



ATSC

ADVANCED TELEVISION
SYSTEMS COMMITTEE

ATSC Candidate Standard: Revision of A/324:2018, Scheduler / Studio to Transmitter Link

Doc. S32-266r56
7 January 2020

Advanced Television Systems Committee
1776 K Street, N.W.
Washington, D.C. 20006
202-872-9160

The Advanced Television Systems Committee, Inc., is an international, non-profit organization developing voluntary standards and recommended practices for digital television. ATSC member organizations represent the broadcast, broadcast equipment, motion picture, consumer electronics, computer, cable, satellite, and semiconductor industries. ATSC also develops digital television implementation strategies and supports educational activities on ATSC standards. ATSC was formed in 1983 by the member organizations of the Joint Committee on Inter-society Coordination (JCIC): the Electronic Industries Association (EIA), the Institute of Electrical and Electronic Engineers (IEEE), the National Association of Broadcasters (NAB), the National Cable Telecommunications Association (NCTA), and the Society of Motion Picture and Television Engineers (SMPTE). For more information visit www.atsc.org.

Note: The user's attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. By publication of this standard, no position is taken with respect to the validity of this claim or of any patent rights in connection therewith. One or more patent holders have, however, filed a statement regarding the terms on which such patent holder(s) may be willing to grant a license under these rights to individuals or entities desiring to obtain such a license. Details may be obtained from the ATSC Secretary and the patent holder.

Implementers with feedback, comments, or potential bug reports relating to this document may contact ATSC at <https://www.atsc.org/feedback/>.

Revision History

Version	Date
A/324:2018 approved	5 January 2018
CS Revision of A/324:2018 approved	23 May 2019
CS Update approved	20 September 2019
2 nd CS Update approved	7 January 2020

Table of Contents

1	SCOPE	1
1.1	Introduction and Background	1
1.2	Organization	2
2	REFERENCES	2
2.1	Normative References	2
3	DEFINITION OF TERMS	4
3.1	Compliance Notation	4
3.2	Treatment of Syntactic Elements	4
3.2.1	Reserved Elements	4
3.3	Acronyms and Abbreviations	5
3.4	Terms	8
3.5	Extensibility	12
3.6	XML Schema and Namespace	12
4	SYSTEM OVERVIEW	13
4.1	Features	13
4.2	System Architecture	16
4.2.1	System Manager	17
4.2.2	Broadcast Gateway	18
4.2.3	Studio to Transmitter(s) Dataflow	19
4.2.4	STL Operation	19
4.2.5	SFN Operation	20
4.3	Central Concepts: CTP	20
4.4	Central Concepts: DSTP	20
4.5	Central Concepts: ALPTP	21
4.6	Central Concepts: STLTP	21
4.6.1	STLTP Payload Data	21
4.6.2	Content Security System	22
4.7	System Time Domains	22
4.8	System Manager Configuration Interface	22
4.9	Real-Time Control Interface	23
4.10	Transmitter Requirements Overview	23
4.10.1	Carrier and Timing Offset for Co-Channel Interference Mitigation	23
5	SCHEDULER DESCRIPTION AND NORMATIVE REQUIREMENTS	23
5.1	Relationship of Broadcast Gateway and Scheduler to the System	24
5.2	Scheduler Functionality	24
5.3	Preamble Construction	25
5.4	Scheduler Management Protocol	26
5.5	Data Source Control Protocol (DSCP)	27
6	COMMON TUNNELING PROTOCOL (CTP).....	28
6.1	Overview	29
6.2	RTP/UDP/IP Multicast Considerations	30
6.2.1	Multicast Addressing	31
6.3	Common Tunneling Protocol Design	31

6.3.1	CTP RTP Encapsulation Example	34
6.3.2	Example SMPTE ST 2022-1 ECC Encoding Process	37
7	DATA SOURCE TRANSPORT PROTOCOL	38
7.1	Overview	38
7.1.1	DSTP Mapping Configuration Description	39
7.2	DSTP Design	42
7.2.1	DSTP Tunneled Packet Information Header Definition	42
7.2.2	Emergency Alert Wakeup RTP Controls	44
8	ALP TRANSPORT PROTOCOL (ALPTP)	45
8.1	Overview	46
8.2	ALPTP Design	46
8.2.1	ALPTP Tunneled Packet Information Header Definition	46
9	STL TRANSPORT PROTOCOL.....	48
9.1.1	Address Assignments	48
9.1.2	Port Assignments	48
9.1.3	Majority Logic	49
9.2	Preamble Data Generator	49
9.2.1	Preamble Data Stream Protocol	50
9.3	Timing and Management Generator	52
9.3.1	Timing and Management Data Stream Protocol	52
9.3.2	Bootstrap Emission Timing and Frame Identification	57
9.3.3	PLP Data Stream	57
9.3.4	Baseband Packet Data Stream Protocol	58
9.4	STL Security Protocol	59
9.4.1	Tunneled Packet Signing Procedures	59
9.4.2	Authentication Key Generation and Utilization	63
9.4.3	Authentication Key Delivery	64
9.4.4	Public/Private Key Generation & Distribution	67
9.4.5	Security Data Stream Protocol	70
9.4.6	Security System Implementation	75
9.5	Studio to Transmitter Link (STL) Transport Protocol (STLTP)	80
9.5.1	STL Transport Protocol Design	80
9.5.2	RTP Encapsulation Example	81
9.6	Channel Bonding in Broadcast Gateways	83
10	TRANSMITTER OPERATION NORMATIVE REQUIREMENTS.....	83
10.1	Timing Manager	84
10.2	Preamble Parser	85
10.3	Other Transmitter Requirements	85
10.3.1	Frequency Accuracy	86
10.3.2	Transmitter Timing Offsets	86
10.3.3	Center-Frequency and Network-Timing Offsets for Co-Channel Interference Mitigation	86
ANNEX A	PHYSICAL LAYER CONTROL	89
A.1	Physical Layer Resources	89
A.1.1	Bootstrap Signaling	89
A.1.2	L1-Basic Signaling	89

A.1.3	L1-Detail Signaling	91
ANNEX B	NETWORK CONFIGURATION EXAMPLES	95
B.1	Example Studio Network Topologies	95
ANNEX C	SCHEDULER FUNCTIONAL DESCRIPTION	98
C.1	Overview	98
C.2	Scheduler Operation	98
C.3	Key Concepts of Scheduler Delivery Metadata	99
C.4	Handling Boundary Conditions	103
C.5	Delivery Order Within and Across Multiple Sessions	104
C.6	Timelines and Deadlines	105
C.7	Concept and Practice of Analyzed Media Duration	105
C.8	Sequence of Required Data and Media Events for Acquisition	108
ANNEX D	UNICAST FOR OUTER TUNNELING PACKETS	110
ANNEX E	BONDED-CHANNEL DISTRIBUTION	111

Index of Figures and Tables

Figure 1.1 In-scope interfaces description.	1
Figure 4.1 High-level overview of system configuration.	14
Figure 4.2 System architecture.	17
Figure 6.1 CTP Tunnel Encapsulation Process	29
Figure 6.2 Common Tunneled Packet Stream	34
Figure 6.3 Data Source Tunneled to Tunnel Packet Relationship	36
Figure 6.4 Example ECC encoding process diagram.	37
Figure 7.1 Multiple Data Source Packet Tunnels Concept	39
Figure 9.1 Example use of Majority Logic for Preamble (denoted as L1) and T&M: (when 3 repetitions of each are applied).	49
Figure 9.2 GCM AEAD and GMAC Authentication of AAD	60
Figure 9.3 STLTP security System Authentication Key generation & distribution: (<i>a</i> , top) Broadcast Gateway Processing Procedure; (<i>b</i> , bottom) Exciter Processing Procedure.	65
Figure 9.4 STLTP Security System Public/Private Key locations after system initialization	69
Figure 9.5 STL Transmission diagram.	80
Figure 9.6 Tunneled Packet packing details.	82
Figure 9.7 Channel Bonding operation modes in a Broadcast Gateway: (a) Sending two outer Tunnel Data Streams, and (b) Sending one outer Tunnel Data Stream that carries two groups of inner Tunneled Packet Streams.	83
Figure 10.1 Example use of carrier and timing offsets for neighboring co-channel stations.	88
Figure B.1 Simple ALP encapsulation.	95
Figure B.2 Multiple PLP example.	96
Figure B.3 Fully redundant routing example.	97
Figure C.1 Cascade of real-time functions involved with Scheduler.	98
Figure C.2 Example depiction of a high-level Scheduler process flow.	99
Figure C.3 Illustration of Earliest and Latest Time with a block interleaver.	100
Figure C.4 Illustration of Earliest and Latest Time options relative to Media Segment play.	101
Figure C.5 Example of MDE and associated Earliest and Latest Times.	102
Figure C.6 Relationship of an MDE in LCT packets to IP encapsulation.	103
Figure C.7 Media is sent at early boundary; e.g., Period start.	103
Figure C.8 Example of ordered delivery across multiple sessions.	104
Figure C.9 Explicit delivery order on PLP with multiple sessions.	105
Figure C.10 Analyzed Media Duration.	107
Figure C.11 Order of events to optimize linear service acquisition w/wo an Application.	109
Figure D.1 RTP/UDP/IP Unicast for outer “Tunneling” packets using STLTP redundancy.	110
Figure E.1 Channel Bond over PLP_ID 0 between Channels A and B.	111
Figure E.2 Channel bonding with two separate STLTPs.	111
Figure E.3 Channel bonding with combined STLTP.	112
Figure E.4 Example of network topology for Plain Channel Bonding with co-located sites and separate STLTP Tunnel Data Streams.	112
Figure E.5 Example of network topology for SNR Averaging with co-located sites and combined STLTP Tunnel Data Stream.	113
Table 5.1 Preamble Parameters and Their Sources	25
Table 5.2 Real Time Control Field Definitions	28

Table 6.1 RTP Header Field Definitions for Common Transport Protocol	31
Table 6.2 Code Values for payload_type	32
Table 6.3 Timestamp Field Definitions for DSTP and ALPTP	33
Table 7.1 Data Source Mapping XML Format	40
Table 7.2 DSTP Tunneled Packet Information Header	42
Table 7.3 type Field Available Values	43
Table 7.4 RTP payload_type Wakeup Control Field Definition	44
Table 7.5 Example Wakeup Bit Controls	45
Table 8.1 ALPTP Tunneled Packet Information Header	47
Table 9.1 Preamble Payload	50
Table 9.2 RTP Header Timestamp Field Definitions	51
Table 9.3 Timing and Management Stream Packet Payload	53
Table 9.4 Initialization Value (IV) Organization	61
Table 9.5 RTP Header Extension	62
Table 9.6 Security Data Stream Packet Payload	71

ATSC Candidate Standard: Revision of A/324:2018, Scheduler / Studio to Transmitter Link

1 SCOPE

This standard specifies the protocol on the Studio-to-Transmitter Link (STL) from studio side infrastructure to a Single Frequency Network (SFN) of Transmitters. It defines delivery protocols for ALP Transport. The document also defines possible interfaces among the studio infrastructure, for example the interconnection of the ATSC 3.0 Link Layer Protocol (ALP) and a Broadcast Gateway. This document specifies certain constraints on the scheduling of content and signaling on the Physical Layer. The described scheduling process enables Preamble generation and emission time management. It specifies certain aspects of Transmitter behavior and certain parameters of Transmitter operation.

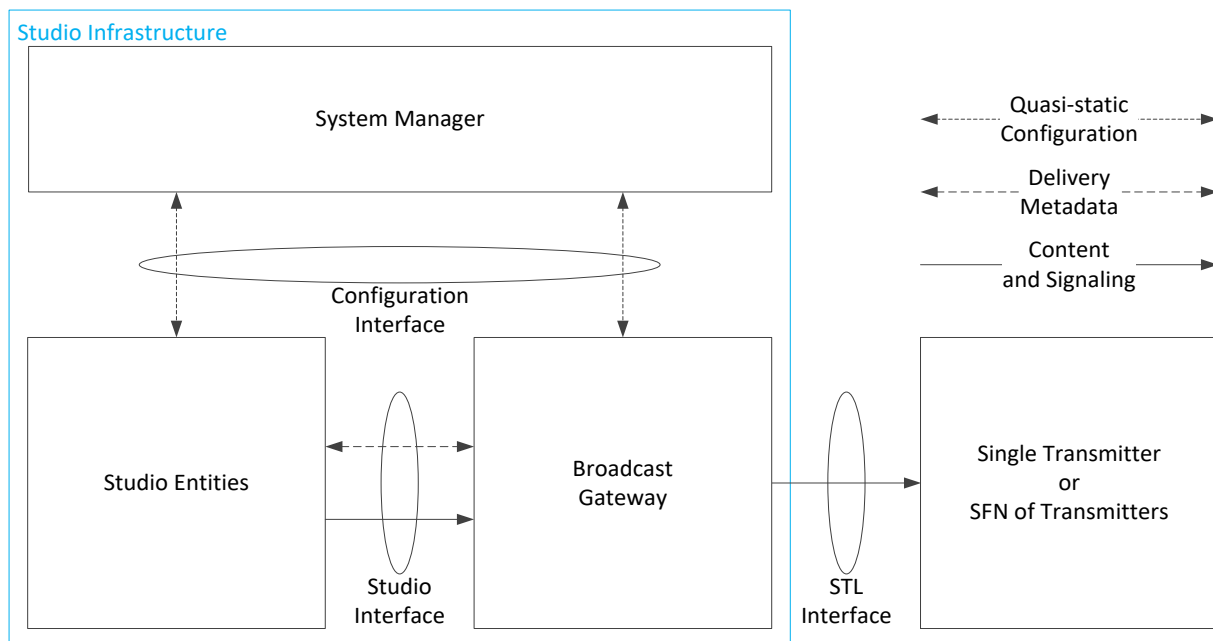


Figure 1.1 In-scope interfaces description.

Figure 1.1 depicts an example configuration and connection of these entities. This document provides no specification of broadband delivery of ATSC 3.0 media content.

1.1 Introduction and Background

The ATSC 3.0 system comprises a number of layers that must be connected to one another to construct a complete implementation. Two of the layers that must be interconnected are the transport layer and the Physical Layer. In addition, the Physical Layer is designed to be implemented partially at the studio or Data Source and partially at one or more Transmitters. To enable the necessary interoperability of the layers and system segments, appropriate protocols are necessary so that equipment from multiple suppliers can be assembled into a working system. This document defines four protocols, the ATSC Link-layer Protocol Transport Protocol (ALPTP), the Studio-to-Transmitter Link Transport Protocol (STLTP), the Data Source Transport Protocol (DSTP) and the Data Source Control Protocol (DSCP), for carriage of data through specific

portions of the system, as well as a number of operational characteristics of the STL and Transmitter(s). Also defined are a Scheduler to manage operation of the Physical Layer subsystems and two protocols used by the Scheduler (1) to receive high-level configuration instructions from a System Manager and (2) to provide real-time bit-rate control information to Data Sources sending content through the transport layer for emission by the Physical Layer.

1.2 Organization

This document is organized as follows:

- Section 1 – Outlines the scope of this document and provides a general introduction
- Section 2 – Lists references and applicable documents
- Section 3 – Provides definitions of terms, acronyms, and abbreviations for this document
- Section 4 – Provides a system overview
- Section 5 – Defines the Scheduler of Physical Layer resources
- Section 6 – Defines the Common Tunneling Protocol (CTP)
- Section 7 – Defines the Data Source Transport Protocol (DSTP)
- Section 8 – Defines the ALP Transport Protocol (ALPTP)
- Section 9 – Defines the Studio to Transmitter Link Transport Protocol (STLTP)
- Section 10 – Describes Transmitter operation
- Annex A – Describes the parameters used for Physical Layer control
- Annex B – Provides Network Configuration examples
- Annex C – Provides the Scheduler Functional Description
- Annex D – Describes Unicast for Outer Tunneling Packets
- Annex E – Describes Bonded-Channel distribution considerations

2 REFERENCES

All referenced documents are subject to revision. Users of this Standard are cautioned that newer editions might or might not be compatible.

2.1 Normative References

The following documents, in whole or in part, as referenced in this document, contain specific provisions that are to be followed strictly in order to implement a provision of this Standard.

- [1] IEEE: “Use of the International Systems of Units (SI): The Modern Metric System,” Doc. SI 10, Institute of Electrical and Electronics Engineers, New York, N.Y.
- [2] ATSC: “ATSC Standard: System Discovery and Signaling,” Doc. A/321:2016, Advanced Television Systems Committee, Washington, D.C., 23 March 2016.
- [3] ATSC: “ATSC Standard: Physical Layer Protocol,” Doc. A/322:2018, Advanced Television Systems Committee, Washington, D.C., 26 December 2018.
- [4] ATSC: “ATSC Candidate Standard Revision: Signaling, Delivery, Synchronization, and Error Protection,” Doc. S33-1-951r14, Advanced Television Systems Committee, Washington, D.C., 22 January 2019.
- [5] ATSC: “ATSC Standard: Link Layer Protocol,” Doc. A/330:2016, Advanced Television Systems Committee, Washington, D.C., 19 September 2016.
- [6] IETF: “RTP protocol,” RFC 3550, Internet Engineering Task Force.

- [7] IETF: “RTP Profile for Audio and Video Conferences with Minimal Control,” RFC 3551, Internet Engineering Task Force.
- [8] SMPTE: “Forward Error Correction for Real-Time Video/Audio Transport Over IP Networks,” Doc. SMPTE ST 2022-1-2007, Society of Motion Picture and Television Engineers, White Plains, NY, 2007.
- [9] SMPTE: “Broadcast Exchange Format (BXF) – Protocol,” Doc. SMPTE 2021-2:2012, Society of Motion Picture and Television Engineers, White Plains, NY, 2012.
- [10] ITU-T, “V.41 Data Communication Over the Telephone Network, Code-Independent Error-Control System,” 1993 (or later, if available).
- [11] IEEE: “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” Doc. 1588, Institute of Electrical and Electronics Engineers, New York, NY, approved 27 March 2008.
- [12] SMPTE: “SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications,” Doc. SMPTE ST-2059-2, 2015, Society of Motion Picture and Television Engineers, White Plains, NY, 2015.
- [13] IETF: “Network Time Protocol Version 4: Protocol and Algorithms Specification,” RFC 5905, D. Mills, J. Martin, J. Burbank, W. Kasch, Internet Engineering Task Force, June 2010.
- [14] ITU-R, “Recommendation ITU-R TF.460-6, Standard-Frequency and Time-Signal Emissions,” Annex 1C “Coordinated Universal Time (UTC).”
- [15] ITU-R, “Recommendation ITU-R TF.460-6, Standard-Frequency and Time-Signal Emissions,” Annex 1B “International Atomic Time (TAI).”
- [16] IETF: “RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed,” RFC 3095, Internet Engineering Task Force, July 2001.
- [17] IETF: “Source-Specific Multicast for IP,” RFC 4607, Internet Engineering Task Force, August 2006.
- [18] SMPTE: “SMPTE Professional Media Over Managed IP Networks (ST 2110),” Doc. SMPTE ST-2110, 2018, Society of Motion Picture and Television Engineers, White Plains, NY, 2015.
- [19] NIST: “Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC,” Doc. NIST Special Publication 800-38D, National Institute of Standards and Technology, November 2007.
https://ws680.nist.gov/publication/get_pdf.cfm?pub_id=51288
- [20] NIST: “Specification for the Advanced Encryption Standard (AES),” Doc. NIST Federal Information Processing Standards Publication 197, National Institute of Standards and Technology, November 26, 2001.
- [21] NIST: “Implementation Guidance for FIPS Pub. 140-2 and the Cryptographic Module Validation Program,” Doc. FIPS 140-2IG, National Institute of Standards and Technology and Communications Security Establishment [of Canada], July 3, 2007.
- [22] IETF: “OpenPGP Message Format,” RFC 4880, Internet Engineering Task Force, November 2007.
- [23] IETF: “Elliptic Curve Cryptography (ECC) in OpenPGP,” RFC 6637, Internet Engineering Task Force, June 2012.
- [24] NIST: “Digital Signature Standard,” Doc. FIPS 186-3, National Institute of Standards and Technology, June 2009.
- [25] IETF: “Advanced Encryption Standard (AES) Key Wrap Algorithm,” Internet Engineering Task Force, September 2002.

- [26] OASIS: “PKCS #11 Cryptographic Token Interface Base Specification Version 2.40 Plus Errata 01,” Organization for the Advancement of Structured Information Standards, 13 May 2016.
- [27] NIST: “Secure Hash Standard (SHS),” Doc. FIPS 180-3, National Institute of Standards and Technology, October 2008.
- [28] W3C: “XML Schema Part 2: Datatypes Second Edition” W3C Recommendation, Worldwide Web Consortium, 28 October 2004. <https://www.w3.org/TR/xmlschema-2/>.
- [29] IETF: “The Secure Shell (SSH) Public Key File Format,” RFC 4716, Internet Engineering Task Force, November 2006.

3 DEFINITION OF TERMS

With respect to definition of terms, abbreviations, and units, the practice of the Institute of Electrical and Electronics Engineers (IEEE) as outlined in the Institute’s published standards [1] shall be used. Where an abbreviation is not covered by IEEE practice or industry practice differs from IEEE practice, the abbreviation in question is described in Section 3.3 of this document.

3.1 Compliance Notation

This section defines compliance terms used in this document:

shall – This word indicates specific provisions that are to be followed strictly (no deviation is permitted).

shall not – This phrase indicates specific provisions that are absolutely prohibited.

should – This word indicates that a certain course of action is preferred but not necessarily required.

should not – This phrase means a certain possibility or course of action is undesirable but not prohibited.

3.2 Treatment of Syntactic Elements

This document contains symbolic references to syntactic elements used in its various subsystems. These references are typographically distinguished by the use of a different font (e.g., restricted), may contain the underscore character (e.g., `sequence_end_code`) and may consist of character strings that are not English words (e.g., `dynrng`).

3.2.1 Reserved Elements

One or more reserved bits, symbols, fields, or ranges of values (i.e., elements) may be present in this document. These elements are used primarily to enable adding new values to a syntactical structure without altering its syntax or causing a problem with backwards compatibility, but they also can be used for other reasons.

The ATSC default value for reserved bits is ‘1’. There are no default values for other types of reserved elements. Use of reserved elements, except as defined in other ATSC Standards or by another industry standards setting body, is not permitted. See individual element semantics for mandatory settings and any additional use constraints. As currently-reserved elements may be assigned values and meanings in future versions of this Standard, receiving devices built to this version are expected to ignore all values appearing in currently-reserved elements to avoid possible future failure to function as intended.

3.3 Acronyms and Abbreviations

The following acronyms and abbreviations are used within this document.

A/V	Audio/Video
AAD	Authenticated Additional Data
AEA	Advanced Emergency Alert
AEAD	Authenticated Encryption with Additional Data
AEAT	Advanced Emergency Alert Table
AES	Advanced Encryption Standard
ALP	ATSC Link-layer Protocol
ALPTP	ALP Transport Protocol
APD	Associated Procedure Description
ATSC	Advanced Television Systems Committee
BCD	Binary Coded Decimal
BCH	Bose, Chaudhuri, Hocquenghem
BBP	Baseband Packet
BMFF	Base Media File Format
BPPS	Baseband Packetizer Packet Set
BRET	Bootstrap Reference Emission Time
BS	Bootstrap
bslbf	bit string, left bit first
BSR	Baseband Sampling Rate
BXF	Broadcast eXchange Format
CAP	Common Alerting Protocol
CRC	Cyclic Redundancy Check
CSRC	Contributing Source Identifier
CTI	Convolutional Time Interleaver
CTP	Common Tunneling Protocol
D/A	Digital to Analog
DASH	Dynamic Adaptive Streaming over HTTP
dB	decibel
DDE	Data Delivery Event
Demux	Demultiplexer
DER	Distinguished Encoding Rules
DSCP	Data Source Control Protocol
DSS	Data Source Signaling
DSTP	Data Source Transport Protocol
EA	Emergency Alert
ECC	Error Correction Coding
ECDH	Elliptic Curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
FEC	Forward Error Correction
FFT	Fast Fourier Transform

FI	Frequency Interleaver
FIFO	First In First Out
Gbps	Gigabits per second
GCM	Galois/Counter Mode
GHASH	GCM-based Hash value
GI	Guard Interval
GMAC	Galois/Counter Mode-based MAC
GNSS	Global Navigation Satellite System
GoP	Group of Pictures
GPS	Global Positioning System
HSM	Hardware Security Module
HTI	Hybrid Time Interleaver
IDR	Instantaneous Decoding Refresh
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IFFT	Inverse Fast Fourier Transform
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IS	Initialization Segment
ISO	International Organization for Standardization
KDF	Key Derivation Function
KEK	Key-Encryption Key
LCT	Layered Coding Transport
LDM	Layered-Division Multiplexing
LDPC	Low Density Parity Check
LLS	Low-Level Signaling
LMT	Link Mapping Table
LSB	Least Significant Bit
MAC	Message Authentication Code
MDCoIP	Media Device Control over IP
MDE	Media Delivery Event
MHz	Megahertz
MIMO	Multiple Input Multiple Output
MISO	Multiple Input Single Output
MITM	Man In The Middle
MMT	MPEG Media Transport
MMTP	MPEG Media Transport Protocol
MPD	Media Presentation Description
MPEG	Moving Picture Experts Group
MSB	Most Significant Bit
MTU	Maximum Transmission Unit
Mux	Multiplexer

NAL	Network Abstraction Layer
NoC	Number of Carriers
NRT	Non Real Time
OFDM	Orthogonal Frequency Division Multiplexing
PAPR	Peak to Average Power Ratio
PHY	Physical Layer
PLP	Physical Layer Pipe
PT	Payload Type
PTP	Precision Time Protocol
QP	Quality Point
RAP	Random Access Point
RDT	ROHC-U Description Table
RF	Radio Frequency
RFC	Request For Comments
ROHC-U	Robust Header Compression UDP
ROUTE	Real-time Object delivery over Unidirectional Transport
RTP	Real-time Transport Protocol
SAP	Segment Access Point
SBS	Subframe Boundary Symbol
SCT	Sender Current Time
SDPS	Security Data Packet Set
SDSP	Security Data Stream Packet
SFN	Single Frequency Network
SISO	Single Input Single Output
SLS	Service Layer Signaling
SLT	Service List Table
SMPTE	Society of Motion Picture and Television Engineers
S-RAP	System Random Access Point
S-TSID	Service-based Transport Session Instance Description
SSM	Source-Specific Multicast
SSRC	Synchronization Source Identifier
STL	Studio-to-Transmitter Link
STLTP	Studio-to-Transmitter Link Transport Protocol
T&M	Timing and Management
T-RAP	Transport-Random Access Point
TAI	International Atomic Time
tcimsbf	two's complement integer, msb first
TxID	Transmitter Identification
UDP	User Datagram Protocol
USBD	User Service Bundle Description
UTC	Universal Time Coordinated
uimsbf	unsigned integer, most significant bit first

3.4 Terms

The following terms are used within this document.

a-millisecond – A time interval approximately equal to one millisecond derived from a binary count of nanoseconds and actually equaling 2^{20} nanoseconds, which represents 1,048,576 nanoseconds (i.e., having a Period of 1.048576 milliseconds).

Advanced Emergency Alert – Provides an emergency notification mechanism in ATSC 3.0 that is capable of forwarding a broad range of emergency data. See [4].

ALP Stream – A sequence of ATSC Link-layer Protocol (ALP) packets processed in the order in which they arrive at a Broadcast Gateway and passed over the STLTP to Exciter(s) for transmission in a particular PLP.

Analyzed Media Duration – The period of time on the media timeline utilized by a Scheduler to map input data Streams to a set of Physical Layer frames.

Authentication – A process in which a hash of certain data is taken at the sending end of a link and the identical hashing process is applied again at the receiving end of the link, with a match between the results of the two hashing processes indicating a successful delivery and a mismatch between the two hashing processes indicating some sort of failure in the delivery.

Authentication Tag – A cryptographic checksum on data that is designed to reveal both accidental errors and the intentional modification of the data.

Baseband Packet – A set of $K_{payload}$ bits that form the input to an FEC encoding process. There is one Baseband Packet per FEC Frame.

Baseband Packetizer Packet Set – A collection of segmented packets representing a single Baseband Packet.

Block – In cryptographic processing, a unit of data that is processed together and to which a Block Cipher is applied. For any given Block Cipher, a bit string the length of which is the block size of the Block Cipher.

Block Cipher – A parameterized family of permutations on bit strings of a fixed length. The parameter that determines the permutation is a bit string called a Key.

Bonded Channels – Two or more broadcast signals on different channels that jointly transmit data or portions thereof representing particular content.

Bootstrap – A defined sequence of symbols that introduces each Physical Layer frame and provides a universal entry point into a digital Transmission signal. Each Bootstrap carries a value that serves as an indicator of the format of an immediately following Preamble symbol.

Bootstrap Reference Emission Time – A time value indicating the instant at which the leading edge of the first symbol of a Bootstrap is to be emitted from the transmitting antenna(s), absent any timing offsets of individual Transmitter(s) in a Network.

Broadcast Gateway – A Broadcast Gateway converts source file objects, for example media, system information, and other opaque files, into SFN baseband description for distribution to Transmitters.

Cell – A single Transmitter of an SFN.

Center Frequency – The point in the spectrum of a Physical Layer signal at which equal numbers of carriers are positioned both higher and lower in the spectrum.

Channel Bonding – A mode of broadcast transmission in which data representing particular content are transmitted on two or more RF channels by separate Transmitters to which portions or all of the data are distributed.

- Cipher Block** – A block of data equal in size to the fixed length of the bit strings to which the parameterized family of permutations of a Block Cipher is applied.
- Cryptographic Token** – A hardware device that is secure from intrusion on both physical and software levels, that executes cryptographic processes, and that provides secure storage for security keys and other data. In this standard, it takes the form of a USB plug-in hardware security module (HSM) and uses the Cryptoki interface protocol.
- Data Consumer** – A device that receives a formatted stream of data and further processes and/or distributes it.
- Data Producer** – A device or process that generates a formatted stream of data.
- Data Source** – An origination point for data to be transmitted as content by the Physical Layer.
- Data Source Signaling** – Information sent from a Data Source to downstream functions to identify specific characteristics of the content of certain packets within the data Stream to provide interlayer communications while avoiding layer violations in the system. Typically, Data Source Signaling information is carried in RTP headers in locations defined in this standard.
- Earliest Time** – A time value that accompanies data sent as input to a Broadcast Gateway to indicate the first instant, as determined using TAI, at which the first byte of related data, including all wrappers and encapsulating protocols, may start emission on the Physical Layer.
- Emission Wakeup Field** – The two Wakeup Bits in a Bootstrap when treated together as a field.
- Exciter** – An element of a Transmitter comprising data processing and signal processing functions including at least framing, error correction coding, waveform generation, modulation, and up-conversion to the output RF channel frequency.
- Fresh** – For a newly generated Key, the property of being unequal to any previously used Key.
- Key** – The parameter of the Block Cipher that determines the selection of the forward cipher function from the family of permutations.
- Latest Time** – A time value that accompanies data sent as input to a Broadcast Gateway to indicate the last instant, as determined using TAI, by which the last byte of related data, including all wrappers and encapsulating protocols, must complete emission on the Physical Layer.
- Majority Logic** – A technique for improving the reliability of delivery of certain data through transmission of an odd number of redundant copies of the data and determination at the point of reception of the data whether more than half of the copies received were the same as one another, thereby validating the data received.
- Man-In-The-Middle** – A type of security attack in which an attacker substitutes content for that which originally was sent.
- Media Segment** – A portion of a data Stream that is treated as a unit by a Scheduler for purposes of analysis and Transmission.
- Multicast** – A set of data sent across, for example, an IPv4 network to many recipients simultaneously.
- Multiplex** – A group of services that are transmitted together over a Network.
- Network** – A group of Transmitters delivering the same Multiplex.
- Nonce** – A value that is used only one time under a particular set of conditions or circumstances.
- Packet Set** – A group of packets carrying segments of a large data structure that has been segmented for the purpose of carriage across a transport connection that is not configured to carry the large data structure.

Packetizer – A process that treats a collection of data (e.g., a portion of a data Stream) by breaking it into segments and wrapping the segments in a header structure, thereby creating packets for Transmission / delivery.

Period – A duration of time.

Physical Layer – A functional protocol that defines the framing, resource allocation, and waveforms of signals emitted for delivery of data and content to receivers.

Plain Channel Bonding – a transmission method in which an ALP Stream is demultiplexed into two or more Baseband Packet Streams and those BBP Streams are transmitted on separate RF channels.

Preamble – The portion of a Physical Layer frame that carries L1 signaling data (see [3]) for the frame.

Preamble Generator – A function within a Broadcast Gateway that accepts instructions from a Scheduler, creates and formats Preamble Packets according to those instructions, and releases the Preamble Packets in the form of a Preamble Stream that can be multiplexed with other data Streams for delivery to the Transmitter(s) under control of the Broadcast Gateway.

Preamble Packet – A packet of data that provides a complete set of information necessary for Transmission following the Bootstrap of a Physical Layer frame to instruct receivers regarding the necessary receiver data processing to permit recovery of the data contained within the frame. The packet also serves to instruct Transmitters with respect to the configuration of Physical Layer frame that is to be emitted so that the Exciter data processing for the frame can be properly configured.

Preamble Parser – A function within an Exciter that receives Preamble Streams, extracts from them Preamble Payload data, and stores the Preamble Payload data until the time for its emission as part of a broadcast signal. The Preamble Parser then makes the Preamble Payload data available to the Exciter control system for use in configuring the data- and waveform-processing of the Exciter, and outputs individual Preamble Packets at the correct times for their inclusion in the emission of the Physical Layer frames the configurations of which they define.

Preamble Payload – The L1 data carried in a Physical Layer frame to define the structure of the frame and to specify the modulation, coding, and other parameters used in delivery of the frame.

Preamble Stream – A data Stream carrying Preamble Packets, comparable to the data Streams carrying data for PLPs and for Timing and Management functionality, that can be multiplexed with other Streams and delivered as a combined data Stream that is tunneled through the STLTP.

reserved – Set aside for future use by a Standard.

Scheduler – A studio-side function that allocates physical capacity to data Streams based on instructions from the System Manager combined with the capabilities of the specific system.

Security Data – A collection of data sent from a Cryptographic Token installed on a Broadcast Gateway to Cryptographic Token(s) installed on Exciter(s) or other Broadcast Gateway(s) to securely deliver keys to be used in the Authentication of data and metadata carried in all of the other Tunneled Packet Streams within an STLTP Stream or to securely deliver keys between Broadcast Gateway(s) through TCP/IP connections between those Broadcast Gateways.

Security Data Stream – A data Stream carrying Security Data Packets, comparable to the data Streams carrying data for PLPs, Preamble Data, and Timing and Management data, that can

be multiplexed with other Streams and delivered as a combined data Stream that is tunneled through the STLTP.

Signing – The process of generating an Authentication Tag for data forming a bit string, attaching the Authentication Tag to the bit string data, and communicating the Authentication Tag together with the bit string data to enable confirmation of the data source and the accuracy of the bit string data at the point of reception.

Single Frequency Network – Multiple Transmitters in proximity to one another radiating the same waveform and sharing a frequency.

