Joint Task Force on Networked Media (JT-NM)
Minimum Viable System Requirements

Of a Sample System Architecture for Live Multi-Camera Studio Production

3 March, 2015

For more info on JT-NM, go to tech.ebu.ch/JTNM or www.videoservicesforum.org/jt-nm

Ownership and Copyright

This work is jointly owned jointly and severally by the European Broadcasting Union (EBU), the Society of Motion Picture and Television Engineers (SMPTE) and the Video Services Forum (VSF), and is licensed under the Creative Commons Attribution-NoDerivs 3.0 Unported License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/3.0/ or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.

Requests for waivers to allow you to use this work in derivative works should be sent to jt-nm-info@videoservicesforum.org.

Minimum Viable System Requirements

Of a Sample System Architecture for Live Multi-Camera Studio Production

Introduction

This report from the Joint Task Force on Networked Media (JT-NM)[1] presents the detailed requirements of the Minimum Viable System (MVS) that was introduced in the "Phase 2 Interim Report December 2014"[2]. MVS is a sample System Architecture that addresses a well-known operational scenario of minimal scope.

This set of requirements could change or be expanded in the future to reflect the development of the JT-NM Reference Architecture.

Context with JT-NM

JT-NM is currently developing a Reference Architecture (RA) for the domain of professional networked media. This Domain RA will provide principles and generalized models that can be reused and tailor made to the specific business context of an organization.

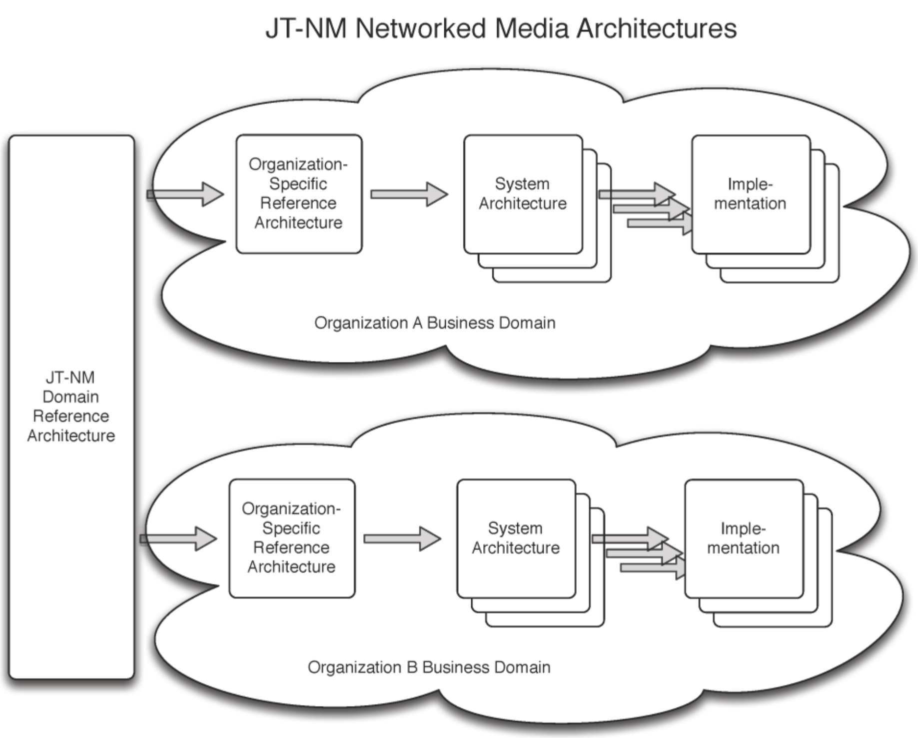

[1] See section "About JT-NM"
[2] See http://videoservicesforum.org/jt-nm/ (Document available to JT-NM participants only.)

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

From an organizational RA, system architects can derive Systems Architectures (SA) that will address specific operational scenarios which in turn can be implemented in real life systems. The aim is that the JT-NM Domain RA will support a broad range of operational scenarios, from conventional live TV production to web-first and cross-media production (for many of Web, mobile, TV, radio, etc.). These workflows include a combination of real-time, near real-time and non real-time workflows. It is hoped that this approach will allow the industry to take advantage of the possibilities offered by networked media architectures.

