Beyond the Digital Conversion

The Integration of Information Technology and Professional Media

31 March 2014

The Convergence of 2 Industries - The Adoption of Information Technology by the Professional Media Industry

Report of the SMPTE Study Group on Media Production System Network Architecture

Society of Motion Picture and Television Engineers®
3 Barker Avenue
White Plains, NY 10601 USA
www.smpte.org

Beyond the Digital Conversion
The Integration of Information Technology and Professional Media

The Convergence of 2 Industries - The Adoption of Information Technology by the Professional Media Industry

Report of the SMPTE Study Group on
Media Production System Network Architecture

Society of Motion Picture and Television Engineers®
3 Barker Avenue
White Plains, NY 10601 USA
www.smpte.org

27 March 2014

Table of Contents

ACKNOWLEDGMENTS

Theodore H. Szypulski, Committee Chair
Senior Director Technology Research & Standards
ESPN, Inc.

It has been my privilege to chair this study group, which performed its work admirably between November 2012 and March 2014. I wish to thank all of the participants for their dedication of time and talent to this effort. While the roster for this group includes over 75 names, there are particular individuals, folks supported by their employers in many cases, who deserve mention. Their names, and their employers, are listed below. This work could not have been completed without the dedication of all the participants.

Special thanks to my fellow officers in this committee, Secretary Thomas Kernen of Cisco Systems, and my colleague at ESPN, and our Document Editor, Steve Posick. Steve has worked untold hours crafting numerous drafts of this report.

I place myself into the category of the traditional broadcast engineer, and one who has worked for over four decades on understanding and deploying new professional media systems with ever increasing proportions of IT based technologies and techniques. It had become clear to me that there existed a dichotomy between my own cadre (broadcast engineers) and my colleagues in the IT industry. I could see there was a lack of full understanding, in both camps, of the performance requirements of professional media systems, and the new technologies being employed therein.

My hope is that folks in both of these groups will develop an appreciation for the knowledge and talents of all participants in this evolution. I also hope that this report will allow the readers to identify areas where they can begin to fill in gaps in their own expertise, improving the future for all of us.

Finally, it is clear to me that SMPTE should investigate what its own role is in this future, taking the steps we recommend in this report to ensure it remains the preeminent leader in worldwide standardization of technologies used in the Motion Picture and Television industries.

I thank the following for their especially untiring effort in the creation of this report:

David Baril	ESPN, Inc.
Paul Briscoe	Unaffiliated
Bob Edge	Bob Edge TV Consulting
Thomas Edwards	Fox
Friedrich Gierlinger	Institut fuer Rundfunktechnik
Brad Gilmer	Gilmer & Associates, Inc.
Thomas Kernen	Cisco
Al Kovalik	Media Systems Consulting
Steve Posick	ESPN, Inc.
Felix Poulin	European Broadcast Union
Edward Reuss	GoPro
Stephen Scott	Skotel Corporation
J. Patrick Waddell	Harmonic Inc.
Marc S. Walker	Grass Valley US LLC
Leigh Whitcomb	Harris Broadcast

Society of Motion Picture and Television Engineers®
3 Barker Avenue
White Plains, NY 10601 USA
www.smpte.org

1. Introduction

The design and architecture of Professional Media Networks (PMN) is becoming increasingly important with the ever-increasing use of shared packet switched networks (PSN). These PMNs are used for applications such as live contribution and production, post-production, and presentation. Typically these are built upon Internet Protocol networks and used for the production of media, including the carriage of media essence (audio & video), metadata, synchronization and control traffic. Where media traffic coexists with other sorts of communication and business traffic (multi-service networks), it is particularly important to balance the requirements of media production network traffic with other types of network traffic. This other traffic may or may not be related to the Professional Media Network workflows.

Although there exist industry accepted best practices for traditional Information Technology (IT) networks, these best practices were developed from the requirements of IT applications and may not account for the requirements and the characteristics of Professional Media Networks. At the time of this report, no such best practices documentation exists for media production networks.

This report is directed to engineering professionals in both the Information Technology and Television/Motion Picture industries. These two groups have been interacting for the past decades in what has often been referred to as the "convergence" of IT and Broadcast Engineering.

The report begins with an introduction to the broadcast plant followed by several typical broadcasting use cases. IT professionals may find these sections useful in order to understand how Broadcasting/Motion Picture professionals have been designing and building systems with legacy technologies. It is important to understand the system characteristics and user expectations with these legacy systems. Following the use case section there is a section on recent successful implementations using IT connectivity and technologies. The next section consolidates the expected characteristics of legacy systems. A discussion of Managed Networks follows. The balance of the report then vets myriad IT technologies and practices and describes how these might be configured to deliver the expected performance and characteristics.

The report concludes with recommended next steps by SMPTE.

A final note to consider before reading this report -- all kinds of technology advance at breakneck speed, and IT is no exception. The reader is reminded that this report was written in 2013 and 2014, and therefore reflects the state of the technology at that time. No doubt, as this report ages it will become less useful for understanding which technologies to consider in designing professional media production networks, and more interesting as a “time capsule” that captured the current state of things. The reader is reminded that it is vitally important to take from this report the knowledge and guidance offered today, and to continually stay abreast of the evolution of IT in the future.

This report includes trade names to illustrate real-world requirements and implementations. These trade names are used solely as examples, and their mention does not imply any degree of endorsement or recommendation.

2. Scope

This report identifies, documents, and describes the requirements and characteristics of well-designed Professional Media Networks and makes recommendations pertaining to the identified requirements. Media-related quality of service may include characteristics such as isochronous streaming, managed latency, and desired faster and slower than real time performance.

This report also makes recommendations to the Society regarding the further study of and the creation of engineering documents pertaining to architectural, design and operational requirements identified for media production systems and their associated networks.

This report does not provide a number of specific cookbook recipes to meet user requirements, but instead enumerates user requirements in specific use cases, and then outlines characteristics of various technologies.

This report also outlines a number of technology challenges that network architects should be aware of in design of PMNs but does not attempt to provide specific solutions. Readers should study the use cases and technologies in order to facilitate their understanding of how to meet expectations of the new system in Professional Media Networks using packet switched network technologies.

3. Introduction to the Professional Media Production Systems

The Professional Media (PM) industry transformation started decades ago with the adoption of the serial digital interface (SDI), and is currently in its digital adolescence. The primary focus of digital to IP conversion projects is the replication of analog behaviors and requirements within the digital domain.

However, the change to all-digital infrastructures offers many opportunities and is changing the way we perceive and organize media. With the encapsulation of signals into streams and files, these signals can now be duplicated without limit or loss and many signals can traverse the same physical network fabric simultaneously. The very way by which media can be controlled has also changed, instead of telling a device to play the media contained within itself, systems are now being designed to instruct the media to play on any device, essentially changing the landscape from controlling devices to controlling media. These changes are improving the efficiency of media systems by allowing for faster than or slower than real-time operations.

The changes are so big that it has taken decades for the PM industry to address reliability, scalability, and interoperability issues. While there has been an influx of professionals from other technology disciplines into the PM industry, bringing additional knowledge, there is still a need to improve the mutual understanding of PM-specific nuances.

3.1. Introduction to the Broadcast Plant (for Information Technology Professionals)

3.1.1. The Typical Broadcast Plant Today

Figure 1 shows a highly simplified block diagram of a typically television broadcast master control. The connection lines with arrows represent coaxial cables carrying HD-SDI video or AES audio streams.

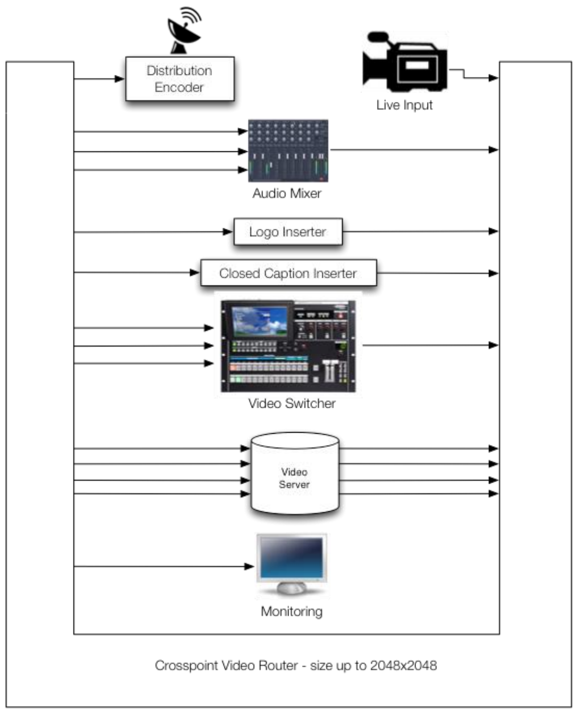

Figure 1: Simplified Schematic of Typical Broadcast Master Control

The large cross-point **video router** is represented as a "U" symbol in this block diagram, but physically it is a monolithic box with hundreds or thousands of coaxial cable BNC connections for audio and video. The video router allows video and/or audio from any input to be routed to any output. Router control panels are found throughout the production plant to control the connections between audio and video sources and destinations. Facilities may have hundreds or even thousands of devices connected to the video router.

Readers should note the collision of terminology with "router" having a very different meaning than within IT. There are numerous overloaded terms used by both industries.

The **video server** is a device for recording and playing back video and audio streams to/from a storage system (typically spinning hard drives). Servers tend to store video using high visual quality compression.

A **video switcher** (also known as a *vision mixer* in Europe) is a device used to select between several different video sources and to composite/mix them together with effects. Some effects include "fade to black", "crossfade" between two video streams, "picture-in-picture", and 3D video warping.

A **Closed Caption Inserter** allows for the insertion of live or recorded Closed Caption data into the video stream.

A **Logo Inserter** is a simple video compositor that adds graphics onto a single video stream.

An **Audio Mixer** performs the mixing, panning and filtering of multiple audio sources, and generally results in a final audio mix that is 5.1 channel and/or stereo. Audio/Video devices may also perform some audio mixing during their video effects (such as a crossfade).

A **Distribution Encoder** compresses the video into a low-bit rate stream for distribution to television stations and/or directly to the viewer. The codec used for distribution is typically MPEG-2 or H.264 and is not designed for a large number of decode/recode cycles. HD video distribution bit rates are typically 15-30 Mbps for MPEG-2.

Not shown in this diagram is the **Automation System**, which has multiple time-driven event lists used to control a large number of devices within the plant, especially video servers, for record and playback, and video switchers. Automation lists can also be manually controlled, for example, in response to a button press that indicates a break in the action of a sporting event, so that an advertising spot can be run. Automation systems previously used RS-422 serial communications to control devices, but the industry has been moving toward IP-based control protocols, with such advancements as the Framework for Interoperable Media Services (FIMS) and the SMPTE ST 2071 standard for Media & Device Control over IP Networks (MDC).

Monitoring includes the displaying of video channels, devices that automatically sense video presence and audio quiet, devices that monitor audio loudness, and Multi-Viewers that display multiple video sources on a single video monitor by combing and scaling the video.

3.1.2. Today's Digital Audio and Video Standards

Although frame rates differ by country high-definition video tends to be found in two primary formats for television viewing, **1080i** and **720p**. 1080i consists of 1920x1080 pixel interlaced frames, while 720p consists of 1280x720 pixel progressive frames. Each frame contains a single image, but interlaced frames, indicated by the "i", are painted as two fields, each field containing either the odd or the even lines of pixels, while progressive frames, indicated by the "p", are painted as a contiguous set of lines of pixels. In 60 Hz power countries that have used the NTSC standard, 1080i has a frame rate of 30000/1001 fps and 720p has a frame rate of 60/1.001 fps. In 50 Hz power countries, 1080i has a frame rate of 25 fps and 720p has a frame rate of 50 fps. Other formats are commonly used in production and consumer media, such as **1080p** at 60/1.001 fps and 60 fps, and **1080p** at 24/1.001 fps and 24 fps; however, the 24/1.001 fps and 24 fps variants are rarely used in broadcast plants.

HD video signals used within production plants are standardized as SMPTE ST 292-1 ("HD-SDI"), a bit-serial data structure with a data rate of about 1.5 Gbps. HD-SDI is usually carried over 75Ω coaxial cable, although there is an optical fiber version (SMPTE ST 297) used mainly for long-distance connections. Uncompressed 4:2:2 chroma subsampled HDTV signals are transmitted on HD-SDI with 10-bit resolution of each component. Extra, ancillary, data space outside the active video area within the HD-SDI signal can carry Ancillary data packets (ANC packets). These Ancillary data packets can include embedded digital audio signals and Closed Captioning data. The HD-SDI channel coding scheme is scrambled NRZI and there is a CRC value calculated for every active line of video.

SMPTE ST 299-2009 defines the embedding of AES3 24-bit PCM digital audio channels, typically sampled at 48 kHz, in the horizontal ancillary data space (HANC) of HD-SDI signal. Embedded audio is particularly convenient as it helps to prevent audio/video lip sync problems that may occur due to timing differences between the audio and video signal paths.

AES3 audio can also be found in the broadcast plant by itself on audio channel pairs using 75Ω coaxial cable (AES3id). Broadcast video routers often have dedicated audio inputs and outputs and some also use AES10 Multichannel Audio Digital Interface (MADI) that can carry up to 64 channels of digital audio on a 100 Mbps network interface. Video routers can have internal audio (de)mux capabilities in order to extract audio channels from or insert audio channels into the video streams.

3.1.3. Uncompressed Video over IP Standards

SMPTE ST 2022-6 defines the transport of HD-SDI over IP using the Real-Time Transfer Protocol (RTP)[RFC 3550]. SMPTE ST 2022-5 defines a Forward Error Correction (FEC) stream that can provide varying levels of protection against RTP packet loss for SMPTE ST 2022-6 streams.

There have been some previous RFCs on HD-SDI carriage over IP (such as RFC 3497), but it is believed that SMPTE ST 2022-6 is an early candidate for the distribution of media essence (audio & video) over IP networks.

3.2. Typical Use Cases

Before a detailed discussion of using packet switched network technologies in Professional Media Systems, it is wise to offer some examples by way of these use cases. While it is impossible to cover every type of use case, understanding some of the more common use cases will be helpful to readers who are not already familiar with Professional Media Systems.

3.2.1. Television Production Control Room User Requirements and Expectations

One example of a facility used extensively in live Television program production is a facility commonly referred to as a “Production Control Room,” or “PCR.” This use case describes a typical Production Control Room at ESPN’s worldwide headquarters in Bristol, CT, U.S.A.

The primary purpose of a PCR is to assemble, in real time, all of the elements needed to broadcast a live program. The control room is typically associated with a Studio or Stage. Often, but not the case at ESPN, an audience is included in the live Studio.

Video elements consist of live camera signals, prerecorded material from video tape recorders and/or media servers. Additional video elements include graphics, remote feeds, images from a still storage device, and the outputs of devices commonly referred to as digital video effects units. Audio elements include live studio microphones; prerecorded jingles, bridges, etc. from audio servers or audio recording devices; audio from the video recorders, remote sites, outputs of audio effects equipment, telephone interfaces and others.

The control room of the PCR suite consists of many (typically ten or more) work positions where people performing the duties of Director, Producer, Technical Director, Graphics Control, Media Playback Control, and various assistant duties. Usually there is a separate, sound proofed room adjacent to the main control room for the control of audio. One to three persons reside in the Audio Control Room (ACR) and perform functions of mixing audio, coordinating with the studio and assistant duties.

Many (typically over 100) video monitors are placed at the front of the control rooms and in consoles in front of the various work positions to allow workers to view current and upcoming video signals. Some positions employ various types of “waveform” monitors to view the technical parameters of the video signals. Some positions also include audio monitors, often used with headphones, which allow the workers to listen to the audio signals. The audio control room includes high quality audio monitoring, in 5.1 channel surround sound, to ensure the outgoing signal is correctly mixed. Supplementary audio cuing and monitoring devices are also available.

All signals that are available in the room and that might be made available to the outputs of the room must be synchronized so that when combined with other signals, or when switched to and from one another, there is no disruption in the cadence of the signal stream. The output stream MUST be a continuous stream at a constant rate. The audio and video shall be synchronized not only with respect to electrical signal rate, but also by always showing the same moment in time in the audio and video signal.

Signals on the various video monitors must always be in "moment of time" synchronization with the audio in the room, within an acceptable margin of error. Many facilities now utilize a composite video monitoring system that combines several images onto one large screen. These systems often introduce a video frame of delay, time needed to process and combine the various images. ESPN has found that the advantages of these systems can be successfully balance with the need for "lip sync" when user expectations are set from the beginning and when the offset is constant and limited to one frame of video.

Often during live productions, interviews between persons take place where one person is remotely located. The industry has dealt with latency issues for years due to satellite delays, causing difficulty for people while in a conversation. It is very undesirable to have this situation at all and therefore the addition of more latency, especially varying latency is to be avoided.

Summary of user requirements for a PCR:

- All video and audio signals bound for air must be synchronized with one another. Video to video timing must be within the input buffer window of the video production devices that combine these signals (typically less than the duration of one line of video). Similarly for audio, the signals must be within the input buffer window of the audio mixing device (this is a digital audio requirement).
- Audio to video "moment in time" match must be less than one frame of video, with the audio never preceding the video. Note: this is often referred to as "lip sync".
- Output signals shall be isochronously streamed with a very low latency from the live event, no more than a total of two frames delayed after passing through the entire PCR, ideally less than one frame delayed. Video shall never freeze nor shall audio ever go silent, unless it is a desired effect. There shall be no visible or audible distortions to, or discontinuities of the media streams due to bandwidth restrictions.
- All signals paths shall have a very high reliability factor. No connection path to the control room, or within the control room, shall be disrupted more than once per year and for a duration of no more than one frame of video. Ideally, no path should ever be disrupted. This restriction applies only to times that the room is being used for production.
- Media signals meant for monitoring shall be, as much as possible, equal in performance to on-air signals. As monitoring signals are often used for quality evaluation and for confidence and continuity purposes, any changes to them, separate from the on-air signals they represent, might be misinterpreted. When absolutely necessary, audio shall be no more than one frame of video later than the video signal it is associated with. Audio shall never precede the video signal.
- Signals that send commands to devices shall be delivered without delay, in a deterministic fashion and with latency in the ballpark of a few lines of video, at most. Reaction to the command is expected by the beginning of the next frame of video.
- When communications amongst workers is delivered via packet switched networks, there shall be no dropped syllables, or any other kinds of distortion that would detract from clear and immediate communications. There shall be no "busy" signals.
- Signal pathways often incorporate physical "patch bays" or "jackfields" so that the room configuration can be changed from its norm for operational, or for emergency bypass purposes. It is required that packet switched networks also accommodate this physical patching.
- While traditional connectivity could never offer it, faster-than-real-time delivery of media in and out of servers is a valuable benefit in live control rooms. Being able to load a recently published element very

quickly, providing time for evaluation and cuing, adds immediacy to the workflow, increasing the value of the finished media production.