SNR Averaging Channel Bonding – A transmission method in which an ALP Stream is demultiplexed into two or more Broadband Packet Streams. The BBP Streams next are data-processed to prepare symbol Streams, data cells then are exchanged between the symbol Streams, and those symbol Streams are transmitted on separate RF channels.

STL Interface – The STL Interface is the origin point for the Studio-to-Transmitter Link Tunneling Protocol (STLTP).

Stream – A sequential set of packets carrying data from a Data Producer to one or more Data Consumers.

Studio Interface – The termination point for the ALP Transport Protocol (ALPTP) and/or the Data Source Transport Protocol (DSTP).

System Manager – A conceptual subsystem outside the Transport and Physical Layers that is responsible for coordinating the functions of at least those two layers and the static and quasi-static configurations of various system aspects, for example definition of PLPs or assignment of IP addresses and port numbers to Services. The System Manager does not manage real-time traffic directly.

System Random Access Point – A location within an emission at which entry points into multiple layers and components of a service align, enabling rapid and efficient acquisition of the service.

Timing and Management Data – A collection of data sent from a Timing and Management Generator in a Broadcast Gateway to the Transmitter(s) under control of the Broadcast Gateway for purposes of controlling the emission times of Physical Layer frames, establishing various Transmitter configurations, and setting various Transmitter parameters, independent of the configurations of the frames and waveforms of the emitted signal itself.

Timing and Management Generator – A function within a Broadcast Gateway that accepts instructions from a Scheduler, creates and formats Timing and Management Data according to those instructions, and releases the Timing and Management Data in the form of a Timing and Management Stream that can be multiplexed with other data Streams for delivery to the Transmitter(s) under control of the Broadcast Gateway.

Timing and Management Stream – A data Stream carrying Timing and Management Data, comparable to the data Streams carrying data for PLPs and for Preambles, that can be multiplexed with other Streams and delivered as a combined data Stream that is tunneled through the STLTP.

Transmission – The signal that is emitted by a Transmitter and the synchronized signal that is emitted by all Transmitters in a Network.

Transmitter – An individual emitter at a specific geographic location on a specific frequency.

Transport Layer – A functional protocol that defines the formatting of data that will be delivered to receivers after its recovery from the formatting required for its physical delivery by the

Physical Layer. It provides logical communication between application processes running on different hosts within a layered architecture of protocols and other network components.

Tunnel Packet – A packet that carries in its payload the contents of one or more multiplexed packet Streams, including the corresponding headers and any other structural elements, i.e., a packet in the “outer” layer of a packet Tunneling system.

Tunneled Packet – A packet within a multiplexed group of packet Streams carried in a Tunnel Packet, i.e., a packet in the “inner” layer of a packet Tunneling system.

Tunneled Packet Information Header – A collection of metadata preceding each Tunneled Packet within the DSTP and ALPTP. Depending on the protocol, each Information Header provides length, routing and priority information to allow recipients of the Tunneled Packets to process them without examining the Tunneled Packet contents.

Tunneled Packet Stream – A Stream of multiplexed inner Tunneled Packets carried within an outer Tunnel Packet, e.g., within an STLTP encapsulation for ECC processing.

Tunneling – A process by which a group of parallel and independent packet Streams is carried within a single packet Stream so that they can be processed, transported, and otherwise treated as a single Stream entity.

Tuple – The combination of Internet Protocol (IP) addresses and port numbers.

Unicast – A set of data sent across, for example, an IPv4 network individually to one recipient.

Wakeup Alert – The occurrence of an AEA message requesting the setting of a non-zero value for the Emission Wakeup Field.

Wakeup Bits – The bits in Bootstrap Symbols 1 and 2 that are used to indicate to receivers that there is high-priority emergency information included in the broadcast. See [4].

3.5 Extensibility

The protocols specified in the present standard are designed with features and mechanisms to support extensibility. In general, the mechanisms include:

- Use of “protocol version” fields
- Definition of fields and values reserved for future use
- Use of XML, which is inherently extensible by means of future addition of new attributes and elements, potentially associated with different namespaces

Receiving devices are expected to disregard reserved values, and unrecognized or unsupported descriptors, XML attributes and elements.

3.6 XML Schema and Namespace

A number of new XML elements are defined and used in this Standard. These elements provide various control elements and attributes defined in this standard (see for example Section 7.1.1, DSTP Mapping Configuration Description). These new XML elements are defined with separate namespaces in schema documents that accompany this standard. The namespaces used by various schemas are described in individual sections of the present document. The sub-string part of namespaces between the right-most two ‘/’ delimiters indicate major and minor version of the schemas. The schemas defined in this present document shall have version ‘1.0’, which indicates major version is 1 and minor version is 0.

The namespace designator, “xs:”, and many terms in the “Data Type” column of tables is a shorthand for datatypes defined in W3C XML Schema [28] and shall be as defined there.

In order to provide flexibility for future changes in the schema, decoders of XML documents with the namespaces defined in the present document should ignore any elements or attributes they do not recognize, instead of treating them as errors.

All element groups and attribute groups are explicitly extensible with elements and attributes respectively. Elements can only be extended from namespaces other than the target namespace. Attributes can be extended from both the target namespace and other namespaces. If the XML schema does not permit this for some element, that is an error in the schema.

XML schemas shall use `processContents="strict"` in order to reduce inadvertent typos in instance documents.

XML instance documents shall use UTF-8 encoding.

In the event of any discrepancy between the XML schema definitions implied by the tables that appear in this document and those that appear in the XML schema definition files, those in the XML schema definition files are authoritative and take precedence.

The XML schema document for the schemas defined in this document can be found at the ATSC website.

4 SYSTEM OVERVIEW

4.1 Features

The STL subsystem exists between the Transport Layer, which creates ATSC Link-layer Protocol (ALP) packets, and the Physical Layer, which formats Streams of ALP packets for Transmission in particular Physical Layer Pipes (PLPs) in an emission configuration specified continuously in detail by a Scheduler. Documents [4] and [5] define ALP and other transport layer protocols. Documents [2] and [3] define the Physical Layer protocols. Figure 4.1 shows a high-level overview of the system configuration with applicable document numbers for the adjoining subsystems not defined herein. As depicted in the figure, data to be transmitted enters a Broadcast Gateway using either ALPTP (defined herein) or DSTP (defined herein) in the lower left of the diagram. Other inputs to the Broadcast Gateway are instructions from a System Manager. A Scheduler internal to the Broadcast Gateway controls the pre-processing functions that occur before delivery of the data and various control information to the Transmitter(s). Delivery of the combined data and instructions to the Transmitter(s) occurs in the lower right of the figure using STLTP (defined herein) with optional ECC applied. At the Transmitter across the top of the figure, the data and instructions from the studio are separated, buffered, and used to control the Transmitter as well as to construct the waveform to be emitted to receivers.

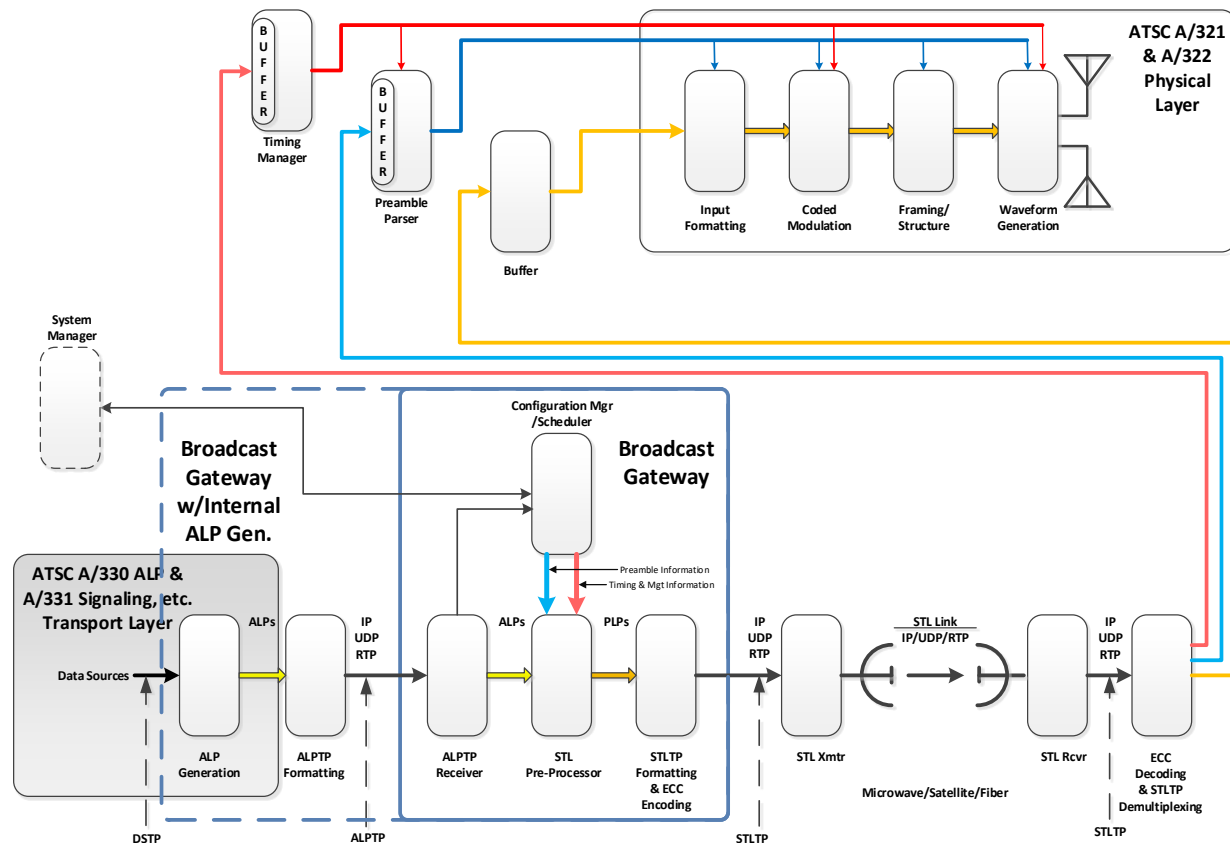


Figure 4.1 High-level overview of system configuration.

There is a one-to-one correspondence between individual Streams of ALP packets and individual PLPs. To prepare ALP packets for Transmission, in the Broadcast Gateway, the ALP packets are encapsulated in Baseband Packets (BBPs), which have defined sizes that are determined by a parameter ($K_{payload}$) related to the specific characteristics of the particular PLP(s) in which they will be carried. The sizes of the BBPs in a given Stream are set to assure that the assigned capacity of the related PLP in a frame is filled by the BBPs derived from an associated ALP packet Stream. ALP packets either are segmented or are concatenated so that they fill the allocated space in the BBPs carrying them as completely as possible without overflowing the available space.

To manage the flow of data through the system, several buffers are required to hold data for purposes of time alignment of data emission. Buffering also is required in certain instances to enable information to be obtained from a data Stream and used to control particular functionality of the system before the corresponding data is processed further. Two specific instances of such buffering exist in the system. The first buffer inserts at least one Physical Layer frame of delay in the STL Pre-Processor to enable sending Preamble information for a given Physical Layer frame to Transmitters prior to arrival of the data to fill that Physical Layer frame. The second buffer accommodates a delay of up to one second to enable synchronization of frame emission timing in the Physical Layer when the delivery delay to each of the Transmitters in a Network is different.

Maintaining the one-to-one correspondence between particular ALP packet Streams and their assigned PLPs through the system requires a method for identifying both the ALP and PLP data Streams. ALP packets are carried between ALP Generators and Schedulers using the ALPTP

which is derived from the Common Transport Protocol (CTP). The binding of ALP packets to PLPs is accomplished using Tunneled Packet Information Headers that are carried for each ALP packet within the ALPTP. Thus, one ALPTP stream can carry ALP packets destined for a single PLP or may multiplex multiple ALP packets for multiple PLPs. For STLTP, RTP/UDP/IP Multicast stacks are used with specific UDP port numbers assigned to particular PLP. Thus, for example, a Baseband Packet Stream derived from an ALP Stream destined for PLP 07 will be carried within an STLTP Stream with a UDP port value also ending in 07. When the emission operates in Single-PLP mode, all of the data to be transmitted will be carried in only a single ALP Stream, and all of that data will be transmitted with the same level of robustness. When multiple PLP Streams are used, each PLP can have a different tradeoff of data rate versus robustness, and data Streams can be assigned to appropriate combinations by the System Manager. If the ALP Stream(s) is (are) created within the same equipment that provides the Broadcast Gateway functionality, use of ALPTP may not be necessary.

Figure 4.1 shows a single path carrying the ALP packet Stream(s) and then the PLP packet Stream(s) on its (their) way(s) from the ALP Generator(s) to the Transmitter(s). In reality, if there are multiple Streams at any point in the system, the processing for each of the ALP packet Streams and/or Baseband packet Streams is applied separately to each of the Streams destined for a different PLP. This separation of processes is diagrammed using parallel paths throughout the remainder of this document, starting with the detailed examination of Figure 4.2.

To manage all of the characteristics of the emission and to coordinate all of the elements of the Physical Layer subsystem with respect to their parameter settings and times of operation, a Scheduler function is included in the Broadcast Gateway. The Scheduler manages the operation of a buffer for each ALP Stream, controls the generation of BBPs destined for each PLP, and creates the signaling data transmitted in the Preamble as well as signaling data that controls creation of Bootstrap signals by the Transmitter(s) and the timing of their emission. To perform its functions, the Scheduler communicates with a System Manager to receive instructions and with the source(s) of the ALP packets both to receive necessary information and to control the rate(s) of their data delivery.

One form of data relationship that the Scheduler must establish is signaling, in the Preamble of any given Physical Layer frame, the presence of Low-Level Signaling (LLS) data in specific PLPs within that frame. To enable the Scheduler to meet that requirement, upstream ALP generators in turn are required to signal to the Scheduler the presence of LLS data in specific ALP packets. When ALPTP is used (i.e., when ALP generators and the Scheduler are in separate equipment units), such signaling takes place in the ALP Tunneled Packet Information Header that precedes each ALP Packet tunneled through ALPTP. This method avoids the layer violation that would occur if the Scheduler had to determine the presence of LLS by inspecting the content of the ALP packets and covers cases in which the content of the ALP packets is not IP packets.

One of the principal functions of the Scheduler is to generate Preamble data for the Transmitter(s) that it controls. Conceptually, as shown in Figure 4.2, the Preamble generation function is assigned to a Preamble Generator, which is part of the Broadcast Gateway. The Preamble Generator outputs the data to be transmitted to receivers to allow their configurations to match the processes and parameters that will be used in Transmission. As the Transmitter(s) process the Preamble data for emission to receivers, the Preamble data also will be used by the Transmitter(s) to set up the Input Formatting, Coded Modulation, Framing/Structure, and Waveform Generation so that the emitted waveform will match what receivers will be instructed by the Preamble to receive. The exact format for the Preamble data is specified in [3].

Similarly, the Scheduler must control the generation and emission of Bootstrap waveforms by the Transmitter(s). To accomplish this, a data structure, similar to the Preamble, is defined in this document to carry Timing and Management (T&M) data to the Transmitters. Conceptually, as shown in Figure 4.2, a Timing and Management Data Generator is included in the Broadcast Gateway and provides the function under control of the Scheduler.

BBP data is carried across the STL as an RTP/UDP/IP Multicast Stream for each PLP. These Streams are multiplexed into a single RTP/UDP/IP Stream for each broadcast emission to enable reliable delivery to the Transmitter(s) of correctly identified and ordered BBPs. Conceptually, the BBP data Streams, as well as the Preamble Stream, the Timing and Management Stream, and the Security Data Stream are encapsulated as inner Streams carried through the outer (or Tunneling) Stream formed by the STLTP. While the inner Streams are Multicast-only, the outer Stream can be Multicast or Unicast. (See Annex B for an example of Unicast use.) The inner Stream provides addressing of BBP Streams to their respective PLPs through use of UDP port numbers. The outer protocol, STLTP, provides maintenance of packet order through use of RTP header packet sequence numbering. The STLTP also enables use of (SMPTE ST 2022-1) ECC to maintain reliability of Stream delivery under conditions of imperfectly reliable STL Networks.

At the Transmitter(s), an input buffer is used for each PLP to hold BBP data until it is needed for Transmission. There also are FIFO buffers for the Preamble Stream and the Timing and Management Stream. The Preamble Stream processing includes a Preamble Parser that collects all of the configuration information for the next and possibly several upcoming Physical Layer frames to use in configuring the Transmitter data processing for those frames.

Preamble data is scheduled to arrive at the Transmitter input at least one full Physical Layer frame Period prior to the first byte of the associated payload to provide time for the Transmitter data processing to be configured properly. Preamble data also can be sent multiple times in advance to enable acquisition of the data with improved reliability. The same considerations also are applicable to the Timing and Management Data; i.e., it is scheduled to arrive at the Transmitter input at least one Physical Layer frame Period prior to the first byte of the associated payload (+processing delay) it describes, and it can be sent multiple times to enable improved reliability of its acquisition.

The ATSC 3.0 system supports in-order delivery of streaming media files, associated metadata, and generic file delivery to the IP layers of receivers. It is imperative that the systems comprising the ATSC 3.0 distribution path maintain in-order sequencing of the respective Data Source streams through the several packet encapsulation and delivery processes so that their data is delivered to receivers in the correct order. The three protocols described in this document that operate within the emission portion of the system, namely, DSTP, ALPTP, and STLTP, utilize RTP to enable order recovery, should it be required. RTP transport need not and does not transit the ATSC 3.0 air interface.

4.2 System Architecture

The Studio to Transmitter Link (STL) interface is typically located between the baseband Packetizer and the Forward Error Correction (FEC) block. There only needs to be one Scheduler and one baseband Packetizer per RF emission. Multiplexing of multiple Services among stations sharing one RF emission can be accommodated on the input side of the Scheduler.

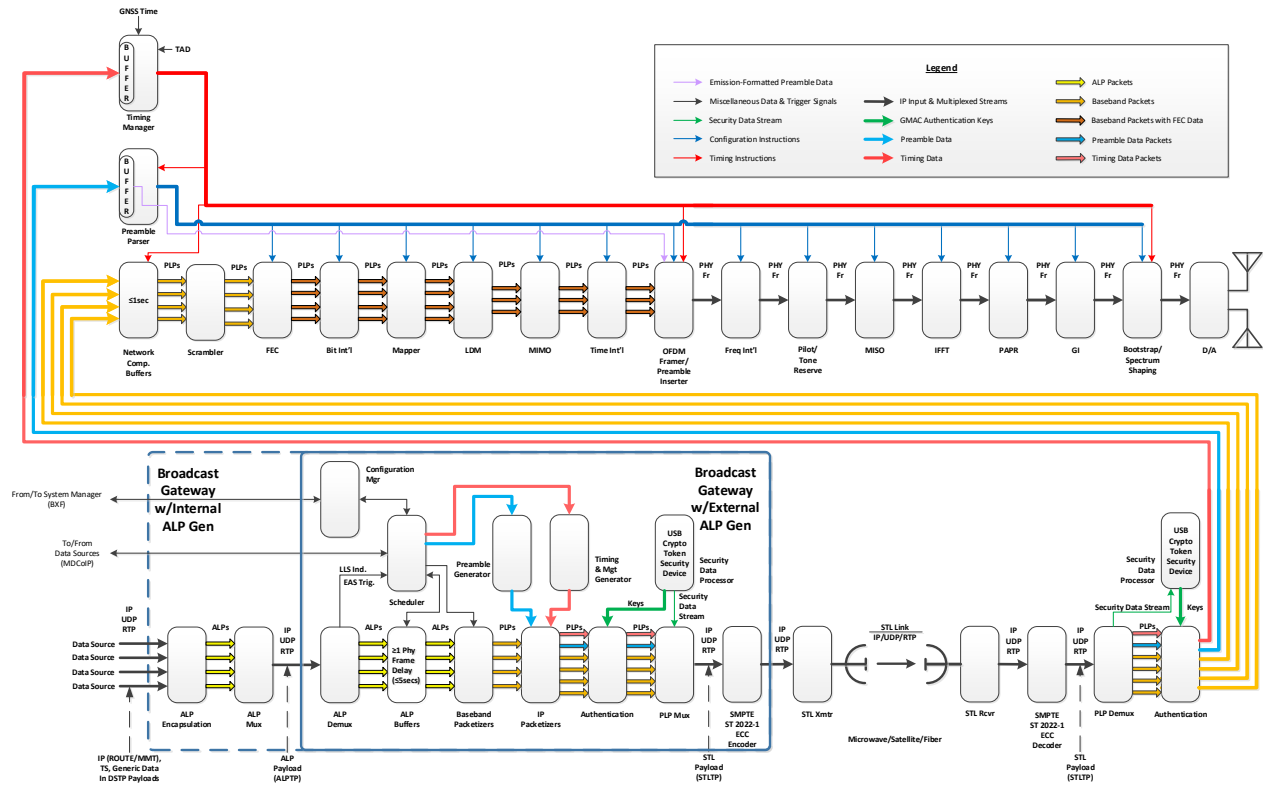


Figure 4.2 System architecture.

Figure 4.2 shows a possible system architecture; other configurations are possible. When considering system configurations, data rates and interfaces between functional blocks must be taken into account in developing practical implementations.

4.2.1 System Manager

Configuration aspects of the overall system are controlled by a single entity called a System Manager, which is represented in Figure 4.2 only as a connection to the Configuration Manager in the Broadcast Gateway. A System Manager can be anything from a web-page based setup screen with manual data entry to a fully automated system; its scope is control of an overall facility or system. The System Manager provides high-level configuration parameters for numerous system functions. The System Manager controls static or quasi-static configurations of the Transmission chain. It controls the Physical Layer configuration with respect to how many PLPs operate and the configurations of the individual PLPs, the Services supplied on those PLPs, and the delivery sessions that support the Services that run in the PLPs. A System Manager can establish a pre-determined schedule for system re-configuration, sending instructions to various system devices and subsystems for delayed execution at specified times. All configuration parameters sent by a System Manager to a Configuration Manager are derived from the Preamble parameter set described in [3] Section 9 and listed in Table 5.1 herein. In response, the Configuration Manager provides feedback of Physical Layer capabilities or structure details of selected Physical Layer frame types. Instructions to the Configuration Manager for these changes in configuration must be sufficiently in advance of the time of emission change. The required minimum advance notice of configuration change time depends on total reconfiguration latency of the combination of the

Configuration Manager and the Scheduler. Further descriptions of Scheduler functionality appear in Sections 4.2.2 and 5.2.

Data sources feeding the Broadcast Gateway are identified to it by the System Manager using IP addresses and port numbers, with one IP address and port combination (i.e., one Tuple) associated with each Data Source or ALP Stream. These Data Sources also can have separate control IP addresses, in which case the IP addresses of the control connections are indicated to the Broadcast Gateway by the System Manager – at most one per ALP Stream. The operation of the Data Source control interface and the messages that cross it are described in overview in Section 4.9 and in detail in Section 5.5. For each input Data Source or ALP Stream, the System Manager assigns a specific PLP number, which is related to the UDP port address used, as described in Section 7.1.1.

4.2.2 Broadcast Gateway

A Broadcast Gateway is an equipment item that incorporates a number of the conceptual functions necessary in processing data prior to its delivery over an STL to one or more Transmitters. Broadcast Gateways can be implemented in two primary ways, differing primarily in the location of one processing function and the type of interface used to deliver data Streams for Transmission by the Physical Layer subsystem. Both configurations of a Broadcast Gateway are shown in Figure 4.2. In the Broadcast Gateway outlined with a solid line, an external function, upstream in the system, generates ATSC Link-layer Protocol (ALP) packets [5] by encapsulating the data that is to be transmitted in specific PLPs. With external generation of ALP packets, those packets are delivered to the Broadcast Gateway using an ALP Transport Protocol (ALPTP) defined herein, and metadata associated with particular ALP packets and necessary for their Transmission is carried with the packets in their ALPTP headers. In the Broadcast Gateway outlined by the combination of the solid and dashed lines, a function within the Broadcast Gateway performs the ALP packet generation, and a slightly different Data Source Transport Protocol (DSTP) is used to carry both content data and related metadata, the latter of which again is carried in the headers of the DSTP packets.

4.2.2.1 Configuration Manager

The conceptual Configuration Manager within a Broadcast Gateway provides a configuration resource to the System Manager that can accept general waveform structure requirements and from them develop a complete description of the frame structure and waveform that is to be emitted by the Physical Layer. It can inform the System Manager of the capabilities of the system and can provide trial frame structure and waveform details to the System Manager for approval. It also can accept from the System Manager an operational schedule based on use of predefined configurations that the Configuration Manager and the System Manager both have accepted and stored for later use with a naming convention for reference to the predefined configurations. The Configuration Manager provides instructions to the Scheduler, at appropriate times, with respect to the detailed frame structure and waveform that are to be adopted starting at a particular instant and continuing until further instructions are given.

4.2.2.2 Scheduler

The Scheduler within a Broadcast Gateway receives an input of configuration instructions from the Configuration Manager and determines both the frame structure and timing and the waveform details for the next and subsequent Physical Layer frames. The Scheduler must look ahead in time in developing the details of a sequence of Physical Layer frames since instructions about the structures and emission timing of future frames must be sent to Transmitters in advance of arrival

of the content data to be transmitted in them. The Scheduler also manages a demultiplexer and a buffer that process data arriving in the form of ALP packets or DSTP packets and any associated metadata in the packet headers. DSTP-delivered data will be encapsulated into ALP packets before reaching the Scheduler, and its associated metadata will be passed to the Scheduler in much the same way as metadata arriving in ALPTP packet headers. The metadata arriving with the content data is needed by the Scheduler for making a number of decisions. Included are decisions to send Wakeup Bits in the emitted Bootstrap of a Physical Layer frame and to signal the presence of various forms of table data (e.g., LLS) in a frame and in a particular PLP within that frame. The Scheduler also controls the Baseband Packetizer blocks, causing them to create, for each ALP Stream, Baseband Packets of one of several fixed sizes, the choice among which depends upon the value of the parameter $k_{payload}$, which is determined according to the characteristics of the specific Physical Layer Pipes (PLPs) into which their Baseband Packets will be sent [3].

Adjuncts to the Scheduler are a Preamble Generator and a Timing and Management Generator. The Scheduler determines the data to be sent to the Transmitter(s) for inclusion in the transmitted Preamble, and it similarly determines the data to be sent to the Transmitter(s) for control of their operation. The Preamble Generator uses instructions from the Scheduler to form and format a packet Stream to be sent in parallel with the PLP packet Streams to the Transmitter(s) for use by the Transmitters in configuring their operation and for emission to receivers to announce the signal configuration to be received and demodulated. The Timing and Management Generator similarly uses instructions from the Scheduler to form and format a packet Stream to be sent in parallel with the PLP packet Streams to the Transmitter(s) for control of their operations. Unlike the Preamble Stream, the Timing and Management Stream will not be broadcast but rather will be consumed by the Transmitter(s) when it is used to control the timing of its (their) emissions and other functions either as a group or individually by Transmitter.

4.2.3 Studio to Transmitter(s) Dataflow

The Studio-to-Transmitter Link (STL) may operate on any of fiber, satellite or microwave links. Reliability of all such links is likely not to be perfect, even when they apply their own error correction coding (ECC); consequently, additional ECC is applied in the STLTP path. In addition, Internet Protocol (IP) is supported on all link types. Because of its exposure on any of the delivery means to potential security breaches, it is important that the data traversing the STL be protected by a security mechanism. Therefore, a system for authentication of the contents of the STLTP data stream is provided.

4.2.4 STL Operation

Broadcasters have a need to send studio-generated data to their Transmitters. Usually those Transmitters are not co-located at the studio. An STL Interface from the studio to the Transmitter(s) is needed. Requirements for such an interface include:

- 1) Support for Real-Time Protocol / User Datagram Protocol / Internet Protocol (RTP/UDP/IP) IPv4 and addressing
- 2) Encapsulation of data for the link
- 3) Providing a synchronization method to TAI for both data and control
- 4) Providing signaling of the Transmitter time synchronization for data and control
- 5) Having measurable maximum latency to allow emission times to be correct
- 6) Applying Error Correction Coding to the data Stream carried
- 7) Applying Authentication of the content of the data Stream carried
- 8) Allowing for redundancy of the delivery channel

4.2.5 SFN Operation

Certain specifications in this standard enable Single Frequency Network (SFN) operations, and the protocol for data carriage from the studio to the Transmitter(s) has the capability to support delivery of certain control data individually to each Transmitter in an SFN. Specifically included are the following specifications and capabilities:

- 1) Inherent synchronization of the data processing functions of multiple Transmitters fed from the same Broadcast Gateway
- 2) Inherent time alignment of the emissions of multiple Transmitters fed from the same Broadcast Gateway
- 3) Specification of the carrier frequency tolerance of Transmitters in an SFN
- 4) Addressing each Transmitter individually and delivering specific data to each
- 5) Providing offsets from the Reference Emission Time individually to each Transmitter in a Network to facilitate Network service shaping

Timing considerations are satisfied by the Timing and Management Protocol described in Section 9.3.1, and carrier frequency accuracy is specified in Section 10.3.1.

4.3 Central Concepts: CTP

The Common Tunneling Protocol (CTP) forms the basis for the three main packet protocols described in this document. The CTP defines an SMPTE ST 2022-1 [8] RTP/UDP/IP *Tunnel* that carries a variety of packets and metadata according to each of the specific protocols. All Tunnel Packets comprising a CTP-based stream may use either multicast or unicast addressing.

Real-time Transport Protocol (RTP) is used with its headers as redefined in Section 6.2.1. Segmentation and reassembly of large Tunneled payload packets and concatenation of small Tunneled payload packets within the Tunnel Packets is performed using RTP header signaling. A segment sequence number within an “outer” RTP header indicating position of a segment within a larger source packet supports segmentation and reassembly, and a value in an “outer” RTP header indicating the offset of the first “inner” packet segment within the payload of the associated “outer” packet supports concatenation. RTP also provides capabilities for ordering packets and determining if any packets have been lost.

If SMPTE ST 2022-1 [8] ECC is in use, any lost packets can be reconstituted. SMPTE ST 2022-1 defines the ECC function and is intended for data rates up to around 1 Gbps, which makes it appropriate for the protocols specified in this document. In addition to the Tunnel Packet Stream, one or two additional Streams of ECC overhead data having related UDP port numbers can be generated by the ECC subsystem and be carried in parallel with the Tunnel Packet Stream.

4.4 Central Concepts: DSTP

The Data Source Transport Protocol (DSTP) provides a solution for transferring content data, associated metadata, and signaling through a typical IP Network from Data Sources to a Broadcast Gateway as shown in Figure 4.2. While Data Source outputs are expected to be conducive to carriage on normal IP Networks, they do not provide the resources for carriage of metadata and signaling related to the content that must be communicated to Broadcast Gateways. Also, Data Source packets must be kept in order, and packet loss cannot be tolerated. To address the needs beyond normal UDP/IP functionality, the DSTP relies on the Common Tunneling Protocol (CTP) to wrap the various data source packets in a convenient RTP/UDP/IP stream. A separate Tunneled Packet Information Header is prepended to each Data Source packet to carry required metadata and signaling that must be communicated from Data Sources to Broadcast Gateways. The DSTP

Tunnel headers and Tunneled Packet Information Headers are discarded by the Broadcast Gateway once the metadata contained in them is processed.

4.5 Central Concepts: ALPTP

The ATSC Link-layer Protocol Transport Protocol (ALPTP) provides a solution for transferring ATSC Link-layer Protocol (ALP) packets [5] through a typical IP Network between two separated devices as shown in Figure 4.2. ALP packets are relatively simple constructs having only a minimal header plus packet length information sufficient just for an emission link layer. These headers are not sufficient for transferring ALP packets through a typical IP Network; consequently, ALPTP is defined to provide necessary additional networking information. Moreover, ALP packets must be kept in sequence for each PLP, and packet loss cannot be tolerated.

ALPTP is based on the Common Tunneling Protocol (CTP). The ALP packets are tunneled through the CTP preceded by a Tunneled Packet Information Header that describes the metadata associated with the ALP packet including its priority and PLP ID. There is a single ALP Stream associated with each PLP and the PLP ID in the ALPTP Tunneled Packet Information Headers allows these packets to be grouped. Note that there are no constraints on how many ALPTP streams can be communicated between an ALP Generator and the Scheduler since metadata regarding binding of ALP packets to PLPs is contained in the Tunneled Packet Information Header accompanying each ALP packet in the ALPTP tunnel.

4.6 Central Concepts: STLTP

The Studio to Transmitter Link Transport Protocol (STLTP) provides a solution for transferring a potentially large number of parallel Baseband Packet (BBP) data Streams carrying content for a like number of PLPs plus Preamble data and Transmitter control signaling through a typical IP Network from a Broadcast Gateway output to one or more Transmitters using Tunneling, as shown in Figure 4.2. Each Tunneled BBP Stream feeding a PLP derives from a corresponding ALP Stream, and PLP IDs provided with the ALP Streams are used to establish the associations. Error Correction Coding (ECC) is specified to maintain reliable delivery of STL data over the links between sites, but the ECC method selected requires use of uniform size packets in its data processing. To accommodate this requirement, the multiple PLP, Preamble, and control Streams are multiplexed together on a packet-by-packet basis and then tunneled through a single resulting RTP/UDP/IP Common Tunneling Protocol data Stream. When the distribution link does not support multicast (e.g., a public IP network), an RTP/UDP/IP Unicast protocol stack also may be used to deliver the multiplexed PLP(s), Preamble, control, and security Streams.

4.6.1 STLTP Payload Data

The STLTP used to traverse STL IP links carries four distinct types of data:

- 1) Baseband Packet Streams, the data from each of which populates an individual PLP,
- 2) Preamble data packets derived from the scheduling process outcome and used both to populate emitted Preambles and to control Exciter configurations, and
- 3) Timing and Management Data packets used to control Transmitter synchronization, to control Bootstrap emission timing, to provide Physical Layer frame identification, and for similar Transmitter management tasks, e.g., TxID control.
- 4) Security packets used to provide security information allowing Exciters to authenticate the various Streams within the STLTP (see Section 4.6.2 below).

4.6.2 Content Security System

The Tunneled Packets within the STLTP may be signed with a cryptographic signing method to allow detection of Man-in-the-Middle (MITM) attacks against the STLTP. The security design is based on the Advanced Encryption Standard (AES) using Elliptic Curve Cryptography and the Galois/Counter Mode (GCM) of AES. The signing functions of GCM produce a Message Authentication Code (MAC) called GMAC. Tunneled Packets within STLTP Streams can be signed with any of four random-number-based Key values or can be left unsigned. Keys for Stream-signing are generated at Broadcast Gateways and communicated to Exciters over Public/Private-key-encrypted and authenticated Security Data Streams also tunneled in STLTP Streams. Keys used both for signing and for encrypted key delivery are generated and retained in hardware Cryptographic Tokens. Private Keys are never exposed outside hardware Tokens. Public Keys are exchanged between and stored only in hardware Tokens. Keys used for signing similarly are exchanged between and stored only in hardware Tokens.

