

SMPTE STANDARD

Seamless Protection Switching of
RTP Datagrams



Table of Contents		Page
1	Scope.....	3
2	Conformance Notation	3
3	Normative References	3
4	Terms and Definitions	4
5	Symbols and Abbreviated Terms	5
6	Creation of Streams for Seamless Reconstruction.....	5
7	Reception of Streams for Seamless Reconstruction	6
Annex A Stream Synchronization (Informative).....		8

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.

SMPTE ST 2022-7 was prepared by Technology Committee 32NF on Network/Facilities Architecture.

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.

This document is a standard for seamless reconstruction of a stream of RTP datagrams based on the transmission of multiple streams of identical content over potentially diverse paths (see Figure 1). The primary reason for this standard is to facilitate interoperability between equipment used to transport media signals composed of RTP datagrams intended for such reconstruction.

Cost effective error free transport of media essence is a constant goal. As costs of redundant network paths decrease, sending redundant streams becomes a more viable option. This standard defines requirements on redundant streams such that a receiver could switch between them on a datagram by datagram basis without impact to the content or the stream.




Figure 1 – Conceptual Diagram of Multiple Paths

1 Scope

This standard defines requirements for multiple redundant streams of RTP packets to allow for the creation of a single reconstructed output stream through seamless protection switching at the RTP datagram level.

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; followed by 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 standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

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>

4 Terms and Definitions

4.1

class HBR stream

a flow of datagrams with a payload bit rate greater than or equal to 270 Mbit/s

4.2

class SBR stream

a flow of RTP datagrams with a payload bit rate of less than 270 Mbit/s

4.3

datagram copies

set of redundant RTP datagrams, perhaps transmitted to a different IP destination or port, but having identical contents in the RTP header and RTP payload

4.4

high-skew

potentially large jitter and very large network transit delay variation due to alternative routes, as might be found on a link between two facilities (an inter-facility link)

Note: The combination of switching-induced jitter and differential path delay in a high-skew link can be greater than 50ms.

4.5

input stream

stream of RTP datagrams in accordance with RFC 3550, which might represent essence or FEC

Note: Some examples of RTP stream specifications are: SMPTE ST 2022-1, SMPTE ST 2022-2, SMPTE ST 2022-3, SMPTE ST 2022-4, SMPTE ST 2022-5, SMPTE ST 2022-6, SMPTE ST 2110-20, SMPTE ST 2110-30, SMPTE ST 2110-40 and AES67.

4.6

low-skew

Jitter and delay due to a modest number of router/switch hops and a reasonably short distance traveled, where the combination of switching-induced jitter and differential path delay is by design less than 10 ms, as might be found on a link within a facility (an intra-facility link)

4.7

moderate-skew

Jitter and delay due to a moderate number of router/switch hops regardless of path length, where the combination of switching-induced jitter and differential path delay in such a link is by design less than 50 ms, as might be found on a short-haul link between two facilities within a region or a small country

4.8

output stream

reconstructed datagram stream that results from processing the multiple copies of the input stream

Note: This output stream might occur at an intermediate stage within a device or system.

4.9

RTP datagram

RTP packet as defined in RFC 3550

4.10**seamless reconstruction**

successful creation of a “reconstructed” output stream based on receipt of two or more (potentially network impaired) input streams meeting the requirements specified in this standard such that the RTP header and payload of the reconstructed stream are identical to the input stream(s)

4.11**ultra low-skew**

a small amount of jitter and delay, as might be found on a link over a small number of LAN switch hops within a facility, where the combination of switching-induced jitter and differential path delay in such a link is by design less than 150 μ s

5 Symbols and Abbreviated Terms

ASI – Asynchronous Serial Interface

DWDM – Dense Wave Division Multiplexing

FEC – Forward Error Correction

HBR – High Bit Rate

MPLS – Multiprotocol Label Switching

RTP – Real Time Protocol (see IETF RFC 3550)

SBR – Slow Bit Rate

VSID – Video Source ID

VPN – Virtual Private Network

6 Creation of Streams for Seamless Reconstruction

The transmitter shall transmit at least two streams, each containing copies of each RTP datagram. The RTP header and the RTP payload shall be identical for each datagram copy. The seamless reconstruction method described herein makes no assumptions about the Ethernet or IP headers of the source streams.

Note: In copies of a SMPTE ST 2022-6 RTP datagram, the VSID field, as part of the RTP payload, will be identical across the datagram copies.

