SMPTE STANDARD

Professional Media Over Managed IP Networks: System Timing and Definitions

Table of Contents

Approved September 18, 2017

Foreword

SMPTE (the Society of Motion Picture and Television Engineers) is an internationally-recognized standards developing organization. Headquartered and incorporated in the United States of America, SMPTE has members in over 80 countries on six continents. SMPTE's Engineering Documents, including Standards, Recommended Practices, and Engineering Guidelines, are prepared by SMPTE's Technology Committees. Participation in these Committees is open to all with a bona fide interest in their work. SMPTE cooperates closely with other standards-developing organizations, including ISO, IEC and ITU.

SMPTE Engineering Documents are drafted in accordance with the rules given in its Standards Operations Manual. This SMPTE Engineering Document was prepared by Technology Committee 32NF.

Intellectual Property

At the time of publication no notice had been received by SMPTE claiming patent rights essential to the implementation of this Engineering Document. However, attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. SMPTE shall not be held responsible for identifying any or all such patent rights.

Introduction

This section is entirely informative and does not form an integral part of this engineering document.

The capability and capacity of IP networking equipment has improved steadily, enabling the use of IP switching and routing technology to transport and switch video, audio, and metadata essence within television facilities. Existing standards such as SMPTE ST 2022-6 have gained some amount of use in this application, but there was a desire in the industry to switch different essence elements separately.

This family of SMPTE standards builds on the work of VSF TR-03 and TR-04, and of AES67, documenting a system for inter-networking various essence streams and capturing the timing relationships between those streams. The system is intended to be extensible to a variety of essence types.

This standard covers the system as a whole, the timing model, and common requirements across all essence types. Subsequent parts of this standard document the individual media essence types and their individual requirements as used within this system.

1 Scope

This family of engineering documents defines an extensible system of RTP-based essence streams referenced to a common reference clock, in a manner which specifies their timing relationships.

This standard specifies the system timing model and the requirements common to of all of the essence streams.

2 Conformance Notation

Normative text is text that describes elements of the design that are indispensable or contains the conformance language keywords: "shall", "should", or "may". Informative text is text that is potentially helpful to the user, but not indispensable, and can be removed, changed, or added editorially without affecting interoperability. Informative text does not contain any conformance keywords.

All text in this document is, by default, normative, except: the Introduction, any section explicitly labeled as "Informative" or individual paragraphs that start with "Note:"

The keywords "shall" and "shall not" indicate requirements strictly to be followed in order to conform to the document and from which no deviation is permitted.

The keywords, "should" and "should not" indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain possibility or course of action is deprecated but not prohibited.

The keywords "may" and "need not" indicate courses of action permissible within the limits of the document.

The keyword “reserved” indicates a provision that is not defined at this time, shall not be used, and may be defined in the future. The keyword “forbidden” indicates “reserved” and in addition indicates that the provision will never be defined in the future.

A conformant implementation according to this document is one that includes all mandatory provisions ("shall") and, if implemented, all recommended provisions ("should") as described. A conformant implementation need not implement optional provisions ("may") and need not implement them as described.

Unless otherwise specified, the order of precedence of the types of normative information in this document shall be as follows: Normative prose shall be the authoritative definition; Tables shall be next; then formal languages; then figures; and then any other language forms.

3 Normative References

The following standards contain provisions which, through reference in this text, constitute provisions of this engineering document. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this engineering document are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

AES67:2015, AES standard for audio applications of networks - High-performance streaming audio-over-IP interoperability

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

Internet Engineering Task Force (IETF) RFC 768 User Datagram Protocol [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc768.txt

Internet Engineering Task Force (IETF) RFC 791 Internet Protocol [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc791.txt

Internet Engineering Task Force (IETF) RFC 2460 Internet Protocol, Version 6 (IPv6) Specification [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc2460.txt

Internet Engineering Task Force (IETF) RFC 3550 RTP: A Transport Protocol for Real-Time Applications [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc3550.txt

Internet Engineering Task Force (IETF) RFC 3551 RTP Profile for Audio and Video Conferences with Minimal Control [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc3551.txt