4.7 System Time Domains

As can be seen in both Figure 4.1 and Figure 4.2, the Broadcast Gateway sits astride the boundary between studio and Transmitters, which also is the boundary between the Transport Layer (ALPs) and the Physical Layer (BBPs). Due to the nature of the data formatting in each of those layers, they operate in different time domains. The Transport Layer and layers above it operate in the UTC [14] domain, matching the time on conventional wall clocks, accounting for leap seconds, and often expressing time in NTP [13]. The Physical Layer, on the other hand, operates in the International Atomic Time (TAI) [15] domain, which has second ticks at the same times as UTC but does not account for leap seconds. This allows the Physical Layer to maintain a constant flow of data, including at leap second insertion instants, without interruption of the Transmission process. Reconciling the differences between the two time domains is a task that falls to the Broadcast Gateway for resolution.

Indication of the value of TAI time on the STL Interface (STLTP) is based on the Precision Time Protocol [11] using 32 bits to represent seconds and 32 bits to represent nsecs.

ATSC Physical Layer Time, which is indicated by the 32 LSBs of TAI seconds and BCD-coded msec, μ sec, and nsec, is delivered to receivers in the Preamble [3].

4.8 System Manager Configuration Interface

As described in Section 4.2.1, a System Manager is a conceptual entity that coordinates and controls all broadcaster facilities necessary to produce a specific desired station output configuration and emission. The configuration interface between the System Manager and the Configuration Manager in a Broadcast Gateway transfers information between the two subsystems to enable the system management process with respect to the Scheduler and its associated functions. The System Manager configuration interface comprises a normal TCP/IP connection between the devices with control functionality structured according to the SMPTE Professional Media Over Managed IP Networks (ST 2110) standards suite [18] and carrying the information described in Section 5.4 herein. The messages are used to negotiate a detailed emission waveform configuration between the System and Configuration Managers and to permit the System Manager to provide instructions to the Configuration Manager with respect to emission requirements and schedules.

4.9 Real-Time Control Interface

As mentioned in Section 4.2.1, the System Manager directs the Scheduler through the Configuration Manager to control data delivery from various Data Sources by providing an address and port number combination (a Tuple) for each Data Source feeding data to the Broadcast Gateway for delivery through a PLP. The Scheduler communicates with the Data Source(s) at the address(es) and port number(s) assigned by the System Manager using a Real Time Control Interface that employs SMPTE Professional Media Over Managed IP Networks (ST 2110) methods [18] and carrying the information described in Section 5.5 herein. Messages that can be communicated between the Scheduler and the various Data Sources are defined in Section 5.5. The Real-Time Control Interface provides such functions as discovery of the capabilities of a Data Source, setting a target bit rate for a Data Source, and managing the speed of data delivery from a Data Source in real time, i.e., updating rates of data delivery as necessary to maintain control of buffer fullness throughout the Physical Layer system.

4.10 Transmitter Requirements Overview

Operation of the ATSC 3.0 Physical Layer system depends on standards-compliant Transmitters meeting certain requirements. Among those requirements are: inclusion of Timing Manager functionality, inclusion of Preamble Parser functionality, conforming to certain buffer models and conditions, maintaining a specified frequency accuracy, supporting emission timing offsets from the Reference Emission Time, and providing for adoption of carrier frequency and emission time offsets as instructed by system Schedulers. These specifications and requirements are detailed in Section 10.

4.10.1 Carrier and Timing Offset for Co-Channel Interference Mitigation

When multiple, geographically neighboring stations are operated on the same RF channel using identical, or even relatively similar, frame and waveform configurations, interference can be caused between the co-channel stations. The interference arises when the pilot locations and characteristics are sufficiently similar in the several signals. Such interference is particularly troublesome when the various stations involved are operating with highly robust emission characteristics having very low signal-to-noise ratios. The mechanism that leads to such interference is addition of pilot energy at each OFDM carrier in an RF channel at which pilots are positioned. Addition of the pilot energy at each pilot frequency will lead to incorrect channel estimation by receivers because the channel estimation will not be valid for either/any of the received signals in the overlap areas of the Transmitters' signals. The same effect will result from both single-Transmitter and SFN operation by the neighboring stations.

To counter the sort of interference described, the carrier frequencies of neighboring stations can be offset by a small amount to avoid pilot super-positioning, which would otherwise result in pilot level addition. In addition, super-positioning of Bootstrap emission times must be avoided to prevent failure of Bootstrap recovery from the individual neighboring stations in regions of signal overlap. Section 10.3.3 herein describes methods for avoidance of super-positioning of both pilots and Bootstraps by neighboring stations, thereby extending their interference-free service areas when they operate with low C/N threshold configurations.

5 SCHEDULER DESCRIPTION AND NORMATIVE REQUIREMENTS

Schedulers are the central subsystems in the Physical Layer that enable movement of data from Data Sources, translation of data formatting from Transport Layer forms to Physical Layer requirements, configuration of the transmitter(s) and the waveforms it/(they) emit, buffering of

data to stage its delivery, timing of emission of Physical Layer frames, and all the other functions necessary to implement the ATSC 3.0 Physical Layer as embodied in [2] and [3]. A Scheduler accepts inputs from a System Manager that give it general instructions as to how the emission is to be configured and scheduled, and it outputs data and instructions to the transmitter(s), controlling precisely how the emissions actually are configured.

5.1 Relationship of Broadcast Gateway and Scheduler to the System

Figure 4.2 shows a conceptual block diagram of a Broadcast Gateway and its associated interfaces. A configuration interface allows provisioning of quasi-static aspects of Physical Layer configuration, such as PLP definitions. A Data Source Interface has two ports, one of which delivers content to the Broadcast Gateway along with Data Source Signaling that carries inter-layer information and the other of which is used for real-time control exchanges between the Broadcast Gateway and one or more Data Sources. The required message interchanges on the control interface are defined in this document. Their semantics and the protocol for their carriage are under development at the time of publication and will be documented in the future either as a revision to this standard or as a separate document. Output from the Broadcast Gateway and its Scheduler are via an STL Interface that communicates using STLTP, which carries a complete description of a Physical Layer instance on a frame-by-frame basis to one or more Transmitters.

5.2 Scheduler Functionality

The functional assets of a Scheduler are defined by data size(s), time(s), and levels of robustness on the Physical Layer. The Physical Layer can deliver defined quantities of data at certain discrete times with particular amounts of resilience to channel impairments. There are a large number of parameters, many having a wide range of settings, under control of the Scheduler, and a major part of its function is to make intelligent choices among settings to most efficiently accomplish the objectives communicated to it by the System Manager.

A Broadcast Gateway receives the following inputs and information:

- 1) Data for transmission, either in the form of ALP-encapsulated content data (via ALPTP) or in the form of content data Streams (via DSTP) that must be ALP-encapsulated within the Broadcast Gateway;
- 2) Inter-layer signaling sent from the Transport Layer to the Physical Layer in the form of Data Source Signaling to identify certain types of content requiring specific actions to be taken by the Scheduler to properly signal or treat the content in the broadcast emission;
- 3) Configuration and scheduling instructions from the System Manager; and
- 4) Capability and status information from the Data Source that feeds each of the ALP/PLP data Streams.

From these inputs, a Scheduler shall create an efficient allocation of the Physical Layer resources conforming to the scheduling and configuration instructions that it receives from its associated System Manager. The Scheduler shall manage buffer fullness throughout the Broadcast Gateway and Transmitter chain, based on current sizes of subframes and frames, maximum delay of the STL delivery Network, available STL channel bandwidth, and requirements for robustness and security across the STL. The Scheduler shall determine Bootstrap Reference Emission Times, create timed control instructions for all Transmitters (collectively and individually), and pass timing and management control messages to the Timing and Management Data Generator so that it can create Timing and Management Data Packets to be sent to the Transmitter(s). As part of managing robustness and security across the STL, the Scheduler shall instruct the Preamble

Generator and the Timing and Management Generator how many repetitions of their respective packets shall be sent to the Transmitter(s) and at what times relative to delivery of the contents of the Physical Layer frames with which they are associated.

The Scheduler shall define all of the Physical Layer parameters, such as frame lengths, subframe sizes, PLP configurations, and modulation and coding type and value selections. It shall provide this information and all other details to be communicated to both the Transmitter(s) and off-air receivers to the Preamble Generator for formation into Preamble Packets. The Scheduler shall assign frame identifiers to all Baseband Packets using the Bootstrap Reference Emission Times of the frames in which they are to be emitted.

5.3 Preamble Construction

The Preamble defined in [3] is used to control the configurations of both Transmitter(s) and off-air receivers so that the signals emitted match the decoding processes in receivers, leading to successful data recovery. The Scheduler shall construct a single Preamble Packet for each Physical Layer frame, for this purpose, which packet shall be sent to the Transmitter(s), used by them for their own configurations, and transmitted as part of the Physical Layer frame with which the Preamble Packet is associated. Multiple copies of the Preamble packet for a particular Physical Layer frame may be sent to the Transmitter(s) over a period of time to improve robustness of the Preamble data delivery to the Transmitter(s), through use of majority logic or similar techniques, as described in Section 9.3.1.

A Preamble Packet consists of two fundamental categories of data: quasi-static and dynamic, with the former changing only at infrequent intervals and the latter changing potentially on every frame. The quasi-static information comes from the System Manager, while the dynamic information is developed within the Scheduler as part of its functionality. Both types of information end up in the Preamble. The various Preamble parameters are shown in Table 5.1, and their sources in either the System Manager or within the Scheduler processes are indicated by check marks in the two columns on the right side of the table.

Table 5.1 Preamble Parameters and Their Sources

Parameters		Instructions from System Manager	Scheduler Generated
Per Frame Data	Channel bandwidth	✓	
	Sampling Rate	✓	
	Major and Minor version values	✓	
	Emergency Alert Wakeup		✓
	L1B Preamble Structure indicator		✓
	Minimum time to next Frame		✓
	Frame length	✓	✓
	Number of subframes		✓
	Number of symbols in the Preamble		✓
	Frame alignment	✓	
	Frame PAPR	✓	
	CRC values (L1B)		✓
	CRC values (L1D)		✓
	Preamble NoC	✓	
	FEC Mode for L1-Detail	✓	
	Additional parity for next frame		✓

	Channel BSIDs involved in channel bonding	✓	
	Center frequency of channels involved in bonding	✓	
	MIMO scattered pilot encoding method	✓	
	Return Channel flag	✓	
	LLS flag		✓
	Time info flag		✓
Per Subframe Data	Subframe size	✓	
	Subframe PLP count	✓	
	Subframe MIMO/MISO/SISO	✓	
	Subframe FFT size	✓	
	Subframe NoC	✓	
	Subframe GI	✓	
	Subframe pilot pattern	✓	
	Subframe pilot boost	✓	
	Subframe boundary symbol flags	✓	
	Subframe FI mode	✓	
Per PLP Data	PLP ID	✓	
	PLP type		✓
	PLP size		✓
	PLP start position		✓
	PLP number of subslices		✓
	PLP subslice interval		✓
	PLP LLS flag		✓
	PLP scrambler type	✓	
	PLP FEC mode	✓	
	PLP position of first complete FEC Block		✓
	PLP LDPC code	✓	
	PLP modulation	✓	
	PLP time interleaver mode	✓	
	PLP CTI depth		✓
	PLP CTI start row		✓
	PLP CTI position of first complete FEC block		✓
	PLP HTI inter-subframe interleaving flag		✓
	PLP HTI number of TI blocks or subframes		✓
	PLP HTI max interleaving FEC blocks per interleaving frame		✓
	PLP HTI number of FEC blocks in the current interleaving frame		✓
	PLP HTI Cell interleaver flag		✓
	PLP LDM layer	✓	
	PLP LDM injection level	✓	
	PLP Channel BSIDs involved in channel bonding	✓	
	PLP Center frequency of channels involved in bonding	✓	

5.4 Scheduler Management Protocol

The interface between the Scheduler and System Management functions shall use SMPTE Broadcast eXchange Format (BXF) protocol as described in [9].

The parameters using this protocol are those listed in Table 5.1 under the parameter column with a check mark in the Instructions from System Manager column. These parameters are quasi-

static in nature and do not routinely change between Physical Layer frames. Scheduler configurations and constraints are set with these parameters, which are allowed to change quasi-statically. An emission schedule for a set of parameters can be similar to a program schedule, in which the parameters can change over the course of a day. Detailed description of these parameters is provided in 5.4.

In addition to providing the quasi-static Preamble parameters, the System Manager shall identify for the Scheduler:

- 1) The Source and Destination Multicast IP addresses of each Data Source associated with each ALP Stream. Only one Data Source shall be selected for a given ALP Stream at a time.
- 2) The Destination Multicast IP address associated with each ALP Stream, where the port number used indicates the particular ALP / PLP pair.
- 3) The Control IP address for each Data Source providing data for each ALP Stream. Redundant Data Sources for any given ALP Stream are possible, but only one can be selected for use at a time.

Messages between a System Manager and a Scheduler shall include:

- Capabilities inquiry to inform the System Manager of the assets and capabilities of the Scheduler.
- Configuration design inquiry to provide a detailed Physical Layer frame design from a Scheduler to a System Manager to meet a set of requirements from the System Manager.
- Current configuration inquiry to report to the System Manager the current state of operation of the Scheduler and any presets available for use.
- System Setup instructions to provide to the Scheduler the fundamental configuration settings for operation in the particular system.
- Scheduled configuration instructions to instruct the Scheduler when to place into operation one or a number of preset configurations stored in the Scheduler and available for use.

Scheduler feedback to the System Manager shall indicate the Physical Layer capabilities and the structure details for each requested frame type. Scheduler to System Manager messages include:

- Capabilities response
- Configuration design response (acknowledgement of preset instructions)
- Current configuration response
- System setup acknowledgement
- Scheduled configuration acknowledgement

5.5 Data Source Control Protocol (DSCP)

The interface between the Scheduler and the one or multiple Data Sources that provide content data for emission shall permit the Scheduler to control operation of the Data Sources so that they deliver the amount of data needed at the times at which it is needed to utilize the capacity of the various PLPs efficiently.

The Scheduler shall use the Data Source Control Protocol to direct Data Sources to provide data in specific target bit rate ranges and to throttle those Data Sources to efficiently fill PLP capacities. Via the DSCP, real-time control messages are sent to Data Source IP control addresses

as specified by the System Manager. Data Sources can be multiplexers, encoders, servers, and the like.

Messages between Schedulers and Data Sources can include:

- Capabilities inquiries and responses
- Target bit range instructions from a Scheduler to a Data Source
- Throttling instructions from a Scheduler to a Data Source to make fine adjustments of data delivery

The control protocol for communication between the Scheduler and Data Sources shall use at least the fields shown in Table 5.2. The Scheduler configures the Physical Layer on a per PLP basis, therefore the DSCP has a link to a Data Source for each PLP. The DSCP can accommodate seamless changes that do not affect service.

Table 5.2 Real Time Control Field Definitions

Syntax	No. of Bits	Format
rt_control () {		
bit_rate_capacity	18	uimsbf
bit_rate_granularity	10	uimsbf
bit_rate_target	18	uimsbf
bit_rate_max	18	uimsbf
}		

bit_rate_capacity field shall indicate to the Scheduler the Data Source bit rate capacity (e.g., minimum and maximum bit rates) in units of kbps.

bit_rate_granularity field shall indicate to the Scheduler the Data Source step size in bit rate in units of kbps. Rate of change in the PLP can be from a few kbps to tens of Mbps depending on input coding rates and PHY frame configurations. The fastest rate of change depends on the shortest allowed frame length, which is 50msec.

bit_rate_target field shall indicate the desired bit rate for Scheduler capability in handling PLP payloads. Range of targets shall be between 1 kbps and 157 Mbps with units of kbps.

bit_rate_max field shall indicate the maximum bit rate value not to be exceeded into the Scheduler in units of kbps.

6 COMMON TUNNELING PROTOCOL (CTP)

The Data Source Transport Protocol (DSTP), ALP Transport Protocol (ALPTP) and STL Transport Protocol (STLTP) all rely on the capability of encapsulating packets of some form specific to the protocol into an RTP/UDP/IP packet stream. The RTP/UDP/IP packets are similar such that a common tunneling protocol can be defined that is common to all three protocols. In each case, source packet streams are combined into a single *Tunnel Packet Stream*. The Common Tunneling Protocol (CTP) hides the details of the various *Tunneled Packets* by treating them and any associated metadata as opaque payload data. The Common Tunnel consists of an RTP/UDP/IP unicast or multicast IPv4 packet stream based on the SMPTE ST 2022-1 [8] standard whose payloads shall contain one or more Tunneled Packet streams. The *Tunneled Packets* are also referred to as *inner* packets while the *outer* packets are referred to as *Tunnel Packets*.

Two of the protocols, DSTP and ALPTP, described in this document specify information headers that precede their respective Tunneled Packets, specifically, Data Source Tunneled

Packets which are UDP/IP packets and ALP Tunneled Packets which are ALP Link Layer packets, respectively. For convenience, this section refers to these combined information header and corresponding packets as simply *Tunneled Packets* with the information headers treated as part of the packet. In fact, this is appropriate since syntax and semantics of the internal packets should not be known to the tunnel except when framing to determine where each Tunneled Packet starts within the RTP/UDP/IP payload of the Tunnel Packet.

6.1 Overview

Tunnel Packets are distributed over UDP/IP multicast or unicast using a modified RTP [6] protocol. UDP/IP multicast is used successfully in many industries to deliver redundant video streams over copper and fiber physical layers. Routing and redundancy can be accomplished using IGMPv3 Source-Specific Multicast (SSM) (RFC 4607) [17]. Examples of various network configuration topologies are described in the Network Configuration Examples found in Annex B.

Since UDP does not guarantee packet order or, in severe cases, delivery, RTP shall be used to maintain and allow reconstruction of the data source tunnel packet order as well as to detect loss in the network. All Tunnel Packets shall be carried with RTP/UDP/IP multicast IPv4 protocol using the RTP header modifications described in Section 6.2.1.

Since the multiple Tunnel Streams can be received through networks of varying reliability, the Common Tunneling Protocol allows SMPTE ST 2022-based ECC to be applied. If ECC is not applied, Tunnel Packets shall still be defined with a fixed packet size to simplify implementations. A depiction of the Tunnel encapsulation is shown in Figure 6.1.

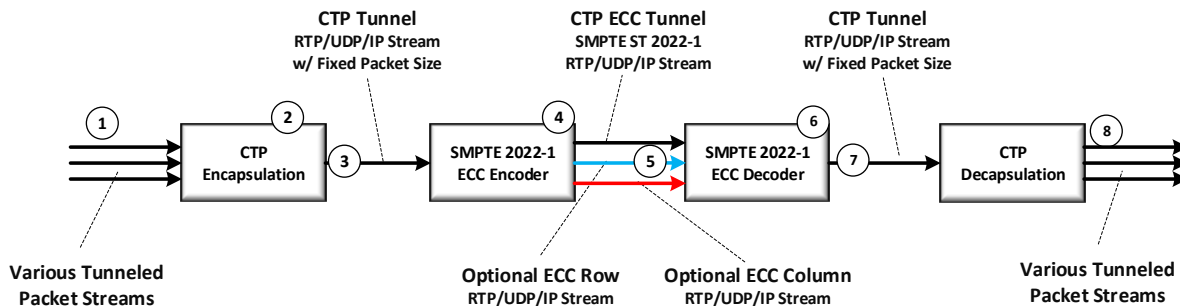


Figure 6.1 CTP Tunnel Encapsulation Process

The following paragraphs describe each of the callouts, (1) through (8), in Figure 6.1. Items (1) through (5) are further detailed in the ECC Encoding Process Section 6.3.1 below.

- 1) The multiple paths represent the Tunneled Packet Streams which have a variety of forms depending on the actual transport protocol. For DSTP, the input is one or more UDP/IP data source streams with metadata information headers. For ALPTP, the packets are ALP packets with metadata information headers. For STLTP, the tunneled packet streams are RTP/UDP/IP Streams that are generated for each PLP as well as the preamble, timing and management and security streams described in Sections 9.2.1, 9.3.1, and 9.4.5, respectively.
- 2) The CTP Encapsulation function is configured to accept packets from multiple packet streams to be tunneled.
- 3) The Tunneled Packets are grouped into fixed-size payloads to accommodate the SMPTE ST 2022-1 ECC process [8]. The CTP RTP fields are defined to allow the Tunneled Packet

Streams to be easily recovered and forwarded (refer to Section 6.3.2 below). The size of the fixed-size ST 2022-1 packets is not defined by this standard and is assumed to be configurable. It is expected that the packet size is within the typical 1500 MTU byte range limit. Note that, the larger the packets and the number of rows and columns defined, the longer the latency when performing ECC calculations. The resultant fixed-size RTP packets with the Tunneled Packet payloads are referred to as the Tunnel Packets.

- 4) The process described in 6.3.2 below buffers multiple fixed-size payloads and performs XOR operations on groups of Tunnel Packets producing additional ECC packets. The level of error correction can be configured, including the number of columns ('L') and the number of rows ('D'). Both Level A and Level B configurations, described in [8], shall be supported, with a preference for use of the 2D scheme (Level B). The minimum packet array size supported shall be 256 packets, which shall be configurable into any valid arrangement according to ST 2022-1 [8].
- 5) The Tunnel Stream and up to two ECC packet Streams are sent to the SMPTE ST 2022-1 ECC Decoder using up to three RTP/UDP/IP streams on three separate ports. If the tunnel Stream is sent on IP address 239.xxx.xxx.xxx, with port N, then the first ECC Stream is sent on the same address 239.xxx.xxx.xxx, with port N+2, and the second ECC Stream is also sent on the same address 239.xxx.xxx.xxx, with port N+4. This is in accordance with SMPTE ST 2022-1 ([8] §8.1, ¶6).
- 6) The Tunnel Stream and the associated ECC streams are processed by the SMPTE ST 2022-1 ECC Decoder. Note that the decoder detects missing packets from the RTP sequence numbers on the Tunnel Stream and reconstitutes them from the ECC Stream(s). The decoder also reorders any packets as they are being received within the constraints of its internal buffering.
- 7) The resultant Tunnel Packet Stream is an in-order collection of fixed-size packets identical to the input Stream described in step (3) above.
- 8) The CTP Decapsulation function extracts each Tunneled Packet Stream from the payloads of the Tunnel Packets. The Tunnel Packet RTP headers contain information allowing the Tunneled Packet Streams to be reframed.
- 9) The reconstituted IP and associated metadata packets, if applicable, then are forwarded to the next stage of the Transmission system.

6.2 RTP/UDP/IP Multicast Considerations

In addition to the context described above, RTP/UDP/IP Multicast with SSM may be used when moving multicast traffic between Data Producers and Data Consumers. Here a 'Data Producer' is defined as a device or process generating an IP multicast, and a 'Data Consumer' is a device or process receiving the IP multicast. Since RTP defines a 'sequence number' field, the various packets can be recovered and re-sequenced within each Data Consumer. This allows packets to be reordered, duplicate packets to be ignored, and packet loss to be detected, possibly avoiding the problems sometimes encountered by UDP datagram delivery.

In redundant systems, the Data Consumer must be aware of the several sources of contribution multicast streams. For each contribution stream, an IGMPv3 SSM 'join' is issued by the Data Consumer for a specific multicast destination IP address, destination IP port number and IP source address. The combination of addresses and port numbers is referenced as a 'tuple'. The router only sends UDP packets matching the specified tuple requested in the IGMP join to the Data Consumer's physical IP connection. If the Data Consumer detects loss or problems with the current

packet stream, it can release that stream and join another using the same IP destination address and port with a different IP source address. The Producers, in this case, operate independently and with no knowledge of the redundant topology.

6.2.1 Multicast Addressing

To use multicast addressing, a CTP stream shall be broadcast using a multicast address in the range 239.0.0.0 through 239.255.255.255, which is within the Organization Local Scope range defined by IETF. The actual addresses used are implementation-dependent and are defined by configuration.

6.3 Common Tunneling Protocol Design

The three protocols described in this document transport the various content, control and metadata that make up an ATSC 3.0 broadcast. As indicated in the introduction to Section 6, each of the protocols relies on the Common Tunneling Protocol (CTP) to hide various aspects of the individual protocol so that they may all be treated in a similar fashion when routing through a broadcast studio. The CTP shall encapsulate these various tunneled streams into an RTP/UDP/IP Tunnel Stream based on SMPTE ST 2022-1 [8]. SMPTE ST 2022-1, in turn, extends the RTP [6] protocol with Generic ECC Payload Format. RTCP and other control Streams are explicitly excluded by SMPTE ST 2022-1 ([8] §7.1, ¶2).

In addition to the constraints defined by [8], this document changes the definition of some of the standard RTP header fields to enable internal payload framing. The **marker (M)** bit along with the repurposed **SSRC_ID** field, defined herein to provide various signaling including the **packet_offset**. These fields and the payload framing mechanism for the Common Transport Protocol RTP header are described in Table 6.1.

Table 6.1 RTP Header Field Definitions for Common Transport Protocol

Syntax	No. of Bits	Format
RTP_Fixed_Header() {		
version (V)	2	'10'
padding (P)	1	bslbf
extension (X)	1	bslbf
CSRC_count (CC)	4	'0000'
marker (M)	1	bslbf
payload_type (PT)	7	uimsbf (Table 6.2)
sequence_number	16	uimsbf
if (PT == DSTP or PT == ALPTP) {		
timestamp_min()	32	Table 6.3
}		
else if (PT == STLTP) {		
timestamp()	32	Table 9.2
} else {		
reserved	32	for (i=0; i<32; i++) '0'

Syntax	No. of Bits	Format
<pre> } protocol_version if (PT == STLTP) { redundancy number_of_channels reserved } else { reserved } packet_offset } </pre>	<p>2</p> <p>2</p> <p>2</p> <p>10</p> <p>14</p> <p>16</p>	<p>Uimbsf</p> <p>Uimbsf</p> <p>Uimbsf</p> <p>for (i=0; i<10; i++) '0'</p> <p>for (i=0; i<14; i++) '0'</p> <p>Uimbsf</p>

Please refer to [6] for base definitions of the fields described in Table 6.1 and refer to [8] Section 7.1 for constraints on those base definitions.

version – Indicates the version number of the RTP protocol. ‘2’ is currently defined by [6].

padding – Indicates, if set to ‘1’, that padding may be included in the payload. No padding is included if the value is ‘0’. Padding shall be supported as specified in [6].

extension – Indicates, if set to ‘1’, that an RTP extension follows the header. If set to ‘0’, no header extension is included. If a header extension is included, all CTP packets must have the same length extension as dictated by Section 7.1 of [8].

CSRC_count – Provides the count of additional contributing source identifier fields. ST 2022-1 [8] requires this field to be set to ‘0’.

marker – Indicates, if ‘1’, that a Tunneled Packet starts within the payload of this RTP tunnel packet. When the **marker** bit is set to ‘1’, the first byte of the first Tunneled Packet starting within the payload shall be indicated by the value of the **packet_offset** field. Subsequent Tunneled Packets can be found by using the length fields within each Tunneled Packet, depending on the protocol, DSTP, ALPTP or STLTP. Typically, bytes preceding the first Tunneled Packet start will be a continuation of the last Tunneled Packet of the previous Tunnel Packet.

payload_type – The values used for the various Tunnel streams defined in this standard shall be as specified in Table 6.2.

Table 6.2 Code Values for **payload_type**

payload_type	Abbreviation	Meaning
81 ('1010001')	DSTP	Data Source Transport Protocol Tunnel Packets and Information Headers
82 ('1010010')	ALPTP	ALP Transport Protocol Tunnel Packets and Information Headers
97 ('1100001')	STLTP	STL Transport Protocol Tunnel Packets

sequence_number – As per [6], the **sequence_number** shall increment by one, modulo 2^{16} or 65536, for each packet with the same IP address and port tuple. The initial **sequence_number**

should be randomized, although this is not important in this setting since the stream will not be restarted frequently nor is it expected to be on a public network. Since UDP/IP does not guarantee ordered delivery and may even duplicate packets, the **sequence_number** can be used to reconstitute any stream as produced and to detect loss, allowing a backup stream to be selected or ECC to be applied if enabled.

timestamp_min() – When the **payload_type (PT)** value is either '81' (DSTP) or '82' (ALPTP), the **timestamp_min** value shall contain the earliest time specified in any of the tunneled packet information headers (see Section 7 and Section 8). If no tunneled packet information headers are present, then this value shall be set to '0'. Providing the earliest time allows the scheduling system to prioritize packets from multiple DSTP input streams. The **timestamp_min** value shall be formatted according to the short-form of NTP specified in RFC 5905 [13] as shown in Table 6.3. Note that this field replaces the SSRC_ID field specified in RFC 3550 [6]. The semantic definitions of the fields may be found in paragraphs following the table.

Table 6.3 Timestamp Field Definitions for DSTP and ALPTP

Syntax	No. of Bits	Format
timestamp_min () {		
seconds	16	uimsbf
fraction	16	uimsbf
}		

seconds shall carry a value equal to the 16 least significant bits (LSBs) of the seconds portion of the UTC time value [14] of the targeted Bootstrap Reference Emission Time.

fraction shall carry a 16-bit fractional seconds value of the UTC time value [14] of the targeted Bootstrap Reference Emission Time – allowing a resolution of approximately 15 microseconds.

The **timestamp_min** field within the RTP header allows a Data Source or ALP encapsulation system to specify schedule constraints on delivery of packets. Methods for determining the timestamp values are outside the scope of this standard, but provision is made herein to enable respecting the constraints without requiring inspection of the contents of the packets coming from Data Sources and being transported to the ALP layer as well as ALP packets being carried to the Scheduling system.

Please note that NTP is based on UTC [14] and, thus, is adjusted for leap seconds. There is no adjustment for leap seconds in emission timing, which is based purely on TAI seconds and fractional seconds [15]. See Section 9.3.2 for more information on Bootstrap emission timing.

timestamp() – If the **payload_type** is '97' indicating an STLTP Tunnel Packet, the **timestamp()** field shall be defined according to Table 9.2 and its subsequent semantic definitions. This value shall indicate the latest time at which the start of the payload may be delivered. A value of '0' shall denote that the packet may be delivered with best effort. Note that this field replaces the SSRC_ID field specified in RFC 3550 [6].

protocol_version – shall indicate the version number of the outer Tunnel Packet protocol. The current version is '01'.

redundancy – When the **payload_type** is set to '97' indicating an STLTP packet, this field shall indicate the number of redundant Streams of outer Tunnel Packets sent over the STL. The value '0' shall indicate that there is no redundancy of outer Tunnel Packets, and values greater

than ‘0’ shall indicate that there is redundancy of outer Tunnel Packet Streams over the STL and the number of such redundant Streams.

number_of_channels – When the **payload_type** is set to ‘97’ indicating an STLTP packet, the **number_of_channels** field shall indicate the number of broadcast channels minus one transported by the outer Tunnel Packet Stream. When the value is set to ‘0’, an outer Tunneling Stream shall contain one group of inner Tunneled Packet Streams, using the port range 30000 – 30066. Note that a group of inner Tunneled Packet Streams is composed of the Preamble, Timing and Management, Baseband Packet, and Security Data Streams required for transmission of one RF channel. When the value is set to ‘1’, an outer Tunnel Data Stream shall contain two groups of inner Tunneled Packet Streams. When the value is set to ‘1’, the port range of one group of inner Tunneled Packet Streams (corresponding to one RF channel) shall be 30000 – 30066, and the port range of the other group of inner Tunneled Packet Streams (corresponding to the other RF channel) shall be 30100 – 30166. See Section 9.6 for an example of use of this field.

packet_offset – If the **marker** bit equals ‘1’, this field shall define the payload offset of the first Tunneled Packet. The field shall indicate the number of bytes *after the RTP header and extension* before the first Tunneled Packet starts within the tunnel payload. If the Tunneled Packet starts immediately after the header and extension (if any), then **packet_offset** shall be ‘0’. Figure 6.2 provides a graphical example of how Tunneled Packets are encapsulated within the fixed-size SMPTE ST 2022-1 CTP RTP Tunnel Packets. Note that only the position of the first packet within the tunnel payload is required since all Tunneled Packet types contain a length field that can be used to calculate the offset of subsequent packets.

6.3.1 CTP RTP Encapsulation Example

Following is an example that shows how encapsulation is done using RTP to create a complete Tunneled Data Stream for SMPTE ST 2022-1 ECC processing. It examines the case of a single stream. Note that the protocols derived from this Common Tunneling Protocol may request multiple, simultaneous Tunnel Streams. The example is broken into two diagrams that show the overall structure of a CTP Stream (in Figure 6.2) and the details of the packing process, including the locations of the several packets of the different layered protocols within the Stream (in Figure 6.3).

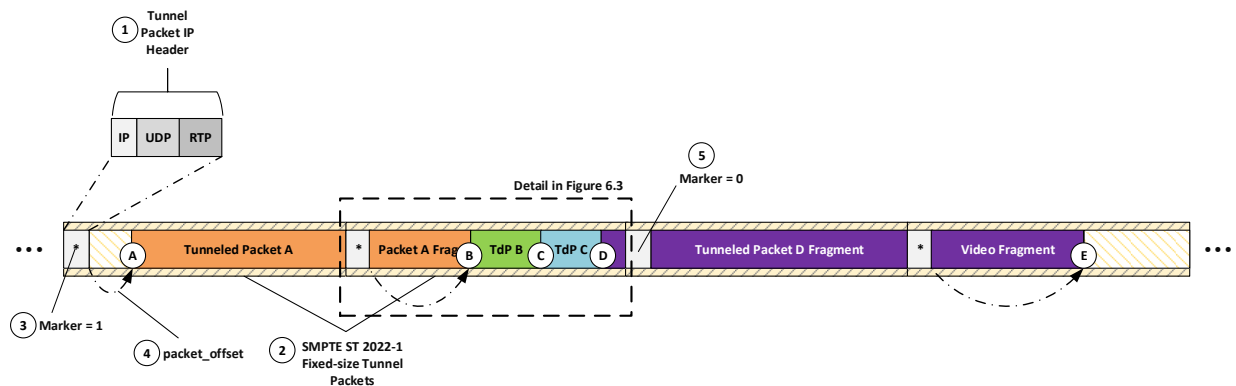


Figure 6.2 Common Tunneled Packet Stream

The following paragraphs describe each of the callouts in Figure 6.2. Note that this is an example of a potential implementation to illustrate how the fixed-size Tunnel Packets are

constructed according to this specification. Actual implementations could differ substantially; the resulting packet Stream, however, would have the same semantic organization.