3.2.2. Remote Production Truck (OB Van)

A Remote Production Truck is essentially a complete Production Control Room (as discussed earlier) on wheels. It is typically a trailer, usually 40 feet (12 meters) in length, towed behind a large diesel powered tractor (truck). It not only carries the full PCR spaces (both video and audio) but also all of the electronics necessary to make the PCR function. Additionally it will also carry camera control units (CCUs), often for over a dozen or more cameras.

To help provide additional space for the humans, most production trucks have one or more expandable sections (called “expando’s”), which often protrude an additional 12 feet (4 meters) beyond the vehicle sides, and sometimes ends.

During major events there will be a number of trucks, hooked together to provide greater capability than any one of them alone could. This will typically include a dedicated audio mix truck and often a dedicated graphics and/or playback truck. For these events, camera control is distributed between those trucks with mounted CCUs.

The working environment within a remote truck is usually noisy, stressful, and cramped. Monitoring, especially audio monitoring may be suboptimal. Despite the hardships, the programs produced within them are typically first class.

Remote trucks are used for most live sports, awards shows, and concerts. There are few venues with this capability built into the building, so a truck is almost always required. While some events backhaul all signals to a central studio facility, the event still requires a large crew to operate the gear and most importantly deal with failures or changes at the site in real-time.

3.2.3. Television Master Control Rooms, User Requirements and Expectations

One example of a facility used extensively in Television is a “Master Control Room,” or “Master Control.” Outside North America, these rooms may be called "Presentation Control Rooms." This use case describes a typical Master Control Room at any large television network.

The primary purpose of a Master Control Room is to merge pre-recorded and live content into a continuous stream which is ready to be sent via satellite, terrestrial transmitter, or by other means to affiliates, cable systems and/or the viewer. The feed is a continuous asynchronous stream with a carefully defined relationship between audio and video content.

Inputs to a Master Control Room can include the output of video servers, audio/video feeds from a live sports, a Production Control Room (a PCR), satellite feeds from remote trucks at remote events, playback from video tape devices, still and moving images and audio associated with graphics and other branding mechanisms, and importantly, inputs from captioning systems and Emergency Alert Systems (EAS). Note that at the time of this report, these last two inputs are the *only* signals U.S. broadcasters are legally required to transmit. Digital video effects boxes may also be available to create effects such as squeezing back the program video so as to introduce branding or “coming-up-next” graphics to the transmitted images. Master control facilities may also have the ability to take an audio-only input such as a phone line in the event of a major equipment failure.

Another critical input to the master control facility is the program schedule, created by the Traffic department. The traffic department takes the program lineup from the program department, combines it with the promo schedule from the promotions department, determines the placement of commercials, and then produces the schedule. The program schedule lists every program, commercial, promotional/interstitial item, in short, every single element that should be combined into the finished program feed, along with the approximate timing for each event. While the schedule is usually viewed on a computer, many media companies also ensure that a printed version is available in master control, in the event of system failure.

You may hear several terms used in relation to lists of content for on-air operations such as schedule, playlist, log, and as-run log. While these may seem interchangeable, they have subtly different meanings. A schedule is a list of content and approximate times at which the content should air. The term Log may be used interchangeably with the term schedule. A playlist is a list of content that should play to air, typically contained in an automation system. A log is a record of what has played and precisely when it played, and an as-run log is a document, or file, which, after certification, contains the record of what played and precisely when it played for billing purposes.

Typically, a schedule is loaded into an automation system at a particular time of day, although some media companies are moving to real-time interaction between traffic systems and on-air automation. Once a schedule is loaded, master control operators check the playlist to ensure that all of the content called for by the list is available in the system. If it is not, they create a "missing list" and then someone is assigned to track down the missing content. The automation system plays the content to air and an as-run log is generated for billing purposes. Some media companies operate their networks according to a strict time-of-day clock. If the schedule says a program starts at 14:02:00, then that is exactly when the program will start. However, other media companies operate their networks according to a floating schedule. Start times on the log are approximate, and operators may adjust schedule timing by inserting interstitial content, dropping items, or the network may insert looping content which runs continuously, the duration of which can be adjusted by sequencing to the next event at the appropriate time. However, even these schedules have certain hard times, which must be met. In some applications, video and audio timing must be extremely closely controlled so that pre-arranged splice points can be met. For example, a main feed may be split off into a regional feed for a certain portion of the day. When the regional feed rejoins the main feed, the timing of that switch must be frame accurate, both in video and audio, or unacceptable video flashes or audio bursts occur.

Since the purpose of a master control operation is to switch between different content streams, the relationship of the timing between the various input streams must be closely maintained in order to avoid video and audio artifacts (flashes, cracks and pops, freezing and other undesirable effects). These timing requirements are identical to those given for a PCR. The original relationship at the source between audio and video streams must be tightly maintained throughout the master control facility. Master control operations frequently put streams on air that originate at remote locations, for example, a truck at a baseball game. The truck in the field is likely *not* using the same time-base as the master control facility. Therefore, the video and audio timing will drift in phase with respect to the master control facility clock. Video synchronizers are used to lock the timing of the incoming feed to the house master reference. These synchronizers delay the video stream. Therefore, audio delay units are inserted into the corresponding audio streams, and these audio delay units are frequently slaved to the video synchronizers, automatically introducing a delay in the audio stream that corresponds with the amount of delay introduced by the video synchronizer.

A number of people typically work in a master control area. There may be a shift supervisor who is responsible for all of the activities in the area, a master control operator who is responsible for the on-air signal, typically for several channels, a tape operator who is responsible for ensuring that long-form content is available for air when needed, and someone working to ingest material, meaning that they are responsible for converting tapes to files and storing them on fileservers for air, or checking on content delivered to the facility as files prior to approving them for air. During live sporting events, there may be a live event coordinator who is in communication with people at the remote event, relaying important information to master control operators. There may also be a utility person who can fill in for any of these positions during breaks or in case of illness. The number of people and their job descriptions vary from company to company. Working in a master control area is akin to working in the operational area of a nuclear power facility. A lot of the work is automated, and the work can be routine. However, when something goes wrong, only highly trained staff can save your company from a highly visible disaster.

The master control operator normally operates the facility through an automation system. His or her job is to check upcoming events to ensure that the automation is properly configured, and to ensure that all content is available. After that, the operator monitors the automation in case something goes wrong. However, during live events, the operator may be integrating commercials into the sporting event. This involves not only playing regularly scheduled commercials during time outs, but also inserting standby schedules in the event of a rain delay, injury or some other unscheduled event. The automation system is responsible for queuing up the content in preparation for air, beginning playback of the content, and then selecting that content, usually through a master control switcher, in order to send the signal to air. In multi-channel facilities, there may not be a master control switcher control head; but instead, switching may be done through a large Serial Digital Interface (SDI) router. The automation system also triggers “secondary events”. These events might be insertion of logos, triggering of Digital Video Effects boxes, or trigger “audio over announcements”, the audio announcements that run over closing credits promoting upcoming programs.

A Master Control Room has sufficient audio and video monitoring to allow the operator to observe the signal being sent from the facility and in the case of satellite or terrestrial transmission, there is usually a return monitor, which is used to verify that the outbound signal is able to be seen by the viewer. There are input monitors that allow the operator to see sources that can be put on-air, and there are also feeds from PCR suites and other control rooms that might need to be taken live to air. One critical issue is the synchronization of audio with video, typically called “lip sync”. Because different monitors have different delays in the time it takes for video to be presented from the input to the display, many broadcasters have gone to one or two very large displays that are used as Multi-Viewers. This allows them to compensate for monitor delay in a way that is consistent over all of the displayed images, something that is difficult or impossible to achieve when using a number of different size displays from different manufacturers with different internal delays. Note that in some cases broadcasters have given up on trying to solve this and simply ensure that the delay is correct at the output of the facility and then let the delays in the master control area be whatever they might be. In this case, lip sync QC is delegated to a special monitoring area where audio/video delay is known to be correct.

Master control designs tend to be relatively straightforward. An absolute premium is placed on reliability and on giving the operator relatively simple options in the case of a system failure. For example, it may be possible for the operator to bypass all master control equipment in order to put a tape machine or fileserver directly to air. Because in many cases a single operator may be responsible for tens of feeds, automated monitoring and signal fault analysis tools are frequently employed. These tools detect extended periods of black video, audio silence, and possibly other more advanced analysis. However, there are some cases where a human being is still a critical part of the quality assurance process, for example in detecting cases where Spanish and Portuguese language tracks have been swapped.

Signal integrity is extremely important in a Master Control Room. What goes in must be the same as what comes out and audio/video synchronization must be tightly controlled. In many cases, master control operators do not have provisions to adjust audio/video levels. Instead, they rely on the ingest and quality control processes to ensure that these levels are set correctly before the content is made available for air. To this end, many facilities embargo content, ensuring that there are areas of fileservers or tape shelves reserved for content which has not completed the QC process. Once content has completed a QC process, it is made available to the automation system for playback.

One very critical point about modern multi-channel file-based master control rooms is that garbage propagates. Once content is made available for air, automated processes may replicate that content not just on servers in the facility, but potentially to remote facilities, which provide business continuation capabilities. So the old adage, "garbage in, garbage out" is actually "garbage in, garbage everywhere" in these operations. QC is a critical process, and automated QC is frequently employed not only at the ingest point, but also at other places in the system, with the goal of catching garbage before it propagates widely.

Another important consideration in master control operations is the importance of a unique identifier. This identifier is sent down from the traffic department in the schedule and is also associated with content as it is ingested into the master control environment. This identifier, historically known as a "house number", is the key link between computer scheduling systems and automation systems playing back content on air. The main issue with house numbers is that, in many cases, they are not guaranteed to be unique. Another issue is that in many master control environments house numbers themselves convey meaning to operators. This has created a very big liability issue for networks and a lot of human effort is expended to ensure that the correct content gets on air. There are a number of proposed solutions to the house number issue; some of them are quite promising. Any new systems or technologies must take all steps necessary to ensure that identifiers are not corrupted or deleted, whether associated with the content, or with the program schedule data.

It is worth mentioning here that networking technology is extensively employed in master control operations. Frequently, files are transferred between video servers using high performance packet switched networks. The master control operation is completely dependent on packet switched networks for the proper operation of the automation systems. Most if not all Master Control Rooms receive some commercial material, which is transmitted to the media company over IP networks, whether terrestrial or via satellite. Most of these operations employ a number of security measures to protect the media company, such as preventing someone from attaching an unauthorized computer to the network or introducing a file infected with a virus into the automation system network. Extreme care is taken in connecting these networks to other networks within the facility, in some cases including an "air gap" between the master control network and other networks, meaning that there is absolutely no connection between master control and other networks within the facility.

Summary of user requirements for a Master Control Room:

- It must work. All the time. "It just works."
- Almost everything about master control systems is driven by the "It just works" requirement. Given that once a particular moment in time is past, any revenue opportunities associated with that moment are lost forever and also given the fact that a commercial on a national network may represent tens or even hundreds of thousands of dollars, a huge amount of effort is expended to ensure that there is always a way to get a signal on the air. People responsible for designing and planning master control facilities are risk-averse, with good reason. Operations are kept simple and backup or recovery scenarios are studied, simplified, and practiced.
- Automated monitoring and QC are required in these facilities, especially multi-channel facilities with centralized ingest operations.
- The facility must have redundant power feeds from the local power company, fed from diverse locations. Facilities are protected by multiple generators, multiple fuel tanks and multiple uninterruptable power supplies and battery banks. Outgoing paths are redundant, both in terms of equipment and physical location. In this way, outbound paths are protected against "backhoe fades", otherwise known as fiber cuts caused by a construction crew.
- All audio and video signals bound for air must be synchronized with one another. Video to video timing must be within the input buffer window of the video production devices that combine these signals,

typically less than the duration of one line of video. Similarly for audio the signals must be within the input buffer window of the audio mixing device.

- Communications is critical to these operations, so redundant telephone connections from the local service provider are required. Intercom communications amongst the master control team and between the master control team and any remote field operations, such as sports trucks, must be extremely reliable and of high quality. The inability to hear a verbal queue from a truck located in the field might mean a missed opportunity for the insertion of an unexpected commercial break. This can directly translate to the loss of tens or even hundreds of thousands of dollars.
- Highly trained and capable maintenance personnel must be available 24 hours per day, 365 days per year. The most critical hours for operations begin about 5:00 PM Eastern Time and continue through “prime time” on the West Coast, ending around 3:00 AM Eastern Time. Your best technical support people need to be available to respond in minutes, if not seconds, during this period of time.
- Regular “switch to protect/repair/restore to main” maintenance protocols are completely unacceptable. In the master control environment there is not a notion of “main” and “protect” for circuits or equipment, or if there is, a restore back to main must be done off hours, usually between the hours of 3:00 AM and 5:30 AM Eastern. It should be noted that the protocol above is frequently baked into IT architectures and results in the unintended consequence of “two hits for the price of one”. The network takes one hit during the initial equipment failure and then take another, unnecessary and sometimes unanticipated, when faulty equipment is repaired and powered back on, causing an automatic switch from the protection equipment back to the main equipment.
- Remote alarms, remote monitoring, and detailed logging are absolute requirements for forensics analysis. Logs must be detailed enough to allow for unambiguous remediation steps to be taken against future failures. This is not an environment where one wants to take a guess at what went wrong and hope you fixed it so it does not happen again. As much as possible, you want to be able to clearly identify the cause of a failure.
- Frequently “shadow” Master Control Rooms are available for operator training and as test beds for new software revisions and hardware upgrades. These systems frequently serve a role in business continuation plans, providing a backup master control room, if required.
- Most media companies and networks are unwilling to put untested software or hardware on air.
- Maintenance Operators must be able to quickly and simply switch between different monitoring points in the transmission chain. They must be able to see an exact copy of the audio/video feeds that are appearing at the inputs of the transmission equipment. Interestingly, delay of these feeds is really not an issue, even a significant delay of up to one second or two is not problematic in most operations, as most operators are used to the satellite return delays.
- Business continuation plans frequently involve the continuous operation of mirroring facilities that can be across town or across country. A lot of effort is put into ensuring that the mirror facility has content that is identical to the main facility and that the mirror is ready for use on a moment’s notice.
- Audio to video “moment in time” match, often referred to as “lip sync”, must be less than one frame of video, with the audio never preceding the video.
- Output signals shall be isochronously streamed with a very low latency between the inputs of the master control room and its output. Video shall never freeze nor shall audio ever go silent. There shall be no visible or audible distortions to, or discontinuities of, the media streams due to bandwidth restrictions.
- Media signals meant for monitoring shall be, as much as possible, equal in performance to on-air signals. As monitoring signals are often used for quality evaluation and for confidence/continuity purposes, any changes to them, separate from the on-air signals they represent, might be misinterpreted. When absolutely necessary, audio shall be no more than one frame of video later than or precede the video signal it is associated with.

- Signals that send commands to devices shall be delivered without delay, in a deterministic fashion, and with latency equivalent to, at most a few lines of video. Reaction to these commands is expected to occur by the beginning of the next frame of video.
- When communications amongst workers is delivered via packet switched networks, there shall be no dropped syllables, or any other kinds of distortion that would detract from clear and immediate communications. There shall be no “busy” signals.
- Signal pathways often incorporate physical “patch bays” or “jackfields” so that the room configuration can be changed from its norm for operational, or for emergency bypass purposes. It is required that packet switched networks also accommodate this physical patching.

3.2.4. Time and Synchronization with IP

At its inception television was a glass-to-glass synchronous environment. With the exception of film, everything was live and there was no color. However there was still the problem of synchronizing various video sources to enable roll-free switching (the introduction of visual artifacts and errors, such as loss of vertical sync).

Although the size of facilities and the types of reference signals has evolved, the need for synchronization and the basic principle of distributing an analog reference signal to devices has remained relatively unchanged.

With the advent of IP technologies for the creation, storage and distribution of media, the time and synchronization landscape is changing. Digital media consists of chunks of data that must be stitched together on a clock-based timeline and the new proposals for the distribution of the synchronization reference consist of time points distributed over IP networks, instead of continuous analog reference signals. The distribution of these synchronization reference time points on packet switched networks pose a significant challenge for the minimization of jitter and the delivery of the absolute time precision required.

The analog reconstruction of audio presents additional challenges for stability (jitter), to prevent sound distortions and other audio artifacts. While the absolute timing precision may be less with audio, there are stringent timing requirements so that the same time points for each audio channel are presented at the same time, to prevent spatial distortion during multi-channel audio presentation.

With the audio and video each undergoing extensive processing via devices with variable latencies, there is the significant task of maintaining audio/video synchronization (lip-sync). To achieve this, there may be a need to bind timing marks to both the audio and video media, so that the media can be reconstructed and presented properly.

3.2.5. Digital Audio Distribution Over IP

Both studio and event audio can involve hundreds of channels, all fed into a mixing console, or a hierarchy of several mixing consoles. Each console may support anywhere from 48 to 190 independent channels. The entire solution may involve several hundred channels. All of these channels result in a huge number of cables routed throughout the facility.

Digital audio networks permit multiplexing several channels onto a single cable. One of the motivations for the conversion to digital audio distribution was to reduce the cabling, especially in touring applications, which can require repeatedly deploying and repacking the entire kit several times a week.

Several types of digital audio networks have been available for years. Most of these are proprietary. The most recent ones are based on IP, such as Dante, Ravenna and AES67. These last two use RTP. They can multiplex several hundred channels together, although there is often a trade-off between the number of channels and the end-to-end latency.

There are two categories of latency requirements for audio. The first is for the talent and the second is for the audience. The performers need to monitor their own individual performance, which requires extremely short latency in order to avoid a distracting echo. (In telephony, this is referred to as "sidetone delay". IEEE 269-2010 requires the sidetone delay to be less than 5 msec.) Some professional musicians claim they can hear a delay on the order of 1 msec. So this monitor path requires extremely short end-to-end delay. Since the sampled audio is conveyed via data packets, the end-to-end latency is directly related to the audio samples per packet. Minimizing this delay requires reducing the packet buffer length, which increases the network load. Minimizing this delay also puts a heavy demand on the packet scheduling capabilities in the managed network routers.

The audience category is more forgiving as the audience is normally further away from the performers. This delay is governed by the speed of sound in air, about 340 m/sec. (~100 f/s, or 1 f/ms) PA systems for live events often include auxiliary speaker systems, either for surround sound or for boosting the volume to the audience in the farthest parts of the venue. These systems include delay elements to synchronize the audio from the auxiliary speakers with the direct acoustic path from the stage. These delay elements may add as much as 300 msec. of delay or more, depending on the size of the venue. However, this delay needs to be controlled to an accuracy within a few milliseconds for best performance.