Internet Engineering Task Force (IETF) RFC 3376 Internet Group Management Protocol, Version 3 [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc3376.txt

Internet Engineering Task Force (IETF) RFC 4566 SDP: Session Description Protocol [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc4566.txt

Internet Engineering Task Force (IETF) RFC 4604 Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc4604.txt

Internet Engineering Task Force (IETF) RFC 7104 Duplication Grouping Semantics in the Session Description Protocol [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc7104.txt

Internet Engineering Task Force (IETF) RFC 7273 RTP clock Source Signaling [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc7273.txt

SMPTE ST 272:2004 Formatting AES Audio and Auxiliary Data into Digital Video Ancillary Data Space

SMPTE ST 299-1:2009 24-Bit Digital Audio Format for SMPTE 292 Bit-Serial Interface

SMPTE ST 299-2:2010 Extension of the 24-Bit Digital Audio Format to 32 Channels for 3 Gb/s Bit-Serial Interfaces

SMPTE ST 2022-7:2013 Seamless Protection Switching of SMPTE ST 2022 IP Datagrams

SMPTE ST 2059-1:2015 Generation and Alignment of Interface Signals to the SMPTE Epoch

SMPTE ST 2059-2:2015 SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications

4 Terms and Definitions

For the purposes of this standard, the following terms and definitions apply.

Note: The terminology used in some of the normative references (particularly those of the IETF) differs from conventional usage within the SMPTE. Readers are reminded that the definitions of this section supersede any similar term in the normative references.

4.1 Device

hardware or software application that can include multiple senders and receivers

4.2 RTP Stream

sequence of IP datagrams constructed in accordance with this family of standards, utilizing the Real-Time Transport Protocol as specified in IETF RFC 3550

4.3 Sender

element within a device which originates one RTP stream into the network

4.4 Receiver

element within a device which terminates one RTP stream from the network

4.5 Network

IP datagram transport mechanism with sufficient capacity to deliver the RTP stream from the sender to the Receiver

4.6 Network Byte Order

convention for the transport ordering of octets within values which require more than one octet to represent, such that the most numerically significant octet is transported first, followed by the remaining octets in order of decreasing significance – an order colloquially referred to as “big-endian”

4.7 Standard UDP Size Limit

maximum limit on the size of the UDP portion of the IP datagrams when operating in the size range specified in section 6.3

4.8 Extended UDP Size Limit

maximum limit on the size of the UDP portion of the IP datagrams when operating in the size range specified in section 6.4

4.9 Image Capture Time

a time instant representative of the scene capture time

4.10 RTP Timestamp

data field specified as “timestamp” in IETF RFC 3550 section 5.1 titled “RTP Fixed Header Fields”

4.11 Media Clock

clock incrementing at a fixed rate as defined for the specific media essence type, with a synchronous relationship to the frame rate or sampling rate of the media

4.12 RTP Clock

clock incrementing at the same fixed rate as the media clock, and potentially offset from the media clock

4.13 Common Reference Clock

network delivered PTP clock, that can be used for synchronization of the device internal clocks across different Devices in the network

4.14 Device Internal Clock

clock internal to a device, that can be synchronized to the common reference clock

5 Textual Conventions

5.1 SDP Parameters and Values

The names and values of SDP media type parameters within the text of this standard are formatted using a monospaced font (such as Courier) except when they appear in section headings.

6 Network Interface Requirements

6.1 General Requirements

Devices compliant to this standard are interconnected at the data-plane by a network. In cases where redundant connections are used, two networks (in parallel) may be used.

The network interfaces specified in this standard shall support IPv4, wherein streams are transported using IP version 4 as specified in IETF RFC 791. Devices should support IPv6 as specified in IETF RFC 2460.

Note: All of the IETF RFC documents which are listed in section 3 (Normative References) are Standards-Track Documents within the IETF; they are however at varying phases of standardization within the IETF process, not all of them have reached the final phase of "Internet Standard". The IETF standardization phases are specified in IETF RFC 6410.

6.2 Real-Time Transport Protocol (RTP)

All of the streams specified in this standard shall use the Real-time Transport Protocol as specified in IETF RFC 3550, and shall conform to the RTP profile specified in IETF RFC 3551, subject to the restrictions and additional requirements specified in this standard and in the specific media-format standards which define each of the media essence types.