- 1) A standard RTP/UDP/IP header precedes each of the CTP fixed-size Tunnel Packets that contain the content to be protected by the ST 2022-1 ECC process. All Tunnel Packet headers are fixed size even if they include an extension header as required by ST 2022-1.
- 2) The example in Figure 6.2 is presented from the decoding standpoint, showing how the Tunnel Packets would be perceived at the consumer of the Tunnel Stream, assuming that the consumer started with an arbitrary packet in a continuous stream of Tunnel Packets. Each Tunnel Packet in the stream (2) is of a fixed size and contains a header and payload that may or may not contain the start of a Tunneled Packet (TdP).
- 3) If the **marker** bit (3) of a Tunnel Packet is set to '1', signified in the diagram by an asterisk in the RTP header box, then **packet_offset** (4) is used to indicate the byte where the first Tunneled Packet starts within the payload of the Tunnel Packet.
- 4) The **packet_offset** is the number of bytes after the RTP header and possible extension where the first byte of the first tunneled packet resides within the Tunnel Packet payload (A). This field is used only if the **marker** bit is set to '1'. Note that when more than one packet starts in the Tunnel Packet payload, only the first is indicated by the **packet_offset**. This is shown in Figure 6.2 at Tunneled Packets (B) and (C). The start of Tunneled Packet B (B) is indicated by the **packet_offset** as being immediately after the end of the remainder of Tunneled Packet (A). The start of Tunneled Packet (C) can be determined using the length of Tunneled Packet (B) found within the Tunneled Packet B. Note that every protocol derived from the CTP has a length field defined somewhere in each Tunneled Packet header. Of course, the start of packet (B) also could be determined by using the total length of Tunneled Packet (A), which started in the previous ST 2022-1 fixed-size Tunnel Packet. Robust decoders likely would use both the length of the previous packet and the **packet_offset** to verify the encapsulation process and that the payload packets were correctly formed. This mechanism of framing the Tunneled Stream Packets allows the consumer to recover lost connections quickly since only one or two fixed-size Tunnel Packets are required to begin extracting the Tunneled Packets.
- 5) The Tunneled Packet, the start of which is labeled (D) in the figure, shows a large packet spanning three ST 2022-1 fixed-size Tunnel Packets. In this case, the middle ST 2022-1 Tunnel Packet does not have the **marker** bit set to '1' and the **packet_offset** field is set to '0' and is ignored.

Figure 6.3 shows the details of the encapsulated Tunnel Packet outlined by the dashed line in Figure 6.2. In the top row of Figure 6.3 are packets from three separate Tunneled Packet Streams. In the center row of Figure 6.3 are the details of the construction of the Tunnel Packet that contains the Tunneled Packets or segments of Tunneled Packets from the top row. As shown by the bottom row, the middle row is a more detailed version of the Tunnel Packet structure shown within the dashed line in Figure 6.2.

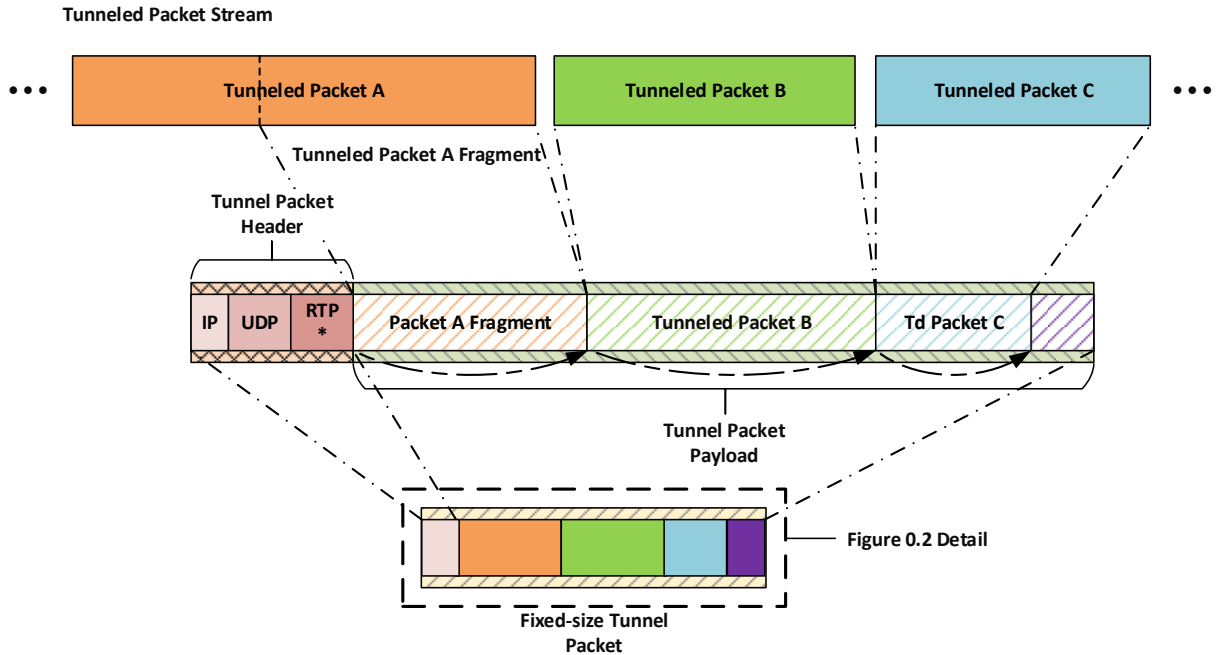


Figure 6.3 Data Source Tunneled to Tunnel Packet Relationship

Starting from the top row of Figure 6.3, the three Tunneled Packet Streams each has its own procession of packets, the details of which are defined by the particular protocol type. Either complete packets or segments of packets from the Tunneled Packet Streams are multiplexed together on the second row to form the payload of a Tunnel Packet, and another header is prepended to the payload to form a new RTP/UDP/IP packet. As shown in the bottom row, that packet is the one in dashed lines in Figure 6.2.

Looking at the middle row of Figure 6.3 in more detail, there is one RTP/UDP/IP header for the Tunnel Packet, two complete Tunneled Packets and two Tunneled Packet fragments. The first fragment, on the left, is the remainder of Tunneled Packet A left over from the previous Tunnel Packet. The second and third blocks represent complete Tunneled Packets B and C that are fully encapsulated within the Tunnel Packet payload. Finally, the beginning of another Tunneled Packet starts immediately after Tunneled Packet C and does not complete so must be continued in the next Tunnel Packet (not shown).

In the sequence of payload packets and segments described within the payload of the Tunnel packet, there is a pointer from the Tunnel Packet RTP header to the start of the first bit of the Tunneled Packet B. The Tunnel Packet RTP header **packet_offset** points there and the **marker** bit is set to '1'. Since the RTP header pointer did not point to the first bit following itself, the implication is that the first data segment within the Tunnel Packet payload is a segment of a Tunneled Packet that was not completed in an earlier Tunnel Packet. When examining the Tunnel Packet RTP header, the consumer can determine both that the first data within the Tunnel Packet payload is a segment of a previously incomplete packet and where the next Tunneled Packet begins within the Tunnel Packet. By examining the length value contained within the next Tunneled Packet, the consumer can locate the next header within the Tunnel Packet payload. This process can continue until the consumer finds a Tunneled Packet header that indicates a length longer than the remaining space available in the Tunnel Packet payload, which implies that there will be at

least a segment of the final Tunneled Packet contained within the next Tunnel Packet having the same port and address as the current Tunnel Packet.

6.3.2 Example SMPTE ST 2022-1 ECC Encoding Process

Figure 6.4 provides a detailed example diagram of how the various Tunneled Packet Streams are encoded into CTP Tunnel Packets with a two-dimensional arrangement of ST 2022-1 ECC packets [8]. This exemplifies some of the steps shown in Figure 6.1. The process shown in Figure 6.4 implements steps 3, 4, and 5 of Figure 6.1. This example shows a combination of the CTP Encapsulation function and the ECC Encoder function operating on a shared memory area. A similar decode process occurs where these packets are gathered into memory, ECC is used to reconstitute missing packets, and the Tunneled Packets are extracted directly from the in-memory buffers. Note that SMPTE ST 2022-1 may also be used in a one-dimensional arrangement, where the ECC encoding operation is applied to a single row of packets.

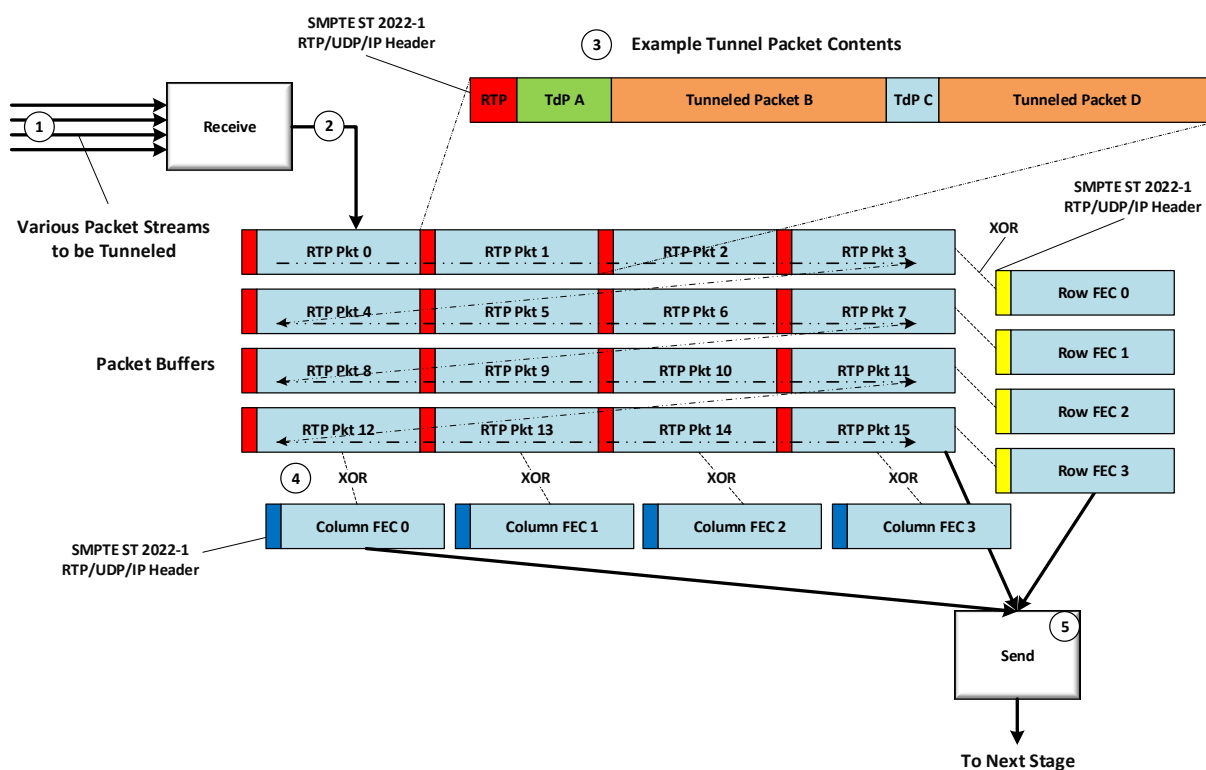


Figure 6.4 Example ECC encoding process diagram.

The following paragraphs describe each of the callouts from Figure 6.4.

- 1) The multiple paths represent the RTP/UDP/IP Tunneled Packet Streams. The content of these streams depend on which protocol is being processed: DSTP, ALPTP or STLTP.
- 2) As Tunneled Packets are received from the input interface, they are placed in memory as they are received. A conceptual memory buffer is preconfigured to contain as many Tunnel Packets as required to meet the ECC packet columns ('L') and rows ('D') configuration with the Tunnel Packet payloads and placeholder RTP headers. The appropriate **marker** bit and **payload_offset** fields will be updated in these headers as the Tunneled Packets are copied into the buffers.

- 3) The “Example Tunneled Packet Contents” diagram a conceptual model of how Tunneled Packets are placed into individual Tunnel Packet buffers. Tunneled Packets are placed into the buffer sequentially as received. If a Tunneled Packet exceeds the length of the current Tunnel Packet buffer, the portion of the Tunneled Packet that fits are placed into the buffer and the remainder placed into the next buffer. The **payload_offset** field of the next Tunnel Packet buffer RTP header is set to indicate the byte of the next Tunneled Packet starts within the Tunnel Packet payload and the **marker** bit is set to indicate that a payload (i.e., Tunneled Packet) start is present within the particular buffer. An example of this packing process is described in detail in Section 6.3.1 above.
- 4) When a column or row is completed, the ECC function (XOR per [8]) is performed creating an SMPTE ST 2022-1 ECC packet for the column or row, respectively. The headers and content of these packets are described in [8].
- 5) The ECC packets along with the Tunnel Packets are sent to the Send function per the guidelines defined in [8]. There are several packet interleaving options defined in [8] that can be configured based on the number of columns and rows defined and the system latency desired.

7 DATA SOURCE TRANSPORT PROTOCOL

The Data Source delivery structure uses the Data Source Transport Protocol (DSTP), which applies between the various Data Sources and an ALP encapsulation function. There may be multiple data sources for each PLP, each carried on a separate packet stream. The packets produced by the data sources are UDP/IP multicast packets intended to be consumed by receivers. Thus, their construction follows rules defined by A/331 [4] including their IP addressing which is not conducive to efficient broadcast studio routing.

The Data Source Transport Protocol hides the details of the Data Source packets by encapsulating them in a Data Source Packet Tunnel derived from the Common Tunneling Protocol (see Section 6). The Data Source Packet Tunnel consists of an RTP/UDP/IP multicast or unicast IPv4 packet stream whose payloads shall contain one or more Data Source packet streams. In addition, the DSTP includes a Tunneled Packet Information Header prior to each Data Source Tunneled Packet that contains information for how to route and schedule each packet without the need to examine the actual packet itself. The Data Source packets along with their associated Information Headers are referred to as the Data Source *Tunneled* Packets or *inner* packets while the *outer* packets are referred to as Data Source *Tunnel* Packets. The details of encapsulation of the Data Source Packets into the Packet Tunnel are described in the Common Tunneling Protocol Section 6.

7.1 Overview

The system diagram in Figure 4.2 provides the functional context of the data source transport protocol. It is possible that an actual implementation may be much more complicated, requiring multiple sources of data source packets. Such sources could be for redundancy or could be external to the studio. To flexibly accommodate various network topologies, there may be multiple Data Source Tunnel Streams sent to an ALP Generator function. The Data Source Tunnel Streams act as transports for carrying one or more Data Source Tunneled Streams from various Data Source generation devices. The ALP Generator function shall accept one or more DSTP-formatted data streams and provide mapping of the tunneled Data Source packets into the appropriate PLP data streams. Figure 7.1 provides a depiction of three Data Source generators supplying multiple Data

Source Streams through three DSTP Packet Tunnel Streams into an ALP Generator. The ALP Generator’s Data Source Mapping configuration, supplied by the System Manager, and the associate tunneled packet information headers dictate how the Tunneled Packets are mapped to the appropriate PLP as shown by the dashed lines within the ALP Generator in the figure. Note that this standard does not constrain the number of DSTP-formatted input streams defined for receipt by the ALP generator. The DSTP descriptions below apply to a single Data Source Tunnel Packet stream.

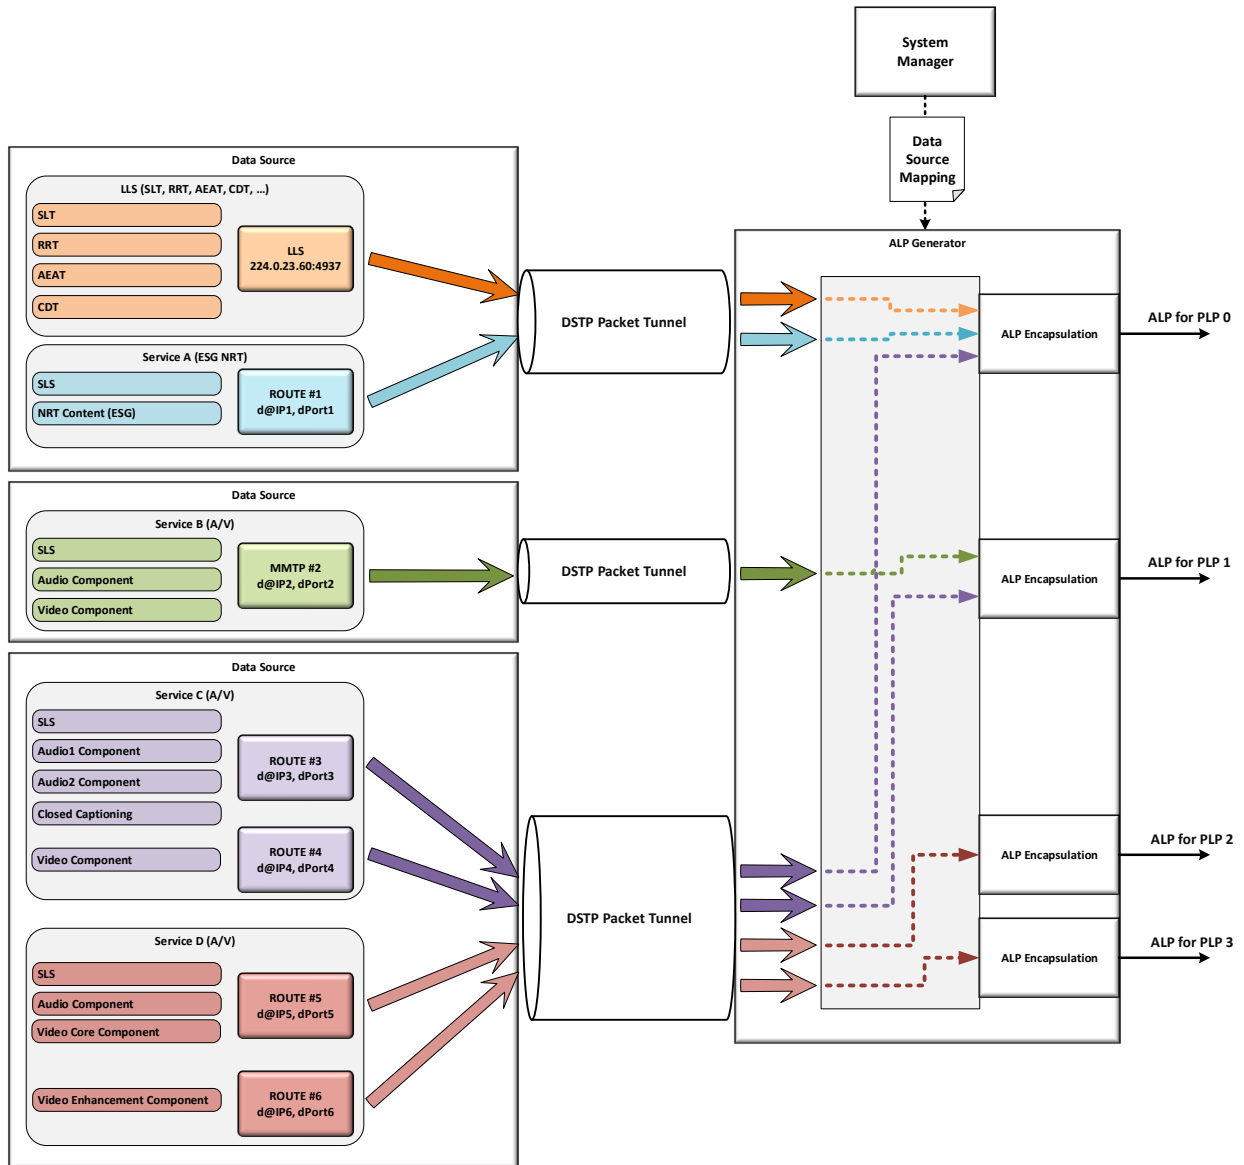


Figure 7.1 Multiple Data Source Packet Tunnels Concept

7.1.1 DSTP Mapping Configuration Description

The DSTP Mapping Configuration provided to the ALP Generator shall be represented as an XML document containing a **DSMapping** root element that conforms to the definitions in the XML schema that has namespace:

tag:atsc.org,2016:XMLSchemas/ATSC3/Delivery/DS_MAPPING/1.0/

The definition of this schema is in an XML schema file, *DSMapping-1.0-2019mmdd.xsd*, accompanying this standard, as described in Section 3.6 above. The XML schema xmlns short name should be "dsm".

The schema provided in Table 7.1 shall be used to capture the Data Source Mapping configuration information. While the indicated XML schema specifies the normative syntax of the **DSMapping** element, informative Table 7.1 below describes the structure of the **DSMapping** element in a more illustrative way. The specifications following the table give the semantics of the elements and attributes.

Table 7.1 Data Source Mapping XML Format

Element or Attribute Name	Use	Data Type	Short Description
DSMapping			Root element of the Data Source Mapping configuration.
DSTunnel	1..N		Defines one or more incoming Data Source Tunnel streams.
@destAddr	1	IPv4address	The destination IPv4 address of the incoming DSTP tunnel stream.
@destPort	1	unsignedShort	The destination IP port of the incoming DSTP tunnel stream.
@srcAddr	0..1	IPv4address	The optional source IPv4 address of the incoming DSTP tunnel stream.
@igmpVersion	0..1	unsignedByte 2..3	The optional IGMP version number. Absence of this attribute indicates IGMP is not in use.
@defaultPLP	0..1	unsignedShort	Optionally, supplies the PLP where any unspecified tunneled packet stream within the Data Source Tunnel should be routed.
DSTBackup	0..N		An optional collection of backup Data Source Tunnel streams.
@srcAddr	1	IPv4address	The source IPv4 address of an optional incoming DSTP backup tunnel stream.
TPS	0..N		Element defining each Tunneled Packet Stream.
@destAddr	1	IPv4address	Defines the IPv4 destination address of a tunneled packet stream within the DSTP tunnel stream. This address is matched with the value found in the Tunneled Packet Information Header (see Section 7.2.1).
@destPort	1	unsignedShort	Defines the IPv4 destination port of a tunneled packet stream within the DSTP tunnel stream. This port is matched with the value found in the Tunneled Packet Information Header (see Section 7.2.1).
@plp	1	unsignedShort	Supplies the PLP where the associated tunneled packet stream should be routed.

The following text specifies the semantics of the elements and attributes in the Data Source Mapping configuration fragment.

DSMapping – The root element of the Data Source Mapping configuration structure.

DSTunnel – The `DSTunnel` element shall contain the configuration of each of the Data Source Tunnel stream to be received and processed by the ALP Generator. Each `DSTunnel` element contains the address of the incoming DSTP packet stream with mapping information for each tunneled stream within it.

@dstAddr – The required destination address of the DSTP Tunnel Stream. This may be a multicast or unicast address.

@dstPort – The required destination port of the DSTP Tunnel Stream.

Note that according to A/331 Section 6.1 [4], all destination addresses and ports must be unique across the broadcast. Therefore, each Data Source packet Stream received by an ALP Generator must have a unique address and may only be routed to a single PLP.

@srcAddr – The optional source address of the DSTP Tunnel Stream. This attribute is required if IGMP version 3 is being used by the routing system to identify the desired source of the DSTP packet stream.

@igmpVersion – The optional `@igmpVersion` attribute allows the version of IGMP to be specified. If '2' is specified, then IGMP version 2 is being used to route the DSTP Tunnel Stream. A value of '3' indicates that IGMP version 3 is being used and Source-Specific Multicast (SSM) is available. If this attribute is not provided, the IGMP is not being used.

@defaultPLP – The optional `@defaultPLP` attribute shall provide the PLP ID to which any unspecified tunneled packet streams should be routed. An unspecified tunneled packet stream is any set of tunneled packets which have a Tunneled Packet Information Header that has `dest_address` and `port_number` values that do not correspond to any of those identified by `TPS` elements within the `DSTunnel` element. Packets from each unspecified tunneled packet stream shall be routed to the PLP specified by this attribute.

DSTBackup – The optional `DSTBackup` element defines a backup Data Source Tunnel stream that may be sourced from an alternate Data Source. The single, required `@srcAddr` attribute specifies the IP source address of the backup Data Source Tunnel.

@srcAddr – The required source address of the backup DSTP Tunnel Stream. This attribute is used in IGMP version 3 to join the backup DSTP Tunnel Stream if the primary Tunnel Stream is not available.

TPS – The `TPS` element contains three attributes that specify the tunneled packet stream and where that stream shall be routed to. Note that if the `TPS` structure is missing then all packets within the tunnel are routed to the PLP specified in the `@defaultPLP` attribute.

@dstAddr – The required IPv4 destination address of the tunneled packet stream within the DSTP tunnel stream. This address is matched with the value found in the Tunneled Packet Information Header (see Section 7.2.1) to determine into which PLP packets from this stream should be placed.

@dstPort – The required destination port of the tunneled packet stream within the DSTP tunnel stream. This port is matched with the value found in the Tunneled Packet Information Header (see Section 7.2.1) to determine into which PLP packets from this stream should be placed.

@plp – The required `@plp` attribute shall provide the PLP ID to which the tunneled packet streams should be routed. Packets from the tunneled packet stream shall be routed to the PLP specified by this attribute.

7.2 DSTP Design

The data source packet streams transport the various signaling, audio, video, captioning and data that make up an ATSC 3.0 broadcast [4]. As indicated in the introduction to Section 7, the data sources create packets with the destination addresses and ports intended to be seen by receivers. To avoid conflicts and readdressing issues, the DSTP encapsulates these Data Source Packet streams in one or more Tunnel Streams based on the Common Tunneling Protocol described in Section 6.

The ALP Generator may receive one or more DSTP streams and differentiates various data sources from within the DSTP Tunnel Packet Stream by a mapping supplied through configuration. To avoid the ALP Generator being required to understand the contents of various Data Source Tunneled Packets, additional delivery metadata is included as a Tunneled Packet Information Header. Placing delivery metadata in a separate header for each packet within the tunnel allows the data source packets to remain opaque to the transmission system. The ALP encapsulation function extracts the Data Source Tunneled Packets and their associated Information Headers from each DSTP tunnel before placing them in the respective ALP streams. Note that the Information Headers are used as part of the mapping and ALP stream creation but are discarded by the ALP Generation function.

7.2.1 DSTP Tunneled Packet Information Header Definition

The DSTP defines a Tunneled Packet Information Header to avoid modification to the Data Source Tunneled Packets. This header provides information to help the ALP Generator route the various packets to the appropriate PLP as well as improve the scheduling performance of the Broadcast Gateway scheduling system when creating physical layer frames. A Tunneled Packet Information Header shall precede every tunneled packet that starts within the DSTP Tunnel Packet.

Table 7.2 provides the field definitions of the DSTP Tunneled Packet Information Header structure. The paragraphs following the table describe the fields within the table.

Table 7.2 DSTP Tunneled Packet Information Header

Syntax	No. of Bits	Format
information_header() {		
dest_address	32	uimbsf
port_number	16	uimbsf
length	16	uimbsf
group	16	uimbsf
type	8	uimbsf
random_access_point	1	bslbf
time_limit_flag	1	bslbf
if (type >= '1' && type <= '5') {		
wakeup_control()	2	Section 7.2.2
}		
else {		
reserved	2	'00'
}		
reserved	4	'0000'
if (time_limit_flag == '1') {		
timestamp()	32	Table 6.3
}		

```

}
|
|
|

```

dest_address – The unsigned, 32-bit **dest_address** shall specify the destination IP address of the tunneled packet following the information header. This value, along with the **port_number** and **group**, are used by the ALP Generator to map incoming tunneled packets to outgoing ALP encapsulated streams. Note that, in practice, this field is a copy of the destination address within the data source tunneled packet. It is provided here to avoid examining the tunneled packet in any way.

port_number – This unsigned, 16-bit value shall define the destination port number of the tunneled packet immediately following the information header within the tunnel packet payload. This value, along with the **dest_address** and **group**, are used by the ALP generator to map incoming tunneled packets to outgoing ALP encapsulated streams. Note that, in practice, this field is a copy of the destination port within the data source tunneled packet. It is provided here to avoid examining the tunneled packet in any way.

length – This unsigned, 16-bit value shall specify the number of bytes occupied by the tunneled packet immediately following the information header. The length is the total size of the tunneled packet, even if it is not fully contained within the current tunnel packet payload. Note that the length here is different than the IP header length contained within the tunneled packet since it also includes the length of any headers contained within the tunneled packet.

group – If the **type** is LLS, this unsigned, 16-bit value shall provide the groupID as defined in [4]. Otherwise, the field shall provide the serviceID associated with the tunneled packet stream. If the serviceID is provided, it shall be taken directly from the service defined within the associated SLT as specified in A/331 [4]. The groupID or serviceID is provided to allow the scheduling functions within the Broadcast Gateway to optimize delivery of various services.

type – This unsigned, 8-bit value shall define the type of packet immediately following the information header within the tunnel packet payload. Table 7.3 defines the types currently supported by the Data Source Transport Protocol.

Table 7.3 type Field Available Values

Type	Title	Description
0	Reserved	
1	LLS SLT	
2	LLS RRT	
3	LLS System Time	
4	LLS AEAT	
5	LLS OSN	
6	LLS CDT	
7-13	Reserved	Reserved for future types
14	LLS Signed Multi-Table	
15	LLS User Defined	
16	SLS	
17-63	Reserved	Reserved for future types
16	Primary Video	
17	Auxiliary Video	
18	Video Enhancement	
19-127	Reserved	Reserved for future types

128	Primary Audio	
129	Auxiliary Audio	
130-191	Reserved	Reserved for future types
192	Primary Application Data	
193	Guide Data	
194	Application Data	
195	AEI Data	Advanced Emergency Information Data
196-254	Reserved	Reserved for future types
255	Unspecified	Indicates that the associated packet stream type is not specified.

random_access_point – If set to ‘1’, indicates that the associated inner packet contains some or all of the payload containing a random-access point (RAP). If set to ‘0’, no random-access point is present within the packet. See Annex C for a discussion of random-access points and their meaning within the system. The **random_access_point** bit shall remain ‘1’ for all of the packets containing the portion of the stream comprising the random-access point. For example, for packets carrying a video stream, all of the packets containing the segment defined as the RAP would have the **random_access_point** bit set to ‘1’.

time_limit_flag – This flag shall be set to ‘1’ to indicate that a **timestamp()** field is present. A value of ‘0’ shall indicate that no **timestamp()** field is present.

wakeup_control() – A 2-bit field that communicates information needed to control the emission wakeup field. See Section 7.2.2 for a definition of this field. The **wakeup_control()** field shall only be interpreted if the tunneled packet type is one of the main LLS types with type values between ‘1’ and ‘5’.

timestamp_min() – This field shall be defined according to Table 6.3 and its subsequent semantic definitions.

7.2.2 Emergency Alert Wakeup RTP Controls

Emergency Alert Wakeup RTP fields in DSTP provide signaling information that indicates to the Scheduler how to manage the **ea_wake_up_1** and **ea_wake_up_2** bits in the Bootstrap. In the text that follows, these two bits are treated together as the “Emission Wakeup Field.” There are two flags provided in the **wakeup_control ()** field. The first, **wakeup_active**, indicates that a currently active Advanced Emergency Alert (AEA) message has the @wakeup attribute set to ‘true’. The flag shall reflect this condition on the DSTP RTP headers of all LLS packets generated by a given source, not just on packets containing an Advanced Emergency Alert Table (AEAT). In contrast, the **AEAT_wakeup_alert** flag shall reflect specific values within the current packet and apply only to packets containing an AEAT.

Table 7.4 RTP payload_type Wakeup Control Field Definition

Syntax	No. of Bits	Format
wakeup_control () {		
wakeup_active	1	bslbf
AEAT_wakeup_alert	1	bslbf
}		

wakeup_active – A 1-bit flag that, when set to ‘1’, shall indicate that the source providing the LLS data is requesting that the Emission Wakeup Field be non-zero. When set to ‘0’, this flag shall

indicate that the source providing the LLS data is requesting that the Emission Wakeup Field be set to a value of zero if no other source is requesting a non-zero value.

AEAT_wakeup_alert – A 1-bit flag that, when set to ‘1’, shall indicate that the packet contains a new or updated AEAT that carries a new Wakeup Alert. When set to ‘0’, it shall indicate that the packet does not contain an AEAT with a new Wakeup Alert. This bit shall only be ‘1’ when the **LLS_table_version** of the AEAT LLS_table () has been incremented and the AEAT contains an AEA element with the @wakeup attribute newly set to ‘true’. In other words, if a new AEA is added with @wakeup set to ‘true’ or a previous AEA has been updated with @wakeup set to ‘true’, then the **LLS_table_version** will be incremented and the **AEAT_wakeup_alert** shall be set to ‘1’.

The **wakeup_active** field controls whether the Emission Wakeup Field is ‘off’ (zero) or ‘on’ (non-zero). **AEAT_wakeup_alert** controls when the Emission Wakeup Field toggles through the ‘on’ (non-zero) states, i.e., 1, 2, and 3. (See Table 7.5 for an example of **wakeup_active** and **AEAT_wakeup_alert** usage.) It is expected that LLS tables, such as the SLT or SystemTime tables, will be updated relatively frequently, allowing timely control. There is no restriction, however, on providing extra LLS signaling to accommodate updating the Emission Wakeup Field. For example, the Data Source could emit an SLT specifically to return the Emission Wakeup Field value to zero.

Table 7.5 Example Wakeup Bit Controls

Time	LLS Table	Wakeup State	wakeup_active	AEAT_wakeup_alert	Emission Wakeup Field
t0	SLT	Off	0	0	00
t1	SystemTime	Off	0	0	00
t2	New AEA @wakeup=true	On	1	1	01
t3	SLT	On	1	0	01
t4	Same AEA	On	1	0	01
t5	SystemTime	On	1	0	01
t6	Updated AEA @wakeup=true	On	1	1	10
t7	SLT	On	1	0	10
t8	SLT	Off	0	0	00
t9	New AEA @wakeup=false	Off	0	0	00

Note that this example is for a single source of LLS. If there is more than one source of LLS, the Scheduler is responsible for combining those controls and setting the Emission Wakeup Field appropriately.

8 ALP TRANSPORT PROTOCOL (ALPTP)

The ALP delivery structure uses the ALP Transport Protocol (ALPTP), which applies between ALP sources and Broadcast Gateway inputs when the ALP encapsulation function is external to the Broadcast Gateway. There may be multiple Data Sources for each PLP, but individual ALP packet Streams carry only a single source data Stream for each respective PLP. Thus, any source multiplexing is done upstream of ALP encapsulation and header compression functions (see Section 7.1).

The ALP Transport Protocol hides the details of the ALP packets by encapsulating them in a ALP Packet Tunnel derived from the Common Tunneling Protocol (Section 6). The ALP Packet Tunnel consists of an RTP/UDP/IP multicast IPv4 packet stream whose payloads shall contain one or more ALP packet streams. In addition, the ALPTP includes a Tunneled Packet Information Header prior to each ALP Tunneled Packet that contains metadata for how to schedule the various packets into baseband packets for the transmission as well as the PLP ID to which the ALP packet is bound. The ALP packets along with their associated Information Headers are referred to as the *ALP Tunneled Packets* or *inner* packets while the *outer* packets are referred to as *ALP Tunnel Packets*. The details of encapsulation of the ALP Packets into the Packet Tunnel are described in the Common Tunneling Protocol Section 6.

8.1 Overview

The system diagram in Figure 4.2 provides the functional context of the ALP Transport Protocol (ALPTP). It is possible that an actual implementation may be much more complicated, requiring multiple sources of ALP packets. Such sources could be for redundancy or could be external to the studio.

For example, consider a Transmitter-sharing agreement where two Physical Layer Pipes (PLPs) are dedicated to two separate stations. The individual studios could produce separate ALP packet Streams that are then delivered for injection into the shared Transmitter. If one of the studios were separated geographically, some sort of robust connection would be required to deliver the ALP packet Stream to the Scheduler / framing subsystem.

The functionality of the ALP standard [5] is sufficient for an emission link-layer protocol but should not be burdened with additional functionality to support intra- and inter-studio routing and distribution.

