivered as RTP data packets in the same manner as the audio and video data.

The current SMPTE 32NF60 SVIP drafting group that is working on SMPTE 2110 will soon address, in a single standard, all of the video, audio, data and other elements that were identified in the VSF TR-03.

The Alliance for IP Media Solutions (AIMS), an alliance of more than 50 members, that includes manufacturers, operators, and integrators, supports TR-03 which will soon become the SMPTE 2110 standard. The AIMS mission statement and roadmap is to implement the TR-03 standard in an effort to promote one open IP standard for using studio video and audio over IP among the AIMS equipment manufacturers.

VSF TR-04

This standard is a derivation of 2022-6 and AES67 using SMPTE 2059-1/2 as the timing mechanism. This may be of particular interest to operators who like the simplicity of 2022-6 but require the need to provide separate AES 67 audio.

Other Approaches

The Evertz's ASPEN system has recently been the choice of some operators who are anxious to incorporate the newer IP technologies and prefer this methodology. Here, the signal type relies on the MPEG-2 transport stream methodology. RDD 37 was recently enacted by SMPTE as a choice for those who like the idea of using IP in this manner. Parts of the RDD 37 standard will be incorporated within the new SMPTE 2110 standard in an effort to allow those operators who began with ASPEN to migrate to the SMPTE 2110 standard.

Newtek's NDI is a proprietary IP video and audio method and it is being promoted as a choice for production applications especially in editing.

Hybrid Baseband Environments

Although it may be easier to design a particular studio or master control room using a single all-encompassing signal type such as uncompressed IP, it may not be very practical to spend tons of cash to do a complete SDI-to-IP conversion given the tremendous constraints of today's CAPEX budgets. And even if it were possible to do a wholesale conversion, certain existing equipment such as monitors, for example, may not need to be modified to accept the IP input signals when a simple throw-down converter may save money as an alternative to replacing the entire monitor. If all the equipment for a given design could be provided to a broadcast architect with uncompressed IP inputs and outputs, it would certainly be easier to design because all of the I/O's would be the same. However, this scenario is very unlikely to occur. As many past transitions have demonstrated, designers will have to salvage some equipment while replacing other systems.

When thinking about the larger broadcast environment, the challenges of using uncompressed IP can be enormous because of how much current equipment uses HD/3G SDI. A large SDI router at the center of most broadcast plants is a good example of the kind of challenge this transition to uncompressed IP will be. If the industry's history as it transitioned from analog to digital and from SD to HD is any indicator, it would be a good bet that the transition to IP will be done in 'islands' rather than full facility-wide conversions given the obvious exceptions of greenfield opportunities. Even greenfield opportunities will most likely be designed with a mix of both SDI and IP equipment. So the question is this: Where in current broadcast facilities will uncompressed IP be first utilized and how will it be used? Most broadcast engineers seem to agree that throwing out all SDI equipment in favor of IP is not only impractical, but downright foolish.

Operational IP Islands

It is logical to expect that most broadcast operators will elect to transition in smaller areas within their plants, where introduction of IP baseband audio and video will likely have the highest impact. One logical area for operators to consider may be video and audio monitoring using uncompressed IP.

There are two considerations that support this idea: Monitoring is considered to be low-risk, and conversion to monitoring over IP can reduce the number of SDI output ports from the current SDI core router. Because monitoring is a key aspect of any SDI plant, converting monitoring signals to uncompressed IP is not a high-risk option. But if this conversion involved a master control room, for example, then it would be a much higher risk to convert. Missing a commercial spot would mean a loss of revenue, and operators are keen to avoid this type of problem as much as possible. When converting SDI to IP for monitoring, operators can take advantage of the opportunity to recovering a lot of SDI output space on the core SDI router. Much of the SDI router is dedicated to monitoring, so converting from SDI to IP would allow operators to recover much of the output — perhaps even half of the SDI router.

SDI router tie-lines are also a logical use of uncompressed IP signals where many operators have to limit SDI connections to and from different SDI routers. Using IP instead of SDI, operators may have a larger selection of signals to choose from using control software to select various IP addresses for conversion to and from SDI sources and destinations.

Another area that operators may consider converting would be smaller live production environments, such as backup control studios. In this scenario, operators would have the ability to get used to the newer IP environments while preserving existing SDI operation rooms. This approach would help to ensure consistent workflow and operational assurance until such time as uncompressed IP technology matures and becomes more familiar to operators. In the early days of HD, this concept was used for production people to become familiar with HD while still doing their daily work in SD.

In any scenario, it's a good bet that SDI and uncompressed IP will coexist as a hybrid environment and allow operators to experiment with IP as a new signal type. As time goes by, operators will get used to using uncompressed IP signals, and they may find that there are some advantages in their workflows that favor IP over SDI. From a historical perspective, this same pattern was seen during the analog video to SDI video transition, as many operators used existing analog composite monitors. The use of these older analog monitors was done mainly because new SDI video monitors were expensive and the old analog video monitors were available and provided the same functionality as the SDI monitor. The other factor was that many new SDI video tape machines, for example, had both SDI and analog outputs so that both formats were available to operators. Having both uncompressed IP and SDI on the same type of machines (like video servers) may lead to a similar situation.