All RTP streams shall use UDP as specified in IETF RFC 768 for transport of RTP. UDP does not guarantee reliable data transport and receivers should be capable of receiving streams with occasional dropped, late or out-of-order packets.

Note: The UDP header checksum is optional in IPv4, and many IPv4 senders populate zero instead of calculating the UDP checksum, as is permitted by IETF RFC 768. In IPv6, IETF RFC 2460 section 8.1 specifically requires the UDP checksum to be calculated, and updates the defined pseudo-header fields.

RTP session multiplexing (on the same multicast group/port) as specified in IETF RFC 3550 section 5.2 shall not be used in the context of this standard.

RTP Control Protocol (RTCP), as specified in IETF 3550 section 6, may be used in the context of this standard. Senders and receivers may implement RTCP, and receivers shall tolerate the presence of RTCP.

All RTP streams shall use dynamic payload types chosen in the range of 96 through 127, signaled as specified in section 6 of IETF RFC 4566, unless a fixed payload type designation exists for that RTP stream within the IETF standard which specifies it.

Unless otherwise specified, all multi-octet numeric values expressed in the RTP header and in the RTP payload headers and payloads shall be expressed in network byte order.

When using redundant streams, the streams shall be generated using the method specified in SMPTE ST 2022-7 and as constrained in section 8.3 of this standard.

6.3 Standard UDP Size Limit

The Standard UDP size limit shall be 1460 octets. The UDP size is reflected in the UDP header, and includes the length of the UDP header (8 octets) and also the RTP headers and data. Senders shall not generate IP datagrams containing UDP packet sizes larger than this limit unless operating conformant to the optional extended UDP size limit specified in section 6.4. Regardless of the presence or size of any RTP header extensions, senders shall adhere to the UDP size constraints.

All receivers shall be capable of receiving UDP packets up to the standard UDP size limit.

Receivers are not required to reassemble fragmented IP datagrams.

Annex A describes the origin of the standard UDP size limit.

6.4 Extended UDP Size Limit

The Extended UDP size limit shall be 8960 octets.

Senders may transmit and receivers may support reception of IP datagrams up to the extended UDP size limit, subject to the constraints of the specific essence transport standard in use.

Senders operating with UDP sizes which exceed the standard UDP size limit shall include a Media Type Parameter `MAXUDP` as specified in section 8.4.

Note: Annex A describes the origin of the extended UDP size limit.

6.5 Unicast and Multicast

Senders and receivers shall support IPv4 multicast transmission and reception (respectively) of streams including IGMP signaling as specified in IETF RFC 3376.

Senders and receivers shall support IPv4 unicast addressing of streams as specified in IETF RFC 791.

Senders and receivers should support IPv6 multicast transmission and reception (respectively) of streams as specified in IETF RFC 2460, including Multicast Listener Discovery Protocol version 2 as specified in IETF RFC 4604.

Senders and receivers should support IPv6 unicast transmission and reception (respectively) of streams as specified in IETF RFC 2460.

7 System Timing Model

7.1 Introduction to the System Timing Model (Informative)

Inter-stream synchronization relies on RTP timestamp values in the RTP packet headers that are transmitted by various Senders originating from a common reference clock. The common reference clock can be distributed to all participating senders and receivers via IEEE 1588-2008 Precision Time Protocol. Synchronization across multiple essence streams at the receiving device is achieved by the comparison of the time offsets (between the RTP timestamp and RTP clock) for each of the various essence streams that share a Common Reference Clock and its associated epoch. If the common reference clock is unavailable, devices may signal use of a device internal clock and epoch, such that streams from that device may still be synchronized with each other.

7.2 Device Internal Clock and Synchronization to the Common Reference Clock

Each device shall maintain a device internal clock. A common reference clock should be provided and distributed on the network by means of IEEE 1588-2008 Precision Time Protocol (PTP). If the device is in PTP slave state, the device internal clock shall be synchronized to the common reference clock. If the device is in PTP master state and hence the device internal clock is the source of the common reference clock, the device internal clock may be synchronized directly to a traceable time source.