For these reasons, the ALPTP is based on the Common Tunneling Protocol carrying ALP packets along with the associated Tunneled Packet Information Headers. The use of the Tunneled Packet Information Header to bind each ALP packet with a PLP means that any number of ALPTP Tunnel streams may be used between the ALP Generator and the Scheduler. However, to avoid issues with timing and scheduling, all ALP packets destined for a single PLP shall be carried in the same ALPTP Tunnel Stream. Furthermore, there is no limit on the number of PLPs that can be supported by a single ALPTP Tunnel Stream. The ALPTP streams may be unicast or multicast consistent with CTP.

8.2 ALPTP Design

ALP describes the link-layer encoding for emitting various types of packets of a broadcast [5]. As indicated in the introduction to Section 8 above, the ALP protocol should not be burdened with routing requirements within the broadcast facility since this routing information is irrelevant to broadcast receivers. To provide ordering, timing and framing information required to route ALP Packet Streams, the ALPTP encapsulates these ALP Packet Streams in a Tunnel Stream for each PLP based on the Common Tunneling Protocol described in Section 6.

8.2.1 ALPTP Tunneled Packet Information Header Definition

The ALPTP defines a Tunneled Packet Information Header to avoid modification or analysis of the ALP tunneled packets. This header provides information to help the scheduler associate the ALP packets with the targeted PLP as well as improve the scheduling performance of the Broadcast Gateway scheduling system when creating physical layer frames. An ALPTP Tunneled

Packet Information Header shall precede every tunneled packet that starts within each ALPTP Tunnel Packet.

Table 8.1 provides the field definitions of the ALPTP Tunneled Packet Information Header structure. The paragraphs following the table describe the fields within the table.

Table 8.1 ALPTP Tunneled Packet Information Header

Syntax	No. of Bits	Format
alptp_information_header() {		
length	16	uimbsf
alp_sid	8	uimbsf
plp_id	6	uimbsf
lls_flag	1	bslbf
lmt_rdt_flag	1	bslbf
random_access_point	1	bslbf
time_limit_flag	1	bslbf
if (lls_flag == '1') {		
wakeup_control()	2	Section 7.2.2
} else {		
reserved	2	'00'
}		
reserved	28	uimbsf
if (time_limit_flag == '1') {		
timestamp_min()	32	Table 6.3
}		
}		

length – This unsigned, 16-bit value shall specify the number of bytes occupied by the tunneled packet immediately following the information header. The length is the total size of the tunneled packet, even if it is not fully contained within the current tunnel packet payload. Note that the length here is different than the IP header length contained within the tunneled packet since it also includes the length of any headers contained within the tunneled packet.

alp_sid – This 8-bit unsigned integer field shall indicate a sub-stream identifier for the ALP packets carrying the multicast identified by the above four fields, src_ip_add, dst_ip_add, src_udp_port and dst_udp_port.

plp_id – This 6-bit unsigned integer field shall be used by the scheduler to map the incoming tunneled packet to the corresponding PLP (defined by L1D_plp_id in A/322 [3]).

lls_flag – A 1-bit flag that shall indicate that the ALP packet payload contains at least one instance of a Low-Level Signaling table in compliance with [4]. A value of '1' shall indicate that the packet contains a Low-Level Signaling table. A value of '0' shall indicate that no LLS is present in the ALP payload.

lmt_rdt_flag – A 1-bit flag that shall indicate that the ALP packet payload contains at least one instance of a Link Mapping Table or a RoHC Description Table. A value of '1' shall indicate that the packet contains an LMT or RDT. A value of '0' shall indicate that no LMT or RDT is present in the ALP payload.

random_access_point – If set to '1', indicates that the associated inner packet contains some or all of the payload containing a random-access point (RAP). If set to '0', no random-access

point is present within the packet. See Annex C for a discussion of random-access points and their meaning within the system. The **random_access_point** bit shall remain ‘1’ for all of the packets containing the portion of the stream comprising the random-access point. For example, for packets carrying a video stream, all the packets containing the segment defined as the RAP shall have the **random_access_point** bit set to ‘1’.

time_limit_flag – This flag shall be set to ‘1’ to indicate that a **timestamp()** field is present. A value of ‘0’ shall indicate that no **timestamp()** field is present.

wakeup_control() – A 2-bit field that communicates information needed to control the emission wakeup field. See Section 7.2.2 for a definition of this field. The **wakeup_control()** field shall only be interpreted if the **lls_flag** is set to ‘1’.

timestamp_min() – Defined in Table 6.3, this value shall indicate the latest time at which the start of the payload may be delivered. A value of ‘0’ shall denote that the packet may be delivered with best effort. Note that this field replaces the **SSRC_ID** field specified in RFC 3550 [6].

9 STL TRANSPORT PROTOCOL

The STL Transport Protocol (STLTP) shall extend the Common Tunneling Protocol (CTP) described in Section 6. The inner, Tunneled Data Streams shall carry sequences of packets comprising an RTP/UDP/IP header plus payload data. All packets delivered over the STL shall use an RTP/UDP/IP IPv4 protocol. The inner layer shall use Multicast Streams only, while the outer layer can use Multicast or Unicast consistent with CTP. See Annex D for more details of the Unicast use case.

9.1.1 Address Assignments

IPv4 packet format and addressing shall be used exclusively on the STL. The Multicast (destination) address range is 224.0.0.0 – 239.255.255.255. Of that range, 239.0.0.0 – 239.255.255.255 are for private addresses and shall be used for both inner Tunneled and outer Tunnel Packets when multicast is used.

When Unicast is used for outer Tunnel Data Packets, assignment of destination addresses and port numbers will depend on Internet Service Providers.

9.1.2 Port Assignments

Each Multicast destination address has usable port numbers from 1 through 65535. Values of 30000 – 30366 inclusive shall be used for this standard. Ports 30000 through 30063 shall be used to identify inner Tunneled Streams having destinations of PLPs 0 through 63, respectively; port 30064 shall be used to identify the inner Tunneled Stream carrying the Preamble data; port 30065 shall be used to identify the inner Tunneled Stream carrying the Timing and Management Data; and port 30066 shall be used to identify the inner Tunneled Stream carrying the Security Data Stream. All the listed Tunneled Packet Streams (i.e., Preamble, Timing and Management, Baseband Packets, and Security Data) required to populate an RF channel comprise a group of inner Tunneled Packet Streams.

The arrangement of UDP port numbers described above for an inner group of Tunneled Packet Streams can be offset in steps of 100 to permit addressing similar sets of Tunneled Packet groups intended for delivery to other Transmitters over the same STLTP. In such cases, an offset value for the port numbers for Transmitter P (where P is a value from 0 – 3) can be obtained from $Offset[P] = P \times 100$. Alternatively, the number P can be used as the middle digit of a 5-digit number, with the first two digits being ‘30’ and the last two digits ranging from ‘00’ through ‘66’. (See Table 6.1 for the semantics of the **number_of_channels** field.) For example, in a case in which a

Broadcast Gateway processes two channels, respectively A and B, the group of inner Tunneled Packet Streams of channel A would use the UDP port range from 30000 to 30066 (offset 0), while the group of inner Tunneled Packet Streams of channel B would use the UDP range 30100 to 30166 (offset 100). All these inner Tunneled Packet Streams then would be carried by the same Stream of Tunnel Packets. See Annex E for descriptions of the use of multichannel STLTP in Channel Bonding cases.

9.1.3 Majority Logic

Majority Logic is a technique for increasing the reliability of delivery of data to which it is applied. Delivery of the Preamble and of the Timing and Management data across the STL are required to maintain synchronization of Transmitters in a Network and possibly to avoid muting of individual Transmitters when there are interruptions of the delivery channels from a Broadcast Gateway to one or more Exciters. By permitting spreading out delivery of the critical data over time, Majority Logic can be useful in improving delivery of Preamble and of Timing and Management data when an STL experiences longer outages than can be corrected by the other reliability-improvement methods included in the STLTP. The amount of improvement that can be obtained with respect to relatively long outages increases with larger numbers of repetitions of the data and with longer time intervals between the repetitions of the data. The number of repetitions and the time between repetitions are separately configurable for the Preamble and the Timing and Management data.

When majority logic is applied to Preamble and/or T&M data, the objective is to spread the repeated information over the longest reasonable time so that disturbances in the STL delivery channel will be overcome through the redundant transmission of the primary synchronization information of the transmission system. The time spread of the repeated packets must be balanced against the latency added to the emission process by a greater spreading time. Therefore, successive redundant packets of Preamble or T&M Data, at a minimum, shall be carried both in separate inner Tunneled Packets and separate outer Tunnel Packets. Figure 9.1 shows an example of how multiple copies of Preamble and T&M can be sent to the Transmitter(s). As depicted in Figure 9.1, the repeated packets of Preamble and T&M may be sent over a duration of multiple Physical Layer frames prior to delivery of the Baseband Packets to be carried in the Physical Layer frames that they control.




Figure 9.1 Example use of Majority Logic for Preamble (denoted as L1) and T&M: (when 3 repetitions of each are applied).

9.2 Preamble Data Generator

Preamble information is constructed in the Preamble Generator according to instructions from the Scheduler. It is output by the Preamble Generator in the form of RTP/UDP/IP Multicast packets, similar to those used to carry Baseband Packets (BBPs) in PLP Streams, so that they form a Stream that can be multiplexed together with the PLP Streams in the STLTP. The Preamble Stream carries a description of the configuration of the Transmitter processing functions and the resulting emitted waveform that is identical to the Preamble data structure sent to receivers.

To set up Transmitter configurations, the Preamble must be sent from the Scheduler to arrive at the Transmitter(s) at least 1 Physical Layer frame in advance of the start of construction by the Transmitter(s) of the frame that the Preamble describes.

To enable receivers to decode transmitted data, the same Preamble data used to configure the Transmitter(s) for a particular frame shall be carried in the emitted preamble waveform immediately following the Bootstrap of that frame.

9.2.1 Preamble Data Stream Protocol

The Preamble data shall be delivered in an RTP/UDP/IP multicast Stream conforming to RFC 3550 [6] with the constraints defined below. The maximum Preamble data structure size can exceed the typical 1500-byte MTU, so a mechanism is defined herein to allow segmentation of the Preamble data across multiple RTP/UDP/IP packets. Note that such segmentation is required only to conform with typical MTU sizes of 1500 bytes. If the local Network allows larger multicast packets, this segmentation may not be needed.

The payload data for each Preamble Stream RTP/UDP/IP packet shall be a fragment of the Preamble Payload data structure described in Table 9.1. To provide validation that the **L1_Basic_signaling ()** and **L1_Detail_signaling ()** structures are delivered correctly over the STL, a 16-bit cyclic redundancy check (CRC) is provided. The CRC is applied to the combined contents of the **length** field, **L1_Basic_signaling ()** and **L1_Detail_signaling ()** and is appended as the last 16 bits of the payload data. The resultant Stream of Preamble Payload packets shall have destination address 239.0.51.48 and destination port 30064 before application of channel number offset of the port number in the case of multichannel carriage within a single STL Tunnel Packet Stream.

Table 9.1 Preamble Payload

Syntax	No. of Bits	Format
Preamble_Payload () {		
length	16	uimsbf
L1_Basic_signaling ()	200	Table 9.2 of [3]
L1_Detail_signaling ()	var	Table 9.12 of [3]
crc16	16	uimsbf
}		

The following paragraphs describe the fields shown in Table 9.1.

The **length** field shall contain the number of bytes in the Preamble Payload data structure following the **length** field excluding the **crc16** bytes. The **length** field allows Data Consumers to avoid knowing the detailed syntax of Preamble data structures and still to reconstruct the overall Preamble Payload data structure.

The **crc16** field shall be the value resulting from application of a 16-bit cyclic redundancy check, as defined in [10], applied to the combined **length**, **L1_Basic_signaling ()** and **L1_Detail_signaling ()** structures in the Preamble Payload immediately preceding the **crc16** field.

The Preamble Data Generator shall form the necessary Preamble Payload data, as detailed in Table 9.1, from the Scheduler configuration and calculated information. Once the data structure has been populated, it shall be partitioned, if necessary, into multiple RTP/UDP/IP packets, each conforming, with the necessary headers, to the local Network MTU size. This process results in creation of a Preamble Payload Packet Set that typically consists of multiple packets of the same

size followed by a smaller remainder packet. Constructing the packets in this way, however, is not normative.

The RTP header fields of the Preamble Payload Packet Set shall be as described below, configured with the **marker (M)** bit of the packet containing the beginning of a Preamble Payload data structure set to '1'. The **marker (M)** bits of remaining packets shall be set to '0'. This allows the Transmission system on the Data Consumer end of the STL to reconstruct the Preamble Payload data structure after any resequencing takes place. The timestamps of the packets of a given Preamble Payload Packet Set shall have the same values. The timestamp values are derived from a subset of the "Bootstrap_Timing_Data ()" appearing in Table 9.3, providing a mechanism to uniquely associate each of the Preamble Payload packets with a specific Physical Layer frame. The format of the timestamp field is described in Table 9.2.

Table 9.2 RTP Header Timestamp Field Definitions

Syntax	No. of Bits	Format
timestamp () {		
seconds_pre	22	uimsbf
a-milliseconds_pre	10	uimsbf
{		

seconds_pre shall carry a value equal to the 22 least significant bits (LSBs) of the **seconds** field of the Bootstrap_Timing_Data (), described in Table 9.3.

a-milliseconds_pre shall carry a 10-bit value identical to the value contained in the 3rd through 12th MSBs of the **nanoseconds** field of the Bootstrap_Timing_Data (), described in Table 9.3. Note that the **a-millisecond_pre** value is used in the RTP Header Timestamp only as an identifier of the Reference Emission Time of the Frame in which its contents belong; consequently, the somewhat longer Period of an **a-millisecond_pre** relative to precisely one millisecond is immaterial for this use.

On reception, the Preamble Payload Packet Set shall be resequenced, when necessary, and extracted into the Preamble Payload data structure as described in Table 9.1. The Data Consumer can accumulate RTP packets until it has received all of the bytes defined by the length field in the first packet. If a packet is missed, as determined by a missing sequence number, or if a packet with the **marker (M)** bit set to '1' is received prematurely, indicating the start of the next Preamble Payload Packet Set, then one or more packets have been lost and the entire Preamble Payload data set has been lost. Any accumulated data shall be discarded.

The RTP header fields shall follow the syntax defined in RFC 3550 [6], with the following additional constraints:

The **Padding (P)** bit shall conform to the RFC 3550 [6] specification.

The **Extension (X)** bit shall be set to '0' to indicate that the header contains no extension and that the packet has not been signed. The **Extension (X)** bit shall be set to '1' to indicate that a signing (or other) extension is present. Refer to Section 9.4 for details regarding packet signing and the definition of the STLTP signing extension.

The **CSRC Count (CC)** shall be set to '0', as no CSRC fields are necessary.

The **marker (M)** bit shall be set to '1' to indicate that the first byte of the payload is the start of the Preamble Payload data. A '0' value shall indicate that the payload is a continuation of the Preamble Payload data from the previous packet.

The **Payload Type (PT)** shall be set to 77 (0x4d), indicating the Preamble Payload type.

The **Sequence Number** shall conform to the RFC 3550 [6] specification.

The **Timestamp** shall be defined as in Table 9.2. The timestamp shall be set to the same value for all of the Preamble Payload packet set.

The **Synchronization Source (SSRC) Identifier** shall be set to '0'. There shall be no other sources of Preamble Payload data carried within an STLTP Stream. Any redundant sources can be managed using IGMP Source-Specific Multicast (SSM) mechanisms.

9.3 Timing and Management Generator

Timing and Management (T&M) information is constructed in the Timing and Management Generator according to instructions from the Scheduler. It is output by the Timing and Management Generator in the form of RTP/UDP/IP multicast packets, similar to those used to carry Baseband Packets (BBPs) in PLP Streams, so that they form a Stream that can be multiplexed together with the PLP Streams in the STLTP. These Timing and Management packets are not emitted. The resulting Timing and Management Stream carries a set of instructions for controlling the emission of Physical Layer frames comprising a Bootstrap, Preamble, and Baseband Packets. Configurations of Bootstraps and certain other components of the Physical Layer frames are carried in the Timing and Management Stream. Also included in the Timing and Management Stream are the emission time of each Bootstrap and, hence, the start of each Physical Layer frame, the offset times of each Transmitter in an SFN from the Bootstrap Reference Emission Times for the Network, and other information used to control the Transmitter(s).

To set up Transmitter configurations, the Timing and Management Data for a Physical Layer frame must be sent from the Scheduler to arrive at the Transmitter(s) at least one Physical Layer frame in advance of the start of construction by the Transmitter(s) of the Physical Layer frame that it describes.

9.3.1 Timing and Management Data Stream Protocol

The Timing and Management Data shall be delivered in an RTP/UDP/IP multicast Stream conforming to RFC 3550 [6] with the constraints defined below. The maximum T&M data structure size may exceed the typical 1500-byte MTU, so a mechanism is defined herein to allow segmentation of the T&M data across multiple RTP/UDP/IP packets. Note that such segmentation is required only to conform with typical MTU sizes of 1500 bytes. If the local network allows larger multicast packets, this segmentation may not be needed.

The payload data for each T&M Stream RTP/UDP/IP packet shall be a fragment of the TMP () data structure described in Table 9.3. To provide validation that the TMP () structure is delivered correctly over the STL, a 16-bit cyclic redundancy check is provided as part of the TMP () data. The resultant Stream of TMP () packets shall have IP destination address 239.0.51.48 and destination port 30065, before application of channel number offset of the port number in the case of multichannel carriage within a single STL Tunnel Packet Stream.

The T&M Data Generator shall form the necessary TMP () data structure, as detailed in Table 9.3, from the Scheduler configuration and calculated information. Once the data structure has been populated, it shall be partitioned into multiple RTP/UDP/IP packets, each conforming, with the necessary headers, to the local network MTU size. This process results in creation of a TMP packet set that typically consists of multiple packets of the same size followed by a smaller remainder packet. Constructing the packets in this way, however, is not normative.

The RTP header fields of the TMP packet set shall be as described below, configured with the **marker (M)** bit of the packet containing the beginning of a TMP () data structure set to '1'. The

marker (M) bits of the remaining packets shall be set to ‘0’. This allows the transmission system on the consumer end of the STL to reconstruct the TMP () data structure after any resequencing takes place. The timestamps of the packets of a given TMP packet set shall have the same values. The timestamp values are derived from a subset of the Bootstrap_Timing_Data, providing a mechanism to uniquely associate each of the TMP packets with a specific Physical Layer frame. The RTP header fields shall follow the syntax defined in RFC 3550 [6] with the following additional constraints:

The **Padding (P)** bit shall be set to ‘0’, indicating no padding is present in the Timing and Management Data packet.

The **Extension (X)** bit shall be set to ‘0’ to indicate that the header contains no extension and that the packet has not been signed. The **Extension (X)** bit shall be set to ‘1’ to indicate that a signing (or other) extension is present. Refer to Section 9.4 for details regarding packet signing and the definition of the STLTP signing extension.

The **CSRC Count (CC)** shall be set to ‘0’, as no CSRC fields are necessary.

The **marker (M)** bit shall be set to ‘1’ to indicate that the first byte of the payload is the start of the TMP data. A ‘0’ value shall indicate that the payload is a continuation of the TMP data from the previous packet.

The **Payload Type (PT)** shall be set to 76 (0x4c) indicating the Timing and Management Data payload type.

The **Sequence Number** shall conform to the RFC 3550 [6] specification.

The **Timestamp** shall be defined as in Table 9.2.

The **Synchronization Source (SSRC) Identifier** shall be set to ‘0’. There should be no other sources of Timing and Management Data carried by the STLTP. Any redundant sources can be managed using IGMP Source-Specific Multicast (SSM) mechanisms.

Table 9.3 Timing and Management Stream Packet Payload

Syntax	No. of Bits	Format
Timing & Management_Packet (TMP) () {		
Structure_Data () {		
length	16	uimsbf
version_major	4	uimsbf
version_minor	4	uimsbf
maj_log_rep_cnt_pre	4	uimsbf
maj_log_rep_cnt_tim	4	uimsbf
bootstrap_major	4	uimsbf
bootstrap_minor	4	uimsbf
min_time_to_next	5	uimsbf
system_bandwidth	2	uimsbf
bsr_coefficient	7	uimsbf
preamble_structure	8	uimsbf
ea_wakeup	2	bslbf
num_emission_tim	6	uimsbf
num_xmtrs_in_group_minus_1	6	uimsbf
xmtr_group_num	7	uimsbf
maj_log_override	3	bslbf
num_miso_filt_codes	2	bslbf

tx_carrier_offset	2	tcimbsf
reserved	6	for (i=0; i<6; i++) '1'
}		
Bootstrap_Timing_Data () {		
for (i=0; i<=num_emission_tim; i++)		
seconds	32	uimbsf
nanoseconds	32	uimbsf
}		
}		
Per_Transmitter_Data () {		
for (i=0; i<=num_xmtrs_in_group_minus_1; i++) {		
xmtr_id	13	uimbsf
tx_time_offset	16	tcimbsf
txid_injection_lvl	4	uimbsf
miso_filt_code_index	2	bslbf
reserved	29	for (i=0; i<29; i++) '1'
}		
}		
Packet_Release_Time () {		
pkt_rls_seconds	4	uimbsf
pkt_rls_a-milliseconds	10	uimbsf
reserved	2	'11'
}		
Error_Check_Data () {		
crc16	16	uimbsf
}		
}		

length shall indicate the number of bytes in the Timing and Management Data packet following the RTP/UDP/IP header structure. Up to 65,535 Bytes can be indicated.

version_major, in conjunction with **version_minor**, shall indicate the version of the protocol used to construct the Timing and Management Data packet. Increments in the value of **version_major** are intended to indicate changes in the structure that are not fully compatible with lower-ordered **version_major** values. The value of **version_major** can range from 0 through 15. Timing and Management packets constructed according to this version of this standard shall have the value of **version_major** set to 0.

version_minor, in conjunction with **version_major**, shall indicate the version of the protocol used to construct the Timing and Management Data packet. Increments in the value of **version_minor** are intended to indicate changes in the structure that are fully backward compatible with lower-ordered **version_minor** values paired with the same **version_major** value. The value of **version_minor** can range from 0 through 15. Timing and Management packets constructed according to this version of this standard shall have the value of **version_minor** set to 0.

maj_log_rep_cnt_pre shall indicate the number of repetitions of Preamble data in the Preamble Stream at UDP port 30064 prior to emission of the Preamble. Permitted values are 1, 3, 5, 7, and 9. Note that value of **LIB_lls_flag** may be correct only in the final copy of the Preamble data sent to Transmitters prior to emission. Consequently, majority logic error correction can be

applied reliably to all portions of the Preamble Stream data except the flag value noted. See Section 10.2 for details of placement of the repeated data.

maj_log_rep_cnt_tim shall indicate the number of repetitions of Timing & Management data in the Timing & Management Stream at UDP port 30065 prior to emission of the next Bootstrap. Permitted values are 1, 3, 5, 7, and 9. Note that values for the **ea_wakeup** bits may be correct only in the final copy of the Timing & Management data sent to Transmitters prior to emission. Consequently, majority logic error correction can be applied reliably to all portions of the Timing & Management Stream data except the **ea_wakeup** values noted. See Section 10.1 for details of placement of the repeated data.

bootstrap_major shall indicate the value of the **bootstrap_major_version** of the Bootstrap symbols that introduce the Physical Layer frame identified by the Bootstrap Timing Data (), which value shall be applied as the root of the Zadoff-Chu sequence of the Bootstrap symbols, as specified in [2].

bootstrap_minor shall indicate the value of the **bootstrap_minor_version** of the Bootstrap symbols that introduce the Physical Layer frame identified by the Bootstrap Timing Data (), which value shall be applied as the seed for the PN sequence of the Bootstrap symbols, as defined in [2].

min_time_to_next shall be the enumerated value indicating the minimum time until the next frame of the same type as defined in [2].

system_bandwidth shall be the enumerated value indicating the bandwidth of the RF Transmission channel as defined in [2].

bsr_coefficient shall be the binary value associated with the baseband sampling rate as defined in [2].

preamble_structure shall be the enumerated value indicating the Preamble configuration as defined in [3].

num_emission_tim shall indicate the number of sequential Bootstrap Reference Emission Times that are contained within the Bootstrap_Timing_Data () ‘for’ loop. Up to 64 values may be indicated. Allowable values shall range from 0 thru 63 and shall be expressed as the number of values carried in the packet minus 1. At least the next Bootstrap Reference Emission Time shall be carried, and it shall be carried in index 0 of the ‘for’ loop.

ea_wakeup shall signal the states of the two EA Wakeup Bits to be included in the Bootstrap signal at the start of the next Frame to be emitted.

num_xmtrs_in_group_minus_1 shall indicate the number of Transmitters minus one to which data is addressed in the Per_Transmitter_Data () ‘for’ loop (i.e., 1 to 64 Transmitters are indexed 0 to 63). The value can be less than the total number of Transmitters in the Network, in which case data addressed to groups of Transmitters shall be sequenced in order across multiple Timing & Management Data packets.

xmtr_group_num shall indicate the ordinal number of a group of Transmitters to which information in the Per_Transmitter_Data () loop is addressed. The value of the field may range from 0 through 127. Only a single value of **xmtr_group_num** shall apply to a given Timing & Management Stream data packet. Information for individual Transmitters shall be organized in groups identified by values of **xmtr_group_num** starting at 0 and incrementing by 1 from one Timing & Management Stream data packet to the next, until the highest-numbered group is reached, at which point the value shall start again at 0 in the following such packet.

Bootstrap_Timing_Data () shall contain a list of the Bootstrap Reference Emission Times of the next and, optionally, successive future frames, the list having a total number of entries equaling

the value of **num_emission_tim**. The values of the Bootstrap Reference Emission Times shall strictly increase from the first entry in the list to the last.

maj_log_override shall indicate that all previous instances of Timing and Management Data and Preamble data for the next and following Physical Layer frames shall be ignored and that the information in the current Timing and Management packet and a subsequent Preamble data packet shall be used to configure the next Physical Layer frame. The non-override condition shall be indicated by a value of '000' in this field. An override condition shall be indicated by a value of '111' in this field.

num_miso_filt_codes shall be set to one less than the number of different MISO filter codes in use within an SFN, as represented by the variable ' N_{TX} ' in Annex J of A/322. For example, when $N_{TX}=2$ is used, **num_miso_filt_codes** would be set equal to 1. The value '0' shall be reserved for future use.

tx_carrier_offset shall indicate the carrier offset of the Transmitter(s) in the frequency domain. The carrier offset shall be expressed in units of a positive or negative integer number of carriers, and it shall be a two's complement signed integer binary number having a range from -1 to $+1$ decimal, representing from -1 to $+1$ OFDM carriers. The carrier offset value shall be equal to the product of the value of **tx_carrier_offset** and the carrier frequency spacing in Hz of an 8K FFT for the value of **bsr_coefficient** indicated in the Structure_Data () for the same frame. Carrier frequency spacing (in Hz) equals BSR (in Hz) divided by 8192. For example, in a system operating with a 6 MHz channel bandwidth and a BSR of 6.192 Mega-samples/second, **bsr_coefficient**=2, the carrier frequency spacing of 8K carriers is 843.75 Hz, and the carrier frequency offset will be -843.75 Hz, 0 Hz, and +843.75 Hz for values of **tx_carrier_offset** of -1, 0, and +1, respectively. **tx_carrier_offset** = -2 shall be reserved for future use.

The **tx_carrier_offset** value also is used by the Scheduler to set reference emission times as described in Section 10.3.3.2 below.

seconds shall carry a value equal to the 32 least significant bits (LSBs) of the seconds portion of the TAI time value [15] of the associated Bootstrap Reference Emission Time, as expressed using the Precision Time Protocol (PTP) defined in [11] and [12].

nanoseconds shall carry a value equal to the nanoseconds portion of the TAI time value [15] of the associated Bootstrap Reference Emission Time. It shall be expressed as a 32-bit binary value having a range from 0 through 999,999,999 decimal.

Per_Transmitter_Data () shall contain information addressed individually to one or a group of Transmitters, with the number of Transmitters for which data is included in the loop equaling the value in **num_xmtrs_in_group_minus_1** plus 1.

xmit_id shall indicate the address of the Transmitter to which the following values are being sent and shall correspond to the seed value used by the TxID code sequence generator of that Transmitter. The value of the address shall be an unsigned integer binary number having a range of possible values from 0 through 8191 decimal.

tx_time_offset shall indicate the emission time offset of the Transmitter to which it is addressed relative to the Bootstrap Reference Emission Times of all frames. The Transmitter time offset shall be expressed in units of positive or negative integer steps of 100 ns and shall be a two's complement signed integer binary number having a range from $-32,768$ through $+32,767$ decimal, representing time offsets from $-3,276.8$ through $+3,276.7$ microseconds.

txid_injection_lvl shall indicate the Injection Level of the TxID signal below the average power of the Preamble symbols emitted by the Transmitter to which its value is addressed. The

Injection Level shall indicate the value in dB listed in A/322 [3] Table N.3.1 for the TxID Injection Level Code included in the **txid_injection_lvl** field (or Off for code value '0000').

miso_filt_code_index shall be set to one less than the specific MISO filter code assigned to the individual Transmitter, as represented by the variable 'h' in A/322 [3] Annex J. For example, when h=1 is used, **miso_filt_code_index** would be set equal to 0. The same value of MISO filter code index shall apply to a particular Transmitter regardless of whether 64-coefficient or 256-coefficient filters are in use.

pkt_rls_seconds shall be the seconds portion of the time of release from the Broadcast Gateway of the specific Timing & Management packet in which the value is found. Its value shall be expressed as 4 bits representing the 4 LSBs of the seconds value of the TAI time [15] when the first bit of the IP header of the T&M packets is released from the Broadcast Gateway.

pkt_rls_a-milliseconds shall be the milliseconds portion of the time of release from the Broadcast Gateway of the specific Timing & Management packet in which the value is found. Its value shall be expressed as 10 bits representing the 3rd through 12th MSBs of the nanoseconds value of the TAI time [15] when the first bit of the IP header of the T&M packets is released from the Broadcast Gateway. Its range will be from 0 to 953 (decimal) as a consequence of the Period of an **a-millisecond** being slightly longer than precisely a millisecond. (See the definition of an **a-millisecond** in Section 3.4).

crc16 shall be the value resulting from application of a 16-bit cyclic redundancy check, as defined in [10], applied to all fields in the Timing and Management Packet payload from the **length** field through the field (and any reserved bits) immediately preceding the **crc16** field.

9.3.2 Bootstrap Emission Timing and Frame Identification

Bootstrap Reference Emission Time shall be used for Physical Layer frame identification. In the RTP headers of the STLTP, the timestamp fields are used to carry number patterns matching portions of a complete Bootstrap Reference Emission Time. The portions to be matched have been selected for unique identification of the Physical Layer frames with which the STLTP packets are associated. The portions of the Bootstrap Reference Emission Time used for identification are the 22 LSBs of the seconds value and the 3rd through 12th MSBs of the nanoseconds value, the latter of which represent approximately millisecond time increments. The selected portions of the Bootstrap Reference Emission Time are sufficient to assure no ambiguity in the association of STLTP packets with Physical Layer frames.

9.3.2.1 Baseband Packet Delivery

The Packetizer functionality accepts ALP packets from the transport layer as described in [5]. All ALP packets are converted to Baseband Packets as described in [3]. The result of this process is a complete description of the baseband data to be radiated within a specific PLP within an identified Physical Layer Frame. The baseband Packetizer then shall encapsulate and segment Baseband Packets into RTP/UDP/IP packets as per Scheduler instructions appropriate to the configuration of the associated destination PLP and Physical Layer Frame.

9.3.3 PLP Data Stream

Each PLP carries parallel data, simultaneous in time and inherently packetized into Baseband Packets. The payload identity must be signaled for correct PLP processing. The port number used to deliver each ALP Stream, as described in Section 6, associates the ALP Stream with a particular PLP number.

9.3.4 Baseband Packet Data Stream Protocol

The Baseband Packet data shall be delivered in an RTP/UDP/IP multicast Stream conforming to RFC 3550 [6] with the constraints defined below. The Baseband Packet data structure size can exceed the typical 1500-byte MTU, so a mechanism is defined herein to segment the Baseband Packet data across multiple RTP/UDP/IP multicast packets. Note that this is required only to conform with typical MTU sizes of 1500 bytes. If the local Network allows larger multicast packets, this segmentation may not be required.

The payload data for each RTP/UDP/IP multicast packet shall be a fragment of the Baseband Packet data structure as defined in [3]. The collection of segmented packets representing a single Baseband Packet is referred to hereafter as the Baseband Packetizer Packet Set (BPPS). The resultant packet Stream shall have IP destination address 239.0.51.48 and IP destination ports 30000 through 30063, corresponding to the PLPs numbered from 0 through 63, respectively, before application of channel number offset of the port number in the case of multichannel carriage within a single STL Tunnel Packet Stream.

The Packetizer, in each path from the corresponding ALP Stream to the target PLP Stream, will form each Baseband Packet from ALP packets in the incoming ALP data Stream as described above. Once a Baseband Packet is ready for emission, it will be partitioned into multiple RTP/UDP/IP multicast packets each conforming, with the necessary headers, to the local Network MTU size. In summary, the resultant BPPS typically will consist of multiple packets of the same size followed by a smaller remainder packet. Constructing the packets in this way is not normative and any size or number of packets is permitted within a BPPS.

The RTP header fields of the BPPS shall be set as described below, with the **marker (M)** bit of the packet containing the beginning of the BPPS set to '1'. The **marker (M)** bit of the remaining packets shall be set to '0'. This allows the Transmission system on the Data Consumer side of the STL to reconstruct the Baseband Packet after any resequencing takes place. The timestamps of the packets of the BPPS shall be set to the same value indicating the Bootstrap Reference Emission Time of the Physical Layer frame. In addition to the **marker (M)** bit being set to '1' on the first BPPS packet, the **Synchronization Source (SSRC) Identifier** field of the first BPPS packet shall be set to the overall length of the Baseband Packet data structure in bytes across the entire BPPS for a given PLP.

On receipt at the Data Consumer, the BPPS shall be resequenced and extracted into the Baseband Packet data structure. The Data Consumer can accumulate RTP packets until it has received all of the bytes defined by the length field in the first BPPS packet. If a BPPS packet is missed as determined by a missing sequence number or if a packet with the **marker (M)** bit set to '1' is received prematurely, indicating the start of the next BPPS, then one or more packets have been lost and the entire Baseband Packet data structure has been lost. Any accumulated data shall be discarded.

Note that, depending on time interleaver modes, a Baseband Packet at the current timestamp may span into the subsequent Physical Layer frame due to the delaying nature of time interleaving. In such a case, the Baseband Packet at the current timestamp would belong to both the current and the subsequent Physical Layer frames.

The RTP header fields shall follow the syntax defined in RFC 3550 [6] with the following additional constraints:

The **Padding (P)** bit shall conform to the RFC 3550 [6] specification.

The **Extension (X)** bit shall be set to '0' to indicate that the header contains no extension and that the packet has not been signed. The **Extension (X)** bit shall be set to '1' to indicate that a

signing (or other) extension is present. Refer to Section 9.4 for details regarding packet signing and the definition of the STLTP signing extension.