Conversion Strategies

In all likelihood, the selection of signal conversion from SDI to IP and from IP to SDI will be dictated by need and by exposure to risk, meaning that most operators will elect to gain experience with the new baseband IP signals before they commit active revenues of daily workflows. As time goes by, operators will gain more confidence with IP sources and destinations and expand them as operators did when HD sources became more prevalent in the case of the SD to HD expansion. Using this approach, operators will have some choices as to how to go about the task of conversion.

External Conversion

Converting signals from one signal domain to another and vice versa utilizing an external frame with separate conversion cards — commonly referred to in the past as 'glue' — has long been a common method for accomplishing the task of conversion. In the case of SDI-to-IP and IP-to-SDI conversion, each signal type is converted separately to give operators the flexibility to pick and choose individual signals that are coming from cameras, satellite receivers, servers and other signal sources.

The idea of putting a number of different specialty conversion cards in a common frame has been around for decades. Typically, the frame can accommodate 10 to 20 cards with a dedicated power supply with an optional backup power supply. This external conversion method also allows systems to be controlled by a common software control environment. While lending operators the ability to monitor and configure equipment to suit the needs of the plant, this model also eliminates the need to set dip switches and perform other laborious manual configuration tasks. That said, the frame and card chosen still must be wired from the signal source, and rack space must be accounted for in the facilities rack or machine room.

Another conversion choice for operators may include conversion as 'throw-down' boxes as this tends to be easier on rack space and allows portability in applications in which permanent installation is neither desired nor required.

Internal Conversion

Another method to consider for incorporation of IP-transported signals is to use existing infrastructure to manage SDI-to-IP and IP-to-SDI conversion. Given that the majority of current SDI environments use core SDI routers, it would make sense to explore the idea of conversion between SDI and IP within the actual router. After all, nearly every SDI signal in current broadcast plants are inside core SDI routers. So it would seem logical that, because these signals are already aggregated in the core router, one would keep them there and use the power of the SDI router to manage selected conversions to and from both SDI and IP sources and destinations.

In this approach, there would be no need for more rack space, additional coaxial cable or additional control cabling along with the bolt-on external conversion; the conversion would be done internally within the router using SDI input and output cards that pass the SDI through their normal router paths and also provide the conversion to IP. In other words, operators could simply replace their normal SDI cards with cards that convert the SDI to IP while maintaining the same SDI signal capacity. If an input card has 12 inputs, then the new card would convert 12 SDI inputs to IP but still pass the SDI on to the SDI outputs. The same is true of the output cards. Furthermore, more IP signals would be placed in the Ethernet fabric switch with the use of denser 40Gbe and 100Gbe Ethernet fiber cables instead of placing a handful of signals using 10Gbe cables.

Broadcast Core Routing Infrastructures

A key operating concept, which began during the SDI transition from analog signals, is the ability to do frame accurate switching from source to source in an XY (source/destination) format. Many vendors of uncompressed IP video have demonstrated the ability to switch IP using an ordinary button-per-source panel. This is a curious observation. In the past when analog video was the predominant signal format, operators had a difficult time making sure that SDI signals could be switched accurately. Back then, vendors built SDI signal features such as 'clean' and 'quiet' so that these newer SDI signals mimicked the analog features when the signals were timed and genlocked through the analog router.

But one of the most compelling demonstrations, when comparing SDI to uncompressed IP, is this same ability to switch cleanly from IP signal source to IP signal source. A common button panel is typically used by manufacturers to demonstrate this switching ability using IP. In essence, a hybrid environment is used to demonstrate a simple function that started from the days of analog video. A number of manufacturers use this common button panel as a means of illustrating the inherent simplicity of switching IP signals just as operators are doing with SDI signals.

At the heart of this demonstration is the ability to show operators a clean switch on a core SDI router between a few SDI source signals using a dedicated button-per-source control panel. But the demonstration uses IP signals being switched within an Ethernet fabric switch. The point here is that demonstrating a switch from IP signal source to another IP signal source could be done using a computer and some graphical user interface, but it is not. This is because operators typically use button panels to switch between signals, and the use of such button panels for IP signal switching helps to assure users that they can maintain the same familiar means of switching video as they transition to the IP world.

In the days of analog video and audio routing, the two analog signal types were kept together so that audio followed the video. When it became necessary to peel off the audio from the video, the concept of switching on different 'levels' became predominant. This allowed operators to have audio on one level and video on another. The level concept continued when baseband stereo audio appeared as well as secondary audio channels, component analog video, SDI, embedded SDI, ASI, MADI, and on and on. Leveling became an easy method for keeping all these disparate signals together, and it also provided a means by which to control them. Uncompressed IP may also follow this same logic in that uncompressed IP may be added to existing router fabric as a separate level for the same reasons of the past.