The configurable PTP dataset member defaultDS.slaveOnly may be set to prevent a device from entering the PTP master state. Devices which are not intended to become the PTP master should be configured with defaultDS.slaveOnly set to TRUE. All devices containing Senders or Receivers shall have a method or control to allow a user to force the device into a defaultDS.slaveOnly equals TRUE state.

All devices conforming to this standard shall support the SMPTE ST 2059-2 PTP Profile. The parameter ranges specified in ST 2059-2 should be constrained to simultaneously meet the parametric limits of the Media Profile as specified in AES67:2015 for systems which expect to interchange audio streams with AES67 compliant devices.

Note: AES has issued AES-R16-2016, a technical report regarding the compatibility of parameter ranges between the AES67 Media Profile and SMPTE ST 2059-2.

7.3 Media Clocks

The media clock shall be frequency locked to the device internal clock specified in section 7.2. The media clock shall advance at an exact rate as specified for the specific media essence type. When referenced to IEEE 1588-2008, the media clock shall have a value of zero at the SMPTE Epoch as specified in SMPTE ST 2059-1.

The sampling clocks used for reconstruction of digital media signals may be derived from the media clock.

If the device cannot lock to the common reference clock (and is thus operating on the local timebase) sender and receiver shall assume the epoch specified in SMPTE ST 2059-1 and utilize their best available source of current time in the generation and interpretation of media clocks, RTP clocks, and RTP timestamps.

7.4 RTP clocks

The RTP clock shall advance at the same rate as the media clock for each specific media essence type.

Notwithstanding the provisions of RFC3550, RTP streams conforming to this standard shall utilize an RTP clock offset of zero from the media clock. Therefore, when conforming to this standard, each RTP clock is identical to its associated media clock.

Note 1: Receivers designed to maintain compatibility with other RTP implementations will need to comply with the RTP clock offset provisions in those RTP standards, specifically the possibility that the RTP clock could be offset from the media clock.

Note 2: The requirement of a zero offset value in this standard allows fast restoration of signals after sender restarts. Eliminating the random offset provision of IETF RFC 3550 allows the receiver to make use of the signal as soon as the packet stream is restored, without waiting for the systemic propagation of a revised SDP object.

7.5 RTP Timestamps

7.5.1 General Provisions

The RTP timestamps of RTP streams are used to synchronize RTP streams from various senders within an overall production environment. RTP packets contain RTP timestamps as specified for the specific media essence type. The RTP timestamps are samples of the RTP clock.

The RTP timestamps, as specified in IETF RFC 3550 clause 5.1, shall reflect the “sampling instant” of the essence samples contained within the RTP packet, subject to additional clarifications in the sections below.

The RTP timestamps of successive video frames shall advance at regular increments based on the prevailing frame rate, truncating to integer values when necessary. For interlaced video, the RTP timestamps of the first field of successive frames shall advance at regular increments based on the prevailing video frame rate, and the RTP timestamp of the second field shall be offset from the RTP timestamp of the first field by one half of the prevailing frame period, truncating to integer values when necessary. For Progressive segmented Frame (PsF) signals, both segments shall have the same RTP timestamp.

Note: The video RTP timestamp is limited in temporal resolution to the values that its rate can convey. Not all frame periods will have an integer relationship with the rate of the video RTP clock. The frame periods (difference between successive video RTP timestamps) might not be exactly constant - for example 60/1.001 Hz frame periods effectively alternate between increments of 1501 and 1502 ticks of a 90 kHz RTP clock.

The audio RTP timestamps of successive audio packets shall advance at regular increments based on the audio RTP clock rate as specified in section 7.4 of this standard, and the audio packet time as specified in section 7.2 of AES67:2015.

7.5.2 RTP Timestamps Generated by Image and Sound Capture Devices

For progressive scan video essence, the RTP timestamp shall reflect the image capture time of the progressive image whose samples are contained within the RTP packets.

For interlaced video essence, the RTP timestamp of the first field shall reflect the image capture time of the first field whose samples are contained within the RTP packet. The RTP timestamp of the second field shall be as specified in section 7.5.1.