The **CSRC Count (CC)** shall be set to '0' as no CSRC fields are necessary.

The **marker (M)** bit shall be set to '1' to indicate that the first byte of the payload is the start of a Baseband Packet data structure. A value of '0' shall indicate that the payload is a continuation of the Baseband Packet data structure from the previous BPPS packet.

The **Payload Type (PT)** shall be set to 78 (0x4e) indicating the Baseband Packet payload type.

The **Sequence Number** shall conform to the RFC 3550 [6] specification.

The **Timestamp** shall be defined as in Table 9.2. The timestamp shall be set to the same value for all packets of a given BPPS.

When the **marker (M)** bit is '0', the **Synchronization Source (SSRC) Identifier** shall be set to '0'. When the marker (M) bit is set to '1', indicating the first packet of the BPPS, the **SSRC** field shall contain the total length of the Baseband Packet data structure in bytes. This allows the Data Consumer to know how much data is to be delivered within the payloads of the BPPS.

9.4 STL Security Protocol

The STL Security Protocol has two primary components: A Signing function that can authenticate the structural data and content of any and all Tunneled Packet Streams other than the Security Data Stream, and a Key Generation and Distribution function that securely generates and delivers the keys used by the Signing function. The Signing function is based on a Symmetric key and process called the Galois/Counter Mode (GCM) [19] of the Advanced Encryption Standard (AES) [20]. When GCM is used exclusively for authentication and not for encryption, as here, the signature data is called Galois/Counter-Mode-based Message Authentication Code (GMAC) [19]. The Key Generation and Distribution function is secured using a Symmetric key wrapped with an Asymmetric Public/Private key and processed in a scheme called OpenPGP [22]. OpenPGP is used to deliver both symmetric keys for the GMAC authentication of Tunneled Packet Streams and Public keys to be used in the OpenPGP processes by new Broadcast Gateways added to the network.

The following subsections treat each of the elements required to implement the entire STL Security Protocol. They cover the signing procedures for Tunneled Packets, the generation of keys for the signing procedures, the delivery of keys used for the signing procedures, the generation and distribution of Public/Private key pairs used for securely delivering the keys for the signing and key-delivery processes, the protocol of the Security Data Stream used to deliver the several keys, and implementation considerations for the Security System.

9.4.1 Tunneled Packet Signing Procedures

The individual packets in Tunneled Packet Streams carrying Baseband Packets, Preamble Packets, and Timing and Management Data packets can be signed for purposes of authentication of the source of the packets. Packets in a Security Data Stream are treated by separate encryption plus authentication methods described in Sections 9.4.3, 9.4.3.1, and 9.4.3.21). Tunneled Packets to be authenticated shall be signed using the Galois/Counter Mode (GCM) [19] of the Advanced Encryption Standard (AES) [20] (i.e., GMAC). Although the GCM is capable of Authenticated Encryption with Additional Data (AEAD), no encryption and only the authentication of Additional Authenticated Data (AAD) shall be applied. See Figure 9.2 for a diagram showing an overview of the GCM processes in a Broadcast Gateway and an Exciter. Keys of length 256 bits shall be applied

in the GCM/AES process, and AES processing Blocks of 128 bits shall be used. The GMAC Tag value shall be 128 bits in length.

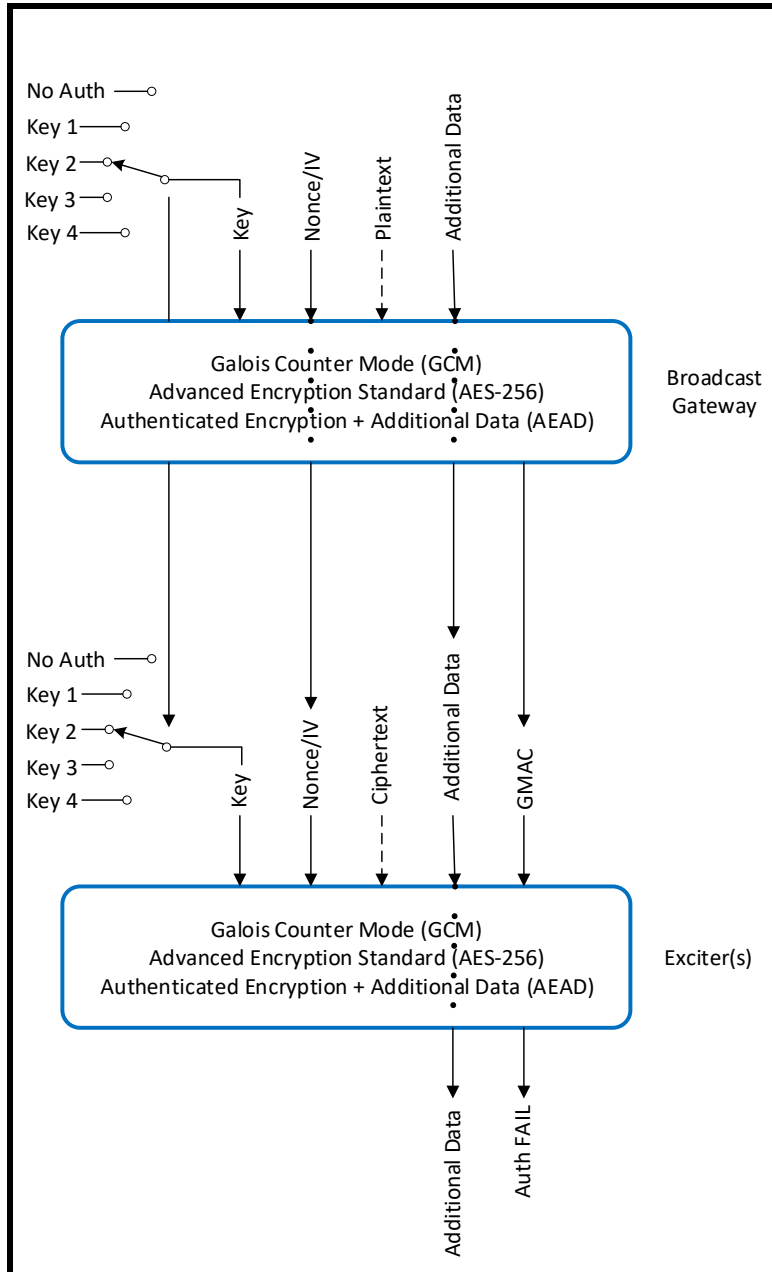


Figure 9.2 GCM AEAD and GMAC Authentication of AAD

The GCM process shall be applied to Tunneled Packets at a Broadcast Gateway and GMAC Tags calculated for all such packets. The GMAC Tags shall be communicated to Exciter(s) in RTP Header Extensions along with their respective Tunneled Packets. At the Exciter(s), the GCM process shall be applied again to the Tunneled Packets and new GMAC Tags calculated for those packets. The GMAC Tags calculated at the Exciters shall be compared with those communicated from the Broadcast Gateway. Any difference between the two GMAC Tags calculated for a

specific packet shall indicate either one or more errors in the delivery of the packet or tampering with the packet.

In addition to the secret GCM keys that must be known to both a Broadcast Gateway and its associated Exciter(s), a separate Initialization Value (IV) also must be known to both types of units. The IV must be a form of Nonce but is not required to be secret. The most important characteristic of the IV is that it not be repeated in conjunction with any particular GCM key value. The required conditions with respect to the IV can be met by concatenating the bit patterns of several values carried in the headers of the packets to be authenticated or nearby in the data Stream. The IVs shall be 96 bits in length, formatted as depicted in Table 9.4 and as described immediately thereafter.

Table 9.4 Initialization Value (IV) Organization

Syntax	No. of Bits	Format
Initialization_Value () {		
bootstrap_time_msb	10	uimbsf
gateway_src_addr_lsb	22	uimbsf
packet_dest_port	16	uimbsf
packet_seq_num	16	uimbsf
timestamp	32	uimbsf
}		

bootstrap_time_msb shall contain the 10 most significant bits of the Bootstrap Reference Emission Time of the Physical Layer frame with which the authenticated Tunneled Packet will be associated.

gateway_src_addr_lsb shall contain the 22 least significant bits of the IP Source Address of the Broadcast Gateway sending the authenticated Tunneled Packet.

packet_dest_port shall contain the UDP port number of the authenticated Tunneled Packet.

packet_seq_num shall contain the RTP sequence number of the authenticated Tunneled Packet.

timestamp shall contain the 32 bits of the timestamp () structure defined in Table 9.2 and in the paragraphs defining the individual timestamp () fields that follow Table 9.2, as carried in the RTP header of the authenticated Tunneled Packet.

Four secret GCM keys shall be stored in Cryptographic Tokens connected to each Broadcast Gateway and to each Exciter, as described in Section 1). The keys shall be communicated from a Broadcast Gateway to all Exciters in a network through use of the Security Data Stream protocol described in Section 9.4.5. The use of a particular one of the four keys for the signing of a specific Tunneled Packet shall be signaled in an RTP Header Extension of the packet as described in the next paragraph. Similarly, the GMAC Tag resulting from GHASH processing of the packet shall be communicated from the Broadcast Gateway to the Exciter(s) in the same RTP Header Extension.

To accommodate indication of the key used for signing of a packet and to support carriage of the GMAC Tag resulting from GHASH processing of the packet, an RTP Header Extension shall be applied to each signed Tunneled Packet, as described in Sections 9.2.1, 9.3.1, and 9.3.4. The RTP Header Extension shall be constructed according to the format defined in Table 9.5 and the following paragraphs. The RTP Header Extension of the Tunneled Packet shall be placed immediately following the SSRC field and the header extension (X) bit shall be set to '1' as described in the previous sections.

Table 9.5 RTP Header Extension

Syntax	No. of Bits	Format
Header Extension () {		
reserved	13	0x1FFF
key_num	3	uimsbf
length	16	0x0004
GMAC_tag	128	uimsbf
}		

key_num shall indicate the index number of the GCM key used in authentication of the packet. It shall have a value in the range from 0x0 through 0x4. A value of 0x0 shall indicate that no authentication processing is applied. Values from 0x1 through 0x4 shall indicate that authentication processing is applied using the key having the index number given. Values 0x5 through 0x7 shall be **reserved**.

length shall indicate the length of the RTP Header Extension following the last bit of the **length** field. Its value shall be the count of the number of 32-bit words included in the Header Extension, excluding the four-byte extension header; thus, zero shall be a valid length. In the instance of packet signing exclusively, **length** shall have a value 0x0004, indicating a length of the RTP Header Extension following the **length** field of 128 bits.

GMAC_tag shall be the 128-bit GMAC Tag calculated for the packet using a combination of the key value indicated by the **key_num** field, the Initialization Value determined as described above in this section, and the GHASH function applied to the structural elements plus the payload of the packet.

9.4.1.1 Procedure for Calculating GMAC Tag

The following conceptual steps are applied when calculating the GMAC Tag for an individual Tunneled Packet.

Broadcast Gateway Procedure – Calculate Signature

- 1) Create a Tunneled Packet including a header extension in the RTP Header portion of the packet.
- 2) Fill the **GMAC_tag** field with 128 '0' (zero) bits.
- 3) Calculate the GMAC Tag value by applying the GHASH function to the entirety of the Tunneled Packet extending from the first bit of its IP Header through the last bit of its payload. Encrypt the GMAC Tag with the Authentication Key selected for use with the Tunneled Packet.
- 4) Substitute the resulting GMAC Tag value into the 128-bit **GMAC_tag** field.
- 5) Transmit the packet through the STLTP Stream.

Exciter Procedure – Check Signing

- 6) On reception of the Tunneled Packet, harvest the **GMAC_tag** value from the packet for later comparison.
- 7) Fill the **GMAC_tag** field with 128 '0' (zero) bits.

- 8) Calculate the GMAC Tag value by applying the GHASH function to the entirety of the Tunneled Packet extending from the first bit of its IP Header through the last bit of its payload.
- 9) Decrypt the **GMAC_tag** value harvested from the Tunneled Packet using the Authentication Key indicated for use with the packet. Compare the GMAC Tag value calculated at the Exciter with the **GMAC_tag** value harvested from the RTP Header Extension upon reception of the packet.
- 10) Any difference between the two renditions of the **GMAC_tag** value shall indicate either errors in transmission of the Tunneled Packet or tampering with its delivery.

9.4.2 Authentication Key Generation and Utilization

Given the parameters selected for this standard, the GCM/GMAC signing scheme requires as one of its inputs a random and Fresh 256-bit key. As described in [19], “The Key shall be generated uniformly at random, or close to uniformly at random, i.e., so that each possible Key is (nearly) equally likely to be generated. Consequently, the Key will be Fresh, i.e., unequal to any previous Key, with high probability. The Key shall be secret and shall be used exclusively for GCM with the chosen Block Cipher.”

As described in Section 9.4.3, security Keys for Signing purposes generally shall be generated using a random number generator in a Cryptographic Token installed in a Broadcast Gateway. Four such Keys shall be generated initially, and they may be applied in any order, as also described in Section 9.4.6. The strength of the security provided by the Keys used in the GCM/GMAC Authentication scheme has a limited lifetime; therefore, the authentication Keys shall be replaced on a routine basis at appropriate times that depend on the patterns of their use. The fundamental limitations on Key longevity are the number of times the Keys are invoked and the number of Cypher Blocks that they process. The number of invocations of a Key shall be limited to 2^{32} . The number of Cypher Blocks processed shall be limited to 2^{64} . Given the largest possible size of the data set to which one invocation of a Key could apply, observing the limit on the number of Key invocations effectively would observe the limit on the number of Cypher Blocks processed with that Key.

In lieu of counting the number of invocations, a conservative method for limiting Key utilization is to limit the time period during which a given Key can be active. The time period limit can be derived from a calculation of the maximum number of packets that possibly can be transmitted over the STLTP per time unit with all parameters set to maximize the packet rate, followed by division of the number of packets per time unit into 2^{32} , yielding the smallest number of time units (i.e., seconds) before the earliest point at which the security capacity of a Key could be fully expended. The time period obtained by such a calculation then could be adjusted using factors relating actual operating parameters with those parameters used to determine the maximum possible packet rate for the system. For the STLTP design described in this standard, the maximum packet rate is approximately 480,000 packets per second, and the maximum safe continuous utilization time for a Key becomes just under 8,950 seconds or about 2½ hours. That value is based on a parameter set including an 8 MHz channel bandwidth, 64 PLPs, and use of MIMO. Adjustment of the time limit for a system with a parameter set including a 6 MHz channel bandwidth, 4 PLPs, and no use of MIMO would result in a time limit per Key of just under 381,900 seconds or about 106 hours (almost 4½ days). If the actual packet rate, based on the Physical Layer parameters in use, were used in the calculations instead of assumed worst-case values, it is likely that the packet rate, and hence the Authentication Key invocation rate, would be substantially lower and the corresponding time before exhaustion of the Authentication Key

security capacity would be proportionally greater. It should be noted that, if Key rotation were used instead of application of a single Key for an extended period, it would be necessary either to track the utilization times of the individual Keys or to apply a further factor to determine Key longevity based on the utilization percentage of each Key.

9.4.3 Authentication Key Delivery

It is a requirement that the Keys used in Signing of all Tunneled Packet Streams other than the Security Data Stream remain secret. Nevertheless, the Keys must be updated from time to time. To facilitate such updates, the Security Data Stream is designed to carry secret data through use of authenticated encryption based on the OpenPGP protocol [22], extended through use of Elliptic Curve Cryptography as described in [23]. OpenPGP uses a combination of symmetric and asymmetric Keys plus standardized hashing algorithms and cryptographic processes. IETF RFC 6637 [23], which describes the addition of Elliptic Curve Cryptography to OpenPGP, applies the latest cryptographic technology and also narrows the choices among many parameters available in OpenPGP. In the next several paragraphs, this standard narrows the choices further to a single interoperable set of processes and parameters.

Delivery of Authentication Keys between Broadcast Gateways and from a Broadcast Gateway to one or more Exciters shall employ OpenPGP [22], extended for use of Elliptic Curve Cryptography [23]. Each GMAC Authentication Key shall be a 256-bit random number, which shall be carried as the payload of an OpenPGP Session. Each Session Key, as specified in [23], shall be derived from a random number and use the prime field curve defined in [24] as “Curve P-384”. The Public Key Algorithm used shall be Elliptic Curve Diffie-Hellman (ECDH), and the Key Derivation Function (KDF) for ECDH shall be that described in [23]. The associated Hash function shall be SHA2-384, and the key wrapping method shall be that specified in [25] with AES-256 as the Key-Encryption Key algorithm (KEK).

The Key delivery process makes use of the relationships between Public and Private Keys in a pair for purposes of Authentication and Encryption/Decryption. The relationships are:

- Data Encrypted with a Public Key can be Decrypted with a Private Key
- Message Authentication Codes (MACs) signed with a Private Key can be verified with a Public Key

The delivery method shall proceed according to the following steps, as diagrammed in Figure 9.3. The procedure for Broadcast Gateway Processing is shown in Figure 9.3(a), and the procedure for Exciter Processing is shown in Figure 9.3(b). In the figures, Keys shown in red are Private Keys; Keys shown in other colors are Public Keys. Keys are labeled as to their types. Small numbers alongside the various blocks and text items tie them to the steps given in Sections 9.4.3.1 for Figure 9.3(a) and 9.4.3.2 for Figure 9.3(b).

In the Security Data Stream, packets carrying GMAC Authentication Keys, as described in this subsection, shall be identified by Payload Type value 80 (decimal), equivalent to 0x50 (hex).

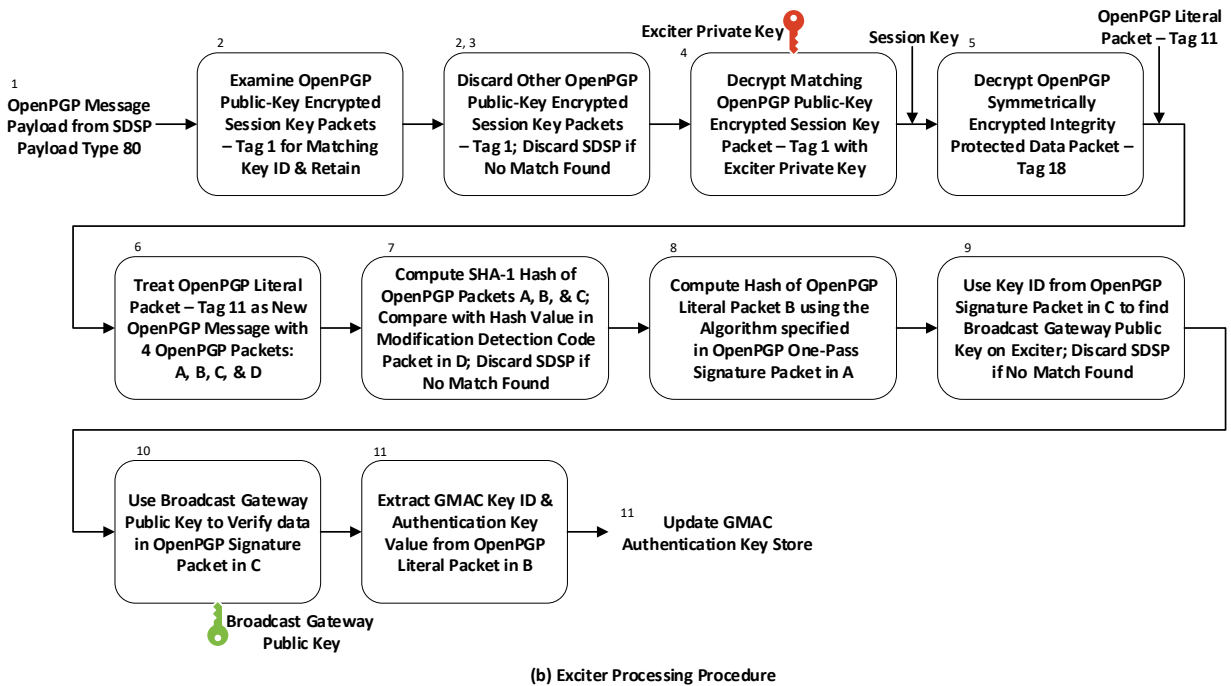
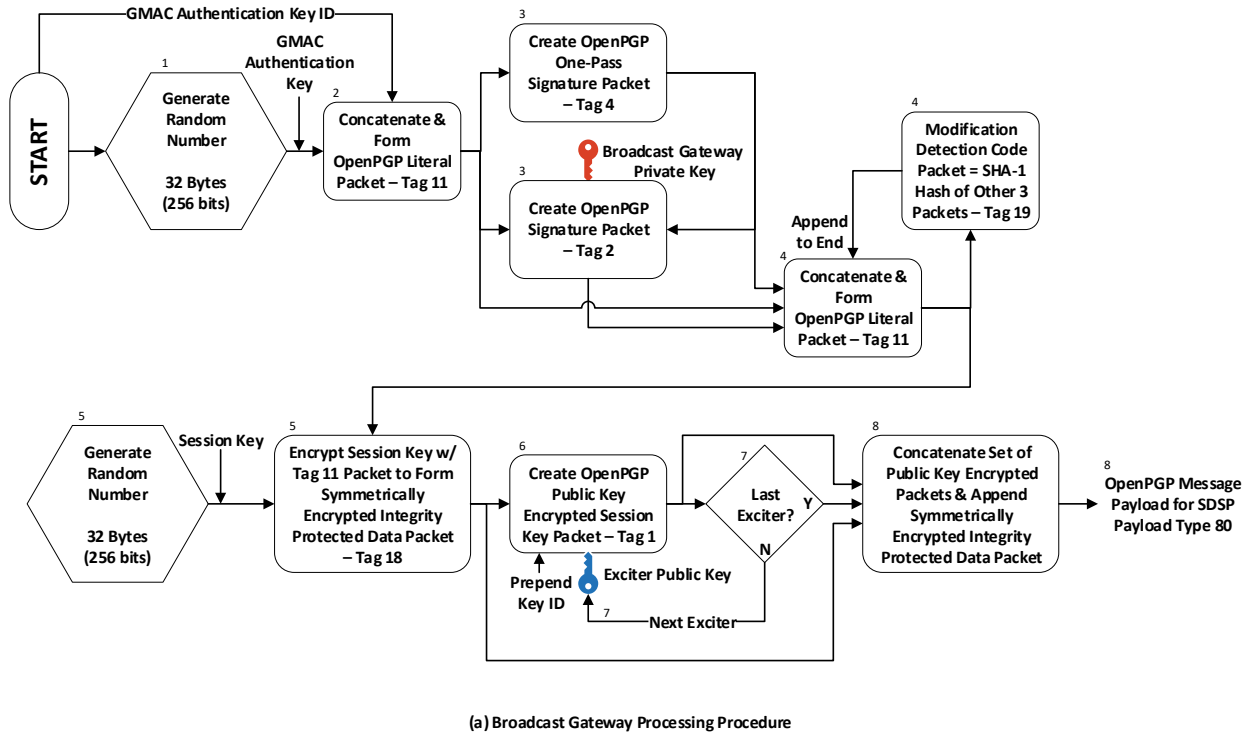


Figure 9.3 STLT security System Authentication Key generation & distribution: (a, top) Broadcast Gateway Processing Procedure; (b, bottom) Exciter Processing Procedure.

9.4.3.1 Broadcast Gateway Processing Procedure

- 1) Generate a random number to be used as a Symmetric GMAC Authentication Key and associate it with one of the four Key Identifiers used for GMAC Authentication Keys.
- 2) Create an OpenPGP Literal Packet (RFC 4880 [22] Tag 11) containing the concatenation of the GMAC Authentication Key Identifier and the GMAC Authentication Key from step (1).
- 3) Select the Broadcast Gateway Key Pair with which to sign the OpenPGP Literal Packet from step (2). Create an OpenPGP One-Pass Signature Packet (RFC 4880 [22] Tag 4). Also create an OpenPGP Signature Packet (RFC 4880 [22] Tag 2) containing the concatenation of information prescribed for the Version 4 Signature Packet Format in RFC 4880 [22], which includes a signature built from a hash of the OpenPGP Literal Packet from step (2) plus some of the previous fields of the OpenPGP Signature Packet as specified in RFC 4880 [22].
- 4) Create a second OpenPGP Literal Packet (RFC 4880 [22] Tag 11) containing the concatenation of the following four elements in the order given: (i) the One-Pass Signature Packet, (ii) the OpenPGP Literal Packet from step (2), (iii) the Signature Packet from step (3), and (iv) a Modification Detection Code Packet (RFC 4880 [22] Tag 19), containing the SHA-1 hash of the preceding elements (i) – (iii) in the concatenation.
- 5) Generate a second random number to be used as the Session Key and encrypt it with the second OpenPGP Literal Packet from step (4) to form a Symmetrically Encrypted Integrity Protected Data Packet (RFC 4880 [22] Tag 18).

Note: Per their description in RFC 4880 [22], Tag 18 packets require use of Cipher Feedback (CFB) mode encryption.

- 6) Select a target Exciter Public Key and create an OpenPGP Public-Key Encrypted Session Key Packet (RFC 4880 [22] Tag 1) containing the encrypted Session Key generated in step (5), prepended with the Key ID of the selected Exciter Public Key. The Key ID shall be the final 8 bytes of the Fingerprint of the Exciter Public Key, which Fingerprint shall be constructed according to RFC 4880 [22] instructions for fingerprints for V4 keys. The timestamp of key creation shall be the 32 bits representing the seconds count of NTP Time per RFC 5905 [13] at the time of creation of the Public Key. The Algorithm-Specific Fields shall be those specified for ECDH keys in RFC 6637 [23].
- 7) Repeat step (6) for each target Exciter.
- 8) Concatenate the set of Public-Key Encrypted Session Key Packets generated in steps (6) and (7) (in any order) and then append the Symmetrically Encrypted Integrity Protected Data Packet created in step (5). This produces an OpenPGP Message Payload that will be distributed to the Exciters in an A/324 Security Data Stream Packet with Payload Type 80 (0x50).

9.4.3.2 Exciter Processing Procedure

- 1) On receipt of a Security Data Stream Packet with Payload Type 80 (0x50), extract the OpenPGP Message Payload and process it as follows.
- 2) Locate each of the OpenPGP Public-Key Encrypted Session Key Packets (RFC 4880 [22] Tag 1) at the beginning of the OpenPGP message and locate a packet that contains an Exciter key Identifier that matches a Private Key installed on the Exciter. Discard any OpenPGP Public Key Encrypted Session Key Packets other than the first one that contains an Exciter key identifier matching a Private Key installed on the recipient Exciter.

- 3) If no matching Exciter key identifier is found, discard the Security Data Stream Packet.
- 4) Decrypt the OpenPGP Public-Key Encrypted Session Key Packet (RFC 4880 [22] Tag 1) matching the Exciter Key Identifier using the corresponding Exciter Public Key to access the Session Key.
- 5) Locate the OpenPGP Symmetrically Encrypted Integrity Protected Data Packet (RFC 4880 [22] Tag 18) and decrypt the contents of the packet with the Session Key to obtain an OpenPGP Literal Packet (RFC 4880 [22] Tag 11).
Note: Per their description in RFC 4880 [22], Tag 18 packets require use of Cipher Feedback (CFB) mode encryption and decryption.
- 6) Treat the contents of the OpenPGP Literal Packet from step (5) as a new OpenPGP Message that contains the following four OpenPGP packets in order:
 - A. An OpenPGP One-Pass Signature Packet
 - B. An OpenPGP Literal Data Packet, and
 - C. An OpenPGP Signature Packet
 - D. An OpenPGP Modification Detection Code Packet
- 7) Verify the contents of the OpenPGP Modification Detection Code packet (6)(D) against the SHA-1 hash of the three packets in (6)(A), (6)(B), and (6)(C). If the integrity protection checks fail, discard the Security Data Stream Packet.
- 8) Calculate the hash of the OpenPGP Literal Packet in (6)(B) using the algorithms specified in the OpenPGP One-Pass Signature Packet in (6)(A).
- 9) Use the Key Identifier from the OpenPGP Signature Packet in (6)(C) to locate the Broadcast Gateway Public Key stored on the Exciter. If the Key Identifier does not correspond with a Broadcast Gateway Public Key known to the Exciter, discard the Security Data Stream Packet.
- 10) Using the Broadcast Gateway Public Key from step (9), verify the signature data contained in the OpenPGP Signature Packet (6)(C) against the hash value calculated in step (8). If the signature does not verify, discard the Security Data Stream Packet.
- 11) Extract the GMAC Key Identifier and the GMAC Authentication Key from the OpenPGP Literal Packet obtained in (6)(B) and update the Exciter GMAC Authentication Key Store.

9.4.4 Public/Private Key Generation & Distribution

To enable the interchange of secured secrets such as Public and Private Key values, Authentication Keys, and the like, it is necessary for a strong set of Keys to be distributed among Broadcast Gateways and Exciters in such a way that all needed communication channels can be protected. The matter is complicated by the fact that Broadcast Gateways and Exciters can be widely separated geographically, and there can be many of them in a network. It also is necessary to support redundant operation of equipment, to allow for growth of networks, and to provide for spare parts and repairs of equipment and components. To meet these needs, resilient security elements must be deployed in semi-permanent installations, and appropriately secure processes deployed to maintain the privacy of the secrets that secure communications between network entities.

All of the foregoing considerations are met in the STLTP Security System design through use of a combination of long-lasting, asymmetric, Public/Private Key Infrastructure with symmetric Key processes that use either replaceable Keys that can be periodically updated or ephemeral Keys that are used only once. The foundation for the system, though, is the Public/Private Key scheme

that is installed in all units on the network. To make that approach work, each unit must have its own Private Key that is well-protected, remains secure, and never is exposed to other entities. As described in Section 9.4.6, Cryptographic Tokens with USB interfaces are employed to process and store the various secrets. As a network is installed, the several Keys needed must be initialized on the Tokens and then installed in the equipment.

To preserve the secrecy of Private Keys, in particular, it is imperative that they be generated on the devices where they will be stored and never exposed to communications between devices or systems where they possibly could be intercepted. Consequently, Public/Private Key pairs shall be generated on Cryptographic Tokens, and the Private Keys shall be retained there, without access from the outside of the devices on which they are generated. Public Keys shall be shared between security devices on the network, but they shall not be distributed to devices outside the network except to back up their values to facilitate generation of replacement or expansion Keys at a later time. As shown in Figure 9.4, each Cryptographic Token shall hold its own Private Key (shown in red) and also its Public Key, shown in different colors for the different applications of the Tokens. For sake of example, three Broadcast Gateways and five Exciters are shown in Figure 9.4; the number of each type can be a value of one or greater. Upon initialization, each Token intended for use in a Broadcast Gateway, in addition to its own Key pair, shall hold copies of the Public Keys of each of the other Broadcast Gateways and of each of the Exciters in the network. Each Token intended for use in an Exciter, in addition to its own Key pair, shall hold copies of the Public Keys of each of the Broadcast Gateways in the network.

In addition to the Public and Private keys stored on a Token, each one shall be given a number unique within the scope of the network, so that it can be addressed directly for delivery of its unique Session Key encrypted with its Public Key by a Token in a Broadcast Gateway. The number associated with each Token and the Public Key value of the Token having that number shall be stored securely in the Broadcast Gateway(s) to permit it/them to properly prepare security data for distribution to the other entities on the network, including both Exciters and other Broadcast Gateways.

When preparing Cryptographic Tokens for a network, it is advisable to populate additional Tokens beyond those necessary just to equip the units that will be installed from the start. Doing so can provide both backup units and units to enable expansion of equipment on the network. When redundant equipment is installed in a network, the same considerations shall apply as to single units, i.e., Tokens for additional Broadcast Gateways shall hold their own Public/Private Key pairs plus copies of the Public Keys of all other Tokens for both Broadcast Gateway and Exciter use, and Tokens for Exciter use shall hold their own Public/Private Key pairs plus copies of the Public Keys of all Tokens for Broadcast Gateway use. When redundant Exciters are used in a Transmitter, each shall be treated as if it were a standalone unit and shall have installed a Token that is identical in the types of Keys it contains to all the other Tokens intended for Exciter use. Since the data to be delivered to all Exciters through the Security Data Stream is identical, any Token set up for Exciter use can be installed in any Exciter; the only requirement is that its number and Public Key be included in the database retained in the Broadcast Gateway(s). Similarly, any Token set up for use in a Broadcast Gateway should be usable in any Broadcast Gateway on the same network, so long as the Broadcast Gateway also contains a copy of the database of unit numbers assigned to Tokens and copies of the Public Keys associated with those numbers.

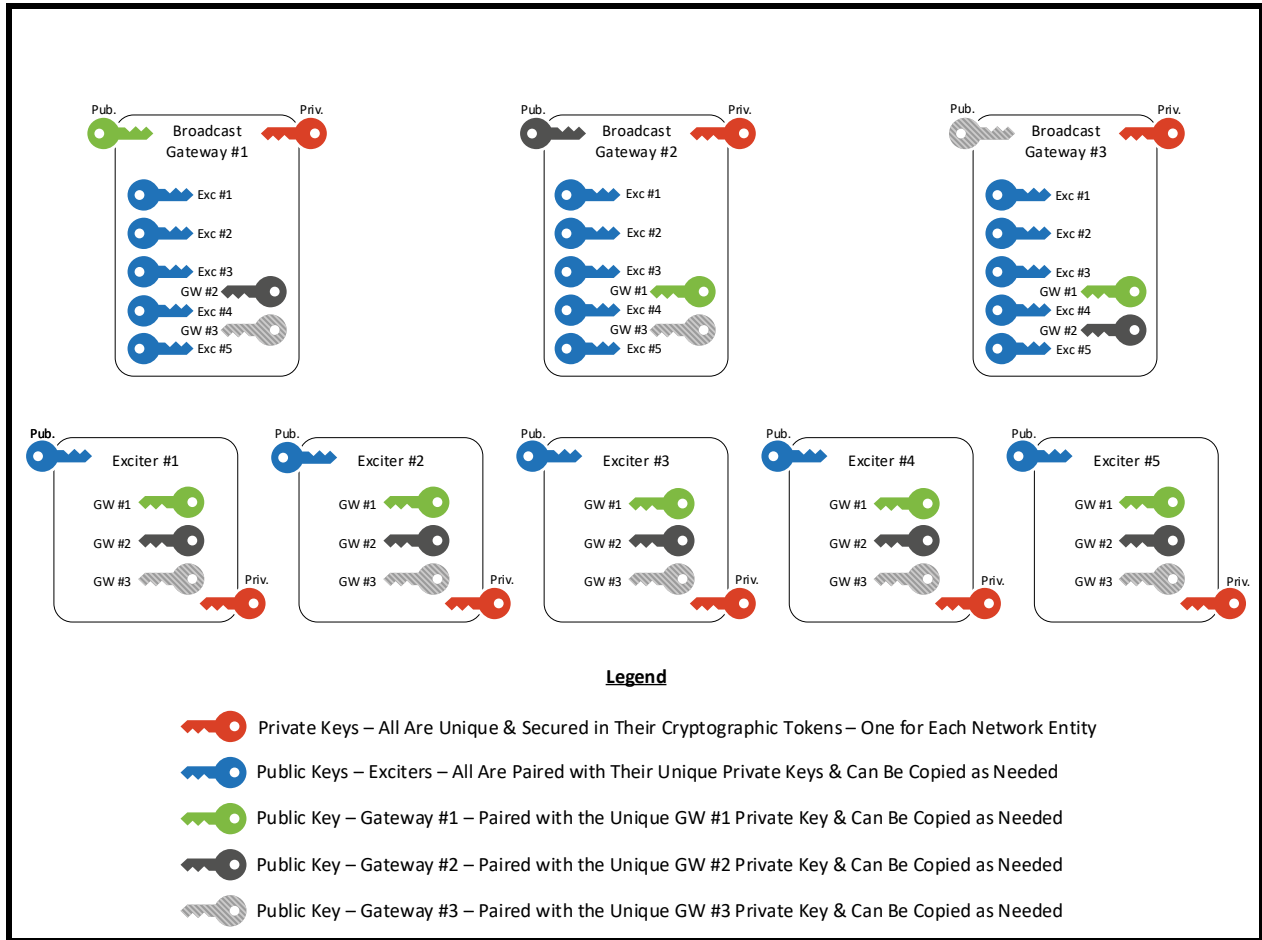


Figure 9.4 STLTP Security System Public/Private Key locations after system initialization

Despite possibly having provided spare Tokens for Broadcast Gateways and/or Exciters, at some point in the life of a network, it may become necessary to prepare additional Cryptographic Tokens and add them to the system. There are two cases that must be considered: Addition of Tokens for use in Broadcast Gateways and addition of Tokens for use in Exciters. As described in Section 9.4.6, it is expected that Token initialization will be carried out using functionality of a Broadcast Gateway or supplemental software. Given such an arrangement, a new Token for Exciter use can be loaded with the Public Keys of the Broadcast Gateway(s), and, conversely, the Public Key of the new Token can be captured and distributed among Broadcast Gateways prior to the new Token being put into service. When a new Token is for use in a Broadcast Gateway, however, it would be inconvenient, perhaps impossible, to bring all the Tokens from perhaps many Exciters to the Broadcast Gateway to load the new Gateway Public Key onto them. For this purpose, provision is made to transfer Public Keys from new Tokens initialized for Broadcast Gateway use to all Exciters on the network in the same way that GMAC Authentication Keys are distributed.