In fact, because there may be more than one common uncompressed IP signal format such as SMPTE 2022-6 and TR-03 (SMPTE 2110 draft), there may be a workflow or practical reason that some operators will assign each IP signal format a different level so that 2022-6 remains in one Ethernet fabric switch and 2110 IP signals remain on another separate Ethernet fabric switch. This was certainly the case when operators chose to put ASI in SDI router hardware and assign it as a separate level.

4K presents different challenges for operators. In a recently published letter to manufacturers, U.S. 4K (UHD-1) truck operators pronounced their desire to utilize SMPTE ST2082 12G single link SDI inputs and outputs. The main reason is that data payloads that exceed 10GB/s are much more difficult to work with in current IP Ethernet fabric. However, in the application of using a traditional SDI router, the UHD-1 format can easily be dedicated as a level so that UHD-1 and traditional 3G SDI signals can coexist. This is not to suggest that 4K and higher resolution formats will never be compatible with IP Ethernet infrastructures, but rather that the bandwidth restriction of 10G hasn't been completely resolved. Some industry developers suggest mezzanine compression as a solution while others support the idea of using 25Gb/sec Ethernet switch ports as a means to address this problem. Either way, there is uncertainty in this market for operators who want to produce live shows for UHD-1 resolutions now. So, it's only natural for truck operators to use something they are completely comfortable with: SDI.

Control of IP Routing Fabric Switches

One of the chief enablers for moving from SDI to IP is the availability of advanced Commercial Off-The-Shelf (COTS) Ethernet switches. These switches have progressed from 10Mb/sec to 100Gb/s in a very short time. Given that these Ethernet switches are used in nearly every sector of the economy, these Ethernet switches also span many different industries. Many industry experts point to the use of these switches as the broadcast industry pushes for greater bandwidths such as 4K and 8K. The other factor is that these IP switches utilize fiber optic cable which is now much cheaper than traditional coaxial copper cable.

The main challenge is controlling this Ethernet switch and its associated Ethernet fabric. Consider a simple switch from one IP signal to another. In the SDI world, this is fairly straight forward. In the Ethernet world it is anything but straightforward. In order to facilitate the IP signal switch, the IP audio video receiver must be told that it will 'tune' to another IP stream. The fact is, this takes a little time to receive the new signal and then switch to it. Some manufacturers refer to this as a make-before-break. In doing so, the receiver must take on two signals for a brief period (thus doubling the bandwidth for this duration) and then breaking the previous connection to receive the new connection. Others vendors use Network Address Translation (NAT) in the Ethernet switch to change the IP address. This method relies on the Ethernet Switch to use a 'table' of pre-defined addresses so that the receiver remains 'tuned' to the same multicast address at all times. The NAT method gets more complicated when there are hundreds of IP signals rather than just a few dozen IP signals.

Nearly 10 years ago, Utah Scientific saw that switching IP signals within an Ethernet switch would be interesting and began developing its own Ethernet switch. After a year and a half of development, Utah Scientific applied for a patent that addressed a method for switching IP signals within a 1Gb/s Ethernet switch and later was awarded a patent for the manner in which this could be accomplished.[1]

The product that was developed in conjunction with this patent didn't go very far simply because Ethernet switches didn't address uncompressed baseband HD — and furthermore, broadcasters weren't too interested in switching from SDI to IP at that time. However, it would appear that this idea was just a bit premature because 10 years later, uncompressed IP is now a subject that is front and center for the broadcast industry.

Many companies feel that providing their own fabric switches isn't a good idea given the enormous number of switches provided by Ethernet switch vendors that are quite active in this market space. Some manufacturers feel that full control of the Ethernet fabric is important and have elected to build their own Ethernet switches. However, many other manufacturers feel that development of these switches and routers is better left to those whose expertise lies in this segment of the IT industry. Audio and video operators can then choose among a number of switch vendors and pick the product that provides cost-effective interoperability. More important is the ability to use COTS hardware, which is more appealing than special or bespoke hardware because it supplies operators with numerous choices as they prepare to migrate their existing SDI infrastructures to IP.

While a full IP infrastructure may be able to offer more signal flexibility in terms of selecting different audio and video signals within the overall Ethernet fabric, it would seem logical that there will be a mix of older SDI and newer IP as a hybrid environment. The question now is how to design this mixed environment to satisfy both current and future workflows.

Conclusion

There are many strategies an operator might utilize to migrate a baseband signal plant to IP, but if history is any guide, the transition will follow the familiar island approach that was adopted in the past with conversion from analog to digital and from SD to HD. It's clear that the transition to IP will not be undertaken in wholesale fashion. Rather, SDI and IP will coexist for years to come simply because operators will not completely abandon prior investments. The scenario in existing facilities will most likely be some mix of old SDI and new IP in some combination that makes technical and economic sense for the broadcast facility.

It's also important to note that during these transitions, workflows change. And one of the chief arguments supporting the move to IP is the ability to manipulate and control the signal more effectively than ever before. While there is much debate across the industry as to the advantages and disadvantages of IP carriage of audio and video, most industry experts concede that the transition from baseband SDI to baseband IP is inevitable.

---

[1] US Patent - US 20080019388 A1; https://www.google.com/patents/US20080019388