The Minimum Viable System presented in this report is a sample System Architecture that was developed by the JT-NM Minimum Viable Approach workgroup[3] (MVA) to address a specific operational scenario of live multi-camera studio production.

This exercise is useful as a bottom-up input to help designing the JT-NM Domain RA by identifying key components and interfaces for a specific operational scenario. It is also used to illustrate how the JT-NM Domain RA can lead to a practical system design. Finally, this set of requirements can be reused by the industry as an input to standardization efforts.

Operational Scenario

The operational scenario to be addressed is the transport of live media within the broadcast plant to support a multi-camera, live studio production; specifically, a live, multi-camera sports half-time show being produced entirely within the four walls of a single facility. It involves switching between a number of cameras, video effects, graphics overlays, live audio production, the playback of pre-recorded pieces, video and audio monitoring, tally lights, intercom, interruptible foldback (IFB), production room monitoring, studio floor monitors, studio audio monitors, and teleprompter.

The scenario involves, switching from a live incoming game feed to the studio show (no interaction with the field is required), where a producer choses from various sources in order to create a live half-time show. The show is fed out live to a network origination facility where it is transmitted nation-wide.

[3] See section "Minimum Viable Approach (MVA) Workgroup"

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Operational Scenario: Live Studio Production - Physical Layout