Synchronization is critical for all of the inputs to the mixing consoles. At 48k samples/sec, the sample interval is about 20.8 microseconds. At 192k samples/sec, it's about 5 microseconds. While not all of the paths require sub-sample synchronization, the delays need to be adjustable to millisecond accuracy or better, even less when used for inter-channel digital signal processing, such as echo cancellation. This requirement is more stringent than the normal lip sync requirements between the audio channels and a video stream with frame sampling intervals on the order of 16 to 42 msec. Several of the audio networks based on IP permit synchronization via IEEE 1588-2008 PTP.

3.2.6. Edit Rooms

One example of a facility used extensively in Television is an "Edit Room". This use case describes a typical Edit Room at any facility. This room can be located in proximity to the systems supporting it or as a remote facility accessing material from systems both near and far.

The primary purpose of an edit room is to merge pre-recorded material from various sources into a continuous stream. This material may be provided in a live streaming format, file transfer, or static images. How the material is ingested into the edit system is not necessarily time dependent. How the material is assembled during the edit process is very time dependent.

Sources utilized in an edit room can include the output of video servers, graphics systems, audio storage devices, live cameras, voice over audio microphones, branding mechanisms, captioning systems, and sometimes, playback from video tape devices. Digital video effects boxes may also be available to create effects such as image manipulation, image color and resolution modification, time manipulation, etc. Audio effects systems may also be available to create effects such as Foley sound effects, alteration of sound, and time manipulation.

Another critical element of the edit facility is the edit decision list (EDL). This is the instruction list that defines how the individual elements are to be assembled into the final program. Most edit systems are completely independent. The EDL is created during the edit session and resides on the edit controller. However, in many cases an "Offline" system is used to create an EDL. In this case the EDL needs to be moved to the "Online" edit system that will assemble the final version of the program.

Both the creation of an EDL as well as the final compilation of the material into a final program requires precise timing of the material from all of the sources used. It is imperative that the content from each source have an associated timing identifier. Since video is a collection of still images, called frames, each individual frame must have its own time stamp. This is called time code. Time code is presented in hours, minutes, seconds, and frames. This time code is continuously running during a moving picture defining exactly which still image is being displayed at that instant of time. Time code is relative to the image it defines. It does not specify an absolute time, but only specifies a relationship between all the frames in the moving picture stream it represents. Edit systems use the time code for each source to identify where an event is to be initiated and terminated for each element of the edit process that involves that source.

Audio is a continuous stream of data. It does not use a collection of static events, and as such, does not break up into frames. However in digital media systems the audio is broken up into time windows. These windows match the frame rate of the video. So, if each frame of video represents 16ms of time. All the audio associated with those 16ms will reside in that window. Many times the audio in a specific window will be placed in the data stream associated with the video frame for that time window. However, if the audio is recorded on a separate device from the video the audio device needs to create its own time code to define where each of the 16ms windows start and end.

Since the purpose of an edit room is to assemble different content streams into a single composite stream, the relationship of the timing between the various input streams must be closely maintained in order to avoid video and audio artifacts (flashes, cracks and pops, freezing and other undesirable effects). These timing requirements are identical to those given for a Production Control Room. The original relationship at the source between audio and video streams must be tightly maintained throughout the edit facility.

Since editing requires the transition between different images, these transitions need to both begin at the start of a video frame and end at the end of a video frame. The processing device in the edit system needs to have knowledge of the frame boundaries in the video. When these transition events occur processing device needs to be able to execute the transition events at specific frame boundaries defined in the EDL.

An edit room has audio/video monitoring sufficient to allow the operator to observe the signal either being evaluated or processed. There are source monitors that allow the operator to see sources that can be processed. There are feeds from, servers, PCR suites, and other subsystems in the facility that might be utilized in the edit session. One critical issue is that of synchronization of audio with video, typically called "lip sync". Many edit facilities have gone to one or two very large displays, which are broken up into multiple monitors with external devices. This adds additional delay to the image display that is not matched in the audio monitors. Care must be taken to assure that combination of all sources of delay to both the audio and video are accommodated in the relation between the video and audio monitors and time code presentation

Networking technology is extensively employed in edit rooms. Frequently files are transferred between video servers using high performance networks. The edit room function operation is completely dependent on data networks for proper operation of machine control. Some edit facilities receive some material, which is transmitted to the edit room over IP networks. Most of these operations employ a number of security measures to protect the networks such as preventing someone from attaching an unauthorized computer to the network or introducing a file infected with a virus into the automation system network. Extreme care is taken in connecting these networks to other networks in the facility, in some cases including an "air gap" between the edit room network and other networks, meaning that there are absolutely no connections between edit rooms and other networks. In other situations a gateway is used between the networks such that the signals pass between networks in a serial digital link that does not allow network based protocol signals to pass.

Since multiple devices are used in the edit process, control of these devices is critical. Often network technologies are used to communicate the machine control requests for the edit controller to the media devices. Device control is usually the only communication the media device receives to define the timing of the signal to be processed in the edit function. As stated above, the media timing in the edit function is extremely important. For this reason the timing of the machine control commands through the network is very time sensitive.

An edit room may consist of one editor or may have a number of people working in cooperation. In addition to the editor there may be a producer who is responsible for all aspects of the program being assembled. There may be a director who is responsible for the creative decisions of the program. There may be a client for whom the program is being created. There may be specialists for color correction, audio and video effects, or other unique skills in the edit room. Sometimes these other functions have their own monitors to view. All monitors in the process must be identical with respect to time and quality, especially colorimetry, to assure the accuracy of the edit process.

Summary of user requirements for an edit room:

- All video and audio signals involved in an edit session must be synchronized with one another. Video to video timing must be within the input buffer window of the video production devices that combine these signals (typically less than the duration of one line of video). Similarly for audio, the signals must be within the input buffer window of the audio mixing device (this is a digital audio requirement).
- Highly trained and capable maintenance personnel must be available during the edit session time. Edit sessions can be either very time sensitive, as in news editing, or very expensive, as during production editing. Extended down time translates into either failed deadline for news systems or lost revenue for high end production systems
- Audio to video “moment in time” match must be less than one frame of video, with the audio never preceding the video. Note: this is often referred to as “lip sync”.
- System output signals shall be isochronously streamed with a very low latency to the monitor system and master recording device. Streaming video shall never freeze nor shall audio ever go silent. There shall be no visible or audible distortions to, or discontinuities of the media streams due to bandwidth restrictions.
- Media signals meant for monitoring shall be, as much as possible, equal in both time and quality performance to on-line signals. As monitoring signals are often used for timing and quality evaluation and for confidence/continuity purposes, any changes to them, separate from the on-line signals they represent can affect incorrect decisions during the creative process. When absolutely necessary, audio shall be no more than one frame of video later than the video signal it is associated with. Audio shall never precede the video signal.
- Signals that send commands to devices shall be delivered without delay, in a deterministic fashion and with latency in the ballpark of a few lines of video, at most. Reaction to the command is expected by the beginning of the next frame of video.
- Signal pathways often incorporate physical “patch bays” or “jackfields” so that the room configuration can be changed from its norm for operational, or for emergency bypass purposes. It is required that packet switched networks also accommodate this physical patching.

3.2.7. Remote Feeds from a Live Event

In live sports program production multiple cameras and microphones are often used to cover the event. These sources are delivered to a video production switcher and an audio mixer to produce a finished program in real time. In some cases there are isolated (ISO) feeds from the individual cameras set up to continuously capture specific views. In one use case, these ISO feeds are recorded locally at the venue to be played back when requested by the producer and inserted into the production, typically a replay of an event, sometimes in slow motion. In another variation of this use case, the live ISO feeds are sent to a remote Production Control Room(s) (see use case above) for remote production of the live event. In yet a third variation of the ISO feed use case, the live feeds are recorded for later use in an editing session, where the editor can select views to use in the finished master, and at a more leisurely pace than in live production.

Packet switched networks are commonly employed for these remote use cases as the cost of providing multiple media channels via terrestrial packet switched network links is significantly lower than the alternate and older approach of using satellite links.

In the remote production use cases, the remote ISO feeds are individually transported back from the venue to the remote broadcast center that may be hundreds or even thousands of miles away. Sometimes there are multiple independent remote production facilities, as is often the case with widely popular or public events. This arrangement is referred to as a “pool feed.” Examples of this are a political debate or ceremony, the Olympics, or other similar worldwide tournaments. By remotely transporting these feeds back to the broadcaster center(s), the production can be performed in a more extensive facility, according to the particular user’s needs, while simultaneously being safely archived for future use. With the correct infrastructure in place, these ISO feeds may be extended to end viewers, allowing them to select their views and replay shots via companion screen applications. Such applications rely on receiving all ISO feeds, having the correct metadata assigned to each feed to ensure that the scripting and annotations allow those downstream applications to trigger the correct feed on demand.

The ISO feeds are often required to be synchronous and timed to one another with each portraying the same moment in time. This is required in remote production of the live event so that the various ISO feeds can be composited. Audio and video must also be in “lip-sync”. Stereoscopic left and right eye views must also meet all of these timing considerations. Since these feeds must be timed in the same manner as if the production was local to the event, the packet switched network from the venue to the broadcast center must be capable of ensuring that the timing information is retained.

3.3. Successful Implementations

What follows are some recent examples of successful implementations of IT-centric professional media systems at the time of this publication.

3.3.1. Use of JPEG 2000 over IP at Big Ten Network

The Big Ten Network (BTN) is a regional sports network in the U.S.A. dedicated to the Big Ten Conference, jointly operated by the conference and Fox Sports.

BTN developed a “Fly Pack” for economical production of live events. This included three HD cameras, a video switcher, and a T-VIPS [at the time of this publication known as Nevion] JPEG 2000 over IP encoder. The IP connection back to BTN was over OmniPoP, a fiber optic network collaboration between participating members of the Committee on Institutional Cooperation (CIC, a consortium of the Big Ten universities plus the University of

Chicago), connecting them to each other and to research hubs worldwide with dual 10 Gbps links. BTN established a 1 Gbps link between BTN and Northwestern University's downtown Chicago campus.

JPEG 2000 video rates were typically 80-100 Mbps. BTN VLANs were created on the network of the venue university to the fly pack location, and a BGP peering session was established between BTN and each university router. VoIP was used for communications between BTN and the production staff at the venues.

Some events were also produced by using four HD camera feeds brought back to production control rooms at BTN studios using 4 T-VIPS JPEG 2000 over IP encoders.

Over 550 live events were produced by IP fly pack in 2012.

3.3.2. Additional Uses of JPEG 2000 over IP

The FOX broadcast network in the U.S.A. uses JPEG 2000 over IP for two use cases. The first is live video backhauls from American football stadiums and other live venues to the FOX Network Center LA (Pico). Connectivity to Pico is through dual diverse paths through the Level 3 network. Three Cisco DCM video gateways per venue are used as JPEG 2000 encoders and IP encapsulators. Pico can receive twenty JPEG 2000 IP feeds at a time.

The second use case is inter-facility connectivity. FOX uses three bi-directional JPEG 2000 IP paths between Pico and Fox Network Center Woodlands, Texas. Cisco DCM video gateways are also used in this case. The JPEG 2000 IP streams are carried over dual private 10 Gbps circuits. These paths are used for many kinds of video services, including live games, studio feeds, tape delayed programs, and commercial feeds.

FOX JPEG 2000 IP streams use a service bit rate of 150 Mbps with 127 Mbps for JPEG 2000 video in streaming MXF in RTP. These feeds include 720p59.94 video with 8 embedded audio channels. "Hitless switching" redundancy for the streams is provided at the IP level by sending two copies of every RTP packet over the dual physically diverse paths. Should an RTP packet not arrive on the primary path, the receiver utilizes the RTP packet from the secondary path. IP RTP packets are multicast, and IGMP is pre-configured on routers.

The use of JPEG 2000 has been a significant improvement over MPEG-2 compression due to reduced latency and the use of hitless switching has reduced visible video errors.

3.3.3. Electronic News Gathering over Packet Switched Networks

One area where packetized network transport of live video is in extensive use today is in the case of Electronic News Gathering, sometimes-abbreviated ENG. This use case describes a typical live ENG application using packet switched network transmission.

In the late Fall of 2012, a hurricane made its way up the Atlantic seaboard of the U.S.A., and came on shore in the state of New Jersey. The northeast side of the hurricane, the area that typically produces the heaviest rainfall and wind conditions, impacted the New York City area, producing record flooding and extreme weather conditions. During the height of the storm, a construction crane high above Columbus Circle near Central Park in the middle of New York City partially collapsed, threatening to fall into other buildings and onto the streets below.

The initial signals put on the air by news organizations covering the crane collapse came from mobile devices such as iPads and various smart phones. Shortly thereafter, networks deployed ENG vans – trucks outfitted with the equipment necessary to produce air-quality live feeds of breaking stories. Networks deployed vans that included bonded cellular technology as the primary means of transmitting live signals from the remote event to the studio. This use case will briefly describe the technology involved in these critical live event transmissions.

Bonded cellular devices have this name because these units typically include cell phone transceivers from multiple carriers, and the devices internally bond the bandwidth available from the carriers into a single pipe which is made available for transmission of the audio/video signals from the van back to the studio. Importantly, these devices provide two-way transmission, although the two-way capability is not typically used to transport video from the studio back to the ENG van. Bonded cellular devices are also sometimes known in the industry as backpack units because the technology has also been deployed as a portable, battery-powered backpack unit.

While these devices are known as bonded cellular devices, they typically can also connect to WIMAX, WiFi and wired Ethernet networks. In most applications today, this bandwidth is used to transmit the ASI output of a compression system. The compression technology most frequently deployed today is MPEG-4 AVC, although MPEG-2 is also used in some cases.

In the typical use case such as the hurricane event described above, an ENG van arrives on the scene, the equipment is turned on, and the device establishes connectivity with cell phone carriers, WIMAX carriers, WiFi devices. The operator confirms connectivity, typically by establishing an IP connection with the computer network at the studio, while simultaneously connecting with video streaming devices provided by the bonded cellular device manufacturer. In turn, this manufacturer provides a live link to the broadcaster. Cameras and microphones are deployed, and the feed of video and audio from the ENG truck to the studio begins.

These units are typically capable of transmitting one HD signal including two channels of audio. This technology also allows the ENG van to transmit video and audio as files, allowing the crew to create a news package (a complete news story) and feed it back to the studio. Also users can access newsroom computers and other devices at the studio through an IP network connection.

Multiple manufacturers have deployed these units for a number of years and they provide hundreds of hours of live feeds per day that appear on stations in major markets all over the U.S.A.

As might be expected, some users report that they do not rely on these units exclusively, but also deploy conventional 2 GHz microwave units, especially in cases where earthquakes or other man-made events might impact data networks. However, these units are being used every day for important news stories in major U.S. markets.

Summary of user requirements for bonded cellular devices:

- Connect to various providers of bandwidth, whether cellular carriers, WIMAX carriers, WiFi units or hardwired Ethernet connections and aggregate that bandwidth together and present it as a single pipe with sufficient bandwidth to allow transmission of live video streams to a broadcast facility
- Dynamically adapt to changing bandwidth availability across any available links
- Provide sufficient bandwidth for transmission of at least one HD feed with associated stereo audio
- Allow for audio/video file transfer when live transmissions are not in progress
- Provide IP connection to broadcast facility computer network while in the field

3.4. The Requirements of Professional Media Networks

Specialized industries have specific computing requirements. These tend to manifest themselves as performance characteristics within the networking and computational systems used by those industries. Many of them were early adopters of digital computing, their workflows and computational requirements evolved along with IT. Unfortunately, this is not true for the PM industry, which is a relatively late IT adopter. However, the PM industry has matured to the point where innovative thinkers are now readdressing the concepts and workflows of old, with a full understanding of the IT world, opening the way for new possibilities.

The traditional, "best effort" nature of general purpose IT networks does not yield the required low latency, high bandwidth, isochronous & temporaneous, and deterministic requirements of PM systems. This does not mean that the IT industry does not already have technologies and methodologies available that address these requirements; it simply means that the identification of these technologies and methodologies can be difficult, as the two industries are not yet speaking the same language.

3.4.1. Real-time media delivery vs. file-based workflows

There are two primary workflows for media essence within a PM system, real-time delivery and file-based delivery. Real-time delivery is when the media essence must be transmitted to one or more clients in an isochronous & temporaneous fashion. Real-time essence delivery requires both high bandwidth and isochronous & temporaneous delivery, meaning that the essence consists of large volumes of data that must be transferred in equal intervals, arriving at the right time. If a data packet cannot be delivered within the temporal limitations set by the real-time presentation rate (frames per second or samples per second), the data packet is no longer valid and should be discarded to prevent issues with the presentation. Real-time delivery is typically used to deliver essence for presentation, such as viewing IPTV and for the recording of real-time input sources, such as cameras and live broadcasts.

Unlike real-time delivery, file-based delivery does not have a temporal limitation; packets can be transmitted at any rate and may be transmitted out of sequence, provided the proper sequence is reconstructed when the file is written or when the file is next read. In file-based workflows, depending on the urgency or priority of the media transmission, files can be transmitted faster or slower than real-time, meaning faster or slower than the temporal component (timeline) associated to the media essence.

For example, media files that are scheduled 24 hours in advance of their broadcast can be transferred at rates slower than real-time. If the media files must be ready for air 12 hours before their broadcast, a 12-hour window in which the media files can be transferred is available. In other cases the media files may have a greater urgency, requiring faster than real-time transmission. For example, if an ingest system possesses 200 hours of essence storage, but ingests 400 hours of essence a day, the media transferred from the ingest system to an external storage system must be transferred at least 2 times faster than real-time. The best-effort transfer rate naturally available on networks may suffice for some applications, but well-engineered PM systems should manage the rate of transfer for media essence, maximizing the available bandwidth for high priority transfers. This can be expressed mathematically as:

$$\boldsymbol{BPS} = \boldsymbol{B} \times (1 \div (\boldsymbol{T} \div \boldsymbol{D}))$$

where

$\boldsymbol{BPS}$ = The Minimum Transfer Rate in bits per second
$\boldsymbol{D}$ = The Duration of the Media File in the same time units as $\boldsymbol{T}$
$\boldsymbol{T}$ = The Amount of Time until the Media File is needed in the same time units as $\boldsymbol{D}$
$\boldsymbol{B}$ = The Constant or Average Bit Rate of the Media File in bits per second

3.4.2. Essence vs. control

Not all network traffic is created equal. Professional Media Networks (PMNs) typically have three classifications of network traffic, control, essence, and everything else. Control traffic is used to access state, manipulate state, and perform action on devices and services that reside upon the network and require deterministic, low latency, behavior. Control traffic typically consists of small, low bandwidth, messages transmitted between two endpoints, but some applications, such as IEEE 1588-2008 Precision Time Protocol (PTP) may transmit control messages to more than one endpoint using multicast messaging.

Essence traffic, on the other hand, typically consists of large data streams requiring high bandwidth and typically utilizes large data packets to optimize performance. All other forms of network traffic on a PMN are ancillary and should not interfere with the transfer of either control messages or essence streams. Of the three traffic classifications control traffic typically has the highest priority, as it controls the flow of essence over the network and is used to execute time sensitive operations.