In the Security Data Stream, packets carrying Broadcast Gateway Public Keys, as described in this subsection, shall be identified by Payload Type value 81 (decimal), equivalent to 0x51 (hex).

9.4.5 Security Data Stream Protocol

The Security Data Stream shall be delivered in an RTP/UDP/IP multicast Stream conforming to RFC 3550 [6] with the constraints defined below. The payloads described in the previous sections can exceed typical MTU sizes, so a mechanism is defined to allow segmentation of the Security Data Stream across multiple RTP/UDP/IP packets. Note that such segmentation is required only to conform with typical MTU sizes. If a local network allows larger multicast packets, such segmentation may not be needed. All packets carrying segments of a single payload shall be sent before starting to send a new payload. In other words, packets from a segmented Security Data payload shall not be interleaved with other packets comprising the Security Data Stream.

The resultant Stream of Security Data packets shall have destination address 239.0.51.48 and destination port 30066.

The RTP header fields of a segmented set of Security Data Packets shall be configured as described below, with the **marker (M)** bit of the packet containing the beginning of a Security Data payload set to '1'. The **marker (M)** bits of remaining packets of the set shall be set to '0'. This allows the Transmission system on the Data Consumer end of the STL to reconstruct the Security Data payload after any resequencing takes place.

On reception, a Security Data Packet Set shall be resequenced and extracted into an SDPS Encrypted Data payload as described in Section 9.4.5.1. This payload shall be processed as described in Section 9.4.5.3 to produce either an Authentication Key payload as described in Section 9.4.5.4 or a Public Key Distribution payload as described in Section 9.4.5.5. A Data Consumer can accumulate RTP packets until it has received a syntactically complete data structure. If a packet is missed, as determined by a missing sequence number or premature reception of a packet with the **marker (M)** bit set to '1', indicating the start of the next Security Data Packet Set, then one or more packets have been lost, and the entire Security Data set has been lost. Any accumulated Security Data shall be discarded. Note that, because packets from different payloads within a Security Data Stream cannot be interleaved, a single segmented payload will appear in consecutive Tunneled Packets.

Security Data Stream RTP header fields shall follow the syntax defined in RFC 3550 [6], with the following additional constraints:

The **Padding (P)** bit shall conform to the RFC 3550 [6] specification.

The **Extension (X)** bit shall be set to '0' to indicate that the header contains no extension. The Security Data Stream is secured using the OpenPGP method [22], described in previous sections, instead of the GMAC Tag signing technique [19] applied to all other types of Tunneled Packets, as described in Section 9.4.1.

The **CSRC Count (CC)** shall be set to '0', as no CSRC fields are necessary.

The **marker (M)** bit shall be set to '1' to indicate that the first byte of the payload is the start of a Security Data Packet Set. A **marker (M)** bit value of '0' shall indicate that the payload is a continuation of the Security Data Packet Set from the previous packet.

The **Payload Type (PT)** shall be set either to 79 (0x4f), indicating the Authentication Key payload type or to 80 (0x50), indicating a Broadcast Gateway Public Key payload type.

The **Sequence Number** shall conform to the RFC 3550 [6] specification.

The **Timestamp** shall be defined as in Table 9.2. The timestamp shall be set to the same value for all packets within a Security Data Packet Set.

The **Synchronization Source (SSRC) Identifier** shall be set to ‘0’. There shall be no other sources of Security Data carried within an STLTP Stream. Any redundant sources can be managed using IGMP Source-Specific Multicast (SSM) mechanisms.

Table 9.6 Security Data Stream Packet Payload

Syntax	No. of Bits	Format
Security_Data_Stream_Packet (SDS) () { Structure_Data () { Length version_major version_minor num_tokens_minus_1 for (i=0; i<length-4; i++) SDPS_Encrypted_Data_payload_byte } } }	16 4 4 8 8	uimbsf uimbsf uimbsf uimbsf bslbf

length shall indicate the number of bytes in the Security Data Stream Packet following the RTP/UDP/IP header structure. Up to 65,535 Bytes can be indicated.

version_major, in conjunction with **version_minor**, shall indicate the version of the protocol used to construct the Security Data Stream Packet. Increments in the value of **version_major** are intended to indicate changes in the structure that are not fully compatible with lower-ordered **version_major** values. The value of **version_major** can range from 0 through 15. Security Data Stream Packets constructed according to this version of this standard shall have the value of **version_major** set to 0.

version_minor, in conjunction with **version_major**, shall indicate the version of the protocol used to construct the Security Data Stream Packet. Increments in the value of **version_minor** are intended to indicate changes in the structure that are fully backward compatible with lower-ordered **version_minor** values having the same **version_major** value. Security Data Stream Packets constructed according to this version of this standard shall have the value of **version_minor** set to 0.

num_tokens_minus_1 shall indicate the number of Cryptographic Tokens for which encrypted Session Keys are included in the Security Data Stream Packet. Its value shall range from 0 through 255 to represent Cryptographic Token counts from 1 through 256, respectively.

SDS_Encrypted_Data_payload_byte shall contain the sequence of bytes that form an OpenPGP message containing the OpenPGP packets described in Sections 9.4.5.1 through 9.4.5.5.

9.4.5.1 SDPS Encrypted Data Processing

The SDS Encrypted Data payload shall contain an OpenPGP message with the following components, in the order specified:

- (1) One OpenPGP Public Key Encrypted Packet for each of the Tokens. The number of these packets shall match **num_tokens_minus_1** in the Security Data Stream Packet plus 1. The relative order of these packets is unspecified.

The content of each of these packets shall be as follows:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 1	The format of the Packet Header shall be as described in RFC 4880 [22] Section 4.2 for New Format Headers
Version	one byte	3	
Key Identifier	eight bytes		The identifier of the Public Key used to encrypt this packet
Encryption Algorithm	one byte	18	Elliptic Curve Diffie-Hellman
Encrypted Data	remainder of packet length		The ECDH algorithm-specific fields described in RFC 6637 [23] Section 10 for Packet Tag 1. The last of these algorithm-specific fields shall contain the encrypted Session Key. See Section 9.4.5.2 for details on encapsulation of the Session Key.

(2) One OpenPGP Symmetrically Encrypted Integrity Protected Data Packet. The content of this packet is as follows:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 18	The format of the Packet Header is described in RFC 4880 [22] Section 4.2 for New Format Headers
Version	one byte	1	
Encrypted Data	remainder of packet length		

The recipient of the SDS Encrypted Data Payload shall perform the following steps:

- 1) Locate each of the OpenPGP Public Key Encrypted packets and select the first packet that contains a Key Identifier that matches a Private Key installed on the Exciter. Discard any OpenPGP Public Key Encrypted packets other than the first one that contains a Key Identifier for the recipient Exciter.
- 2) If no matching Key Identifier is found discard the Security Data Stream Packet.
- 3) Decrypt the OpenPGP Public Key Encrypted packet matching this Exciter Key Identifier with the corresponding Exciter Private Key to access the Session Key.
- 4) Locate the OpenPGP Symmetrically Encrypted Integrity Protected Data Packet and decrypt the contents of this packet with the Session Key to obtain an OpenPGP Literal Packet which contains an OpenPGP message with the SDS Authenticated Key.

The SDS Authenticated Key shall be processed as described in Section 9.4.5.3.

9.4.5.2 Session Key Encapsulation

The Session Key shall be a randomly generated 256-bit (32-byte) value that is encapsulated in the manner specified in RFC 6637 [23] Section 8 to form a 40-byte value that is input to the key wrapping method. The content of the encapsulated Session Key shall be as follows:

Field Name	Size	Value	Description
Algorithm Identifier	one byte	AES-256= 09	The AES algorithm used for the Session Key
Session Key	32 bytes		A randomly generated 256-bit number
Checksum	2 bytes		Sum of the 32 Session Key bytes modulo 65536. (Note: The Algorithm Identifier byte is not included in this calculation.)
Padding	5 bytes	each byte = 05	

9.4.5.3 SDS Authenticated Key Processing

The SDS Authenticated Key is an OpenPGP formatted message with the following four components:

(1) An OpenPGP One-Pass Signature Packet

The content of this packet shall be as follows:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 4	The format of the Packet Header shall be as described in RFC 4880 [22] Section 4.2 for New Format Headers
Version	one byte	3	
Signature Type	one byte	0	Signature of a binary document
Hash Algorithm	one byte	9	SHA-384
Public Key Algorithm	one byte	19	ECDSA
Key Identifier	eight bytes		The identifier of the Broadcast Gateway Public Key used to authenticate this packet
Flag	one byte	1	Indicator that this is the last (only) One-Pass Signature Packet

(2) An OpenPGP Literal Packet

This packet may contain either an Authentication Key Payload (in the case that the Security Data Stream RTP header is type 80 (0x50)) as described in Section 9.4.5.4 below or a Public Key Distribution Payload (in the case that the Security Data Stream RTP header is type 81 (0x51)) as described in Section 9.4.5.5 below.

(3) An OpenPGP Signature Packet

The content of this packet shall be as follows:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 2	The format of the Packet Header is described in RFC 4880 [22] Section 4.2 for New Format Headers
Version	one byte	4	
Signature Type	one byte	0	Signature of a binary document
Public Key Algorithm	one byte	19	ECDSA
Hash Algorithm	one byte	9	SHA-384
Hashed Data Length	two bytes	16	
Creation Time Subpacket Length	one byte	5	

Creation Time Subpacket Type	one byte	2	
Signature Creation Time	four bytes		Number of seconds since midnight 1 January 1970 UTC.
Issuer Subpacket Length	one byte	9	
Issuer Subpacket Type	one byte	16	
Issuer	eight bytes		The identifier of the Broadcast Gateway Public Key used to authenticate this packet
Unhashed Data Length	two bytes	0	
Truncated Hash	two bytes		The leftmost 16 bits of the signed hash value
Signature Data	remainder of packet length		The value of the ECDSA signature as specified in RFC 6637 [23] Section 10.

(4) An OpenPGP Modification Detection Code Packet

The content of this packet shall be as follows:

Field Name	Size	Value	Description
Packet Header	two bytes	Packet Tag = 19	The format of the Packet Header shall be as described in RFC 4880 [22] Section 4.2 for New Format Headers. For this packet the packet length is always encoded in a single byte.
SHA-1 Hash	twenty bytes		The SHA-1 hash of the preceding three packets – One-Pass Signature Packet, Literal Packet, and Signature Packet – in that order

The recipient of the SDS Authenticated Key shall perform the following processing steps:

- 1) Verify the contents of the OpenPGP Modification Detection Code packet against the SHA-1 hash of the three preceding packets. If the integrity protection checks fail, discard the Security Data Stream Packet.
- 2) Calculate the hash of the OpenPGP Literal Packet using the algorithms specified in the OpenPGP One-Pass Signature Packet.
- 3) Use the Key Identifier from the OpenPGP Signature Packet to locate the Broadcast Gateway Public Key stored on the Exciter. If this Key Identifier does not correspond with a Broadcast Gateway Public Key known to the Exciter, discard the Security Data Stream Packet.
- 4) Using the Broadcast Gateway Public Key from (3), verify the signature data contained in the OpenPGP Signature Packet against the hash value calculated in (2). If the signature does not verify, discard the Security Data Stream Packet.
- 5) Process the contents of the OpenPGP Literal Packet as specified in Section 9.4.5.4 (for Authentication Key Payloads) or Section 9.4.5.5 (for Public Key Distribution Payloads).

9.4.5.4 Authentication Key Payload

The Authentication Key Payload shall be formatted as an OpenPGP Literal Packet with the following content:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 11	The format of the Packet Header shall be as described in RFC 4880 [22] Section 4.2 for New Format Headers
Data Format	one byte	'b'	Binary Data
File Name string length	one byte	'1'	Length of following "file name" field
File Name (= Authentication Key Number)	one byte	'1' through '4'	The index number to which the following Authentication Key Value is assigned.
Date/Time	four bytes	= RTP Timestamp	Carries RTP Timestamp format from Table 9.2 indicating time of first use of Key
Authentication Key Value	32 bytes		The value of the Authentication Key

A recipient of the Authentication Key Payload shall update the HSM Token containing the Authentication Keys with the new Authentication Key Value for the specified Authentication Key Number.

9.4.5.5 Public Key Distribution Payload

The Public Key Distribution Payload shall be formatted as an OpenPGP Literal Packet with the following content:

Field Name	Size	Value	Description
Packet Header	one to five bytes	Packet Tag = 11	The format of the Packet Header shall be as described in RFC 4880 [22] Section 4.2 for New Format Headers
Data Format	one byte	'b'	Binary Data
File Name string length	one byte	'0'	Length of following "file name" field
Date/Time	four bytes	= RTP Timestamp	Carries RTP Timestamp format from Table 9.2 indicating time of creation of the Key
Public Key Identifier	eight bytes		The value of the Network-unique Key Identifier used to access this Public Key
Public Key Data	Remainder of the packet length		The DER encoded value of the Public Key

The recipient of the Public Key Distribution Payload shall update the HSM Token containing the Broadcast Gateway Public Keys with the new Public Key Data for the specified Public Key Identifier.

9.4.6 Security System Implementation

Certain aspects of Security System implementation involve physical security and interoperability in addition to interoperability of data interchange protocols. In particular, it is critical that certain Security System secrets not be exposed to potential interception by an adversary attempting to carry out a Man-In-The-Middle (MITM) attack. Primary among such secrets are the Private Keys of all of the equipment on the network. They should not be exposed through communicating them anywhere, and even exposing them within the equipment that must use them can lead to a security breach. For these reasons, protection of the Private Keys of all equipment as well as of other secrets such as Authentication Keys is entrusted to hardware-based Cryptographic Tokens. The Tokens operate using a common interface defined in [26] and related documents referenced in the cited work. They have both internal processing capability and non-volatile storage within their hardened physical instantiations. The Tokens connect to equipment with which they work through a USB

interface. Physical security for the Tokens themselves also is a consideration in overall network security.

9.4.6.1 Token Initialization

To establish a secure STLTP network, it is necessary to initialize a Cryptographic Token for each equipment unit (Broadcast Gateway or Exciter) on the network. It is anticipated that Broadcast Gateways will include the necessary facilities for managing the initialization process. It is assumed that the Broadcast Gateway used for initialization will have at least two USB ports for connecting Tokens. The initialization process will include the following steps, which assume redundant Broadcast Gateways and multiple Exciters:

- 1) Install a Token into the 1st USB port of the Broadcast Gateway to be used for initialization purposes
- 2) Generate a Public/Private Key pair within the Token of step 1.
- 3) Install a Token to be used in another Broadcast Gateway into the 2nd USB port of the Broadcast Gateway of step 1 (if none needed, continue with step 7)
- 4) Generate a Public/Private Key pair within the Token of step 3
- 5) Copy the Public Key from the Token of step 3 to the Token of step 1
- 6) If another Token for use in a Broadcast Gateway is needed, return to step 3; otherwise, continue with step 7
- 7) Install a Token to be used in an Exciter into the 2nd USB port of the Broadcast Gateway of step 1
- 8) Generate a Public/Private Key pair within the Token of step 7
- 9) Copy the Public Key from the Token of step 7 to the Token of step 1
- 10) Copy the set of Public Keys from all the Tokens of steps 3 – 5 for use in Broadcast Gateways from the Token of step 1 to the Token of step 7
- 11) If another Token for use in an Exciter is needed, return to step 7; otherwise, continue with step 12
- 12) Reinstall a Token from steps 3 – 5 for use in a Broadcast Gateway into the 2nd USB port of the Broadcast Gateway of step 1
- 13) Copy the Public Keys of all other Tokens of steps 3 – 5 for use in Broadcast Gateways from the Token of step 1 to the Token of step 12
- 14) Copy the Public Keys of all Tokens of steps 7 – 11 for use in Exciters from the Token of step 1 to the Token of step 12
- 15) If another Token for use in a Broadcast Gateway remains to be completed, return to step 12; otherwise continue to step 16
- 16) Remove the last Token for use in a Broadcast Gateway from the 2nd USB port of the Broadcast Gateway of step 1
- 17) Install a Token processed in steps 3 – 5 for use in a Broadcast Gateway into another Broadcast Gateway and repeat until all Broadcast Gateways have Tokens installed
- 18) Install a Token processed in steps 7 – 10 for use in an Exciter into an Exciter and repeat until all Exciters have Tokens installed
- 19) Save initialized but unused Tokens as spares

9.4.6.2 Security in a Redundant Environment

Broadcast transmission systems often are constructed with redundant equipment to improve the reliability of those systems. Redundant operation must be considered when configuring an STLTP Security System. Fortunately, most of the considerations have been taken into account in the design of the Security System itself.

In Transmitters, multiple Exciters frequently are installed. In the STLTP Security System, Exciters are treated independently, whether they are installed in individual Transmitters or installed together in a single Transmitter. Thus, nothing needs to be done to facilitate redundant Exciters in a Transmitter other than installation of an initialized Cryptography Token in each Exciter.

With respect to redundant Broadcast Gateways, somewhat more is required. Because they serve as the sources of the content and of the security data and Keys, they generally will operate in a switched configuration in which one unit will be on line at a time, with additional units performing the same functions in a synchronized way and ready to be switched on line at any time when a downstream monitor senses a problem in the unit currently on line. When no security functionality is applied, the redundant operation adds only a relatively small increment to the complexity of operation. When security functionality is active, and especially if a failure occurs at a critical time, considerably more complexity in the security system is required to assure a smooth and reliable transition between Gateways.

Two specific, potential conditions must be considered with respect to switching between redundant Broadcast Gateways. First, since monitors typically will examine the outer, Tunnel Packet Stream and switch on a packet boundary in that Stream, and since Gateways can be set to rotate or randomize their use of the different Authentication Keys that are available on the network, it becomes necessary to synchronize the selection of the particular Key in use for each Tunneled Packet and to assure that the same one is used by all Gateways in the redundant set. A further implication is that all Gateways must have available the same set of Keys at all times, requiring additional synchronization in the Key update process. Another complication can occur when the Gateway that is on line is in the process of updating Keys on the Exciter(s) when it fails. In that case, the redundant Gateway that goes on line next must continue the process of downloading the same Key to the Exciter(s) until the new Key is fully installed. At the current time, the requirement for such synchronization between Broadcast Gateways points to the need to obtain them from a common source for any particular redundant installation.

9.4.6.3 Data Processing and Storage

Cryptographic Tokens have the capability to provide both data processing and non-volatile data storage. For interoperability between Tokens and the host systems in which they operate, it is necessary that Tokens selected for use in the STLTP Security System have certain characteristics and support certain algorithmic functions. Functions needed by the Security System but not provided by Cryptographic Tokens must be provided by the host systems. To minimize and normalize the security workload on the host systems, Tokens shall have the following features:

- Support for USB 2.0 or greater using a USB Type A connector
- Non-Volatile Memory of at least 80 kB
- Support for OASIS PKCS #11 Cryptographic Token Interface API Version 2.40 [26]
- Data processing for NIST AES-256 [20]
- Data processing for SHA2-384 Hash [27]
- Data processing for the OpenPGP Message Format [22]

- Data processing for Elliptic Curve Cryptography in OpenPGP [23]
- Data processing for the Elliptic Curve Diffie-Hellman KDF [23]
- Data processing for the Elliptic Curve Digital Signature Algorithm [24]
- Support for NIST “Curve P-384” [24]

Cryptographic Tokens found in the marketplace as of the date of this standard do not directly support GCM/GMAC data processing. Consequently, for the GMAC Authentication function described herein, the data processing shall be divided between the Token and the host system. Since GCM/GMAC is a mode of operation of the Advanced Encryption Standard (AES) and since data processing for AES is included in available Cryptographic Tokens meeting the specifications above, application of the underlying AES block cipher process shall take place in a Token, while the remaining processes involved in forming and applying GMAC shall take place in the host. Easing the processing burden and speed requirement on the Token is the factor that only one application of the AES-256 cipher process is required for each invocation of GHASH and GCTR, which comprise GMAC Authentication and are applied once per authenticated packet. Additional processing is required of the host, however, as it performs the remainder of the GHASH and GCTR function processing.

Data stored within a Cryptographic Token can take one of two forms: Secrets and related data generated within the Token (i.e., the Public and Private Keys and any Authentication Keys generated by the Token), and secrets and related data generated by other Tokens in the network and delivered to the Token as part of one of the security or Authentication processes (i.e., Public Keys and Authentication Keys generated by other Tokens). All such Keys and related data shall be stored in the Cryptographic Token installed on each host and shall not be stored on the host itself. Rather, when a host has need of any data stored on a Token or a processed version of such data, the host shall withdraw a copy of the data or processed data from the Token for one use only, and the host shall withdraw another copy of the data or processed data when it is needed again. One use shall constitute the processing of one Tunneled Packet of an STLTP Inner Stream at both Broadcast Gateways and Exciters; the data withdrawn shall be a Hash subkey for use in the authentication processing of one Tunneled Packet of an STLTP Inner Stream. All other STLTP Security System processing shall take place within the Tokens.

It should be noted that, through the storage of all the security data and Keys on Tokens, it becomes relatively easy to replace equipment while retaining the information necessary to resume operation immediately without requiring any refresh process. Further, since all of the security data is tied to a number identifying a Token and not the equipment in which the Token is installed, it also is possible to freely replace a Token, if necessary, with another Token, so long as the Token has been initialized for the correct type of application, i.e., for a Broadcast Gateway or for an Exciter. In the case of replacing a Token with one that was initialized at an earlier time, it may be necessary to update Authentication Keys and possibly Public Keys in the replacement prior to its installation if any such values have been updated in the system while the replacement was not in service.

9.4.6.4 Public Key File Format

In addition to the key delivery methods described in §9.4.3, to enable interchange and storage of Public Keys for system initialization and key value retention, a format is specified for placing required information into text files for these purposes. The text files shall contain copies of the Public Key data itself, of the Fingerprint of the Public Key data, and of the Creation Time of the Public Key. Only one Public Key and its associated data shall be stored in each text file of the

type described in this section. Note: The Public Keys referenced in this section are Public Keys derived using the Public Key Algorithm described in §9.4.3.

The text files shall be formatted with ASCII headers and footers as well as metadata descriptions, as described below. The text files shall contain only ASCII text and data expressed in ASCII characters.

The first line shall begin with a “Title: ” description followed by title information. The second line shall begin with a “Comment: ” description, followed by any comment information. The third line shall begin with a “Creation Time: ” description, followed by the ASCII character representation of the decimal NTP seconds count at the time when the Public Key was created. The fourth line shall be blank.

The fifth line shall contain a fixed header value of “-----BEGIN PUBLIC KEY-----”. Note that the header line begins with five dashes before the text and ends with five dashes immediately following the text. After the header row, the next line shall begin the Public Key identifiers and data, which shall be expressed in the Base64 format, which uses only ASCII characters.

The Base64-formatted text shall begin with a sequence of two Object Identifier (OID) values, one that identifies an Elliptic Curve Public Key (OID = 1.2.840.10045.2.1), and one that identifies the P-384 elliptic curve used in the STLTP Security Data scheme (OID = 1.3.132.0.34). The OIDs shall be followed by the Public Key stored in the file. The Public Key shall constitute 776 bits of data.

Following the end of the Base64-formatted text, the next line shall contain a fixed footer value of “-----END PUBLIC KEY-----”. Note that the fixed footer line begins with five dashes before the text and ends with five dashes immediately following the text.

The fixed footer shall be followed by a blank line.

After the blank line following the fixed footer, the next line shall begin with “Fingerprint: ”, which shall be followed by the Fingerprint data expressed in hexadecimal values represented in ASCII characters, with two characters in the range from 0 – 9 and a – f per byte of Fingerprint data. The Fingerprint value shall be computed according to the procedure and using the component data elements specified for V4 fingerprints in the section on Key IDs and Fingerprints in [22]. The Creation Time value used in calculating the Fingerprint shall be the value included on the third line of the file content.

An example of the text file format follows:

Title: This is a title.

Comment: This is a comment.

Creation Time: 1570618441

-----BEGIN PUBLIC KEY-----

MHYwEAYHKoZIzj0CAQYFK4EEACIDYgAEWN6ZQhII4jbS7Kph0x4QoX4kXn0+/k53
18At30WaO3umME3hMZRoMqklzA8auRoxeVj1OWdnjGMizrrkXfQDp8BEJm6MdaTiw
qtIDNBDbcvEMcuY6cVFTR2Rgi0H6pb

-----END PUBLIC KEY-----

Fingerprint: 5587FDC6F0792985AF59DA4AB2262BD4B32B0CB6

9.4.6.5 Physical Security

For full security of STLTP delivery, the Security System itself must be secured. While the components used are secure electronically, they also must be secured physically. Physical security, in this instance, involves protecting the Cryptographic Tokens from removal, substitution, or reprogramming. Physical designs of Broadcast Gateways and Exciters that permit installing the Tokens inside the equipment or behind lockable doors are among the ways in which physical security of the Tokens can be obtained.

9.5 Studio to Transmitter Link (STL) Transport Protocol (STLTP)

The STL Transport Protocol (STLTP) shall be an RTP/UDP/IP Stream based on the Common Tunneling Protocol defined in Section 6 of this standard.

9.5.1 STL Transport Protocol Design

Figure 9.5 provides a detailed diagram of the portion of the broadcast PHY layer chain described by this section. Refer to Figure 4.2 for a complete diagram of the system architecture.

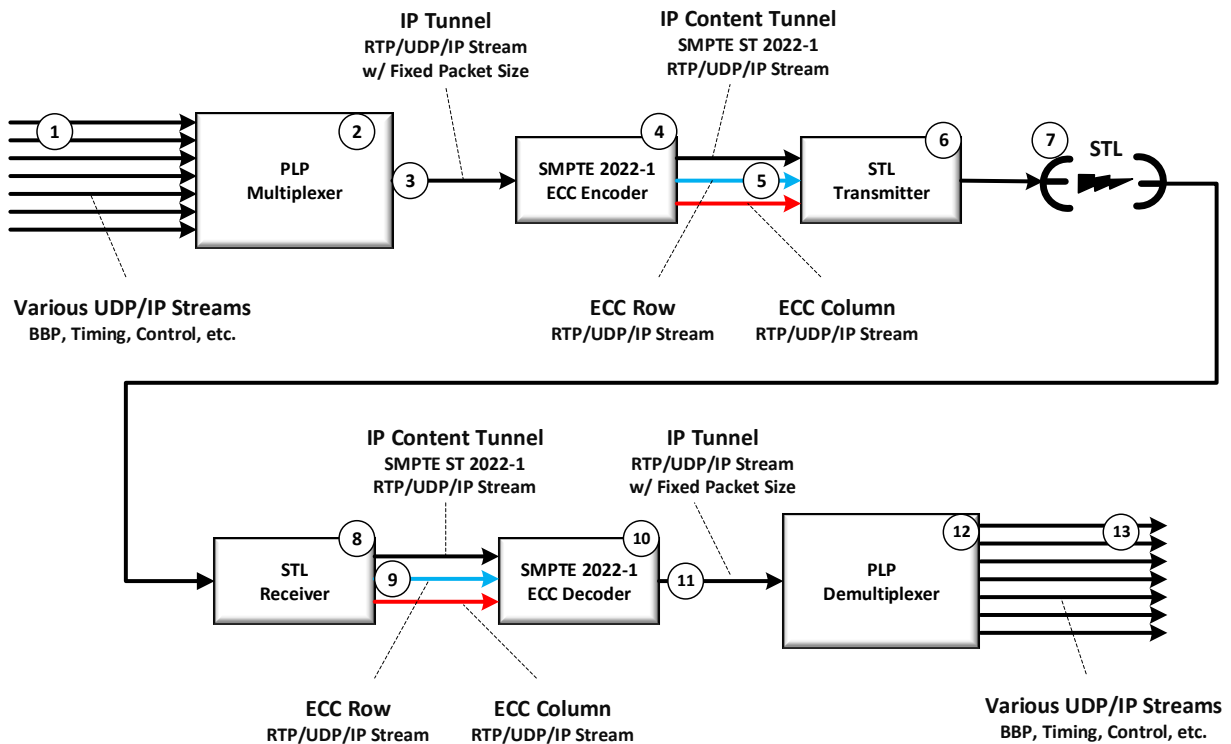


Figure 9.5 STL Transmission diagram.

The following paragraphs describe each of the call-outs, (1) through (13), in Figure 9.5. These are an extension of the figure shown in the Common Tunneling Protocol specifically for the STL use case. Steps (1) through (5) and (9) through (13) correspond to the steps described in Figure 6.1 CTP Tunnel Encapsulation Process.

- 1) The multiple paths represent the RTP/UDP/IP Streams that are generated for each PLP as described in Sections 9.2, 9.3, and 9.3.4. These are referred to as the Tunneled Packet Streams.

- 2) The PLP Mux is configured to accept packets from multiple RTP/UDP/IP multicast Streams to be tunneled.
- 3) The Tunneled Packets are grouped into fixed-size payloads to accommodate the SMPTE ST 2022-1 ECC process [8]. See the CTP Section 6.3.2 for more details.
- 4) The process described in [8] buffers multiple fixed-sized payloads and performs XOR operations on groups of Tunnel Packets producing additional ECC packets.
- 5) The tunnel Stream and up to two ECC packet Streams are sent to the STL Transmitter using up to three RTP/UDP/IP on three separate ports.
- 6) The three Streams are processed by an IP-capable STL Transmitter where they are
- 7) sent via the STL to
- 8) the STL Data Consumer where they are extracted into
- 9) up to three RTP/UDP/IP Streams; i.e., the Tunnel Packet Stream and up to two ECC Streams.
- 10) These Streams are processed by the SMPTE ST 2022-1 decoder. Note that the decoder detects missing packets from the RTP sequence numbers on the tunnel Stream and reconstitutes them from the ECC Stream(s).
- 11) The resultant Tunnel Packet Stream is an in-order collection of fixed-size packets identical to the input Stream described in step (3) above.
- 12) The PLP Demux extracts each packet Stream from the payloads of the Tunnel Packets. The Tunnel Packet RTP headers contain information allowing the packet Streams to be reframed according to the CTP Section 6.3.1.
- 13) The reconstituted IP packets then are forwarded to the remaining stages of the Transmission system.

9.5.2 RTP Encapsulation Example

Following is an example that shows how encapsulation is done using RTP to create a complete Tunneled Packet Stream for SMPTE ST 2022-1 ECC processing, specifically for the STLTP case. Please refer to the Common Tunneling Protocol (CTP) Section 6.3.1 for an example of the general encapsulation process. It examines the case of a Single-PLP emission, specifically, the details of the packing process, including the locations of the several headers of the different layered protocols within the Stream (in Figure 9.6).

Figure 9.6 shows the details of the encapsulated Tunnel Packet. In the top row of Figure 9.6 are packets from three separate Tunneled Packet Streams. In the center row of Figure 9.6 are the details of the construction of the Tunnel Packet that contains the Tunneled Packets or segments of Tunneled Packets from the top row. As shown by the bottom row, the middle row is a more detailed version of the Tunnel Packet structure shown within the dashed rectangle in CTP Figure 6.2.

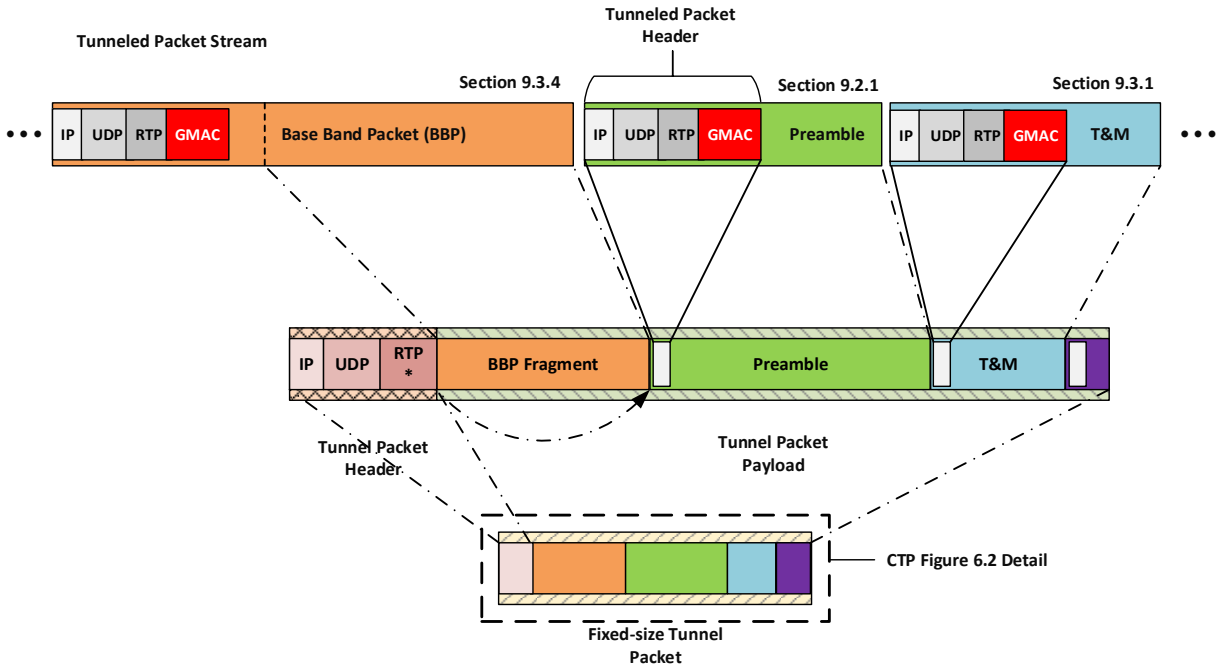


Figure 9.6 Tunneled Packet packing details.