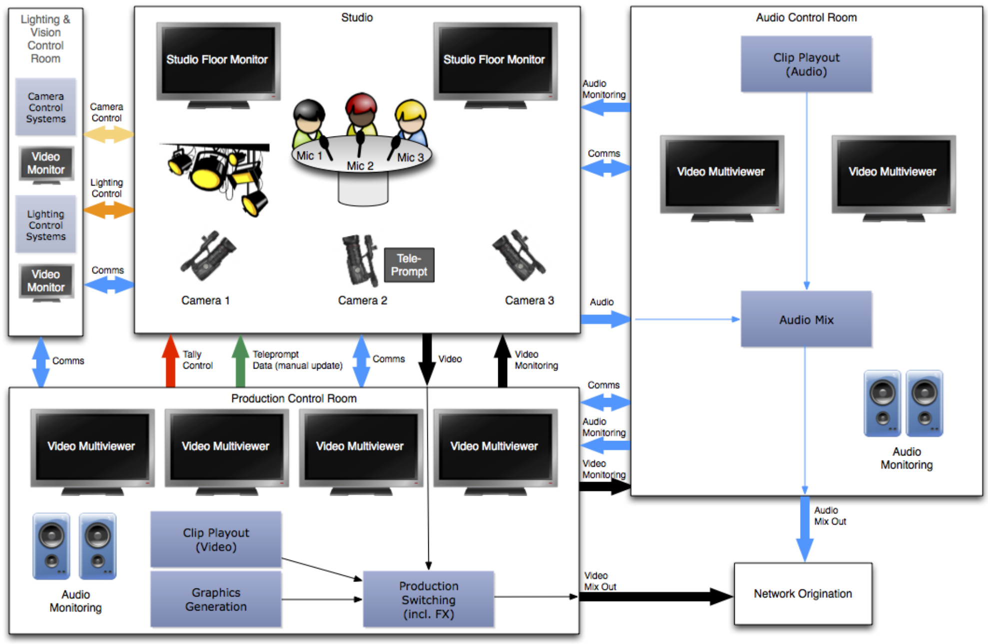

Operational Scenario: Live Studio Production - Signal Flow

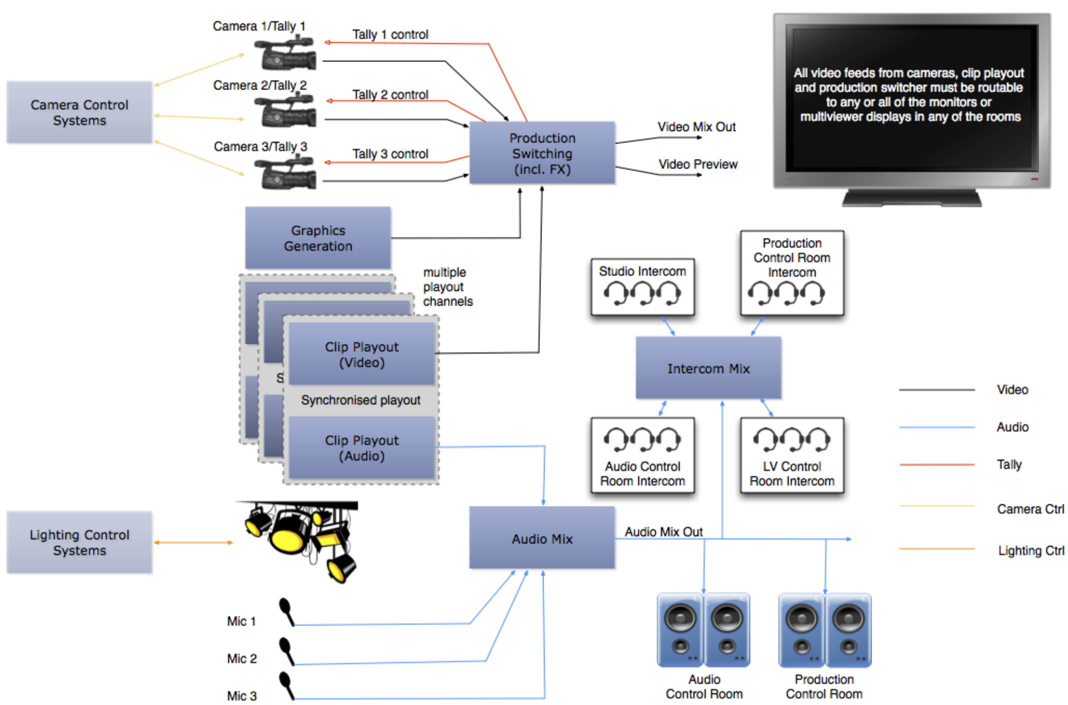

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

This scenario does not include interaction with talent in the field, and does not include interaction with other production facilities, either in the facility or at remote locations. It also does not include commercial integration – this is done in the network origination facility.

This first iteration of MVS requirements only addresses the movement of Live Media within the broadcast plant. Live Media shall be defined to include uncompressed video (including key & fill), uncompressed audio, and Ancillary Data as defined by SMPTE.

Definitions

Flow – A sequence of packets of a single type of essence from a single source.

Bundle – A logical association of related flows. The association may be for simultaneous presentation (such as audio and video). The type of association of the flows in a bundle shall be defined by the bundle description. The architecture may need to support a hierarchy of bundles (bundles of bundles).

Source - The originator of one or more flows.

Destination - The receiver of one or more flows.

Source Device - A device with one or more Sources.

Destination Device - A device with one or more Destinations (sinks).

Node – Generic term for a source device or destination device.

Lock Up Time - The time from beginning of physical connection or flow reception to the usability of the media; this should be specified in a solution.

SDI – Where the term "SDI" is used, the group refers to all three of "SD-SDI" (SMPTE 259M), "HD-SDI" (SMPTE ST 292-1), and "3G-SDI" (SMPTE ST 424).

Requirements

The requirements include notes indicated by italicized text within brackets such as [*this is a note]*. These notes are not part of the MVS requirements, but the authors felt that the group's thinking on these matters was worth making a note of for those working on a solution to these requirements, or in some cases those working on future requirements documents beyond the MVS scenario and functionality.

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Fundamental Precept

The solution should be developed using existing standards where possible, or using existing standards with the least amount of amendment or extension. An amendment or extension should be fully defined and described.

Payload

Payload is the data representation of uncompressed live media & metadata for transport. A key requirement is that elementary essence types (e.g., video, audio, and associated data essence) shall be carried as separate flows. This is to facilitate processing, modification, and add/drop of essence flows using network switching without the need to perform the SDI mux/demux operation. Flows are constrained to maintain a constant type of payload, they are required to have data boundaries that are octet-aligned, and they should be free of additional octet padding when possible.

Flow Payload Consistency – An individual flow shall maintain a constant data type and description (i.e. there shall be no changes to frame rate, resolution, etc.). If there is a need for a change in the flow data type, this shall be supported by using a different flow.

[Note: The ability to have changes in flow data type may be useful for future iterations, especially OTT distribution. There is also a general assumption that flows coming into the facility will be "normalized" to specific flow types within the plant for MVS purposes.]

Elementary Payload Flows - Elementary essence types (e.g. video, audio, ancillary data) should be carried as separate flows.

[Note: The reason for having elementary essence flows is to facilitate processing and modification of the essence flows, for instance to add or subtract elementary essence flows, especially in the software domain without the need to perform an SDI mux/demux operation.