3.5. Characteristics of Professional Media Systems

Most professional media systems share a common set of characteristics. While not all systems require all of these characteristics, it is common to strive for as many of these requirements as feasible.

- Robust and Resilient System – A minimal set of features must be available 100% of the time. Systems must be designed to resist both equipment and human errors.
- Isochronous & temporaneous streaming – When presented for viewing/listening at any point it is required that the media be presented in a smooth and contiguous fashion.
- Low Latency – Especially in live production, and in many other cases, it is important to ensure that the delay in presenting media be minimized.
- Market Leverage – Businesses centered on the production and delivery of professional media, while large, are not of significant size as compared to the Enterprise market. Therefore, the ability to influence the direction of new networking products and technologies is much diminished as compared to influence upon the professional media market.
- Data loads – Modern production quality video is quite large in file size when stored, and in bandwidth requirements and duration when streaming.
- Interconnectivity – Whatever forms of media transport are chosen, there must be a means to move the media amongst them, whether within the system or to other systems.
- Source alignment – The compositing of media mandates that multiple sources be aligned in frequency and phase. In common vernacular, they must be genlocked and timed.
- “Event” alignment – Production of media involving multiple sources, audio video, and stereoscopic views mandates that the multiple media sources portray the same moment in time. In common vernacular, audio and video out of alignment is referred to as out of “lip sync”.
- Operational flexibility – Systems must be designed with an appropriate balance of flexibility and security. Protection mechanisms must be applied judiciously.
- Tolerance of media data errors – Media data streams are unlike other types of data, such as financial data. When media is displayed or listened to, some errors are nearly imperceptible. Often having even an error prone media stream is preferable to having none at all.
- Maintenance – One must have the ability to know how media is moving through the system at any moment, and route around defective equipment.

4. Managed Networks

This section covers the salient aspects and advantages of a managed network system.

What is a Managed Network?

It is a network system (or individual elements) that is actively controlled and monitored, all or in part, across the spaces of faults/warnings, configuration/provisioning, performance, accounting, and security. Managed networks provide predictable and assured performance across a wide range of requirements and fault scenarios.

What is an Unmanaged Network?

Devices and systems are initially installed but not actively managed. This is an "install and forget" approach. Any required changes to the configuration are done reactively not proactively. Of course, this may seem ideal from the standpoint of simplicity.

What is a self-managed network?

A self-managed or plug-and-play network requires little or no configuration. Small networks may be built with unmanaged switches, which have minimal or no configuration capability. Some plug-and-play network technologies are self-configuring. Examples include audio video bridging (AVB) extensions for Ethernet and peripheral networks such as USB and Thunderbolt.

4.1. Managed Networks Basics

Networks and transmission links can be divided across four main "Connectivity Domains". Figure 2 outlines instances of these domains. The remainder of this section will consider these domains under active management.

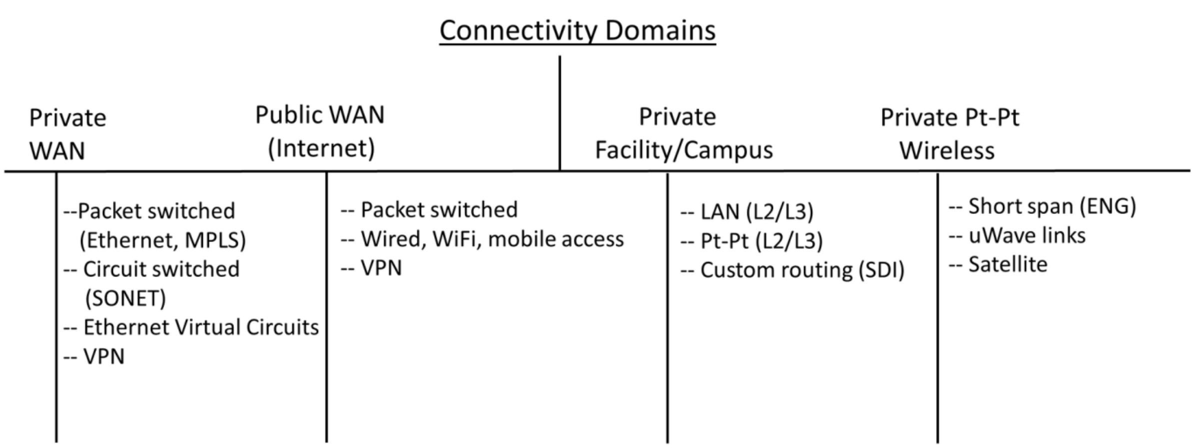

Figure 2 – Generic Connectivity Domains and Options per Domain

It is beyond our scope to discuss each entry in Figure 2. Individually, depending on capability, the listed options may be used to move media intra facility, inter facility or between end-points. Each element, network segment, or entire network can be actively managed with associated benefits.

Elements can be managed separately, holistically as one system or as multiple sub-systems. Using the "divide and conquer" principal it is prudent to decide when to manage an element alone or in association with a system of network components.

4.2. Main Features of Managed Networks

High-level network management methods are divided across 3 related areas; **models, processes,** and **tools**. There are several available **models** but the fault, configuration, accounting, performance and security (FCAPS) model is well respected, mature and defined by ITU Rec M.3010. **Processes** are used to realize the model's goals. **Tools** are used to implement the processes. An example will be given below. First, what is the FCAPS model? Figure 3 outlines the five elements and usage examples for each.

Fault Management	**C**onfiguration Management	**A**ccounting Management	**P**erformance Management	**S**ecurity Management
Fault/warnings/alarms; detection, notification, isolation, correction	Connectivity provisioning; in band, out of band, scale, reach, reliability, automation	Usage tracking per user, net, link	Performance data collection; Key QoS parameters (uptime, rate, latency, loss, jitter)	Selective access
Logging	Auto discovery	Billing	Non-intrusive QoS monitoring	Performance profiles adherence
In service testing and monitoring	Inventory, database tracking	Auditing	User and link profiles	Security alarms, lockout
Automatic and manual repair methods	Configuration status, dashboards		SLA conformance	Audit trail
3rd party support services	User profiles of resource		Status dashboards	

Figure 3 – FCAPS Model Outline for Connectivity Domains

The FCAPS model can be applied to any system-wide element not just network-related elements. For example, FCAPS is used daily to manage apps, servers, networks, storage, appliances and more. However, for our scope, the model is limited to network management. The well-managed network uses many, but not necessarily all of the FCAPS features. Element faults/alarms are dealt with promptly and, if possible, automatically repaired using spare capacity and resources. For link transmission, the five QoS parameters of uptime/MTBF, rate, latency, loss, and jitter are paramount. To assure these metrics are within bounds, design and monitoring steps should be established to assure the parameters are within operational range.

As mentioned, **models, processes** and **tools** are needed for a complete managed network. Let's look at an example based on the FCAPS model;

- Link failure occurs, reported by monitoring method, notifications and alarms generated (F in FCAPS **model**)
- Documented **process** describes how to resolve this particular failure (manual or auto)
- **Tools** used to monitor, re-configure or repair faulty component. Alarm cleared.

4.3. Administration Aspects

The well-managed network needs people and processes to assure compliance to the desired system-wide QoS and scale. Management personnel may be local IT staff. This is especially true for local facility networks and some regional networks. However, for many WANs, the transmission capacity is purchased from a "telecom service provider". Examples of this are MPLS networks, private lines using the services of Metro Ethernet (or Carrier Ethernet) providers and other point-to-point connectivity.

The service providers write a Service Level Agreement (SLA) that defines the extent of service provided with possible fee refunds based on non-compliance. A typical transmission related SLA will define QoS parameters including uptime, Mean Time Between Failure (MTBF) and Mean Time To Repair (MTTR), monitoring techniques and provisioning methods. Some facilities offer link and network capacity as a service to individual departments. These often have an associated SLA. The SLA is of little worth if the end-user doesn't test for compliance on a regular basis. If you can't prove conformance, you may not get what you are paying for.

Not all WAN providers offer SLA metrics that meet the needs for professional real time video transport. For example, using the wide-open Internet to move video streams from a venue to a media facility may require sophisticated forward error correction (FEC) to conceal unspecified packet loss. Each case will be different. Bottom line; purchase a service QoS sufficient to meet your needs (or can afford), monitor it, and augment with methods to mitigate problems when metrics extend beyond the preferred ranges.

If you manage a system under your direct control, it's a good policy to write a virtual SLA and monitor so it performs according to your requirements. At the core, your network long-term performance will be strongly correlated to the abilities of your management methods.

5. The Convergence of Two Industries

The Information Technology (IT) industry is ubiquitous within nearly every organization and with driving factors, such as Moore's Law, commoditization, globalization, and general customer demands, the industry is advancing at an unprecedented rate. Like any other industry, the IT industry is also seeking to expand into new markets, such as Professional Media, offering new technologies, such as Virtualization, Software/Hardware as a Service, Service Oriented Architectures, and more. These offerings sound very appealing because they offer a reduction in cost, an increase in available resources, human and technical, and minimize vendor lock-in. But how is the Professional Media Systems Architect to know which of these offerings best fit their needs, which do not function as promised, or which can negatively impact the durability of their existing systems? This section focuses on the technologies available from the IT industry that are most useful to the Professional Media Systems Architect and will likely be most useful for builders of traditional systems, who desire to have a better understanding of the benefits and risks when employing these new technologies.

This is not meant to be an all-inclusive list of every technology where Professional Media Production System designers should be fluent. Neither is it an exhaustive explanation of each topic, nor do we make specific recommendations as to use and configuration. The reader is encouraged to find detailed information elsewhere regarding any topic where they are not currently fluent. The creators of this report have, however, endeavored to include many topics of import to professional media production system designers.

5.1. Differences between Professional Media and Information Technology

Historically, there has been an appreciation of the differences between the IT and Professional Media industries. These differences were primarily different views on the management and treatment of data during transmission and storage. Today, with increased network and computing performance, coupled with the availability of media essence compression techniques, these differences are disappearing and IT-class network and computational elements can now readily accommodate Professional Media application requirements.

The primary differences between the Professional Media and IT industries are:

- Professional Media workflows typically require isochronous and temporaneous data presentation, meaning that data must be presented in equal intervals, in the right order, at the right time. IT workflows rarely require isochronous and temporaneous data presentation.
- Limited data loss, primarily in essence streams, is usually tolerable within Professional Media workflows, while data loss is rarely tolerable within IT workflows.
- Raster Based: Streaming video signals (classically) are built upon scans (or in the post CRT world, read-outs) of the video signal, divided line by line, and not packet-by-packet. Another portion of the signal called "vertical interval" represents the top of the image and both dividing points need to be maintained in order and at the right time.
- Media signals are non-retransmittable, as opposed to most data signals, which can safely and easily be retransmitted and reassembled at the receiver.

The Professional Media industry relies on error concealment for essence streaming. Error concealment in this context is generally the repeating of the last known good frame of video until a new good frame of video is received. The expectation is the perception of an error free transmission and if only 1 to 2 frames are repeated, the consumer will rarely notice the error. This is not true in many IT workflows, where even the slightest loss of data can have a huge impact on the end results.

The compression techniques used by both industries also illustrate the data loss differences. Most compression methods used within the Professional Media industry, for compressing media essence, provide controlled loss of essence. The assumption is that not all of the audio-visual data is required to recreate a suitable representation of the original. The IT industry however relies mainly on lossless compression techniques, because even the slightest loss of data could have a significant impact on the final result.

5.2. Common Ground

Despite historical differences, both the Professional Media and IT industries share a lot in common. Both industries share the same concerns and except for the treatment of audio-visual essence, both industries have similar requirements and expectations from their users. The IT industry also has specialized workflows where the treatment of data is nearly identical to the treatment of data within Professional Media workflows.

High Availability – Reliability, Resilience, and Robustness

Reliability, Resilience, and Robustness are three key fundaments of Professional Media workflows. Both the Professional Media and IT industries share this same desire, to ensure their systems are reliable, resilient, and robust, to the greatest degree technically feasible for the budgetary constraints.

Service Level Agreements (SLAs)

Service Level Agreements (SLAs) are agreements between two parties, the service provider and the service consumer, which guarantee a minimum level of performance and availability for a service or set of services. Service Level Agreements are used by both the Professional Media and IT industries and typically specify penalties or consequences service levels are not maintained.

Ease of System Maintenance

Network and computing resources have become highly complex and require an orchestration of hardware, software, and network components to fulfill workflow requirements. The design and construction of such an orchestrated system can be very complex, requiring volumes of configuration. Both the Professional Media and IT industries have a strong desire to make the process of configuring and maintaining these highly orchestrated systems as easy as possible.

Low Total Cost of Ownership (TCO)

Total Cost of Ownership (TCO) is a concern in most industries, including the Professional Media and IT industries. The Total Cost of Ownership ultimately controls the profitability of an operation and is one of the key constraints for system designs.

The Use of Commercial Off-The-Shelf, Commodity Resources

The use of off-the-shelf or commodity hardware is desired by both the Professional Media and IT industries for numerous reasons, such as interoperability via the use of standards and helps to prevent vendor lock-in.

Amongst its other benefits, the use of off-the-shelf or commodity hardware also helps in the Total Cost of Ownership (TCO) of a system. The use of commodity hardware allows equipment to be mass manufactured, reducing the price of the equipment and also reduces the cost in human resources required to maintain the equipment.

The Efficient Use of Resources

Waste can be a large factor in the Total Cost of Ownership. In many cases dedicated hardware remains idle for a significant amount of the time. Both the Professional Media and IT industries are looking into ways to reduce these wasted computational cycles. Reducing the amount of hardware required to build a system, reduces the electricity and cooling costs and can also reduce the human resources required to maintain the system.

6. Information Technology Solutions for Engineered Networks

There is a tremendous wealth of technology, resources, and knowledge available within the Information Technology (IT) industry that pertains to the treatment and conveyance of data. Nearly every conceivable data requirement has been addressed in one form or another; the trick is finding the right solution for a specific need, but before one can even begin the search for a solution, they must have a basic understanding of basic networking principles, the technologies that are available, and the dialect commonly used within the IT industry.

6.1. Network Models and Stacks

A network stack is an implementation of a computer network protocol suite. While a network model is a conceptual model of the functional levels of a network stack. The two are related, as the design of a system is related to its implementation.

6.1.1. OSI Model

The OSI Model was defined by the ISO as a 7 layer abstract model of networking and set of specific protocols. The OSI model is not concrete, but is critical for network related communications. The layers of the OSI Stack are illustrated below:

OSI Model			
	Data Unit	Layer	Function
Host Layers	Data	7. Application	Network to Application Interface, API
		6. Presentation	Representation, Encryption, Abstraction
		5. Session	Session management between clients
	Segments	4. Transport	End-to-End Connections, Flow Control
Media Layers	Packet/Datagram	3. Network	Logical Addressing, Path Determination
	Frame	2. Data Link	Physical Addressing, Data Framing
	Bit	1. Physical	Hardware and Binary Transmission

Figure 4 – OSI Networking Model

1. Layer 1 – Physical Layer: The physical layer defines the electrical and physical specification for the devices.
2. Layer 2 – Data Link Layer: The Data Link layer provides the data framing including the physical addressing
3. Layer 3 – Network Layer: The Network layer packetizes the datagrams to allow for logical addressing and path determination.
4. Layer 4 – Transport Layer: The Transport layer establishes the connection between the endpoints and provides reliability and flow control.
5. Layers 5 – Session Layer: The Session layer provides inter-host communication and session management between applications.
6. Layer 6 – Presentation Layer: The Presentation layer provides the platform independent data representation, translating platform dependent data to platform independent data.
7. Layer 7 – Application Layer: The Application layer connects the network stack and the application.

To the Professional Media Architect it is important to understand that the OSI model is a description on how the signal is broken down and communicated across the system. Each layer defines an interface that is used by the higher layers, thus allowing low layers to be swapped out with different implementations without impacting the higher layers.

6.1.2. Internet Protocol (IP) Stack

TCP/IP model does not match the OSI stack 100% since the concepts are quite different. There are no strict layers as in the OSI model. TCP/IP does recognize 4 layers of functionality that are:

1. Layer 1 – The Link Layer is equivalent to Layers 1 & 2 of the OSI model. Providing data framing and access to the physical hardware.
2. Layer 2 – The Internet Layer (IP) translates the physical addresses of the hardware to the logical addresses used by the IP protocol. In addition to address virtualization, the Internet Layer provides the translation between datagrams/packets to hardware frames and is used to establish paths between network hosts for internetworking.
3. Layer 3 – The Transport Layer provides host-to-host communications.
4. Layer 4 – The Application Layer contains all of the protocols for data communications on a process-to-process level.

The IP Stack allows system components to communicate over the network without the component having to know the actual hardware address of the destination device. Each device is assigned an IP address. The IP address is a human readable address that is mapped by the system to the MAC address of the hardware. Being a virtual address, the device identified by the IP address can be changed without impacting the clients of that device.

6.1.3. Network Stacks (Conclusion)

The OSI and IP Network Stacks provide an abstraction over the physical hardware, allowing for one wire to transmit multiple signals, each independent from the other. In contrast to protocols such as SDI that only have 2 defined layers, Physical and Data Link.

OSI Model	Protocols				IP Model
Application	AVB (API)	Telnet	HTTP	FTP	Application
Presentation		ASCII/Binary/SSL			
Session	AVTP	gPTP	Sockets	NetBIOS	
Transport			UDP/TCP		Host to Host
Network			IP		Internet
Data Link	Ethernet/Infiniband/?				Network Interface
Physical	Physical				

Figure 5 – OSI to IP Stack Comparison

6.2. IP Addressing

IP addresses are abstract and virtual; each device obtains an IP address, mapped to its MAC address, from either static configuration, system services, such as DHCP or via Router Advertisement. Before 2 network devices can communicate the IP addresses must be translated to the MAC addresses. This translation is performed using the ARP protocol. The IP Address defines a network host, but provides the abstraction required to be able to swap out hardware of the device without affecting the clients or other devices within the network.

6.2.1. Subnetting

In the 1980's and 1990's the concept of IP address classes was introduced. Originally address classes were used to lay out address blocks for the Internet, but as addresses became a commodity the concept of subnet classes was thrown out and replaced by a more dynamic method of subnetting, using the bit-boundary of the subnet mask. By the early 1990s, it was recognized that the classful model not scale and in 1993 Classless Inter-Domain Routing (CIDR) was introduced. This enabled all bit boundaries to be used to define the network subnet mask and therefore break away from the rigid classful model initially designed. Subnetting is used to keep traffic confined within a physical or virtual location, but too much subnetting can become difficult to manage. Subnetting generally lends to routing, please refer to the Routing section for more details about routing.