7 Reception of Streams for Seamless Reconstruction

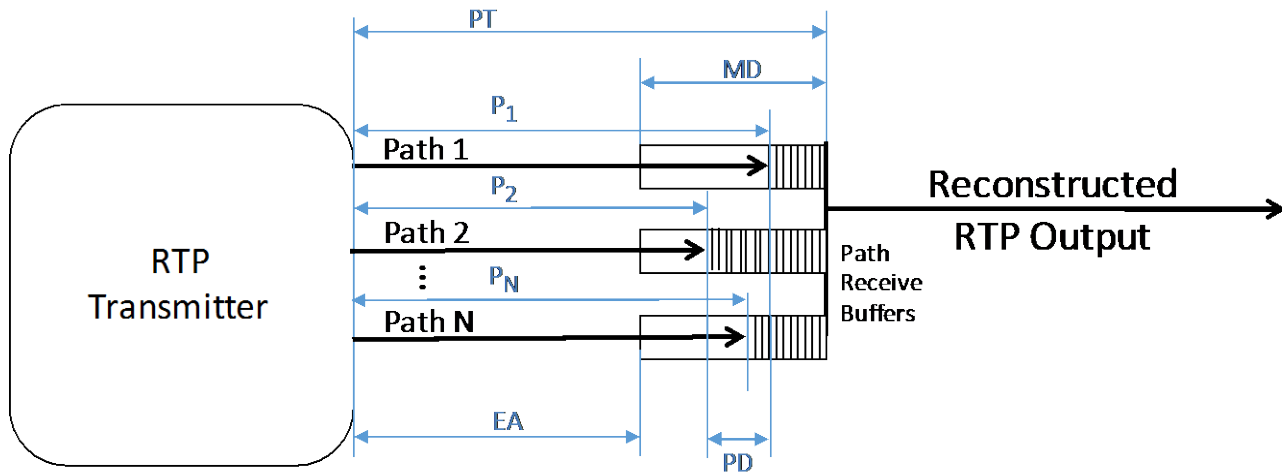


Figure 2 – Key Timing Points for Seamless Reconstruction

The key timing points for seamless reconstruction of RTP streams are illustrated in Figure 2.

P_n (where $n \in 1 \dots N$) is the instantaneous latency from transmission to reception of datagrams on path number n of N paths, inclusive of any network jitter.

PT is the latency from transmission to the final reconstructed output. It is also the latest time that a packet could arrive at the receiver to be part of the reconstructed output.

EA is the earliest time that a packet could arrive at the receiver to ensure seamless reconstruction.

MD is the maximum differential and is the difference of PT and EA.

$$MD = (PT - EA)$$

PD is the instantaneous path differential, and is equal to the maximum of the absolute value of the differences between instantaneous latencies.

$$PD = \max_{i,j \in 1 \dots N} |P_i - P_j|$$

The latency PT is established by the receiver at startup. After startup, P_n might change due to changes in network routing and latency, but only to the extent that PD remains within the bounds specified below.

A compliant receiver for Classes A, B, C or D shall support Seamless Reconstruction from streams that maintain a PD value of less than or equal to those specified in the Table 1 for its classification.

Receiver Classification	Use Case (example)	Class SBR Streams	Class HBR Streams
Class A: Low-Skew	Intra-Facility Links	PD \leq 10ms	PD \leq 10ms
Class B: Moderate-Skew	Short-Haul Links	PD \leq 50ms	PD \leq 50ms
Class C: High-Skew	Long-Haul or special circumstance Links	PD \leq 450ms	PD \leq 150ms
Class D: Ultra Low-Skew	Physical Layer LAN Redundancy	PD \leq 150 μ sec	PD \leq 150 μ sec

Table 1 – Receiver Classifications

Note: The Class SBR and Class HBR defined in this document were loosely modeled on class 6 and class 7 as described in Recommendation ITU-T Y.1541. However, Y.1541 describes their use for Access Distribution infrastructure, whereas the application of this standard is for contribution links.

Only paths with a latency P_i that is greater than EA and less than PT can be used to source packets to potentially be used in the output streams.

As long as there are at least two paths with latencies P_i and P_j that are greater than EA and less than PT then seamless reconstruction is able to recover from a packet loss in one of those two paths and create a successful output stream.

A path may include a mechanism of temporal offset redundancy.

Note: The exact the method of reconstruction is left to the implementer.

Annex A describes a mechanism to correlate stream times. Annex B provides suggestions on management and monitoring. Annex C describes use cases for seamless protection switching.

Annex A Stream Synchronization (Informative)

By the method envisioned herein, the RTP sequence numbers can be used to correlate multiple packet streams. Since the sequence numbers roll over relatively frequently in Class HBR streams, in high bandwidth applications the RTP timestamps can be used in addition to the sequence numbers to ensure proper matching of the incoming streams, when they are present.

In the case of a 1080p/60Hz video stream, the HD-SDI bit-rate is 2.970 Gbits/sec. Each packet of ST 2022-6 carrying this video stream conveys 1376 octets (or 11008 bits) of the signal, generating 270 packets (approximately) each millisecond.

In the case of a class C receiver and an HBR stream, when the PD value is zero at startup, the receiver must plan to accommodate future path delay excursions of up to 150ms in either direction on either stream; establishment of a 300ms window ($PT = P_1 + 150\text{ms}$, $EA = PT - 300\text{ms}$) fulfills this requirement. During this 300ms window for class HBR streams, there would be approximately 81,000 packets in the window in the case of 1080p60Hz HD-SDI. There are only 16 bits used to convey the RTP sequence number in ST 2022-6, rolling over every 65,536 packets. Therefore it is impossible to tell by sequence number alone the relative values of P_n while maintaining the buffers implied for a class C receiver.