There also have been suggestions that the different color components of an image could be their own elementary essence flows, or that samples from specific color components could be segregated into contiguous blocks in the payload. This may provide advantages in video processing by simplifying the loading of component planes into a frame buffer without having to decompose every pixel into its separate color components.]

Payload Octet Alignment - Data in a packet of a flow should be octet-aligned when possible.

[Note: This requirement is to make processing easier on typical computer systems.]

No Payload Packing Required - The use of padding should be avoided. Fixed size packets are not a requirement.

Video

Video Format Capability - The payload shall at least support what today's SDI transport supports. This should include all the digital formats currently in use, as well as those formats

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

that are likely to be used in the near future. Although the MVS scenario does not require any use of SDI, the solution should be capable of providing "transparent transport" of SDI payload bit streams over the network (with "transparent" meaning that the input SDI bit stream to an SDI to IP converter shall match the SDI bit stream egressing from an IP to SDI converter).

[Note: The MVS effort is not intended to replicate SMPTE ST 2022-6 (SDI over IP), but instead specifies that legacy SDI signals be decomposed into elementary essence streams over IP, and then recomposed back to the SDI mux if necessary for legacy SDI receivers.]

[Note: There were concerns expressed about the ability to route a bundle as if it were a single flow. There also was a feeling that individual flows should be routable via typical networking addressing (L2/L3/L4). These concerns may be in conflict as the solution is specified. Further thought may be required to resolve this issue.]

Video Format Resolution – The system shall be capable of carrying video payload of any resolution up to the size of UHDTV2 (7680 x 4320).

Video Image Rate – The system shall be capable of carrying video payload of any frame or field rate up to 300 Hz, and shall be capable of carrying NTSC style fractional frame rates at multiples of (1000/1001).

Video Sample Depth - The system shall be capable of carrying video payload sample depths of 10 or 12 bits.

Video Chroma Sampling - The system shall be capable of carrying a video payload of 4:2:2 or 4:4:4 chroma sampling.

Alpha Channel – The system shall be capable of carrying a video payload that contains an Alpha Channel (a component that represents transparency).

Color Spaces - The system shall be capable of carrying video payload in the color space of ITU-R Rec. BT.601, ITU-R Rec. BT.709, and ITU-R Rec. BT.2020.

[Note: The group felt that HDR spaces should be considered for future requirements. The ability to carry 12 bit video sample depths may provide some capability for this. It would also be a good design goal to allow incremental support for enhancements in video format capability in the future without requiring a complete plant redesign, if possible.]

Image Scanning Types - The payload shall support both progressive and interlaced scanned images.

Video Payload Frame Alignment - Data carried in a single packet shall contain data from a single frame.

[Note: This requirement is to allow for frame accurate switching at packet level without decomposing the packet.]

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Video Payload Line Alignment - Data carried in a single packet should contain data from a single line.

[Note: This requirement is to allow for line-based processing without decomposing the packet, but the group noted that this might not always be efficient for the carriage of blocks of one color component.]

Video Payload Sample Alignment - In the case where a flow contains multiple color components, a packet should contain all samples from the pixels contained within the packet.

[Note: This requirement is so all pixels in a packet can be obtained without having to decompose a second packet to extract additional color samples for that pixel.]

Video Payload Pixel Alignment - The order of pixels on a line shall be reflected in the order of pixels within a packet, but packing techniques to more efficiently process samples that are not octet sized shall be allowed.

[Note: For example, one may take a set of 10-bit samples and transmit the most significant octets first, followed by a packed set of the 2 least significant bits from each sample.]

Video Payload Boundaries - The occurrence in a packet of the Start of frame, End of frame, Start of line, and End of line shall be identified.

Video Payload Location - The line number and start pixel number (x,y) position in the video raster of video data within that packet shall be contained within every single packet.

[Note: This enhances the ability to handle video switching without context, and to handle packet loss or out-of-order packets gracefully.]

Video Payload Length Determination - The number of pixels in a packet shall be determinable.

Video Payload Sequence Numbering - Frame/field sequence numbers [with a short roll-over length] should be available in each packet.

[Note: Some of the potential reasons behind this requirement include synchronization of flows such as stereoscopic and hitless failover at the frame level. It is possible that timestamps (with a flow ID) could provide this, but the group is uncertain if a specific header metadata item is required or not. Potential counter lengths discussed included 8 bits, 10 bits, or trying to identify a number that fits a certain time length such as 100ms, etc.]