6.2.2. IPv4 and IPv6

IPv4 is a 32-bit addressing model that was introduced in the early 1980s. There was no concern about address space exhaustion since the model was designed to fit the purpose of the DARPA's original mandate of testing networking concepts. By the end of the 1980s, the growth in allocated addresses was such that subnetting wasn't going to be sufficient to solve the address space management problem. Work was carried out to design a new scheme completed by 1996. The most visible difference is the 128-bit addressing model and the primary purpose of the new scheme. Many other changes related to protocol behaviors are included, some of which are discussed later in this document.

6.2.3. Address Management

Being abstract and virtual, IP addressing requires management. Some categories of devices such as those originating and terminating services (e.g., video servers, IP to SDI gateways, video mixers) require a means to consistently obtain the same IP address; other devices may not always require the same identical address but, insteadone from a pool of addresses assigned to a specific device category. Addresses tend to be a limited commodity and carved up into pools for specific applications and therefore must be managed to ensure consistency and resource management.

Such address assignments may be provisioned and statically assigned to an interface by means of configuration within the device itself, or dynamic via an external process such as DHCP (Dynamic Host Configuration Protocol) for IPv4 hosts, DHCPv6 for IPv6 hosts or SLAAC (Stateless Address Auto-configuration) for IPv6 hosts. Note that a dynamic provisioning mechanism doesn't imply the address has to be different for each request. It can be mapped to the MAC address of an interface and always provide the same address for that specific request.

While DHCP and DHCPv6 rely on the response from a specific request initiated by the host to which the DHCP server or proxy will respond to, SLAAC relies on NDP (Network Discovery Protocol) whereby the router for the local subnet announces the network prefix to be used as the first portion of the address and the 2nd part is derived from the MAC address of the physical interface.

A DHCP server doesn't need to be collocated on the same physical medium as the interface requesting an IP address. Many routers and switches support a DHCP proxy mechanism that can forward a DHCP request originating from an interface that is directly connected to them to a remote DHCP server. This provides scalability for DHCP management and is a common mechanism in large DHCP deployments.

In the case that an interface requests is not fulfilled or that none of the dynamic mechanisms are available, a Link-Local mechanism enables communications within a network segment between hosts without the need to rely on an external entity to provide IP addresses. These IP addresses are self-generated by the host and therefore routers should never forward packets originating from them.

6.2.4. In-band and Out of band management

For ease of management, many devices have multiple interfaces whereby one is dedicated to managing the device via terminal sessions, remote desktop, web interface, SNMP and/or any other management protocol. Such an interface is typically connected to a management network. On the other hand, all data sessions related to the primary purpose of the device is sent and received by one of multiple other interfaces that are connected to the main network. This is called out of band management vs. in-band management when all the management traffic and data sessions reside on the same interface(s).

Traffic separation models ease the administration of devices by separating the roles for each component of the device, helps with enforcing security by isolating the traffic and is typically viewed as best practice in many IT and Service Provider deployments.

Some equipment extends such concepts with Lights Out Management (LOM) and/or serial interfaces that can be used even if the main system is powered off or has crashed. These options offer future remote management capabilities and are found on many servers and embedded systems.

6.3. Layer 2 – Data Link Layer

6.3.1. Ethernet

Ethernet is comprised of a series of standards that define the physical and data-link layers of the OSI networking model. The term "Ethernet" is commonly used to refer to the data-link layer portion of these standards. Ethernet was first introduced in 1980 and standardized in 1983 as IEEE 802.3. Most LANs in use today utilize Ethernet as the primary data-link layer protocol on which layer 3 protocols such as IP are carried. The physical layer specifications that make up Ethernet include many wiring specifications such as 10BASE-2, 10BASE-5, 10BASE-T, 100BASE-T, and 1000BASE-T. Originally, Ethernet predominately used coaxial cables to create serial runs, but this was superseded by the use of switches and hubs to form a star based network architecture. Ethernet is lossful, meaning that there is no guaranty that data will be delivered, therefore elements at the higher layers of the OSI stack must be developed to guaranty the delivery of data. In recent years Ethernet data rates have increased to 100 gigabits per second.

6.3.2. IEEE 802.1 Audio Video Bridging (AVB)

In 2006, the IEEE 802.1 LAN and Bridging Architecture Working Group started a special task group devoted to efficiently managing time sensitive media streams, called IEEE 802.1BA-2011 IEEE Standard for Local and metropolitan area networks - Audio Video Bridging (AVB) Systems, based on a groundswell of interest in supporting high-quality, plug-and-play AV streaming over Ethernet. This group has been active and growing ever since, and recently changed their name from "Audio Video Bridging" (AVB) to "Time Sensitive Networking" (TSN) to address a wider scope of critical real-time applications beyond media transport, but otherwise has the same charter.

This group developed two new standards:

- IEEE 802.1BA-2011 Audio Video Bridging (AVB) Systems
- IEEE 802.1AS-2011 Timing and Synchronization for Time-Sensitive Applications in Bridged LANs (gPTP).

Along with two amendments to IEEE 802.1Q-2011 Media access control (MAC) Bridges:

- IEEE 802.1Qat-2010 Stream Reservation Protocol (SRP)
- IEEE 802.1Qav-2009 Forwarding and Queuing Enhancements for Time-Sensitive Streams (FQTSS).

The two amendments were eventually reincorporated into the latest maintenance release of IEEE 802.1Q-2011 Media Access Control (MAC) Bridges, as new clauses 34 and 35. IEEE 802.1Q-2011 is the definitive standard that defines how an Ethernet bridge should operate, and it is widely accepted and deployed today. The AVB amendments are thus expected to become standard features in Ethernet switches, just as VLAN features, Attribute Registration, and Tagged Frames have become over the past decade. At present, it is difficult to project how long this will take; surely it will depend on market demand for AVB as perceived by the switch vendors.

The above standards are available for free download at http://standards.ieee.org/about/get/ (for the 802.1Qat and 802.1Qav standards, just download 802.1Q-2011).

IEEE 802.1BA-2011 describes the overall architecture of an AVB time-sensitive network, defining media sources as "Talkers", media destinations as "Listeners", and AVB-aware bridges as "AV Bridges". It also describes the requirements on the operation of each of these elements and how they can be interconnected in order to meet the performance requirements for media streams.

The media network management system uses IEEE 802.1Qat-2010 SRP works dynamically to verify that every link in the media stream path from Talker to Listener has adequate remaining capacity to support a new media stream before the stream is set up, and that future stream requests will not oversubscribe any of the links. This is called "admission control" because the AV Bridge will reject any attempt to establish a new SRP stream that would over-subscribe any of the links in the stream path from the Talker to the Listener. Once a stream reservation is successful, the network infrastructure defends its reserved bandwidth (except for emergency interruptions which are supported via a higher priority code that usurps AVB traffic priority).

IEEE 802.Qav-2009 (FQTSS) amends the IEEE 802.1Q-2011 MAC Bridge & VLAN standard by adding a traffic-shaping feature for AV Bridges that controls the queuing of time-sensitive media stream packets over each of the AVB-aware ports in the bridge. The bridge's packet queuing algorithm uses either priority-based queuing, or else the new credit-based traffic shaper used to reduce the chance of packet "bunching" at any node in the network. The AV Bridge uses either priority or credit-based traffic shaping defined in IEEE 802.1Qav-2009, along with the stream admission control features in the Stream Reservation Protocol (SRP) described in IEEE 802.1Qat-2010, to ensure that the AV Bridge will not over-allocate data on any port beyond its capacity.

These two protocols work together to ensure that critical real-time streams (e.g. media) and best-effort traffic (e.g. file transfer, web, email, etc) can coexist on the same network without disturbing the flow rate and latency of the streams. Latency is kept to a minimum (typically 1 to 2 msec), and is bounded due to AVB's deterministic treatment (queuing and forwarding) of reserved stream packets. This convergence of what has traditionally been two separate networks offers a lot of promise for simplification and cost reduction of facility infrastructure.

The AVB standards form the basis for time-constrained media networking at the MAC layer (Layer 2 of the ISO model). In today's world, these services are not broadly available across the generic routed networks because they rely on non-routable protocols. Instead, they are typically deployed within the boundaries of the Local Area Network (LAN) itself. Nevertheless, these AVB services can support many audio and video studio applications very efficiently since they do not rely on higher layer protocols, which can add some amount of overhead.

Size Limitations of AVB networks

As of today, AVB systems have been tested with reliable operation up to about 1000 concurrent streams on the same LAN. Limitations exist within the Stream Reservation Protocol's ability to pack large numbers of attributes into its protocol messages, and also in the number of multicast MAC addresses that a bridge can support. A worst-case situation can limit the size to 318 talkers sending one stream each; however it is typical that talkers send multiple streams having common attribute values (e.g. Stream Class, Bandwidth, and MSRP Accumulated Latency), and in this case a talker can send many streams without hitting a limit. It is also possible (and perhaps typical) for multiple talkers to share common stream attributes, in which case the number of talkers can also scale well beyond 318. Two known AVB bridges on the market today can support up to 1024 and 4096 multicast MAC addresses, respectively. Although unicast destination addresses (physical MAC addresses) are not recommended for AVB, they can be used if necessary to overcome a maximum stream count imposed by these table sizes in certain bridges.

Of course, the bandwidth capability of any given link also imposes an upper bound, which can impact system scalability. The topology of the network, and the choice of 100M, 1G, 10G, 40G, or 100G port speeds, are critical factors to consider in provisioning an AVB network for deployment. While AVB supports a default of 75% of a given links bandwidth allocated to AVB traffic, this parameter is programmable and can be set as high as perhaps 90% depending on the use case and traffic mix expected.

Integration of RTP streaming and AVB services

Many streaming media applications use RTP over UDP/IP as their fundamental “lingua franca” for streaming. To reuse this large base of RTP media applications, active and concerted efforts are underway in the IEEE, IETF, and AVnu Alliance (see below), to define how RTP streams will employ AVB services in an IP-based (routed) network.

The IEEE 1733-2011 standard was created in effort to define integration of RTP media streams with IEEE 802.1 AVB protocols. However, in the past few years, new efforts and standards developed within the IETF have made IEEE 1733-2011 obsolete. These new IETF RFC’s have adopted a more favorable approach to utilizing AVB services with RTP streams, for example by using RTP header extensions. A survey of key players in the media streaming and networked AV equipment industry (including those who developed IEEE 1733-2011) confirms that there has been no commercial adoption of the IEEE 1733-2011 standard (and none is expected in the future).

The AES67-2013 – AES Standard for audio applications of networks – High-performance streaming audio-over-IP interoperability – Annex C – “AVB network transport”, Annex D – “Interfacing to IEEE 802.1AS clock domains”, and the IETF draft-ietf-avtcore-clksrc – “RTP Clock Source Signaling” provide suggestions beneficial for the carriage of RTP over AVB and for the synchronization of RTP streams with the AVB clock.

AVnu Alliance

The AVnu Alliance is an industry consortium, established in 2009, dedicated to the advancement of professional-quality networked audio and video through promoting the adoption of the IEEE 802.1 AVB standards, as well as the IEEE 1722-2011 layer-2 transport protocol and IP-based (layer-3) media transports operating on AVB infrastructure. Its mission is similar to what the Wi-Fi Alliance has done for 802.11-based products.

The Alliance is focused on applications of AVB technologies in the Professional AV, Consumer Electronics, and Automotive markets. As of August 2013, the membership of AVnu comprises 55 companies including prominent audio, video, and IT equipment manufacturers, automotive OEMs, silicon vendors, software vendors, and design services firms.

The primary mission of AVnu is to develop and deploy a compliance and interoperability (C&I) certification program for AVB end station and bridge products. As of August 2013, real products are in the C&I lab undergoing testing, and the test processes are being validated and refined. Certified products are expected to be on the market by 2014. Testing is currently based at the University of New Hampshire's InterOperability Lab (UNH-IOL), which has been a leader in testing data communications equipment and proving out network-related technologies since 1988.

More information on the AVnu Alliance can be found at www.avnu.org.

6.3.2.1. IEEE Standard for Layer 2 Transport Protocol for Time Sensitive Applications in a Bridged LAN

IEEE 1722-2011 IEEE Standard for Layer 2 Transport Protocol for Time Sensitive Applications in a Bridged Local Area Network defines several specific Ethernet frame formats for transporting audio and video using the services provided in IEEE 802.1 AVB networks. . The latest revision, IEEE 1722a, is still in development, but it will add uncompressed HD video (SDI payloads) and several compressed video modes. IEEE 1722a also adds support for real-time control streams, clock reference streams, and encrypted payloads.

IEEE 1722-2011 utilizes a presentation timestamp embedded in each media stream header to ensure synchronized "playout" (or other treatment by the listener device) of the media relative to the IEEE 802.1AS-2011 clock. This allows any streams originating from synchronized source (or sources) to remain synchronized after being transported across the network (and subjected to uncertain delivery time). Since AVB guarantees an upper bound on network-imposed latency, the presentation timestamp merely needs to be offset from the transmit time of the media by this upper bound (or greater); this will ensure that the media is available for "presentation" within the listener at the prescribed time.

Since IEEE 1722-2011 a non-IP based protocol, it is constrained to operate across IEEE 802-based local area networks only. It cannot operate over IP-routed networks spanning multiple subnets.   Since many studio networks use a single bridged LAN or subnet, this is usually not a problem.   It may be worth noting that discussions are recently underway in the IEEE AVB/TSN working group to define a means for tunneling IEEE 1722-2011 payloads inside IP datagrams. This would allow 1722-based devices on different subnets to stream across a WAN, for example, albeit without the strict performance guarantees provided within the same subnet.

6.3.3. Data Center Bridging

Another LAN-based (Layer-2) technology that has been identified for possible use in providing QoS for media streams is called "Data Center Bridging" (DCB). Like AVB's QoS solution, DCB is standardized as a set of amendments to IEEE 802.1Q-2011 (notably IEEE 802.1Qbb-2011 and IEEE 802.1Qaz-2011). The DCB umbrella includes Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS) and Data Center Bridging Capabilities Exchange Protocol (DCBX). This technology is supported today in enterprise-class switches from Cisco and other vendors.

DCB compares with AVB in that both technologies are LAN-based and have mechanisms to ensure that no packets from "critical traffic" are dropped. AVB achieves this via reservations, while DCB uses its PFC and ETS mechanisms to prevent packet loss but do not put an upper bound on latency. DCB enforces fairness based on Traffic Class (TC) assignments placed on all traffic feeding into the network; DCB scales back the bandwidth allotted to a given sender, for a given traffic class, based on pausing of specific traffic classes (PFC) as network congestion occurs. AVB, by contrast, locks down reservations and does not change their parameters (bandwidth or latency) over time. Thus, AVB seems most suitable for data streams which have steady and constant flow rates, while DCB is most suitable for traffic having time-varying flow rates and intolerant of packet loss (such as storage traffic and high-performance cluster computing).

6.4. Name Resolution

Name Resolution is the process by which a name is translated into meaningful information, such as name to IP address translation. In IT centric implementations Name Resolution tends to be dynamic, with little importance placed on the fixidity of network addresses to which a particular name refers. However, Professional Media centric Name Resolution implementations tend to be less dynamic, placing significance on the fixidity of network addresses to which a particular name refers. When names are used within a PM systems, the Name Resolution services become critical aspects of the system, requiring low-latency and extremely high up times, relative to IT implementations. Names within the PM space also tend to be fixed and relate to a device, for example "Camera 5" or "Server 1". The names being typically based on what a device "is" or what a device "does", with a much greater likelihood that the device represented by the name will be changed, without changing the name by which it is referenced.

6.4.1. Static Name Assignment

The static assignment of names can be configured using the Host Table available within the Operating System (OS) of a networked node, as prescribed in the IETF RFC 952. The Host Table is typically represented within the OS as a file, known as the Hosts File (/etc/hosts for POSIX systems, %SystemRoot%\system32\drivers\etc\hosts for Microsoft™ Windows™ systems). While the Hosts Table provides a means for a networked node to contain statically configured names and a high performance method of Name Resolution, the use of the Hosts Table is difficult to maintain and requires significant administrative coordination to effectively manage. Due to its inherent administration complexities the Hosts Table was abandoned by ARPANET, the predecessor of the Internet, as the primary means of Name Resolution in favor of the Domain Name Services (DNS).

6.4.2. Domain Name Services

The Domain Name Services (DNS) was developed to provide a means of Name Resolution for networked nodes without the administrative complexities associated to the Hosts Table. The DNS was developed as a hierarchical database in which names can be assigned to networked node, each name residing within a namespace, known as a Zone or Domain; the name being resolved to zero or more IP Addresses. Although the primary purpose of the DNS was for name to IP Address resolution, the DNS has grown to encompass a wide variety of extensions, allowing for far more complicated tasks, such as aliasing, Service Discovery and metadata storage.

Zones (aka Domain)

Zones, also known as Domains, are defined within the DNS to provide for the distribution of the Name Resolution services to local authorities, in which the names are defined. Each name is associated to a Zone/Domain and each networked node is assigned a name. When a networked node wishes to resolve a name to an IP Address, or other information, it asks its locally configured DNS server for the information related to the name. If the local DNS server is not responsible for the Zone/Domain associated to the name, the DNS server passes the request to its configured parent server. The process repeats itself until either the name is resolved or the parent DNS server knows which other DNS server is the Source of Authority (SOA) for the Zone/Domain. The SOA is then asked to resolve the name. If the SOA can resolve the name it does so, but if the name belongs to a domain that is assigned to a child DNS server representing a sub-domain, the request is sent to the child DNS server for resolution. This process of distributed Name Resolution relies on the hierarchy inherent to the Zone/Domain naming convention, with a number of root level servers defined for the whole of the Internet, by which clients can use to resolve names over the Internet. Within the DNS each name is locally stored and managed by a server/service level administrator.

Aliases

Aliases are defined in various ways within the DNS, for example CNAME, DNAME, and PTR resource record types. Each Alias type has distinct behavioral characteristics, but they are all means by which one name may be mapped to another name. For example the name “Some Service” can be aliased to multiple hostnames, e.g. “host1”, “host2”, and “host3”, for the purpose of providing a round-robin load-balanced service or for fault tolerance. Aliases may also be used to associate meaningful names to less meaningful hostnames, such as mapping “Playout Server 1” to “ER-1-R-12-S-4”, where the actual hostname, “ER-1-R-12-S-4” in this case, is cryptic and contains location information. Note that in this example if the physical hardware is relocated, the hostname will also change, but the alias that describes its purpose will not. Thus, clients using aliases for Name Resolution will be isolated from changes to the underlying host names.

6.4.3. Dynamic Addressing

Dynamic Host Configuration Protocol (DHCP)

The Dynamic Host Configuration Protocol (DHCP) provides a means by which networked nodes may dynamically obtain an IP Address and other network centric configuration, such as the IP Address for the local DNS server. To facilitate the resolution of names referring to network nodes with dynamically assigned IP Addresses, the ability to update information pertaining to a domain name was introduction to the DNS. This update capability allows networked nodes to inform the local DNS server of its new IP Address, thereby allowing clients to locate the node by name.