The RTP header timestamp values in ST 2022-6 are a required field. They are specified in terms of a 27MHz free-running clock, and 32 bits are used, leading to an RTP timestamp roll-over period of approximately 159.07 seconds ($2^{32}/27,000,000$). In the perfectly smooth case, there are about 100 ticks of the 27MHz RTP timestamp clock for every 1080p60Hz ST 2022-6 datagram.

The multiple redundant streams, by normative provision of this document, have identical RTP timestamps. These timestamps by themselves can indicate the relative skew (time) between P_n 's quite accurately, and the timestamp values themselves could be used to match up the packets between the two streams. Alternatively, the monotonically increasing sequence numbers can be augmented with a function of the difference between the timestamps to correctly correlate the arriving packets.

In the general case, the receiver does not know the absolute values of P_n 's. It only knows the difference between them. Individual implementations can take different approaches to establishing a startup delay PT, however, barring any a-priori knowledge of the network, a class C receiver, to account for worse case, could start up as follows:

For Class C receivers (high skew) processing SBR streams — Set PT (the latency from transmission to the final reconstructed output) to 450 ms after the earlier of the two streams, and set EA to 900 ms less than PT.

For Class C receivers (high skew) processing HBR streams — Set the PT to 150 ms after the earlier of the two streams, and set EA to 300 ms less than PT.

Annex B Management and Monitoring (Informative)

Management and monitoring are important capabilities of any video network. A receiver could allow the tracking of received and lost datagrams per stream.

A receiver could expose the matching buffer setting to the users and/or provisioning/monitoring software so that it can be observed and adjusted to compensate for potential delay problems between the streams, or to allow configuring the receiver for specific path delay situations about which the operator might have a-priori information.

A receiver might expose indicators and publish notifications to provisioning/monitoring software as the receiver transitions between a protected and non-protected state so that users are made aware of the protection state at all times.

Annex C Use Cases (Informative)

Potential network points to implement seamless protection switching

- Video encoder/decoder with network port
- ASI network receiver/transmitter
- SMPTE ST 2022 network receiver/transmitter
- AES67 network receiver/transmitter
- SMPTE ST 2110 network receiver/transmitter
- Network router or switch
- Network interface card
- Dedicated network attached seamless processor

Network Use Cases

There are many types of network situations that can be augmented with seamless protection switching. The following cases define some common network types. These network types could be used for one or both of the streams.

Satellite transmission can add substantial latency to the stream transmission. The minimum requirements of this standard might not support combinations of satellite and fiber links with instantaneous path differentials higher than that of Class C receivers in this document. A device with capabilities to support instantaneous path differentials higher than that of Class C could support this use case.

When using MPLS or Internet connectivity with or without a VPN, network latency can change over time as paths reroute due to failures and traffic grooming (the latter being the process of grouping many small telecommunications flows into larger units, which can be processed as single entities, such as in a single wavelength). The required class of seamless protection switching can be chosen to accommodate this.

The main variants of protected private line are protected SONET and protected DWDM wavelengths. Protected networks can offer a higher level of reliability, but also can complicate a seamless network. If the two protected circuits share a segment of the same physical path there can be a situation where the stream is not actually protected from failures along that segment.

Unprotected SONET and unprotected DWDM wavelengths are also available. Unprotected circuits are generally much less expensive than a protected circuit. They offer complete visibility to both paths allowing the user to know the status of both paths of the network at all times.

Another potential use case differs in that it is not a diverse path network type, but is instead formed by offsetting in time the primary stream from the secondary stream; this arrangement offers protection from network errors smaller than the offset. This method allows some form of protection while not requiring two network paths.

Bibliography (Informative)

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

Recommendation ITU-T Y.1541, Network Performance Objectives for IP-Based Services

SMPTE ST 2022-1:2007, Forward Error Correction for Real-Time Video/Audio Transport Over IP Networks

SMPTE ST 2022-2:2007, Unidirectional Transport of Constant Bit Rate MPEG-2 Transport Streams on IP Networks

SMPTE ST 2022-3:2010, Unidirectional Transport of Variable Bit Rate MPEG-2 Transport Streams on IP Networks

SMPTE ST 2022-4:2011, Unidirectional Transport of Non-Piecewise Constant Variable Bit Rate MPEG-2 Streams on IP Networks

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

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

SMPTE ST 2110-10:2017, Professional Media Over Managed IP Networks, System Timing and Definitions

SMPTE ST 2110-20:2017, Professional Media Over Managed IP Networks, Uncompressed Active Video

SMPTE ST 2110-30:2017, Professional Media Over Managed IP Networks, PCM Digital Audio

SMPTE ST 2110-40:2018, Professional Media Over Managed IP Networks, SMPTE ST 291-1 Ancillary Data