Audio

Audio Format Capability – The audio payload shall support 24-bit uncompressed audio as per SMPTE ST 299-1 (i.e. the audio data derived from AES3) that may contain linear PCM audio or non-PCM data formatted according to SMPTE ST 337.

[Note: It is unclear to the group if metadata bits in AES3 need to be carried, as they are not carried in AES67, SMPTE 302, or EBU ACIP TECH3326.]

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Audio Format Sampling – The audio payload shall support 48 kHz and 96 kHz sampling.

Audio Format Phase Coherency – The audio payload shall support the ability to deliver a number of related phase coherent channels (e.g. stereo, 5.1, 22.2, etc.).

[Note: There is ongoing work in the SMPTE, EBU, and AES regarding audio channel signaling which could be valuable for an MVS solution. As with enhancements of video format capability, it would be a good design goal to allow for audio format capability enhancements in the future without requiring a complete plant redesign, if possible.]

Ancillary Data

Ancillary Data Format Capability – The ancillary data payload shall be capable of carrying all non-audio SMPTE ANC data packets that may be found in HANC and VANC space in SDI as per SMPTE ST 291-1.

Ancillary Data SDI Transparency - The ancillary data payload shall carry metadata required to reconstruct the original line placement of ANC packets within SDI if the original source that is being packetized is SDI.

[Note: Some legacy devices have limited capability for ANC data carriage or requirements for ANC data to be on specific lines.]

Non-Located Ancillary Data Signaling - If ancillary data is to be carried without any specification as to the original line placement, that fact shall be signaled.

Additional Payloads

Additional Flow Payloads - The solution shall not preclude the addition of other flows of essence or metadata that are not defined in these requirements.

Transport

Transport provides network carriage for payloads. It provides for sequence numbers to ensure re-ordering of out-of-sequence packets, to allow detection of lost packets, and to allow for "hitless" failover of redundant flows. Transport shall work on typical commercial-grade off-the-shelf Ethernet switches with functionality typically available from multiple vendors, and shall not require Forward Error Correction (FEC).

Sequence Numbers Loss & Re-Order – The transport shall provide sequence numbers with enough bits to ensure re-ordering of out-of-sequence packets and to allow detection of lost packets.

[Note: The maximum number of lost packets to be detected accurately by sequence numbers should be determined by the group developing the MVS solution, and should reflect actual expected loss scenarios for different kinds of flows.]

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Sequence Numbers "Hitless" Failover - The transport shall carry sequence numbers with enough bits to allow for hitless failover using redundant packet transmission [e.g. SMPTE ST 2022-7].

Transmission Type - The transport shall support point to multipoint transmission.

Single Connection - The transport shall support point-to-point transmission over a single physical connection.

[Note: The group felt that it was important that the MVS solution should work over a single direct Ethernet link, say between a camera and a video recording system, without any additional hardware (such as servers or switches).]

COTS Capability - The transport shall work on typical commercial-grade off-the-shelf Ethernet switches with functionality typically available from multiple vendors.

No FEC Required - The transport should not require forward error correction (FEC) or packet retransmission for reliability against packet loss of the transport, but will not preclude these mechanisms.

[Note: The group feels that FEC may be necessary in some cases, for example carrying flows over a WAN, but is likely out of the scope of this effort.]

[Note: It was unclear to the group if timestamps referring to the time of transmission of transport packets is required. RTP does not have this, nor does MPEG-TS. MVS does not include a requirement for timestamps referring to the time of transmission of transport packets. See the timing section regarding more timestamp discussion.]

Identity

Identity provides unique identification of nodes, sources, destinations, flows, and bundles. The unique identifier shall be the link to a discovery mechanism for technical and descriptive metadata about the entity.

[Note: The group felt that urn:uuid might be a good unique identifier.]

Uniqueness – Identifiers shall allow all nodes, sources, receivers flows & bundles to be uniquely identified.

Key to Identification – The unique identifier (ID) shall be the link to a discovery mechanism for technical and descriptive metadata about the entity.

Flow Identity – The unique identifier of a flow shall be discoverable from metadata carried in every packet of the payload of that flow.