For audio capture devices, the RTP timestamp of audio RTP packets shall reflect the sampling instant of the first sample of the audio signal within the audio RTP packet.

7.5.3 RTP Timestamps Generated by Devices Which Emit Synthetic Essence Including Playback

In the case of video playback from storage devices, or generation of video essence via synthetic means, the RTP timestamp of a frame (or of the first interlaced field within a frame) shall be defined as the time point when the first sample of the frame is presented for transmission on the sender’s transmission interface. This provision is subject to the general provisions of section 7.5.1 including the “regular increments” provisions. The RTP timestamp of the second field if applicable shall be as specified in section 7.5.1.

For the case of audio playback or synthetically generated audio essence, the audio RTP timestamp shall be a sample of the audio RTP clock at the time point when the audio RTP packet is presented for transmission on the Sender's transmission interface.

7.5.4 RTP Timestamps Generated from a Serial Digital Interface (SDI) signal

In the case of encapsulation of the essences contained within an SDI signal, the RTP timestamps shall be determined as follows:

For a progressive video frame, or for the first field of an interlaced frame, the video RTP timestamp shall be a sample of the value of the video RTP clock at the alignment point of the SDI signal, as specified in SMPTE ST 2059-1 or the appropriate media reference standard. If applicable, the video RTP timestamp of the second field shall be as specified in section 7.5.1.

For Progressive segmented Frame (PsF) video signals on SDI, the segment data shall be treated as a single progressive image, and the RTP timestamp of the second segment shall be the same as the first segment.

For audio essence embedded in SDI as specified in SMPTE ST 299-1, ST 299-2, or ST 272, the effective sampling instant of the first audio sample of each audio channel related to a frame of video shall be contemporaneous to the video frame RTP timestamp as determined above, offset by an amount determined from the audio embedding process. The effective sampling instant of subsequent audio samples for each audio channel shall increase monotonically with each sample. The audio RTP timestamp of each audio RTP packet shall be the effective sampling instant of the first audio sample contained within the audio RTP packet.

Note: SMPTE ST 299-1 for HD and 3G signals describes the timing relationship between embedded audio and SDI including the phase offset information. For SD signals, SMPTE ST 272 applies, and some additional skill-of-art is needed to infer the phase information.

8 Session Description Protocol (SDP)

8.1 General

Devices which contain one or more senders shall construct one SDP object per RTP stream as specified in IETF RFC 4566. These SDP objects shall be made available through the management interface of the device.

All stream descriptions shall have a media-level `mediaclk` attribute as per IETF RFC 7273 section 5.2. The `direct` reference should be used.

An example SDP is shown in Annex B.

8.2 Reference Clock

All stream descriptions shall have a media-level `ts-refclk` attribute as specified in IETF RFC 7273 section 4. Devices which are referenced to IEEE1588-2008 shall use the `ts-refclk:ptp` form, signaling either the grandmaster clockIdentity and domain, or signaling that the PTP is traceable (if it is). Devices which are not referenced to IEEE1588-2008 shall signal using the extended form shown below, indicating the MAC address of the sender using the token `localmac`. Receivers may assume that different streams which signal the same value for `localmac` are synchronized together.

```
a=ts-refclk:localmac=<Ethernet MAC address of sender>
```

The following examples show the PTP form and the localmac form:

```
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:37

a=ts-refclk:ptp=traceable

a=ts-refclk:localmac=7C-E9-D3-1B-9A-AF
```

Note: The first PTP example above signals that the sender is using a PTP clock conforming to IEEE 1588-2008, the clockIdentity of the grandmaster is `39-A7-94-FF-FE-07-CB-D0`, and the domain number is 37. The second PTP example indicates that the PTP source is traceable as specified in IETF RFC 7273 section 4.7. The case of the local timebase signals an Ethernet MAC address of the Sender in EUI-48-format. ClockIdentity is expressed in EUI-64 format, which is a sequence of hexadecimal values, while the PTP domain number is expressed as a decimal number.

Additional requirements of the SDP object may be specified in the media essence-specific requirements documents.