IPv6 Stateless Address Autoconfiguration

The IETF RFC 4862 – IPv6 Stateless Address Autoconfiguration provides a means by which networked nodes may dynamically configure network interfaces on IPv6 networks. This configuration includes obtaining the link-local address, generating the IP Address, and the prevention and detection of duplicate IP Addresses. IPv6 Stateless Address Autoconfiguration does not provide the tight constraints provided by DHCP, but is a software service of the network hardware, e.g., switches and routers.

6.4.4. The Handle System

The Handle System is a distributed Name Resolution system designed to provide an efficient, extensible, and secure means by which global names can be resolved into meaningful data associated to globally unique names. The Handle System is designed and built with many of the same principles as the DNS; however the Handle System is designed to provide a finer granularity of administrative control.

Within the Handle System each name and its associated data may be administered by one or more entities other than the server and/or service administrator. Thus, allowing for registrants to maintain the information that they have registered, without having the privileges to alter the information registered by others. The server and/or service administrator must have the required privileges to maintain the configuration of the Handle service, but does not have the necessary privileges to view and/or modify the information registered within the service. As with the DNS, Names are configured and managed locally, within a Local Handle Service, and the Handle System uses the hierarchical namespace associated with each name to facilitate the distribution of Name Resolution requests to the appropriate Local Handle Service.

The Handle System was designed to act as a repository of data, with minimal search capabilities, and with fine-grained access control provisions. This differs from the DNS model, which provides some mechanisms for search, but does not provide the required access granularity. With the growth of the Internet, the demand on the DNS has increased, causing organization such as the IETF to apply great caution when developing general-purpose uses for the DNS infrastructure. The critical nature of the DNS for network operations and the requirement that general-purpose name to data resolution processes should not negatively impact network operations led to the creation of the Handle System.

6.5. Switching

Layer 2 & 3 Switching refers to the layers in the OSI model pertaining to IP. Layer 2 switching refers to the switching of frames at the Data Link layer (such as Ethernet), using hardware addresses, where Layer 3 switching is the switching of the packets at the Networking Layer (IP). Layer 2 switching uses the frame headers and the MAC addresses of the network devices.

Both approaches have their pros and cons. Layer 2 switching can be optimized to take only 100s of nanoseconds to switching Ethernet frames between source and destination ports (even less in some specific configurations), it uses a model where each switch must know the entire set of MAC addresses on the entire network. This works fine in smaller environments, but doesn't scale to larger networks such as a campus or beyond. The Layer 3 approach relies on routing (discussed later in this section), which aggregates ranges of IP addresses and provides information about which interface must be used to reach the next hop in the network. This permits scaling the infrastructure to a much large number of devices, as seen with the public Internet or corporate networks. Layer 3 interfaces with low data rates will be taxed by a serialization delay; this is the maximum wait time for an IP packet before it can be sent onto the wire. With the introduction of higher interface data rates, this delay has become negligible. At 1Gbit/s a 1500 bytes IP packet takes 12µs to be processed while at 10Gbit/s the same operation takes 1.2µs.

The term "Layer 3 switch" refers to a class of Layer 3 routing devices that do not rely on the more traditional design based on centralized microprocessors to make IP packet forwarding decisions but one that is analogous to Layer 2 switches with hardware based forwarding that can run up to the interface line rate. This is a general classification that vendors use according to features they are promoting. As with any product being evaluated for a specific purpose, all aspects need to be understood and qualified as part of the process.

6.6. Routing

Switching operates at the Data Link Layer of the OSI network model (Layer 2) and uses the Data Link Identity (e.g. MAC address) within the Data Link header to determine the path by which the network frame will take to the destination. While routing operates at layer 3 of the OSI network model (Network) and uses the Network Identity (e.g. IP Address) to determine the route the packet is to take to the destination.

6.6.1. Static Routing

Static Routing is the static configuration of routes and routing information. Static routing is typically used when a fixed set of routes are desired for packet traversal.

6.6.2. Dynamic Routing

The basic purposes of network routing protocols are to "learn" of the available routes that exist on a network, build the routing tables, and make routing decisions. The differences between the various routing protocols reflect differences in the methods of learning, persistence, and decision-making relating to discovery of the available routes, level of detail and complexity of the routing tables, and the sophistication of the decision-making algorithms.

Routing Information Protocol (RIP)

RIP is a multi-vendor interior gateway protocol using the distance vector method. Originally created in 1988, it is outdated, having been superseded by OSPF for IP networks, and IS-IS for OSI networks.

RIP v1 was a multi-vendor routing protocol created in 1988 for routing IP and Novell IPX. RIP v1 was a classful routing protocol, with routing updates that did not carry subnet information, and lacked support for variable length subnets. In RIP v1, all subnets in network class must have the same size. There was no router authentication, leaving the protocol vulnerable to attacks. RIP v1 distributes the full routing table to all direct-connected neighbors using broadcast, incurring overhead in non-router nodes to ignore the broadcast full route tables. RIP v1 supports a maximum of 15 hops.

RIP v2 was initially developed in 1993, with last standard update in 1998. RIP v2 supports Classless Inter-Domain Routing (CIR) route updates including subnet information. RIP v2 uses multicast to distribute the routing table efficiently on multicast address 224.0.0.9, and uses MD5 authentication. Route tags are used to distinguish internal and external redistributed routes from external gateway protocols. RIP v2 supports a maximum of 15 hops.

Interior Gateway Routing Protocol (IGRP)

IGRP is a Cisco proprietary interior gateway protocol using the distance vector method. It is designed to overcome the confines of RIP v1 in small and medium Cisco networks. IGRP will route IP, IPX (Novell), Decnet, and AppleTalk. IGRP is somewhat more scalable than IP-only RIP, and supports a maximum hop count of 100. IGRP only advertises routing updates every 90 seconds (lessening overhead) and uses a composite of five different metrics to select best path destination.

IGRP is a compromise. IGRP advertises less frequently (90 seconds) than RIP, but converges much slower than RIP since it is 90 seconds before IGRP routers are aware of changes.

Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP is an enhanced version of Cisco IGRP. EIGRP is an interior gateway protocol using a hybrid combination of the distance vector and link state methods. It is designed to lessen routing unsteadiness after topology change, lessen bandwidth usage, and lessen processing power.

Open Shortest Path First (OSPF)

OSPF is a multi-vendor interior gateway protocol using the link state method. It is designed for routing IP (only) across large multi-vendor networks. Version 2 was introduced in 1998 and was designed for IPv4 (only). OSPF needed to be completely rewritten for IPv6, and version 3 was introduced in 2008. To support both IPv4 and IPv6, you need to run both OSPF v2 and OSPF v3.

OSPF is a widely used interior routing protocol for large (but not the largest) IP networks. OSPF uses a hierarchy with assigned areas that connect to a core backbone of routers. Each area is defined by one or more routers with established adjacencies. OSPF uses shortest-path-first Dijkstra algorithm to determine optimum path.

OSPF and IS-IS are similar, with the largest ISPs favoring IS-IS for efficiency and unified support for both IPv4 and IPv6.

Exterior Gateway Protocol (EGP)

EGP is the name of the generic exterior gateway protocol type and the proper name of its first implementation in 1982. This original Exterior Gateway Protocol (EGP) used a tree-structured approach, and did not use distance vectors or path vectors. It is mostly obsolete, having been replaced by the Border Gateway Protocol (BGP)

Border Gateway Protocol (BGP)

BGP is multi-vendor exterior-only gateway protocol using the path vector protocol. BGP is the core routing protocol of the Internet. BGP is an enhancement on EGP to permit completely decentralized routing. Version 4 was introduced in 1994 and included Classless Inter-Domain Routing (CIR) and route aggregation to reduce the size of routing updates. Version 4 was finally codified in 2006, having gone through over 20 draft standards. Most Internet Service Providers must use BGP to establish routing between one another, especially if they are multi-homed. BGP routes IP traffic only.

Intermediate System to Intermediate System (IS-IS)

IS-IS is a multi-vendor interior gateway protocol using the link state method, similar to OSPF. The protocol was defined in defined in ISO/IEC 10589:2002 as an international standard within the Open Systems Interconnection (OSI) reference design. However, originally an ISO standard, the IETF republished the protocol as an Internet Standard in RFC 1142.

While OSPF is natively built to route IP and is itself a Layer 3 protocol that runs on top of IP, IS-IS is natively an OSI network layer protocol. In effect, IS-IS runs over layer 2, side-by-side with IP. IS-IS is at the same layer as OSI's Connectionless-mode Network Service (CLNS).

IS-IS supports a "Single Topology" mode for IPv6 routing. This allows an IPv4 and IPv6 network to use the same shortest-path-first calculation. This reduces the complexity and resources needed to run both protocols. With OSPF, you have to run both OSPFv2 and OSPFv3 to support IPv4 and IPv6 separately.

6.7. Multicast

In contrast with a unicast mechanism, where a unique flow from any given source to any destination must to be set up for each and every destination, multicast offers the ability to distribute a single flow from a source to any number of destinations (known as receivers). By doing so, the same amount of data is sent from the source independently of the number of receivers that wish to subscribe to the flow. The IP network nodes, such as switches and routers, perform the flow replication where needed along the path between sources and receivers.

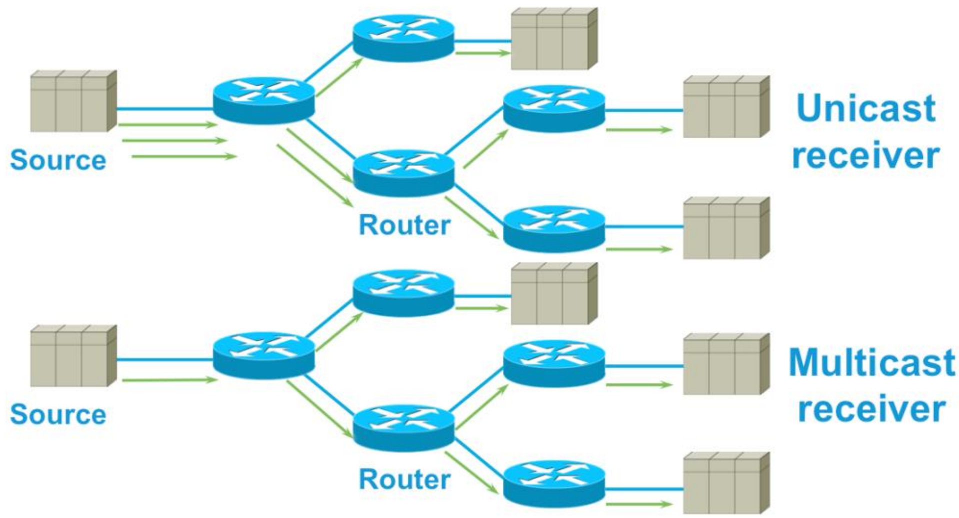

Figure 6 – Unicast and Multicast Network delivery

Any-Source Multicast (ASM)

In the original Any Source Multicast (ASM) model the receiver is only aware of the multicast group (G) that he wishes to subscribe to, and not of the source (S) transmitting that group. The receiver relies on his directly attached (designated) router (DR) to discover which sources (S) are sending to the multicast group (G) of interest. The receiver relies on the IGMPv2 (or MLDv1 for IPv6) protocol to request the (*,G) membership. In order for the router to discover the source of the group in question, there is a need for a Rendezvous Point (RP), to allow for the source discovery to occur. The RP will handle the initial request from the receiver, set up the multicast tree, known as the shared tree, and forward the flow from any source via the RP to the receiver.

Once the receiver's DR knows the source, it will set the direct multicast tree from the source to the receiver, and switch over the active stream to the new path known as the source or shortest path tree (SPT). The shared path between the receiver and the source (via the RP) is then torn down. This means that all new flows will initially transit the RP before the source path is set up.

Source-Specific Multicast (SSM)

Source Specific Multicast (SSM) greatly simplifies the above process versus the ASM model, especially for architectures with a large number of flows from a limited number of sources to a large number of receivers as typically found in a broadcasting environment. It allows for a receiver to subscribe directly to a specific source by specifying both the source (S) and the multicast group (G) it wishes to receive from, hence the name Source Specific. No changes are required at the source, while the receiver and directly connected router must support IGMPv3 (or MLDv2 for IPv6) to allow for (S, G) membership reports. The receiver application must also be written to send source (S) and group (G) to a well-defined API that is present in all modern IP protocol stacks. Therefore, only the source/shortest path tree needs to be set up; this is done directly without the need for a Rendezvous Point. SSM also provides additional security whereby, in case multiple multicast sources are distributing content to the same group address, there is no stream overlap, since the source address remains unique. IETF RFC 4607 standardized SSM in 2006.

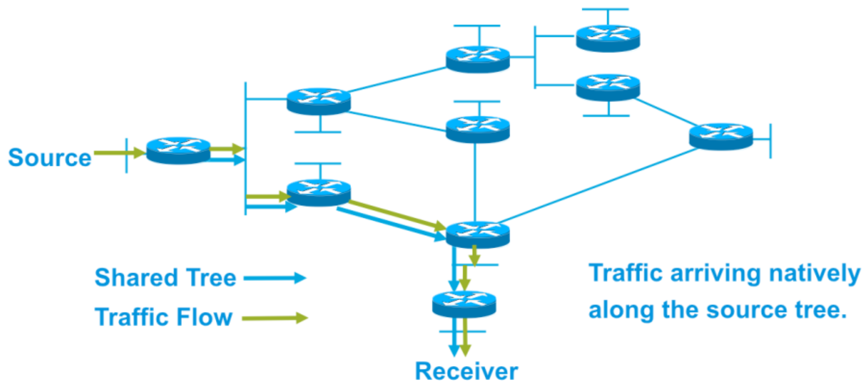

Figure 7 – Source-Specific Multicast (SSM) Routing and Traffic Flow

SSM Source Mapping

For legacy receivers that may not be able to support SSM, due to a lack of IGMPv3/MLDv2 in the IP network stack, or the previously mentioned API calls at high layers, SSM Source Mapping offers a migration solution.

SSM Source Mapping works by deploying SSM at the network edge close to the multicast receivers. IGMPv2/MLDv1 (*,G) membership reports from IGMPv2/MLDv1 receivers are mapped to specific source IP addresses. This mapping can be provided statically (from the local router configuration file), or dynamically (using a well-known address lookup mechanism known as DNS). This enables managed IP networks to offer SSM services to both IGMPv2/MLDv1 and IGMPv3/MLDv2 receivers, and provides a smooth transition path from ASM to SSM, without the need to implement Rendezvous Points for legacy receivers.

6.8. Anycast

Anycast is a network addressing and routing methodology that routes datagrams to a single member of a group of potential receivers that all share the same destination address. Where broadcast and multicast is one sender to many destinations, Anycast is one sender to only one of many potential destinations.

6.9. Quality of Service (QoS)

Quality of Service is the concept of using queuing and queue prioritization to guarantee a level of performance for specific network protocols. QoS is typically implemented as a series of rules that define the queuing policy for network packets as they are processed network equipment.

6.9.1. Explicit Congestion Notification (ECN)

Explicit Congestion Notification (ECN) defines mechanisms by which network nodes may notify data senders of network congestion and imminent packet loss. The ECN specification defines an IP header field, known as the ECN field, which consists of the 2 bits immediately following the Differentiated Services Code Point (DSCP). Combined the DSCP and ECN fields replace the deprecated IPv4 Type of Service (ToS) and IPv6 Traffic Class Octet (TCO) header fields.

The ECN specification defines 4 values, Not-ECT, ECT(1), ECT(0), and CE (binary 00, 01, 10, and 11 respectively). The ECT(0) and ECT(1), or ECN-Capable Transport, values indicate that the network end-points support ECN markings.

While the CE, or Congestion Experienced, value indicates that packet loss is imminent at the current data transmission rate.

6.9.2. Data Center TCP

DCTCP is an extension to TCP/IP leveraging Explicit Congestion Notification (ECN), defined in RFC 3168, and provides further support to ECN marking to gauge the extent of the congestion. ECN enables end-to-end notification of network congestion before a packet drop occurs. This is unlike the traditional behavior of TCP that detects network congestion after a packet drop. Dropped packets can in some cases lead to long TCP timeouts and therefore loss of throughput. This loss will impact the incriminated TCP flow but is likely to affect other flows as well. Instead of dropping a packet, DCTCP is used to signal impending congestion to the sender, who in turn can slow the rate of traffic before a packet drop occurs. This is useful for both short and long-lived flows as found in modern media workflows that may be relying on TCP as part of their transport design.

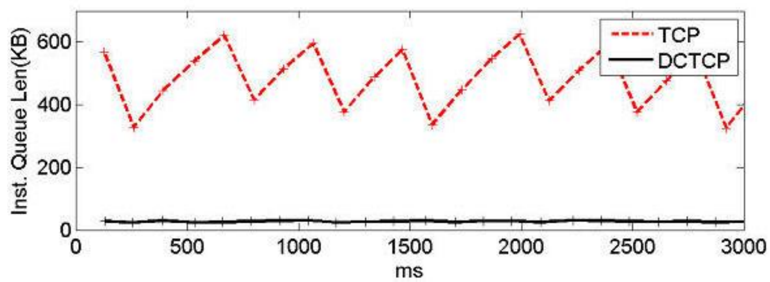

Figure 8 – DCTCP vs. TCP Queue Length

As illustrated above, DCTCP specifies a single value ***K*** at which if ***K*** elements exist within the queue the ECN bits of the IP packet header are set. Based on the probability of further packet transmissions causing an increase in the number of elements within the queue, DCTCP dynamically adjusts the TCP window.

6.9.3. Congestion Exposure

The Congestion Exposure (ConEx) protocol defines an IPv6 mechanism by which networked nodes may notify data senders of congestion and packet loss. Similar in function and purpose to DCTCP, ConEx provides additional information pertaining to the number of packets and bytes that have experienced congestion or have been dropped. Unlike DCTCP, ConEx is implemented by the source node and does not require ECN, nor does it require a change to network hardware. The IETF ConEx Working Group is currently studying implementation requirements and related modifications for supporting this effort.

6.9.4. Priority Queuing

The queue scheduler within the network equipment may service queues in a round-robin or weighted round-robin fashion, which may result in the loss of packets, high latency, and/or jitter. To mitigate this issue a priority queue solution may be implemented.

Priority queuing ensures that the queue scheduler services the traffic with the highest priority first, the priority being indicated in the Differentiated Services Code Point (DSCP) of the packet's IP header. If the DSCP field of the packet is assigned a value of 46, the Expedited Forwarding (EF) Per-Hop Behavior, as described in RFC 3246, the packet will be delegated to the priority queue, also referred to as the "strict priority queue." Traffic within the strict priority queue is processed first, in its entirety, before the queue scheduler services other traffics.

Any packet assigned a DSCP value of 46 (the EF PHB) that encounters congestion within a network node will be moved to the beginning of the queue, also referred to as the "front-of-the-line", and will serviced in a strict priority manner, hence the term priority queuing.