[Note: The group wanted to be clear that the unique flow identifier could either be indirectly discoverable or directly discoverable. An example of indirect discovery would be the SSRC in an RTP packet pointing to an identity carried in an RTCP SDES packet.]

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Flow Source Identity - The identity of the source of a flow shall be discoverable through reference based on the unique identifier associated with that flow.

Hitless Failover of Flows - Identity of flows should be a key to whether flows are "coherent" for hitless failover (e.g. SMPTE ST 2022-7).

Programmatic Assignment - The identifiers shall be capable of being generated uniquely via a programmatic method.

Source Device Identity - The identity of a source device shall be discoverable through the use of the unique identifier associated with a source.

Registration

Registration is a mechanism to provide the registration of identities.

Registration of Unique Identifiers – The registration mechanism shall support the registration of unique identifiers and the linking of metadata items to that unique identifier.

Manual or Automatic Registration - The registration mechanism shall support either manual or automatic registration.

[Notes: The group felt that even flows that were not currently transporting essence data should be able to have an identity registered. The group also felt that a metadata element linked to the identity of the flow should provide information as to whether the flow is "active" and transporting essence data or "inactive" and not transporting essence data.]

Centralized or Distributed Registry - The registration mechanism should support either a centralized registry or a distributed registry.

Discovery

Discovery is a mechanism to discover nodes, flows, and bundles. Mechanisms include peer-to-peer, pre-configured, pushed or pulled, centralized and decentralized.

Peer-to-Peer Discovery – The discovery mechanism shall allow for two or a small number of devices connected via a "single cable" or attached to a networking switch to discover each other.

Flow Description Discovery – The discovery mechanism shall allow for the ability to infer media flow description without access to external servers.

Pre-configured Discovery – A node should allow for the manual or other direct loading of discovery information.

Pushed Discovery – A node should be able to listen "promiscuously" for announcements to achieve discovery.

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Pulled Centralized Discovery – A node should be able to query a central registry to achieve discovery.

Pulled Decentralized Discovery - A node should be able to direct a query at a node or set of nodes to achieve discovery.

Discovery Capability – The discovery mechanism shall allow for discovery of the following from nodes: their sources, their network address, their flow reception capabilities, their flows, and their flow descriptions.

Connection Management

Connection management is a mechanism to manage the establishment & tear down of flows.

Loose Coupling of Sources and Destinations – Connection management shall allow for sources and destinations to be location independent within a network, to be able to be moved to different hosts. Hosts addresses need not be hard-coded in applications.

Point to Multipoint – Connection management shall support a flow being received by multiple destinations.

Non-Interruptability – Connection management shall allow for the addition or removal of destinations without interruptions of flows to other destinations.

Detect Disconnects – The connection management should be able to detect the disconnection of destinations from flows.

Publish/Subscribe – Connection management should make use of a publish/subscribe model.

Timing/Clocks/Sync

This group of requirements provides support for synchronization and time stamping of media in flows. There are specific requirements for tolerances of audio/video synchronization, multi-channel audio imaging, and multiple camera synchronization.

Clock Synchronization - Synchronization of clocks on devices shall be achieved through network transmissions.

Legacy Timing Support - Legacy synchronization signals (such as black burst and timecode) shall be allowed as one option of a source of network time through an appropriate gateway.

Slave Synchronization - Slave clocks shall be synchronized to the master with an accuracy of 1µs within 5 seconds of connection to the network.

Timestamp Time Source - Timestamps of captured media shall be derived from networked distributed time.

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Audio/Video Synchronization - The synchronization method shall allow for synchronization of audio and video within the bounds of EBU TR R37-2007. This means that synchronization accuracy at each stage should lie within audio 5 ms early (sound before picture) to audio 15 ms late (sound after picture), and at any output intended for emission, the difference in the relative timing of the sound and vision components should lie within the overall (end-to-end) range of ≤40 ms (sound before picture), ≥ 60 ms (sound after picture). The preferred goal is for all processing subsequent to capture is zero additional drift from A/V synchronization if possible.

Multi-Channel Audio Imaging Synchronization – The synchronization method shall allow for synchronization of multiple channels of audio intended for synchronized audio imaging to within 10 µs.