8.3 SDP for Duplicate RTP Streams

Duplicate RTP streams meeting the requirements of SMPTE ST 2022-7 may be used for redundant transmission to achieve higher system availability. Senders which transmit these duplicated RTP streams, using the mechanisms of separate source addresses (IETF RFC 7104, Section 4.1) or separate destination addresses (IETF RFC 7104, section 4.2) shall signal the RTP duplication using the session level `group` attribute of IETF RFC 5888 and the duplication grouping `DUP` semantics of IETF RFC 7104.

Redundant streams shall not use both identical source addresses and identical destination addresses at the same time.

Note: SMPTE ST 2022-7 allows the above methods of specifying duplicate RTP streams, but also allows for RTP streams with duplicated source and destination addresses (on separate physical networks); such a mechanism cannot be represented with SDP, and therefore the use of duplicate source and destination addresses is not supported by this Standard.

8.4 UDP Datagram Size

As specified in sections 6.3 and 6.4, Senders may choose to operate within the standard UDP size limit or the Extended UDP Size Limit.

Senders operating with UDP Sizes which exceed the Standard UDP Size Limit shall include a Media Type Parameter `MAXUDP` with a decimal value indicating the largest UDP Datagram Size (in octets) that might be present in the stream.

If the `MAXUDP` parameter is not present, Receivers shall assume the Standard UDP Size Limit specified in section 6.3.

Note: the media type parameter refers to the maximum UDP datagram size – the Ethernet MTU size also includes the IP header in addition to the UDP packet itself.

Annex A Datagram Size Limits (Informative)

This family of documents defines RTP essence formats, and mandates their transport over UDP and IP. The most common method for transporting IP datagrams within facilities is Ethernet, as specified in the IEEE 802.3 family of standards.

The length of an IP datagram is limited only by the representable values within the IPv4 or IPv6 header fields. However, in practice the underlying transport mechanism imposes more significant limits on the datagram sizes.

For technical reasons relating to certain older variants of the Ethernet system, the payload of the Ethernet frame is officially limited to a maximum of 1500 octets within the IEEE 802.3 family of standards. When IP datagrams are transported over Ethernet, this limits the size of the IP datagram, including the IP, UDP, and RTP headers and data, to a total of 1500 octets.

The IPv4 standard header is 20 octets long, while the standard IPv6 header is 40 octets long. In order to accommodate either standard, and to simplify the case of in-network mappings between the two IP standards, the larger of the two values is assumed. Thus the “standard” UDP datagram size limit specified in section 6.3 of this standard is 1500-40 = 1460 octets.

While not strictly allowed within the IEEE 802.3 family of standards, support for so-called “Jumbo” Ethernet frames have been a common feature of Ethernet networking equipment for many years, with an industry consensus value of 9000 octets as the maximum payload length within the Ethernet “Jumbo” frame.

The 8960 octet limit in section 6.4 is based on a 9000 octet Ethernet “Jumbo” frame payload size, and accommodates IPv4 or IPv6 headers.

Ethernet frames have a minimum payload size of 46 octets, and in the rare case of an IP datagram smaller than 46 octets the payload is zero-padded in the process of mapping IP datagrams into Ethernet frames. There is no need to pad up the IP datagram to 46 octets in the RTP protocol.

When contemplating the potential throughput of network interfaces, it is important to remember that the nominal bit-rates of the interfaces, such as “Ten Gigabit Ethernet” or “Twenty-Five Gigabit Ethernet” refer to the bit rate of the Ethernet Frames.

Each Ethernet frame contains the following “overhead” which must be accounted for in any calculation of the IP throughput of an Ethernet connection:

18 octets	Layer-2 Ethernet frame header and FCS (without an 802.1Q VLAN tag)
4 octets	Optional 802.1Q VLAN tag if present
20 octets	Fixed preamble, start-of-frame delimiter, and minimum inter-packet gap

As an example, the throughput of Ethernet payload data for maximum “standard” sized UDP datagrams specified in section 6.3 (assuming IPv4 and an 801.2Q tag) on a 10 Gigabit Ethernet link is:

Ethernet-Frames/sec = 10,000,000,000 / (8 * (1460 + 20 + 18 + 4 + 20)) = 821,287.7792

IP bits/sec = Ethernet-Frames/sec * IP bits/frame = 821,287.7792 * (1480 * 8) = 9,724,047,306

The above calculation assumes minimum inter-packet gap, an 801.2Q tag, and assumes that all of the IPv4 datagrams are 1480 octets (1460 payload + 20 IPv4 header).