Queue Starvation

Queue starvation is a risk of QoS and Queuing techniques. Queue starvation occurs when higher priority frames prevent lower priority frames from being processed. For example, without proper precautions, due to the sheer volume of traffic, it would be easy for media essence traffic to prevent other forms of traffic from being forwarded to their destinations. When implementing queuing techniques for QoS, caution should be taken so that all traffics can be processed within a reasonable amount of time. Please note that "reasonable" is dependent on the traffic type.

Implementation Differences

There are few standards governing how QoS and Queuing techniques are implemented. Therefore, implementations differ from one another, both in the algorithms used and in the methods by which the QoS strategy is configured. Each vendor defines their own Queuing techniques, algorithms, and methods of configuring those options. While designing, building, and maintaining an engineered network it is important to take note of these differences and account for them in design and purchase decisions.

6.10. Issues for the Carriage of SDI over IP using SMPTE ST 2022-6

SMPTE ST 2022-6 provides for encapsulation of the payloads of a variety of SMPTE serial digital video standards, including SMPTE ST 292-1 ("HD-SDI"). The standard carries 1376 octets of video payload into each RTP datagram. There is not an integral number of RTP datagrams per line in SMPTE ST 2022-6. Also the last RTP packet in a frame in SMPTE ST 2022-6 contains zero padding once it has exhausted the SDI sample data from the frame.

Ethernet frames carrying SMPTE ST 2022-6 tend to be 1450 to 1458 bytes long (depending on whether the High Bit Rate Media Video Timestamp and High Bit Rate Media Header Extensions are present or not). This includes 8 octets of Ethernet preamble and start of frame delimiter, 14 octets of Ethernet header, 20 octets of IP header, 8 octets of UDP header, 12 octets of RTP header, and 8 to 16 octets of High Bit Rate Media header, with 4 octets of Ethernet frame check sequence (Ethernet and IP headers may be longer with various extensions such as the 802.1Q tag). In addition to this, Ethernet frames generally require an Ethernet interframe gap period of 5 to 12 octets depending on Ethernet standard being used. Total header overhead is thus 5.4% to 5.9% for SMPTE ST 2022-6 datagrams that carry 1376 octets of SDI data.

The last datagram in a frame in SMPTE ST 2022-6 is zero padded to fill out all 1376 octets of datagram payload. For 4:2:2 10-bit sampling, there is 874 octets of zero pad per frame for 720p formats (0.028% overhead), and 372 octets of zero pad per video frame for 1080i formats (a 0.006% overhead).

There are three key metadata elements in SMPTE ST 2022-6 that must be kept valid during any processing of the stream:

- The 16-bit RTP sequence number, which must increment by one for each RTP datagram sent.
- The 32-bit RTP timestamp, a 27 MHz clock which reflects the sampling instant of the first octet in the RTP datagram.
- The High Bit Rate Media Payload Header 8-bit frame count. This counter must increment to a new value for the next RTP datagram immediately after the RTP datagram with the marker bit set to indicate the end of the video frame.

- The High Bit Rate Media Payload Header 32-bit video timestamp. This field is optional, and is left zero in some implementations.

Some of these metadata fields (such as RTP sequence number and frame count) may be analyzed by implementations of SMPTE ST 2022-7 and other stream redundancy schemes to detect a loss of data indicating an invalid stream.

Conceptually, one does not need any timestamps at all to decode an SMPTE ST 2022-6 stream. The payload data carried in SMPTE ST 2022-6 is a bit-for-bit replacement of every single bit in the equivalent SDI stream. Therefore, an SMPTE ST 2022-6 to SDI converter can recreate the original SDI timing just from the payload data alone, with help from the sequence number and the frame count.

When processing SMPTE ST 2022-6 carriage of HD-SDI video (for logo insertion, for example), it is important to remember that any changes in the active video must be reflected in a revised line CRC.

Video stream switching is an important part of any broadcast operation. The goal of video stream switching is to accomplish a clean switch between two video streams with a minimal amount of latency. SMPTE RP 168 defines a switching point in the video raster such that the effects of any signal discontinuity in the processing chain is minimized. The line designated for the switching point is chosen to be after the vertical sync, but early in the vertical blanking interval. This is to ensure that ancillary signals transmitted during the vertical interval remain with the video frame with which they are associated. Also ancillary data is excluded from the line following the switch line to allow for receivers to synchronize to EAV/SAV before data becomes present to avoid its loss. This exclusion is reflected in several SMPTE ancillary data standards. Finally, the recommended switch area on the switching line is a region relatively within the middle of the active video to provide some tolerance for inaccurate timing and also to provide a reasonable “guard area” to protect the SAV and EAV data flags. Some of these issues are of less concern for packetized video transmission, but conservative implementations may wish to switch during the RP 168 area. Switching during the vertical interval also generally means the active video line of the switch point always contains all black pixels (if there is no VANC on the line), so that the line CRC does may have to be adjusted if there is a switch during the active video area.

As mentioned in “SDI Over IP – Seamless Signal Switching in SMPTE ST 2022-6 and a Novel Multicast Routing Concept”, one can simply count a number of SMPTE ST 2022-6 RTP datagrams into the frame to find the appropriate switching point between datagrams that is located as per RP 168. The last RTP datagram of a frame in SMPTE ST 2022-6 has the RTP marker bit set, thus the next RTP datagram after that is the first datagram of its frame.

For 720p formats, the boundary between RTP datagrams 20 and 21 of the frame is on line 7 at 738 pixels from SAV end, which is within the RP 168 recommended switching area on the recommended switching line. For 1080i formats, a similar boundary can be found in the RP 168 switching area between RTP datagrams 26 and 27 of the frame in field 1 (on line 7) and between RTP datagrams 2272 and 2273 in field 2 (on line 569).

6.11. Managing Oversubscription

Oversubscription is the process by which the allocated potential consumption of a resource is greater than the physical limitations of the resource. Oversubscription relies on the fact that most dedicated network resources are not utilized to their full potential, remaining idle or underutilized a significant portion of time. By assigning the resource to multiple consumers, the resource can be utilized more efficiently. The gained efficiency reduces the total cost, but when oversubscribing a resource, care must be taken to calculate the true peak utilization of the resource or the resource will be over utilized, causing performance issues and/or latency. Please note that with today’s technology it is possible to build a network that is not oversubscribed.

6.11.1. Network Oversubscription

Network oversubscription is a common technique used in almost every network. With network oversubscription the uplink or interconnect speeds are lower than the aggregate of the individual links connected to a networking device, such as a switch or a router. As with the general concept of oversubscription, the link speeds should be carefully calculated, as if the link speeds are too low for the actual network usage, latency and dropped frames will be introduced into the network.

Please refer to section IEEE 802.1 Audio Video Bridging (AVB) for details pertaining to AVB and oversubscription.

6.11.2. Hardware Oversubscription

The concept of oversubscription is also applied to hardware servers and various hardware based service providers, such as network based storage services or virtual machines. Virtual Machines are instances of operating systems that are isolated from the hardware, allowing for multiple operating system instances to run on the same physical hardware at the same time. Each virtual machine is isolated from the other instances on the physical hardware. Since the operating system instances are not physically tied to the hardware, they can be moved between different hardware platforms, and completely backed up, including the contents of the physical memory / run time state. Each virtual machine is assigned its own unique address and its address remains the same regardless of which physical hardware the virtual machine is located.

6.12. Network Architecture

6.12.1. Network Elements

Modern Ethernet and Internet Protocol based networks are constructed using many devices. These devices are typically switches and routers, but may also be hubs, firewalls, etc... The interconnection between each of these devices with a given network path forms what is referred to as a Network Hop. Frequently, the devices that make up the Network Hops in any given communication channel are transparent to the network devices utilizing that channel, but each Network Hop is significant, as each device in a communication path adds latency and is a potential point of failure.

How Many Points of Failure are There Between Two Points?

As mentioned above, many times the devices that network packets must traverse in a given communication channel are transparent, but there are tools available that allow a network administrator or professional media engineer to determine how many devices physically reside between 2 or more points in the networks. These tools include, but are not limited to “ping” and “traceroute”.

What is the Impact of Failure?

As mentioned above, each device within a communication path, also known as a network hop can be a potential point of failure. Since each point of failure will have a negative impact on the communication channels traversing it, it is important to understand the impact each of these points of failure will have upon the network as a whole and localized to the section of the network where the failed device resides. Careful mapping of the impact of device failures should be performed, so that the professional broadcast engineer can predict the impact of a device failure and the types of problems that will occur upon a failure.

6.12.2. Virtual LANs (VLAN)

A Virtual LAN (VLAN) is a logical partitioning of network endpoints to create a logical LAN, which is mutually isolated from other LANs within the overall network. A VLAN may be comprised of a subset of network endpoints connected to the same physical LAN and/or a superset of network endpoints that are connected to multiple physical LANs. Being configured in software, VLANs are flexible and dynamic, allowing for the formation of virtually connected devices independent of the physical network infrastructure.

6.12.3. Software Defined Networking (SDN)

Software Defined Networking (SDN) is a recent network organization technique where a well-defined Application Programming Interface (API) separates the control and data planes of the networking equipment. This separation of the control plane from the data plane allows for the consolidation of the control plane into SDN Domain Controllers, which in turn allows for the system to make more intelligent decisions pertaining to the routing, Quality of Service, security, and other critical functions of the network. In addition to simplified management and more intelligent control the introduction of SDN Domain Controllers also adds a layer by which software applications and virtualization products can utilize to dynamically adjust the network resources they require and to adjustment of routes and VLANs for connectivity. This added dynamic control of the network provides new capabilities for virtualized environments, as network routes and parameters can be adjusted on the fly as virtual machines migrate between hardware platforms.

6.12.4. High Availability

High Availability is a process that increases the availability of a network node, path or system by adding additional functionalities beyond the common routing protocols.

Bidirectional Forward Detection (BFD)

Bidirectional Forwarding Detection (BFD) is a network protocol used to detect faults between two network nodes, connected by a link. It provides a low overhead detection of faults for many types of physical and virtual media links. If configured, protocols can rely on BFD to receive faster notification of failing links than they would with their own keep-alive mechanisms.

Multicast Only Fast Re-Route (MoFRR)

Multicast Only Fast Re-Route (MOFRR) is an IP solution that minimizes packet loss in a network when there is a link or node failure. It works by making simple enhancements to multicast routing protocols like Protocol Independent Multicast (PIM). MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to Reverse Path Forwarding (RPF) checks. When a failure is detected on the primary path, the failure is worked around by changing the interface on which packets are accepted to a secondary interface. Because this work-around is local, it is faster; greatly improving the convergence times in the event of a link or node failure on the primary path.

Restrictions for MoFRR:

- The Equal Cost Multipath Protocol (ECMP) feature is a requirement in order for the MoFRR feature to function.
- MoFRR works only for Specific Multicast (SM) S, G, and Source Specific Multicast (SSM) routes.
- MoFRR is applicable to only IPv4 Multicast, not IPv6 Multicast.
- Both primary and secondary paths should exist in the same multicast topology.

Loop Free Alternative (LFA) Fast Re-Route (FRR)

Loop Free Alternate (LFA) Fast Re-Route (FRR) creates a backup route, pre-computed using the dynamic routing protocol, whenever a network fails. The backup routes (repair paths) are pre-computed and installed on the router as the backup for the primary paths. Once the router detects a link or adjacent node failure, it switches to the backup path to avoid traffic loss.

Restrictions for IPv4 Loop-Free Alternate Fast Reroute

- Load balance support is available for FRR-protected prefixes, but the 50ms cutover time is not guaranteed.
- IPv4 multicast is not supported.
- IPv6 is not supported.
- IS-IS will not calculate LFA for prefixes whose primary interface is a tunnel.
- LFA calculations are restricted to interfaces or links belonging to the same level or area. Hence, excluding all of the neighbors on the same LAN, when computing the backup LFA, can result in repairs being unavailable in a subset of topologies.

Hitless Switching (Live - Live)

Hitless switching, also referred to as "Live - Live" or "1+1" switching, is switching that relies upon redundant signals that are transmitted over redundant paths. This data redundancy allows for rapid, lossless, recovery from transmission failure, whether the failure is a long-term failure in the data path or is a short-term failure. This rapid, lossless, recovery requires significantly more bandwidth than other techniques of lossless recovery, such as Forward Error Correction, requiring double the bandwidth of the transmitted signal, but does not have the additional computational and temporal costs.

Seamless Protection of SMPTE ST 2022 Streams

The SMPTE ST 2022 suite of standards describes how to map compressed and uncompressed professional audio-visual content into IP networks, for the purposes of essence transfer. The standard also contains a companion, optional, Forward Error Correction (FEC) mechanism. SMPTE ST 2022 is widely deployed worldwide and is used to transport a variety of high-value live events, from the venue back to the studio. Given the impact even momentary outages have on events, such as the Olympics and national sporting events, a method was needed to ensure, as much as possible, continuous transmission, even in the face of a total link failure. SMPTE ST 2022-7 provides hitless protection of SMPTE ST 2022 streams by providing a way to switch seamlessly between two streams, which may be carried over diverse physical paths. Hitless switching is guaranteed, provided the streams meet the timing requirements described in SMPTE ST 2022-7 and the mapping requirements described in the base SMPTE ST 2022 mapping document.

As the cost of network paths decreases, redundant streams become even more viable, especially for sporting events with a very high value. The transmission of redundant streams allows for network errors to be eliminated without requiring an error threshold or FEC, while preventing the delays incurred by packet retransmission.

6.13. My Network, Your Network, and Everything in Between (The Cloud)

The Cloud is a hot topic in the IT industry and most of us are already using it daily. So exactly what is the Cloud? The Cloud, also known as Cloud Computing, is the use of computing resources as on-demand Services over a network. Cloud Computing is multi-faceted and includes Infrastructure, Platform, Software, and even Network "as a Service", IaaS, PaaS, SaaS, and NaaS respectively. The primary advantage of Cloud Computing is its scalability, scaling from Private Clouds to Public Clouds and every variant in between; Hybrid Clouds and Community Clouds are examples of Clouds that share elements from both Public and Private Clouds. Cloud Computing regardless of the variant, consists of loosely coupled resources arranged as subscription based Services and trust in the

Service/Cloud provider's ability to meet the required Service Level Agreement (SLA) is paramount. Cloud Computing also has network implications, requiring additional network resources for state transmission.

6.13.1. What does the Cloud mean for the Media industry?

The Cloud provides on-demand scalability and maximizes the use of available computational resources. The cost for this scalability and optimization depends on whether the Cloud architecture is private, public, community, or hybrid. With Private Cloud architectures, full control of the data is maintained at all times, however the cost savings are minimalized compared to other Cloud solutions, as the owner of the Cloud must purchase all of the computing resources and commission the network bandwidth. With Public or Community based Cloud architectures the cost of the computing resources is distributed across multiple entities or to another entity altogether. This provides for greater cost savings opportunity, as you pay only for the resources you use. Hybrid Cloud solutions allow for the balancing of control over cost.

In Cloud Computing any computing resource can be a Service; Infrastructure, Network, Platforms, and Software can all be exposed as a Service. Which Services are chosen is entirely up to the implementer. Hybrid solutions can be built where some aspects of the system are Services, while others remain static. For example, the implementer may choose to store all of their essence locally, but delegate transcoding of the essence to an external Service Provider. In a Community or Public Cloud based solution, the entire system may also be relegated to an external Service Provider with the essence and all functions pertaining to the essence being implemented and maintained by the Service Provider.

- **Infrastructure as a Service (IaaS)** – Offers computers, physical or virtual, as a Service. IaaS can also include additional physical resources such as firewalls, storage, and load balancers.
- **Platform as a Service (PaaS)** – Offers platforms, such as operating systems, databases, programming environments, web servers, and the like.
- **Software as a Service (SaaS)** – Offers software applications as a Service.
- **Network as a Service (Naas)** – Offers network connectivity and inter-cloud connectivity.

Advantages

- On-demand resource utilization, scalable up, down, and out to handle bursts of activity, providing elasticity to the system in a pay as you go manner
- For private Cloud architectures, the Cloud resources could be shared with other non-critical workflows, reducing hardware redundancy and allowing for better utilization of the available computational cycles
- Service Oriented Architecture (SOA), Web Services, Representational State Transfer (REST)
- Use of commodity resources over the Internet, such as Web Applications
- CapEx to OpEx transfer
- Added choices for insourcing or outsourcing

Disadvantages

- Dependence upon a 3rd party and their ability to keep their system operating optimally.
- Dependence on the Internet and your ISP ability to provide the Service Levels required.
- For private Cloud architectures, the building out of a server farm capable of handling burst utilization
- Lack of security certifications
- The "Trust us" model for SLA and Security

6.14. Time and Sync Distribution

Professional media systems are typically synchronous by nature, requiring all equipment to be provided with reference signals, by which they derive their output signal timing, including video and audio reference signals and timecode. These reference signals in turn require distribution infrastructures, which are installed as part of the foundational infrastructure of a facility. These reference signal distribution infrastructures are hard-wired and nailed-down systems of cables and distribution amplifiers and typically one such distribution infrastructure exists for each reference format required. Despite containing many single points of failure, these reference signal distribution infrastructures provide reasonable robustness and required reliability, but unfortunately are inflexible and have a significant cost associated with their installation and maintenance. Therefore, it is highly desirable to move these reference signal distribution infrastructures to a technology that offers all references with one infrastructure, which provides comparable or improved robustness, reduces cost, provides the flexibility to accommodate ongoing facility evolution, and provides a foundation for future synchronization signal formats. Better still would be a technology that offers the ability to be used for synchronization in the both file-based and live-media networked domains.

6.14.1. IEEE 1588-2008 Precision Time Protocol

Internet Protocol (IP) networks pervade every aspect of a professional media production system, and one by one, applications, workflows, and media are moving to IP infrastructures. Most equipment, even the simplest, such as DAs, offers IP connectivity. Using IP connectivity for a universal reference signal infrastructure is very attractive and in conjunction with the IEEE 1588-2008 Precision Time Protocol, IP networks meet the desired requirements described above. IEEE 1588-2008 delivers a timekeeping service over IP networks, which allow slave devices to maintain locked timekeeping with the master, even in a network with thousands of slaves. This timekeeping allows for the delivery of date and time, spanning potentially many hundreds of years, from which an equivalent of SMPTE ST 12 timecode can be derived by slave equipment. IEEE 1588-2008 also provides sub-microsecond accuracy, allowing for the generation or emulation of any streaming reference signal used today.

Since IP networks are built from commodity IT hardware and an IEEE 1588-2008 reference signal infrastructure is piggy backed onto an existing IP infrastructure for inter-machine communications, it has a significant reduction in cost when compared to a traditional dedicated reference signal infrastructure. IP-based reference signal infrastructures also remove the need for separate distribution infrastructures by allowing all reference signal types to be transmitted on the same infrastructure. With IEEE 1588-2008, multiple master timing servers can be utilized for redundancy and the same resources that manage the entire IP infrastructure can be used to manage the reference signal infrastructure. The use of IEEE 1588-2008 in large-scale systems with fine-grained timing accuracy requirements require network switches that are aware of the IEEE 1588-2008 protocol. However, small systems or systems with low timing accuracy requirements can function properly without IEEE 1588-2008 protocol aware network switches.