[Note: This requirement is based on the finding that a just noticeable difference in interaural time difference down to 10 µs is found for sinusoids below 1500 Hz in normal hearing listeners, from Yost, W.A., "Discrimination of interaural phase differences," J. Acoust. Soc. Am. 55, 1299–1303, 1974. Also a suggestion of allowable audio sample skew between stereo channels within the boundaries of ±11 µs from R. Steinmetz, "Human perception of jitter and media synchronization", IEEE J. Select. Areas Commun., vol. 14, pp. 61-72 1996.]

Number of Phase Aligned Groups - The synchronization mechanism shall provide for multi-channel audio imaging synchronization for up to 8 mono channels.

Multiple Camera Synchronization Requirement – The Synchronization system should support the ability to have the capture of video from multiple cameras with the synchronization of raster sampling within ±2 µs.

[Note: This is provided so that in live capture, a switch or fade between cameras pointed at the same object is undetectable. The ±2 µs timing is based on the length of the SMPTE RP 168 switch point for 720p59.94, which is the shortest RP 168 switching area (4.37 µs). The group notes that this requirement may be overly stringent, but without additional information it is difficult to relax this constraint.]

Ancillary Data Synchronization - Ancillary data shall be able to be synchronized to the video frame with which it is associated.

SMPTE 12M Timecode – SMPTE 12M time code shall be able to be derived from networked distributed time.

Synchronization Timestamp Meaning – The synchronization timestamp for video shall be the time of the capture of the top of the field (interlaced) or the top of the frame (progressive), and for audio shall be the time of the first sample of a packet payload.

Synchronization Timestamp Use – The synchronization timestamp shall be used to achieve synchronization between flows.

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

[Note: If flows are recorded, there may be a need to have a separate capture timestamp that is a persistent record of the original capture time for index to the storage, and a synchronization timestamp that is reflective of the time a recorded flow is re-emitted on playback.]

Single Timestamp Time Base - Timestamps should be based on a single time base.

[Note: The group agrees that there are certain benefits to placing the timestamp in the payload including editing of streams where timing modifications are not necessary.]

Flow Description

The characteristics of flows are required to be described. This includes items such as resolution and frame rate for video, sampling frequency and bit depth for audio. Flow descriptions are required to be linked to the flow unique identifier.

Description - There shall be a mechanism to describe the characteristics of the flow data (for instance, the video is 720p/59.94, the audio is 24-bit 48 kHz, etc.).

ID Linkage - The description shall be linked to the unique identifier of the flow.

[Note: The decision to carry all or some of the flow description along with the data transport "in band" or to have all or some of it carried "out of band" should be considered by implementers - also whether the description is "pushed" or "pulled" - the solution may depend on whether a "single cable connection" between two can be supported.]

Video Flow Description – A flow description shall include:

- Horizontal and vertical active image size;
- Frame/Field rate (including 1000/1001 fractional frame rates);
- Sample depth in bits;
- EOTF/Gamma;
- Chroma sampling;
- Alpha Channel presence;
- Color space.

Audio Flow Description – An audio flow description shall include:

- Sampling frequency;
- Sample depth in bits;
- Channel configuration (number of channels, packing, channel mapping (L/R/C etc.)).

Availability – The flow description shall include information required for network availability and transport information of flow.

Bundle Description

Bundle description provides a description for logical grouping of flows, for example associating a particular audio flow, video flow, and ANC data flow.

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Bundles – A bundle shall be able to provide an association between any combinations of any number of audio, video, or ancillary data flows.

ID Linkage – The bundle description shall be linked to the unique identifier of the bundle.

Association – The bundle description shall provide a description of the association between flows in the bundles (e.g. simultaneous presentation).

Big Issues that We Don't Know

The JT-NM process has revealed much information about the potential Reference Architecture of professional network media, but it has also highlighted major issues that we don't know, and are likely only to be determined through further experimentation and implementation. These include:

Best Payload Formats – Especially for video, we have many choices for payload formats. Today, the "lingua franca" of the broadcast plant is uncompressed 10-bit 4:2:2 chroma subsampled video. This format was standardized when the cost in both computational power and latency of compressed formats was much higher than today. We now have the opportunity to trade off a few lines of latency for a 2:1 compression, or perhaps even a frame of latency for 10:1 compression.

Moreover, even for the uncompressed case, the existing uncompressed "sample serial" digital format of SDI was built around the need for scanning CRTs. A Professional Networked Media (PNM) plant could send samples of a single color component channel per packet to simplify the loading of color channel planes in memory for video processing (at the expense of a few packets of added latency before an entire pixel is available for processing). There are also potentially computationally more efficient ways to pack 10-bit or 12-bit uncompressed digital samples into an octet-based network packet than simply running the bits of the sample together.