Annex B SDP Example (Informative)

The following example SDP for a video stream including the fields described in this standard.

```
v=0
o=- 123456 11 IN IP4 192.168.100.2
s=Example of a SMPTE ST2110-20 signal
i=this example is for 720p video at 59.94
t=0 0
a=recvonly
a=group:DUP primary secondary
m=video 50000 RTP/AVP 112
c=IN IP4 239.100.9.10/32
a=source-filter:incl IN IP4 239.100.9.10 192.168.100.2
a=rtpmap:112 raw/90000
a=fmtp:112 sampling=YCbCr-4:2:2; width=1280; height=720;
exactframerate=60000/1001; depth=10; TCS=SDR; colorimetry=BT709;
PM=2110GPM; SSN=ST2110-20:2017;
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:37
a=mediaclk:direct=0
a=mid:primary
m=video 50020 RTP/AVP 112
c=IN IP4 239.101.9.10/32
a=source-filter:incl IN IP4 239.101.9.10 192.168.101.2
a=rtpmap:112 raw/90000
a=fmtp:112 sampling=YCbCr-4:2:2; width=1280; height=720;
exactframerate=60000/1001; depth=10; TCS=SDR; colorimetry=BT709;
PM=2110GPM; SSN=ST2110-20:2017;
a=ts-refclk:ptp=IEEE1588-2008:39-A7-94-FF-FE-07-CB-D0:37
a=mediaclk:direct=0
a=mid:secondary
```

This SDP reflects a session ID of 123456 and a session version of 11. The session name is indicated in the s= clause, with additional session information in the i= clause. The source addresses of the primary and secondary streams are indicated in the a=source-filter clauses for each stream.

The t=0 0 clause indicates that this session is permanent (has no begin or end time).

The a=group:DUP clause is as specified in section 8.3 of this standard, indicating two RTP streams are sent, tagged primary and secondary inside this SDP.

The first m= section describes the primary RTP stream, which is transmitted via IPv4 on group 239.100.9.10 and UDP port 50000 from source address 192.168.100.2. Each m= line signals the start of a "media-specific section" within the SDP.

The a=fmtp clause contains a number of media-specific parameters specified in the media-specific document.

The a=ts-refclk clause is as specified in section 8.2 of this standard. The a=mediaclk:direct=0 clause signals that the media clock is directly referenced to the clock in the ts-refclk clause, and the offset of 0 is as mandated in section 7.4 of this standard. The a=mid:primary section tags this media section as the "primary" stream within the grouping semantics.

The second media section (starting with the second m= line) documents the same information for the secondary stream. This secondary stream is on multicast group 239.101.9.10, UDP port 50020, and originates from source address 192.168.101.2.

Note that when utilizing SMPTE ST 2022-7 hitless reconstruction, the RTP headers and RTP payloads need to be identical between the two RTP streams, because packets can be selected from either stream at any time.

Bibliography (Informative)

AES-R16-2016, AES Standards Report - PTP parameters for AES67 and SMPTE ST 2059-2 interoperability

Internet Engineering Task Force (IETF) RFC 6410 Reducing the Standards Track to Two Maturity Levels [online, viewed 2017-08-10] Available at https://www.ietf.org/rfc/rfc6410.txt

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

Video Services Forum (VSF) TR-03 Transport of Uncompressed Elementary Stream Media over IP [online, viewed 2017-08-10] Available at http://www.videoservicesforum.org/download/technical_recommendations/VSF_TR-03_2015-11-12.pdf

Video Services Forum (VSF) TR-04 Utilization of ST-2022-6 Media Flows within a VSF TR-03 Environment [online, viewed 2017-08-10] Available at http://www.videoservicesforum.org/download/technical_recommendations/VSF_TR-04_2015-11-12.pdf