6.14.2. IEEE 802.1AS – Generalized Precision Timing Protocol (gPTP)

IEEE 802.1AS defines a “profile” for the application of IEEE 1588-2008 PTP services across IEEE 802-based local-area networks. IEEE 802.1AS-2011 (gPTP) provides precise, LAN-wide, time synchronization. The network management system uses the PTP services defined in IEEE 802.1AS-2011 to synchronize each element in the network, including end stations and bridges, to a “grandmaster” clock which is elected using a simple “best clock master algorithm” (BCMA). 802.1AS works by characterizing the path delays through each element in the network—including the cables—and performs a simple calculation within each node to compensate for these delays as the grandmaster’s “announced time” is propagated through the network. If a grandmaster change occurs, either by choice or by failure, a new grandmaster is elected and the network resynchronizes within a fraction of a second.

The precision of 802.1AS is ±500 nanoseconds (ns) between any two nodes within 7 network hops from each other, using 100Mbps links and 100ppm crystals in each node. With higher link speeds and/or more accurate crystals, precision can be even tighter. The precision between gPTP clocks degrades linearly as the number of hops between gPTP nodes increases beyond 7.

6.14.2.1. SMPTE ST 2059 Family of Standards for Timing Reference

The SMPTE Time and Synchronization work leverages IEEE 1588-2008 PTP to create a SMPTE IEEE 1588-2008 PTP-based synchronization specification, adding the necessary SMPTE pieces. Key to this work is the establishment of the SMPTE Epoch. The SMPTE Epoch is a point in time (midnight, Jan 1st, 1970) where the alignment of all signals is defined. From that point in time, and forward, it is possible to determine the phase of any signal that has been Epoch-aligned. On this basis, independent systems, which are locked to the GPS timing reference, can generate signals that are locked together in frequency and time, eliminating the drop / repeat behavior of frame synchronization. The SMPTE Timing Reference can also accommodate any video frame rate or audio sample rate, has no issues with fractional frame rates like the fractional frame rate 1/1.001, even when used in concert with non-fractional systems, such as 1/1 systems.

The deployment of a SMPTE ST 2059-1, SMPTE ST 2059-2, and IEEE 1588-2008 PTP-based synchronization system also removes the need for independent reference infrastructures and leverages the evolving migration of control, configuration, management, and most importantly live-media to IP networks. Whether it is delivered on a private connector to the equipment, or on a connector shared with these other services, it brings convergence to the systemization of equipment and the integration of 'systems of systems'. It enables deterministic timing without the need to time everything manually.

6.14.2.2. AES67

The AES67 standard for high-performance streaming audio-over-IP interoperability provides all necessary services related to audio distribution over a single network connection. Delivery of synchronization signals is achieved using the IEEE 1588-2008 precision time protocol. Standard IP networks and networks with IEEE 1588-2008 capability are both supported. IEEE 1588-2008 features a profile concept such that the operation of the precision time protocol can be customized to meet the needs of different applications. AES67 defines a Media Profile that is recommended when using standard IP networks. The Media Profile offers improved start-time and accuracy compared to the IEEE 1588-2008 default profiles, supports interfacing to the digital audio reference signals defined in AES11, and is designed for compatibility with the SMPTE ST 2059 suite of standards on timing and synchronization.

6.15. IP Studio Framework – "Flows and Grains"

BBC Research & Development is developing a framework, which they refer to as "Flows and Grains", to investigate going beyond traditional technologies such as SDI towards an IP-based infrastructure, using widely adopted networking protocols and techniques. The framework treats items of content, such as video frames, and data events generated during production, as uniquely identified and time-stamped objects. For live working, these are multicast locally at high quality or for remote access, published as web feeds; they can also be recorded for immediate use by downstream devices and processes.

At the heart of the IP Studio framework is the idea that everything of interest generated during production should be captured and made available for immediate access, with provisions for later access. This goes beyond just audio and video, and typical "clip-oriented" metadata, to include events: time-related data objects that contain information such as: camera settings, actor positions, logging entries and AV quality analysis data. In the project's content model, events and frames, or sections of video, and audio are all treated as individually identifiable grains. Grains are nanosecond-precision time-stamped objects within flows of time-sequential information coming from one or more sources. An application can access grains in real time, as they are created, or it can access them later using the timestamp.

Each source and flow is uniquely identified (with a UUID), allowing a system to determine the origin of any grain. Information about the relationships between sources, and with other things such as cameras and people is also recorded. This data allows complex queries to be constructed. As an example, a user might want to view what was recorded from each camera when a particular line of dialogue was spoken. This could be achieved by: querying a database to find a particular (speech recognition) event grain; finding its corresponding source; finding related sources; and finding video grains with matching timestamps. Because sources, flows and other objects are uniquely identified, they can be addressed using URIs.

Flows of grains can be processed in several ways. For example a video flow may be encoded into different formats, multiple flows may be composited, and audio or video flows may be analyzed to produce event flows. Such processing operations are combined into pipelines that run on computing instances, or nodes. Pipelines start and end by receiving and sending external flows from and to a network port, a pre-existing interconnect such as SDI, or local storage on the node itself.

The concept of Flows and Grains, introduced by BBC Research, was found to be of particular value and warrants further study. The concept of Flows and Grains ties very neatly with packetized network technologies, and allows for different types of essence and metadata to move across the network in a dynamic fashion, while retaining the association between hundreds of thousands or even millions of individual packets that are associated to a piece of content that is intended for presentation as a contiguous whole. Each snippet of essence and metadata, called a grain, is time-stamped with high precision, uniquely identified, and is associated to a particular flow, which is also uniquely identified. The combination of unique grain identity, unique flow identity, and high precision time stamping allows for systems to capture flows within databases, so that they may later be deconstructed, queried and used to construct entirely different flows.

7. Conclusion and Further SMPTE Action

This report has been prepared for the SMPTE Standards Committee and the general SMPTE Membership. It describes the increased deployment of Information Technology (IT) methodologies, techniques, and technologies within the Professional Media (PM) industry.

This report is not a tutorial, but is intended to introduce concepts that readers may pursue further to improve their knowledge of both IT and PM methodologies, techniques, and technologies.

The migration to new IT-centric methodologies, techniques, and technologies by the PM industry is driven by many advantages, including lower equipment cost, faster technology advancement, a greater talent pool with technical expertise, and greater flexibility in workflow and system design. However, it is evident that many individuals from the IT industry lack sufficient familiarity with the techniques and performance expectations of PM systems and that many individuals within the PM industry lack sufficient familiarity with the applicable IT technologies. This discontinuity can lend to the creation of unreliable, inefficient, or overly complex PM systems.

This Study Group has noted that many early attempts to implement technological change have suffered from a limited view, constrained by existing/legacy technologies. This is evident in many of the early attempts to create tapeless workflows, which replicated the behavior of physical tape, even though more advanced file-based workflows existed that allowed for new capabilities, such as the ability to transfer and/or play out media essence files while those files are being recorded. The Study Group has also noted that the advancement of network technologies and network-based time and synchronization techniques, such as those afforded by the SMPTE Time, Label, and Synchronization Working Group, are making it possible to create deterministic IT-centric transport networks with dynamic bandwidth properties. These efforts have facilitated the transition from legacy SDI-based architectures to packetized network technologies, such as the "Flows and Grains" work developed by BBC Research.

An additional purpose of this report is to suggest to the SMPTE Standards Committee how it can better understand its new role in the development of IT-centric standards, engineering guidelines, and recommended practices for the PM industry. Although SMPTE will remain the pre-eminent source of specifications for the PM industry; standards bodies and other industry groups already exist for the development of standards and best practices for the IT industry. Consequently, SMPTE's process for the development of standards should evolve so that SMPTE can tap into the larger pool of IT expertise that already exists within these groups and the SMPTE should make its PM expertise available to these organizations, to support the standards development process wherever the IT and PM industries intersect. This study group feels that a faster and more pro-active liaison process be devised in order to better represent the PM industry's needs within the IT space. Such an approach will allow a wider range of industries to benefit from the resulting specifications and make possible their integration into commercial-off-the-shelf products used by these industries.

Therefore, this Study Group believes the PM industry would be better served if SMPTE were to enhance its liaison and standards activities to proactively and continuously communicate and interact with these organizations. Additionally, the formation of a new Working Group focused on the long term analysis of the application of IT centric technologies within the PM industry will likely **drive the transition** from SDI-based architectures to IT-centric network based technologies. Finally, once specific PM requirements have been identified; these requirements can be brought to the appropriate organization by official delegations from SMPTE, where applicable.

In conclusion, this Study Group recommends that the SMPTE Standards Committee promptly devise the means to implement the recommendations within this report.

8. Terms and Acronyms

A

Asynchronous – Not occurring at the same time; of or pertaining to operation without the use of fixed time intervals.

Audio Video Bridging – A set of technical standards developed by the IEEE Audio Video Bridging Task Group.

AVB – Acronym for Audio Video Bridging

B

Bit Rate – The number of bits that are conveyed or processed per unit of time.

C

Circuit switched – A methodology of implementing a network in which two network nodes establish a dedicated communications channel (circuit) through the network through which the nodes can communicate.

Connectivity provisioning – Ensuring that the required connectivity is present between specified locations for the desired communication.

D

Differentiated Services Field – A 6-bit field in the Internet Protocol (IP) packet header intended to provide a framework to enable the deployment of scalable service discrimination in an Internet Protocol Network. Please refer to RFC 2474 for greater detail.

DS – Acronym for Differentiated Services Field.

E

ENG – Electronic News Gathering, television producers, reporters, and editors making use of electronic video and audio technologies for gathering and presenting news from remote locations, often from a truck or other vehicle with a microwave or satellite antenna to transmit media back to the station, or sometimes using bonded cellular data connections. The term ENG was developed to differentiate from previous film based news gathering.

Ethernet – A family of networking technologies for local area networks (LANs), standardized by IEEE 802.3 Ethernet Working Group, along with key bridging technology from IEEE 802.1 Higher Layer LAN Protocols Working Group.

Ethernet virtual circuits – An end-to-end representation of a single Layer 2 Ethernet service through a provider network, often implemented through the use of 802.1Q-2011 VLAN tags.

F

FEC – Acronym for Forward Error Correction

Forward Error Correction – A technique for reducing data errors in noisy communication channels through the addition of redundant data to the transmission, typically involving the use of an error-correcting code.

G

Genlocked – Short for "sync generator locking", the process of timing the emission of video from a device based on a shared sync reference signal.

H

I

Ingest – The action of capturing a media stream into a media file.

Internet Protocol – The Internet Protocol is designed for use in the interconnected systems of Packet Switched Networks. The Internet Protocol provides for transmitting blocks of data, called datagrams, from sources to destinations where the sources and destinations are host systems identified by fixed length addresses. The Internet Protocol is limited in scope to provide the function necessary to deliver a package of bits (a datagram) from the source to the destination. The Internet Protocol is the base protocol for all Internet protocols, such as User Datagram Protocol and Transmission Control Protocol. Please refer to RFC 791 for greater detail.

IP – Acronym for Internet Protocol**.**

Isochronous – A sequence of events is isochronous if the events occur regularly, or at equal time intervals relative to the cadence of real time. Isochronous timing differs from synchronous timing, in that the latter refers to relative timing between two or more sequences of events, while isochronous events occur at equal time intervals; having a uniform period of vibration or oscillation.

J

Jitter – In packet switched networks, the term "jitter" is synonymous with packet delay variation, the variability of packet latency across a network measured over time.

K

L

Latency – The time delay experienced in a system. In a packet switched network, latency may be the one-way delay from the source transmission of a packet until its destination, or it may be a round-trip delay of a packet from the source to a destination and then back to the source.

Loss – In a packet switched network, packet loss is the dropping of a packet by a network element such that the packet does not arrive at its desired destination.

M

Metro Ethernet – A Metropolitan Area Network (MAN) based on Ethernet standards, generally referring to services defined by the Metro Ethernet Forum.

Mobile access – Access to a network from a non-fixed position, typically via some kind of wireless connection.

MPLS – Multiprotocol Label Switching, standardized by RFC 3031, is a mechanism that directs data between network nodes based on short path labels rather than long network addresses, avoiding complex lookups in a routing table and also allowing precise traffic engineering in a network.

N

Non-Synchronous – Synonym for *Asynchronous*

O

P

Packet – A block of binary data that is grouped into a suitably size for transmission over a data network

Packet switching – A digital network communications method that groups transmitted data into packets

Packet Switched Network – A digital networking communications method that groups all transmitted data into blocks called packets, which are transmitted over a shared network that allocates transmission resources using statistical multiplexing or dynamic bandwidth allocation.

Performance profiles – Criteria that define different levels of the performance capabilities of a device.

Playout – The emission of a media stream from a device using media data from a media file or tape.

PMN – Acronym for Professional Media Network.

Professional Media Systems – Systems operated by professionals to acquire, manipulate, edit, process, and distribute media.

Professional Media Network – A network infrastructure to support some or all of the activities of a Professional Media System.

PSN – Acronym for Packet Switched Network.

Q

QoS – Acronym for Quality of Service.

Quality of Service – Aspects of a data network that allow for the transport of certain classes of traffic (such as media streams) that have special requirements for reliable transport. These transport requirements could include zero packet loss, low packet delay, or low packet delay variation.

R

Representational State Transfer (**REST**) - REST is a style of software architecture for distributed systems such as the World Wide Web. The REST architectural style was developed by W3C Technical Architecture Group in parallel with HTTP/1.1. RESTful Web Services differ from SOAP based Web Service in the fact that RESTful Web Service rely upon the HTTP command set and HTTP resources for the transfer of state, where SOAP based web services redefine these function using the SOAP container. Where SOAP based Web Services use the "soapAction" and contents of the SOAP envelope to route requests, RESTful Web Services work more like standard HTTP client, utilizing the requests "path" for request routing.

S

Selective access – The restriction of access to a resource or network to a select group of devices and/or users.

Service Level Agreement – (SLA) A term in a service contract where the performance of a service is formally defined. An SLA for network services may often include maximum allowable times of outages, data bandwidth throughput, packet loss, packet latency, and packet latency variation.

Service Level Agreement Metrics – Metrics such as availability, packet loss, latency, packet delay variation, and bandwidth specified in an SLA term in a service contract.

Service Provider – An external or third party supplier of services.

SLA – Acronym for Service Level Agreement

SLA Metrics – refer to Service Level Agreement Metrics

SONET – Synchronous Optical Networking, a protocol that transfers multiple digital bit streams over optical fiber using light. The SONET standard was defined by Telcordia GR-253-CORE and ANSI T1.105.

Synchronous – Occurring at the same time; coinciding in time; going on at the same rate, exactly together, and exactly in phase with one another. Synchronous refers to the relationship between two or more signals or things.

T

TCP – Acronym for Transmission Control Protocol.

Telecom Service Provider – A communications service provider that provides telephone and similar services, such as Internet connectivity.

Timed –

Time Sensitive Networking – The name of an IEEE 802.1 Task Group to produce new standards that build upon and enhance the existing AVB standards.

Traffic Shaping – Also known as "Packet Shaping", is a network traffic management technique in which delays are introduced to some or all of the datagrams transmitted across a network to bring the datagrams into conformance with some desired traffic profile.

Transmission Control Protocol – The Transmission Control Protocol is an extension of the Internet Protocol intended to be a highly reliable host-to-host protocol for the transmission of data between hosts on a Packet Switched Network. The Transaction Control Protocol guarantees the delivery and prevents the duplications of messages. Please refer to RFC 793 for greater detail.

TSN – Acronym for Time Sensitive Networking

U

UDP – Acronym for User Datagram Protocol.

Uptime – A measure of a time a service has been working and available, the opposite of downtime.

User Datagram Protocol – The User Datagram Protocol is an extension of the Internet Protocol defined to allow for the applications to transmit of messages between network nodes with a minimum amount of overhead using the Internet Protocol. User Datagram Protocol is transaction oriented and does NOT guarantee delivery or provide duplicate protection. Please refer to RFC 768 for greater detail.

User profiles – A set of characteristics associated with a particular user to help enable their effective use of a system.

V

Virtual Private Network – Mechanisms to extend a private network across a public network, often maintaining the functionality, security, and management policies of the private network.

VPN – Acronym for Virtual Private Network

W

WAN – Acronym for Wide Area Network

Wi-Fi – As defined by the Wi-Fi Alliance, any wireless local area network (WLAN) based on IEEE 802.11 standards.

Wide Area Network – A network that covers a broad area and that may be used to connect Local Area Networks (LANs) together.

X

Y

Z

9. References

IEEE 269-2010, IEEE Standard Methods for Measuring Transmission Performance of Analog and Digital Telephone Sets, Handsets, and Headsets.

IEEE 802.1BA™-2011, IEEE Standards for Local and Metropolitan Area Networks: Audio Video Bridging (AVB) Systems.

IEEE 802.1AS™-2011, IEEE Standards for Local and Metropolitan Area Networks: Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks.

IEEE 802.1Q™-2011, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks.

IEEE 802.1Qat™-2010, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks, Amendment 14: Stream Reservation Protocol (SRP). Recent works have incorporated IEEE 802.Qat into IEEE 802.Q.

IEEE 802.1Qav™-2009, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks, Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams.). Recent works have incorporated IEEE 802.Qav into IEEE 802.Q.

IEEE 1722™-2011, IEEE Standard for Layer 2 Transport Protocol for Time-Sensitive Applications in a Bridged Local Area Network.

IEEE 1722a, Unpublished draft revision of IEEE 1722.

IEEE 1733™-2011, IEEE Standard for Layer 3 Transport Protocol for Time-Sensitive Applications in Local Area Networks.

IEEE 1588™-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.

SMPTE ST 2022-5:2013, “Forward Error Correction for Transport of High Bit Rate Media Signals over IP Networks (HBRMT)".

SMPTE ST 2022-6:2012, “Transport of High Bit Rate Media Signals over IP Networks (HBRMT)".

SMPTE ST 2022-7:2012, “Seamless Protection Switching of SMPTE ST 2022 IP Datagrams".

SMPTE RP 168:2009, “Definition of Vertical Interval Switching Point for Synchronous Video Switching".

M. Laabs, “SDI Over IP – Seamless Signal Switching in SMPTE ST 2022-6 and a Novel Multicast Routing Concept,” EBU Technical Review, 2012 Q4.

P. Brightwell, J. Rosser, R. Wadge, P. Tudor. “The IP Studio,” White Paper WHP 268, BBC Research & Development, Sept. 2013. http://www.bbc.co.uk/rd/publications/whitepaper268