Latency Requirements – Over the SDI coaxial cable and across digital circuit crosspoints (crosspoint switches), the broadcast industry had the luxury of near zero serial data latency. Network switches have latency that is generally on the order of a few microseconds. Moreover, practical compressed payload representations can have significant latency, up to a frame or so. What is the real requirement in our broadcast plants for interconnection latency? Unfortunately, it is difficult to determine how much latency is too much until large, complex network-based broadcast plants are actually designed.

Synchronization of Transport – In the days of analog, broadcast plants needed raster synchronization down to color subcarrier phase to achieve reliable switching. In digital, SMPTE RP 168 decreased the required raster synchronization to around four microseconds. However with the availability of fast and cheap RAM for buffering, there is no longer a clear need for the transport of video in the broadcast plant to be raster synchronous at all. A plant could operate with the equivalent of a "frame sync" at the input of every device that mixes, combines, or switches media. This would not mean abandoning synchronization of audio channels for multichannel audio imaging or audio/video synchronization, but those synchronizations would

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

only need to be achieved at nodes that combine or present media. Theoretically, such a non-synchronous system could even have less overall latency than a synchronized plant. However, the stability and practicality of a large, non-synchronized PNM broadcast plant needs to be examined, as re-entrant flows could cause latency oscillations in a PNM system that is not well engineered.

Heterogeneity versus Standardization – A PNM plant could very well operate with a mix of the above-mentioned choices. There could be some uncompressed video, and some compressed video. Some elements of the plant could operate with low latency, some with higher latency. Some nodes of the plant can transmit synchronously, others asynchronously. All this is possible with the carriage of media over generic network protocols.

However, there likely would be substantial industry benefits from some level of standardization, especially if there are operating profiles that fit a large amount (perhaps 80%) of use cases. Part of the challenge of PNM will be determining how much to standardize, and how much flexibility to allow in those standards.

About JT-NM

The Joint Task Force on Networked Media (JT-NM) was formed by the European Broadcasting Union, the Society of Motion Picture and Television Engineers and the Video Services Forum to address issues that have arisen in the transition from purpose-built broadcast equipment and interfaces (SDI, AES, crosspoint switchers, etc.) to IT-based platforms (Ethernet, IP, servers, storage, cloud, etc.); a transition that is ongoing in the professional media industry. It was set up to foster a discussion among subject-matter experts to drive the development of an interoperable network-based infrastructure that will encompass live media production and file-based workflows. It brings together broadcasters, manufacturers, standards bodies, and trade associations. The JT-NM comprises more than 295 participants from 175 organizations.

This is a critical activity, since the dynamics of the industry are rapidly changing, with new players vying for a share of the revenue pie, and with continually evolving viewer consumption habits. Not only do media companies need to be more flexible in order to respond quickly to new opportunities to monetize content, but they have a strong desire to take advantage of economies of scale by leveraging the massive and ongoing investments currently being made in the IT industry.

At a high level, the JT-NM is seeking to make it easier to build professional media infrastructures by promoting interoperability between IT-based media components. The effort extends beyond this, however. It includes helping the industry to develop a shared understanding about how to describe these systems, developing a common vocabulary, and promoting best practices for system designs.

(cc) BY-ND This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported

Minimum Viable Approach (MVA) Workgroup

This report presenting the Requirements for a Minimum Viable System (MVS) was prepared by the Minimum Viable Approach workgroup (MVA).

Chair: Thomas Edwards (Fox)

Participants: Merrick Ackermans (Turner Broadcasting System), Bryan Balcombe (London Live), Michael Bany (DVBLink), Peter Brightwell (BBC), Peter Chave (Cisco), Magnus Danielson (Net Insight), Mike Ellis (BBC), Paul Gardiner (Sony), Rajive Joshi (RTI), Toshiaki Kojima (Sony), Al Kovalick (Media Systems Consulting), Eric Pohl (NTC), Steve Posick (ESPN), Christian Adell (CCMA), Andrew Rayner (Nevion), Pat Waddell (Harmonic), Robert Wadge (BBC).

This work is licensed under a Creative Commons Attribution-NoDerivs 3.0 Unported