Starting from the top row of Figure 9.6, each of the three Tunneled Packet Streams has its own procession of packets, complete with IP/UDP/RTP header structures, including the header extension containing the GMAC Tag resulting from the packet signing process (see Section 9.4.1). Either complete packets or segments of packets from the Tunneled Packet Streams are multiplexed together on the second row to form the payload of a Tunnel Packet, and another header is prepended to the payload to form a completely new packet. As shown in the bottom row, that packet is the one in a dashed rectangle shown in Figure 6.2 in CTP Section 6.3.1.

Examining at the middle row of Figure 9.6 in more detail, it can be seen that there are four IP/UDP/RTP header sequences. The first, on the left, is the header sequence for the Tunnel Packet. The second header sequence is that of the complete Preamble packet, shown in green, that is fully encapsulated within the Tunnel Packet. The third header sequence is that of the complete T&M packet, shown in cyan, that is fully encapsulated within the Tunnel Packet. The fourth header sequence leads into an initial segment of a Baseband Packet, shown in purple. Also included in the packet described in the middle row is a segment of a Baseband Packet, shown in orange, that does not have its own header because it is a continuation from an earlier Tunnel Packet that contained the header sequence for that Tunneled Packet. Note that the packet headers from the Tunneled Packets may include a header extension that contains the GMAC Tag resulting from the Tunneled Packet signing process.

In the sequence of payload packets and segments described within the payload of the Tunnel Packet, there is a pointer from the Tunnel Packet RTP header to the start of the first byte of the header sequence (i.e., the IP header) that introduces the Preamble Packet. Since the header of the Preamble Packet is the first packet to appear within the Tunnel Packet, the Tunnel Packet RTP header pointer points there. Since the RTP header pointer did not point to the first byte following itself, the implication is that the first data segment within the Tunnel Packet payload is a segment of a Tunneled Packet that was not completed in an earlier Tunnel Packet. A Data Consumer of the

STLTP Stream, upon examination of the Tunnel Packet RTP header, can determine both that the first data within the Tunnel Packet payload is a segment of a previously incomplete packet and where the next Tunneled Packet begins within the Tunnel Packet. By examining the length value contained within the header sequence of that next Tunneled Packet, the Data Consumer can locate the next header within the Tunnel Packet payload. This process can continue until the Data Consumer finds a Tunneled Packet header that indicates its length to be longer than the space remaining available in the Tunnel Packet, implying that there will be at least a segment of the final Tunneled Packet contained within the next Tunnel Packet having the same port address (i.e., going to the same PLP) as the current Tunnel Packet.

9.6 Channel Bonding in Broadcast Gateways

When Channel Bonding is used, as specified in Section K of [3], a Broadcast Gateway shall include a Stream partitioning block that allocates Baseband Packets of Channel-Bonded PLP(s) to two RF channels, as depicted in Figure 9.7. For Channel Bonding, STLTP packet Stream generation from the Broadcast Gateway shall be enabled in one of two operation modes, which can be selected by setting **number_of_channels** as described in Table 6.1. When **number_of_channels** is set to ‘0’, a Broadcast Gateway shall generate two outer Tunnel Data Streams, each of which shall carry one group of inner Tunneled Packet Streams corresponding to one RF channel, as depicted in Figure 9.7 (a). A group of inner Tunneled Packet Streams comprises the Preamble, Timing and Management, Baseband Packet, and Security Data Streams required for one RF channel. When **number_of_channels** is set to ‘1’, a Broadcast Gateway shall generate one outer Tunnel Data Stream that carries two groups of inner Tunneled Packet Streams, one each for the two RF channels, as depicted in Figure 9.7 (b). When **number_of_channels** is set to ‘1’, the two groups of inner Tunneled Packet Streams shall be identified by different port assignments, such that one inner Tunneled Packet Stream group appears on ports 30000 – 30066, and the other group appears on ports 30100 – 30166. These two STLTP configurations are applicable to all the Channel Bonding modes defined in [3] (i.e., plain Channel Bonding and Channel Bonding with SNR averaging). Further details with examples of different Channel Bonding modes are described in Annex E.

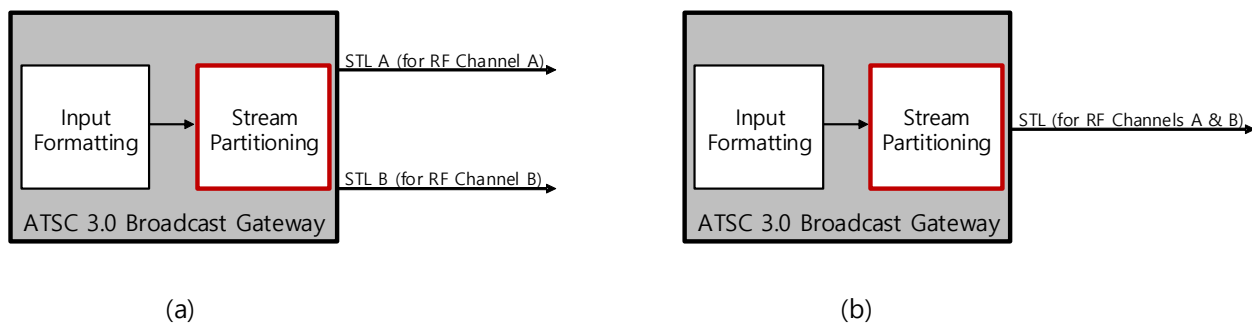


Figure 9.7 Channel Bonding operation modes in a Broadcast Gateway: (a) Sending two outer Tunnel Data Streams, and (b) Sending one outer Tunnel Data Stream that carries two groups of inner Tunneled Packet Streams.

10 TRANSMITTER OPERATION NORMATIVE REQUIREMENTS

There are a number of requirements placed on Transmitter(s) that facilitate their use of the data arriving over the STL from a Broadcast Gateway and that enable operations of Transmitters in

SFNs. The requirements relate to timing of emissions, buffer management, frequency control, offsetting of carrier frequencies for interference mitigation, and the like.

10.1 Timing Manager

A Timing Manager is required functionality in a Transmitter to permit use of the Timing and Management Data sent to Transmitters by a Broadcast Gateway. The Timing Manager shall receive, buffer, and process T&M Data arriving in the form described in Section 9.3 from the STL PLP Demux, as shown in Figure 4.2 (heavy, light-orange line). T&M Data must arrive at a Transmitter in advance of the time for its application. This data shall control the overall operation of the Transmitter and the times at which the numerous processes in the Transmitter are executed. The T&M Data that the Transmitter receives is not intended for emission but rather for control purposes only.

The Timing Manager also shall receive time information directly from a highly accurate source such as the Global Navigation Satellite System (GNSS) or communicated through a data Stream capable of precise time information transfer such as the Precision Time Protocol (PTP). The time information shall be used to accurately control the emission time of the leading edge of the first Bootstrap symbol of each Physical Layer frame. The same timing information used for control of emission timing also shall be capable of use for derivation of a precise frequency reference for control of the Transmitter emission frequency.

The T&M Data shall be stored in a buffer to permit Transmission of Bootstrap Reference Emission Times by the Scheduler to the Transmitter(s) well in advance of those emission times. The size of the buffer should be determined from the maximum advance of delivery time of the Bootstrap Reference Emission Time information permitted by the specification of the Timing and Management Data Stream protocol, as described in Section 9.3.1.

The primary data carried by the Timing and Management Stream is the Bootstrap Reference Emission Time at which each Physical Layer frame is to be emitted. Multiple instances of the Timing and Management Data for a given frame may be sent to Transmitter(s) by the Scheduler. Those instances will be identified as to the Physical Layer frame to which they apply by having a matching time value indicating the Bootstrap Reference Emission Time of the frame. The multiple instances can be held in a buffer and used by the Timing Manager to improve reliability of the T&M Data through application of majority logic or other processing of the redundant copies of the data.

The number of copies of the T&M Data that are sent by the Scheduler to the Transmitter(s) for each Physical Layer frame is indicated in the Timing and Management Data Stream itself, as specified in Table 9.3. Similarly, the number of copies of the Preamble data that are sent by the Scheduler to the Transmitter(s) for each Physical Layer frame also is indicated in the Timing and Management Data Stream. The number of Preamble copies sent by the Scheduler shall be communicated by the Timing Manager to the Preamble Parser so that it can utilize similar processing of the redundant data for purposes of improving reliability of the data.

Other functions of the Timing Manager are:

- Control of the configuration of Bootstrap symbols, waveforms, and the data carried
- Control of the emission times of Bootstrap symbols
- Compensation of the Transmitter-to-antenna emission delay (TAD)
- Control of buffer delays in the Transmitter data processing path
 - Including compensation for STL Network delivery delays in SFNs

- Control of Preamble data insertion following Bootstrap emission
- Parsing of Timing and Management Data for individual Transmitters in SFNs
- Control of Transmitter emission time offsets in SFN operation
- Control of TxID insertion
 - Including insertion level and off

10.2 Preamble Parser

A Preamble Parser is required functionality in a Transmitter to permit use of Preamble data from a Scheduler both to set up the modulation, coding, and other processes in the Exciter and to communicate to receivers the configuration of the signals emitted by the Transmitter. The Preamble Parser shall receive from the STL PLP Demux Preamble data in the form described in Section 9, as shown in Figure 4.2 (heavy, light-blue line) and shall buffer that data. Preamble data will arrive at a Transmitter in advance of the time for its emission. The Preamble data shall be used in two ways: (1) to set up the Transmitter data and signal processing, as controlled by the Scheduler through the delivered Preamble data, and (2) to deliver to the Preamble Inserter block the Preamble data Stream at the correct time for its emission in the transmitted signal.

To achieve the two required uses of the Preamble data, the Preamble Parser shall provide two primary functions. The first function shall be a buffer that holds the Preamble data Stream until it is needed for emission. The output of the buffer shall release the Preamble data Stream to the Preamble Inserter (thin, light-violet line in Figure 4.2) at the appropriate time, as determined by the Timing Manager, so that emission of the Preamble data occurs immediately following emission of the Bootstrap. Both the **length** data at the start and the **crc16** value at the end of the Preamble Payload data packet shall not be emitted by the Transmitter and can be deleted after any use by the Transmitter for error detection. The second function shall be parsing of the data in the Preamble into separate instructions for each of the data- and signal-processing stages of the Transmitter and delivery of those instructions to the processing stages at the appropriate times, as determined by the Timing Manager (thick, dark-blue line in Figure 4.2).

Advance arrival of the Preamble data at the Preamble Parser is for the purpose of providing time for the Exciter to set up its processing stages to the configuration specified by the Scheduler for the waveform, in advance of processing of the next Physical Layer frame. During that setup Period, the Preamble data Stream itself is held in the buffer until it is needed for emission in the appropriate frame.

Multiple instances of the Preamble data for a given frame may be sent to Transmitter(s) by the Scheduler. Those instances will be identified as to the Physical Layer frame to which they belong with a time value matching the Bootstrap Reference Emission Time of the frame. The multiple instances can be used by the Preamble Parser to improve reliability of the Preamble data through application of majority logic or other processing of the redundant copies of the data. The number of copies of the Preamble that are sent by the Scheduler to the Transmitter(s) is indicated in the T&M Data Stream that goes to the Timing Manager. The Timing Manager is expected to communicate that information to the Preamble Parser (thin, red line in Figure 4.2) when it is intended to process the redundant data for improved data reliability.

10.3 Other Transmitter Requirements

To enable use of Transmitters in Single Frequency Networks (SFNs) and to minimize interference between stations, certain additional requirements not covered elsewhere in this standard are placed

on Transmitters. Those requirements relate to frequency accuracy, timing offsets for network management, and co-channel interference minimization through application of frequency and timing offsets.

10.3.1 Frequency Accuracy

The Center Frequency of the transmitted signal shall be either the frequency of the middle carrier in the signal's spectrum when an odd number of carriers is in use or shall be the frequency half-way between the two middle carriers when an even number of carriers is in use.¹ To enable both SFN operation and the co-channel interference mitigation methods discussed in Section 10.3.3, the Transmitter Center Frequency shall be maintained at the nominal Center Frequency $\pm 0.5\text{Hz}$, with a long-term-averaged error of zero. Use of GNSS as a time base and frequency reference can enable such precision.

10.3.2 Transmitter Timing Offsets

To cause Transmitters in a Network to emit their waveforms at times needed for Network shaping, offsets may be applied to emission time values individually at each Transmitter. A Transmitter timing reference with high precision, low jitter, and zero long-term drift can be derived from a single reference and may be delivered to each site or derived at each site. Each Transmitter must be locked to a source equal to the GNSS Time interval. It can be GPS time or TAI, but it must be independent of leap seconds.

A frequency reference with high precision, low jitter and zero long-term drift can be derived from the timing reference. GPS-like sources can be available simultaneously at multiple locations at the studio and Transmitter to enable such an accurate reference.

10.3.3 Center-Frequency and Network-Timing Offsets for Co-Channel Interference Mitigation

The ATSC 3.0 Physical Layer protocol [3] supports low-SNR operation, which can extend coverage areas beyond those served by individual broadcast stations using earlier digital transmission systems. When there are two or more neighboring stations operating with low-SNR configurations on the same RF channel while transmitting different data, reception failure can occur due to channel estimation mismatches in strong co-channel interference environments. For example, when two co-channel stations use the same pilot pattern, FFT size, and guard interval length, after undergoing different propagation channel distortions, the frequency responses of the two arriving signals will be different. In areas where the pilot patterns of the two signals from the co-channel stations overlap, linear addition of the powers of the individual pilots will occur at receiving locations, resulting in creation of a channel model that does not represent the frequency response of either received signal. Receivers at those locations will attempt to correct the channel impairments of the resultant, combined "channel" represented by the summed pilots and consequently will receive neither of the desired signals arriving at the receiver input.

Similarly, when Bootstrap signals from two neighboring co-channel stations are overlapped in the time dimension, receivers may not detect whether such overlap is due to multipath distortion of a single station or to interference from a co-channel station. To enable receivers to separate such interfering signals without requiring design changes, carrier and timing offsets, as described in the following sub-sections, shall be applied to Transmitters so that ATSC 3.0 receivers can successfully receive low-SNR signals in locations where strong, co-channel interference exists. The offsets must be assigned to triads of adjacent stations in such a way that each station in a triad

¹ In the current version of A/322 [3], all possible ATSC 3 Physical Layer frame configurations contain only an odd number of carriers.

has a different offset value. Given that most stations will belong to more than one triad, offsets must be organized on a large-area basis, considering all of the stations on each channel across an entire region.

10.3.3.1 Center Frequency Offset

Pilot pattern overlap always occurs when neighboring co-channel stations use the same FFT sizes, guard intervals, and pilot patterns. Even when neighboring co-channel stations use some different Physical Layer parameters (FFT sizes, guard intervals, or pilot patterns), there may be partial pilot overlap, depending on the choices of pilot patterns. Such pilot pattern overlap (both full and partial overlap) causes the sorts of channel estimation errors described above, which result in reception failure or performance degradation. To avoid overlapping pilot patterns, offset of Transmitter Center Frequencies shall be used. A Center Frequency offset parameter, having values of -1 , 0 , and $+1$ OFDM carriers, is carried in the **tx_carrier_offset** field of the Timing and Management Data. The actual offset frequency value to be applied to each Transmitter shall be the carrier frequency spacing in an 8K FFT for the channel bandwidth in use. The offset frequency value in Hz is defined in the semantics for **tx_carrier_offset** found following Table 9.3. The Baseband Sampling Rate is indicated by the **bsr_coefficient** field of the Timing and Management Data. When signals from two or more neighboring, co-channel stations are receivable in an area, different carrier offset values shall be assigned to each of the co-channel stations. For a group of Transmitters that emits the same waveform (i.e., an SFN), the same carrier offset value shall be applied.

10.3.3.2 Network Timing Offset

Bootstrap signal overlap can occur when neighboring co-channel stations use the same frame size or integer multiples of a particular frame size, which is most likely to occur when frame sizes related to 1-second periods are used. To avoid such time-domain overlaps of Bootstraps from multiple stations, Bootstrap Reference Emission Time offsets of the station waveforms shall be applied by the respective Schedulers. Bootstrap Reference Emission Time offsets shall be applied when multiples or submultiples of Physical Layer frame lengths occur at integer multiples of 1-second periods. Directions of the Bootstrap Reference Emission Time offsets shall be the same as indicated by the **tx_carrier_offset** values sent by Schedulers to their Transmitters in the Timing and Management Data. Positive values of **tx_carrier_offset** shall indicate delays, and negative values of **tx_carrier_offset** shall indicate advances, in the Bootstrap Reference Emission Times sent by Schedulers to their Transmitters. When the assigned offset for a station is zero, its Bootstrap Reference Emission Time shall be 0 ± 1 msec of a Second Tick of TAI time [15]. When the assigned offset for a station is non-zero, its Bootstrap Reference Emission Time shall be $\geq \pm$ (Bootstrap Length + 1 msec) up to \pm (Bootstrap Length + 10 msec) of a Second Tick of TAI time [15], with the sign of ' \pm ' determined by the sign of **tx_carrier_offset**. Bootstrap symbols are 0.5 msec in length, and the minimum number of Bootstrap symbols is 4, making Bootstrap Length 2 msec or greater, depending upon whether the Bootstrap is extended per [2].

An example of using Center-Frequency and Network-Timing offsets among neighboring co-channel stations is shown in Figure 10.1.

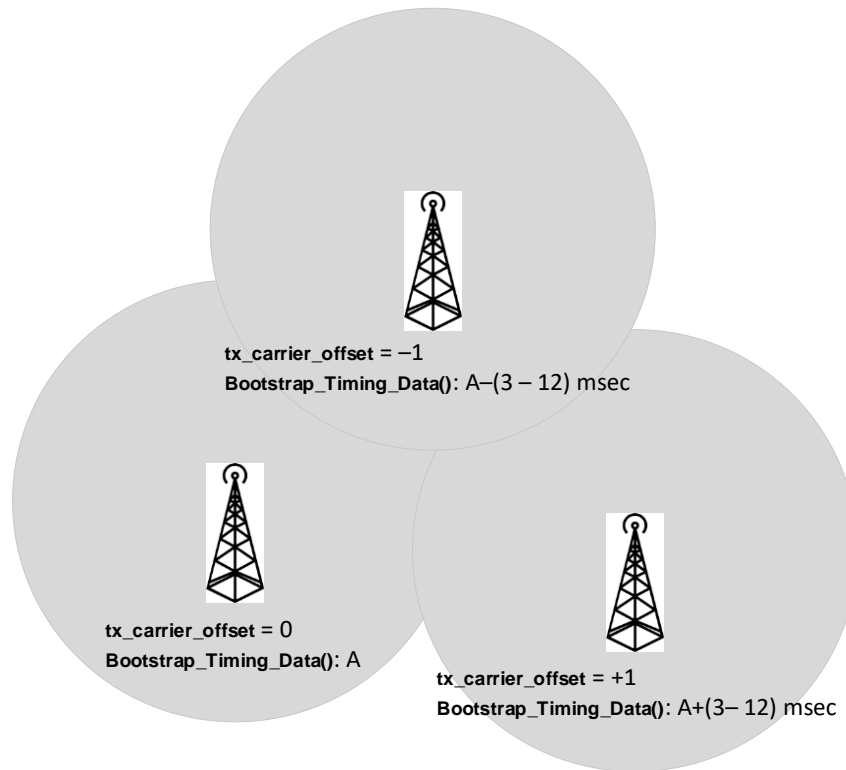


Figure 10.1 Example use of carrier and timing offsets for neighboring co-channel stations.

Annex A Physical Layer Control

A.1 PHYSICAL LAYER RESOURCES

System Manager pre-determined schedule consists of at least the control parameters listed in Section 9 of [3]. The control inputs to the Physical Layer are listed in the following sub-sections. Dependency of some parameters may be based on others from a formula. Those are identified when possible.

A.1.1 Bootstrap Signaling

bootstrap_major_version – This parameter is defined by [2] and constrained by [3]. Intent is to signal a structure change that is not backwards compatible with existing version(s) of the Bootstrap.

bootstrap_minor_version – This parameter is defined by [2] and constrained by [3]. Intent is to signal a structure change that is backwards compatible with existing version(s) of the Bootstrap.

ea_wake_up – This parameter is defined by [2] to signal the presence of emergency alert information.

min_time_to_next – This parameter is defined by [2] to signal when the next similar Bootstrap type (a Bootstrap that matches the same major and minor version number as the current Bootstrap) will be available in an emission.

system_bandwidth – This parameter is defined by [2] to signal the channel bandwidth of the post-Bootstrap portion of the current Physical Layer frame.

bsr_coefficient – This parameter is defined by [2] and constrained by Table 9.1 in [3] with broadcaster-intended Baseband Sample Rate.

preamble_structure – This parameter is defined by [2] and constrained by Table H.1.1 in [3] with settings for FFT size, guard interval and scattered pilot patterns for L1 Basic Preamble symbols chosen by the broadcaster. Preamble symbols should be at least as robust as the most robust payload settings.

A.1.2 L1-Basic Signaling

L1B_version – This parameter is defined and constrained by [3]. It is minor version signaling of the Preamble L1-Basic signaling structure.

L1B_mimo_scattered_pilot_encoding – This parameter is defined by [3] and signals MIMO pilot encoding scheme used by any MIMO subframes.

L1B_lls_flag – This parameter is defined by [3] to signal if Low Level Signaling (LLS) is available in the current frame. LLS is defined by [4] and is the starting point for finding which services are available on a given broadcast channel.

L1B_time_info_flag – This parameter is defined by [3] to signal the presence of Physical Layer timing information in the current frame.

L1B_return_channel_flag – This parameter is defined by [3] to signal the presence of a dedicated return channel.

L1B_papr_reduction – This parameter is defined by [3] to signal the technique used to reduce the peak to average power ratio within the current frame.

- L1B_frame_length_mode** – This parameter is defined by [3] to signal that the current frame is time-aligned with excess sample distribution to the guard intervals of data payload OFDM symbols or if the current frame is symbol-aligned with no excess sample distribution.
- L1B_frame_length** – This parameter is defined by [3] to signal the time Period measured from the beginning of the first sample of the Bootstrap associated with the current frame to the end of the final sample associated with the current frame when time-aligned frames are used. Frame length has many considerations like the longest segment length as defined in [4], or time interleaving depth that broadcasters must choose depending on desired content and robustness. Sizes are constrained by [3] to be between 50 msec and 5000 msec in 5-msec increments.
- L1B_excess_samples_per_symbol** – This parameter is defined by [3] to signal the additional number of excess samples included in the guard interval of each non-Preamble OFDM symbol of the post Bootstrap portion of the current frame when time-aligned frames are used.
- L1B_time_offset** – This parameter is defined by [3] to signal the number of sample periods between nearest preceding or coincident millisecond boundary and the leading edge of the frame when symbol-aligned frames are used.
- L1B_additional_samples** – This parameter is defined by [3] to signal the number of additional samples added at the end of a frame to facilitate sampling clock alignment when symbol-aligned frames are used.
- L1B_num_subframes** – This parameter is defined by [3] to signal the number of subframes present within the current frame.
- L1B_preamble_num_symbols** – This parameter is defined by [3] to signal the total number of OFDM symbols contained within the Preamble, not including the first Preamble symbol.
- L1B_preamble_reduced_carriers** – This parameter is defined by [3] to signal the number of control units of carriers by which the maximum number of carriers for the FFT size used for the Preamble is reduced. It applies to all Preamble symbols except the first one of the current frame which uses the maximum possible carrier reduction.
- L1B_L1_Detail_content_tag** – This parameter is defined by [3] and signals new information is available in L1 Detail.
- L1B_L1_Detail_size_bytes** – This parameter is defined by [3] to signal the size (in bytes) of the L1-Detail information.
- L1B_L1_Detail_fec_type** – This parameter is defined by [3] to signal the FEC type for L1-Detail information protection. It is a combination of 16K LDPC code length with variety of non-uniform constellations and code rates.
- L1B_L1_Detail_additional_parity_mode** – This parameter is defined by [3] to signal the Additional Parity Mode which gives the ratio (K) of the number of additional parity bits for the next frame's L1-Detail that are carried in the current frame to half of the number of coded bits for the next frame's L1-Detail signaling.
- L1B_L1_Detail_total_cells** – This parameter is defined by [3] to signal the total size (specified in data Cells) of the combined coded and modulated L1-Detail signaling for the current frame and any modulated additional parity bits for L1-Detail signaling of the next frame.
- L1B_first_sub_mimo** – This parameter is defined by [3] to signal whether MIMO is used for the first subframe of the current frame.
- L1B_first_sub_miso** – This parameter is defined by [3] to signal whether MISO is used for the first subframe of the current frame.

- L1B_first_sub_fft_size** – This parameter is defined by [3] to signal the FFT size associated with the first subframe of the current frame.
- L1B_first_sub_reduced_carriers** – This parameter is defined by [3] to signal the number of control units of carriers by which the maximum number of carriers for the FFT size used for the first subframe of the current frame is reduced.
- L1B_first_sub_guard_interval** – This parameter is defined by [3] to signal the guard interval length used for the OFDM symbols of the first subframe of the current subframe.
- L1B_first_sub_num_ofdm_symbols** – This parameter is defined by [3] to signal a value equal to one less than the total number of data payload OFDM symbols, including any subframe boundary symbol(s), present within the first subframe of the current frame.
- L1B_first_sub_scattered_pilot_pattern** – This parameter is defined by [3] to signal the scattered pilot pattern used for the first subframe of the current subframe whether it is SISO or MIMO.
- L1B_first_sub_scattered_pilot_boost** – This parameter is defined by [3] to signal the amplitude of the scattered pilots used for the first subframe of the current frame.
- L1B_first_sub_sbs_first** – This parameter is defined by [3] to signal whether or not the first symbol of the first subframe of the current frame is a subframe boundary symbol.
- L1B_first_sub_sbs_last** – This parameter is defined by [3] to signal whether or not the last symbol of the first subframe of the current frame is a subframe boundary symbol.

A.1.3 L1-Detail Signaling

- L1D_version** – This parameter is defined by [3]. It signals the minor version of the Preamble L1-Detail structure for the current frame.
- L1D_num_rf** – This parameter is defined by [3] to signal the number of BSIDs involved in channel bonding of the current ATSC 3.0 system, not including the BSID of the present channel.
- L1D_rf_id** – This parameter is defined by [3] to index the implicit IDs of the other RF channels involved in channel bonding.
- L1D_bonded_bsid** – This parameter is defined by [3] to signal the Broadcast Stream ID of a separate RF channel that is channel bonded with the current RF channel.
- L1D_time_sec** – This parameter is defined by [3] to signal the seconds component of the time information.
- L1D_time_msec** – This parameter is defined by [3] to signal the milliseconds component of the time information.
- L1D_time_usec** – This parameter is defined by [3] to signal the microseconds component of the time information.
- L1D_time_nsec** – This parameter is defined by [3] to signal the nanoseconds component of the time information.
- L1D_mimo** – This parameter is defined by [3] to signal whether MIMO is used for the given subframe.
- L1D_miso** – This parameter is defined by [3] to signal whether MISO is used for the given subframe.
- L1D_fft_size** – This parameter is defined by [3] to signal the FFT size associated with the given subframe.
- L1D_reduced_carriers** – This parameter is defined by [3] to signal the number of control units of carriers by which the maximum number of carriers for the FFT size used for the given subframe is reduced.

- L1D_guard_interval** – This parameter is defined by [3] to signal the guard interval length used for the OFDM symbols of the given subframe.
- L1D_num_ofdm_symbols** – This parameter is defined by [3] to signal the value equal to one less than the total number of data payload OFDM symbols, including any subframe boundary symbol(s), present within the given subframe.
- L1D_scattered_pilot_pattern** – This parameter is defined by [3] to signal the scattered pilot pattern used for the given subframe, whether it is SISO or MIMO.
- L1D_scattered_pilot_boost** – This parameter is defined by [3] to signal the amplitude of the scattered pilots used for the given subframe.
- L1D_sbs_first** – This parameter is defined by [3] to signal whether or not the first symbol of the given subframe is a subframe boundary symbol.
- L1D_sbs_last** – This parameter is defined by [3] to signal whether or not the last symbol of the given subframe is a subframe boundary symbol.
- L1D_subframe_multiplex** – This parameter is defined by [3] to signal whether the given subframe is time-division multiplexed / concatenated in time with adjacent subframes.
- L1D_frequency_interleaver** – This parameter is defined by [3] to signal whether the frequency interleaver is enabled or bypassed for the given subframe.
- L1D_sbs_null_cells** – This parameter is defined by [3] to signal the number of null cells in the subframe boundary symbols of the current subframe.
- L1D_num_plp** – This parameter is defined by [3] to signal a value equal to one less than the total number of PLPs used within the given subframe.
- L1D_plp_id** – This parameter is defined by [3] to signal a value equal to the ID of the given PLP, with a range from 0 to 63, inclusive.
- L1D_plp_lls_flag** – This parameter is defined by [3] to signal whether the given PLP carries the Low Level Signaling (LLS) defined by [4].
- L1D_plp_layer** – This parameter is defined by [3] to signal a value equal to the layer index of the given PLP.
- L1D_plp_start** – This parameter is defined by [3] to signal a value equal to the index of the data Cell that holds the first data Cell of the current PLP in the given subframe.
- L1D_plp_size** – This parameter is defined by [3] to signal a value equal to the number of data Cells allocated to the given PLP.
- L1D_plp_scrambler_type** – This parameter is defined by [3] to signal the choice of scrambler type for the given PLP.
- L1D_plp_fec_type** – This parameter is defined by [3] to signal the Forward Error Correction (FEC) method used for encoding the given PLP.
- L1D_plp_mod** – This parameter is defined by [3] to signal the modulation used for the given PLP, whether SISO or MIMO.
- L1D_plp_cod** – This parameter is defined by [3] to signal the code rate used for the given PLP.
- L1D_plp_TI_mode** – This parameter is defined by [3] to signal the time interleaving mode for the given PLP. The Scheduler is responsible for selecting the appropriate time interleaving mode to use for each Core PLP configured for an RF channel. An Enhanced PLP follows the time interleaver mode of the Core PLP(s) with which it is layered-division multiplexed.
- L1D_plp_fec_block_start** – This parameter is defined by [3] to signal the start position of the first FEC Block that begins within the current PLP during the current subframe.

- L1D_plp_CTI_fec_block_start** – This parameter is defined by [3] to signal the position, after the CTI, of the first cell of the first complete FEC Block, before the CTI, for the current PLP in the current or subsequent subframe.
- L1D_plp_num_channel_bonded** – This parameter is defined by [3] to signal the number of RF channels, not including the present RF channel, involved in channel bonding of the given PLP.
- L1D_plp_channel_bonding_format** – This parameter is defined by [3] to signal the channel bonding format for the given PLP, whether plain channel bonding or SNR averaging channel bonding.
- L1D_plp_bonded_rf_id** – This parameter is defined by [3] to signal the RF IDs of the channels that are used for channel bonding with the given PLP.
- L1D_plp_mimo_stream_combining** – This parameter is defined by [3] to signal whether the Stream combining option of the MIMO precoder is used in the given PLP.
- L1D_plp_mimo_IQ_interleaving** – This parameter is defined by [3] to signal whether the IQ polarization interleaving option of the MIMO precoder is used in the given PLP.
- L1D_plp_mimo_PH** – This parameter is defined by [3] to signal whether the phase hopping option of the MIMO precoder is used in the given PLP.
- L1D_plp_type** – This parameter is defined by [3] to signal when the given PLP is non-dispersed (i.e., all data Cells of the current PLP have contiguous logical addresses and subslicing is not used for the current PLP) or when the current PLP is dispersed (i.e., not all data Cells of the current PLP have contiguous logical addresses and subslicing is used for the current PLP).
If **L1D_plp_type** has a value of one, the number of subslices and subslice interval are set.
- L1D_num_subslices** – If the given PLP is dispersed, this parameter is defined by [3] to signal a value equal to the one less than the actual number of subslices used for the given PLP within the given subframe.
- L1D_subslice_interval** – If the given PLP is dispersed, this parameter is defined by [3] to signal a value equal to the number of sequentially-indexed data Cells measured from the beginning of a subslice for a given PLP to the beginning of the next subslice for the same PLP.
- L1D_plp_TI_extended_interleaving** – This parameter is defined by [3] to signal whether extended interleaving is to be used for the given PLP.
- L1D_plp_CTI_depth** – This parameter is defined by [3] to signal the number of rows used in the Convolutional Time Interleaver.
- L1D_plp_CTI_start_row** – This parameter is defined by [3] to signal the position of the commutator at the start of the given PLP within the current subframe.
- L1D_plp_HTI_inter_subframe** – This parameter is defined by [3] to signal the hybrid time interleaving mode for the given PLP.
- L1D_plp_HTI_num_ti_blocks** – This parameter is defined by [3] to signal either the number of TI Blocks per interleaving frame, N_{TI} , when **L1D_plp_HTI_inter_subframe**=0 or the number of subframes, P_I , over which Cells from one TI Block are carried when **L1D_plp_HTI_inter_subframe**=1 in the given PLP.
- L1D_plp_HTI_num_fec_blocks_max** – This parameter is defined by [3] to signal a value one less than the maximum number of FEC Blocks per interleaving frame for the given PLP.
- L1D_plp_HTI_num_fec_blocks** – This parameter is defined by [3] to signal a value one less than the number of FEC Blocks contained in the current interleaving frame for the given PLP.
- L1D_plp_HTI_Cell_interleaver** – This parameter is defined by [3] to signal whether the Cell Interleaver is used in the given PLP.

L1D_plp_idm_injection_level – This parameter is defined by [3] to signal the Enhanced PLP injection level relative to the Core PLP.

L1D_bsid – This parameter is defined by [3] to signal the Broadcast Stream ID of the current RF channel.

Annex B Network Configuration Examples

B.1 EXAMPLE STUDIO NETWORK TOPOLOGIES

The following discussion provides various Network topology examples to clarify how systems may be deployed to achieve various capabilities and robustness. It is informative only. Each deployment will be different and may change over time as new capabilities are added.

For environments with a single ALP Stream Producer and Data Consumer, simply connecting the two units with a cross-over cable would suffice. The Network connections would have to be on the same subnet but would require no further setup. Typically, systems are connected using hubs or switches as is depicted in Figure B.1. Note that Figure B.1 also shows using RTP/UDP/IP multicast to move various signaling, NRT data, audio and video components into a conceptual multiplexer. The multiplexer would remove the RTP headers prior to the ALP encapsulation stage, resulting in {4} in Figure B.1, since broadcast receivers do not support RTP as part of the emission standard. Using RTP allows the multiplexer to manage the order and timing of the contribution source Streams which are UDP/IP multicast as defined in [4]. The Data Sources may also supply timing information in the RTP Streams to enable reconstruction of packet spacing within the ALP Stream.

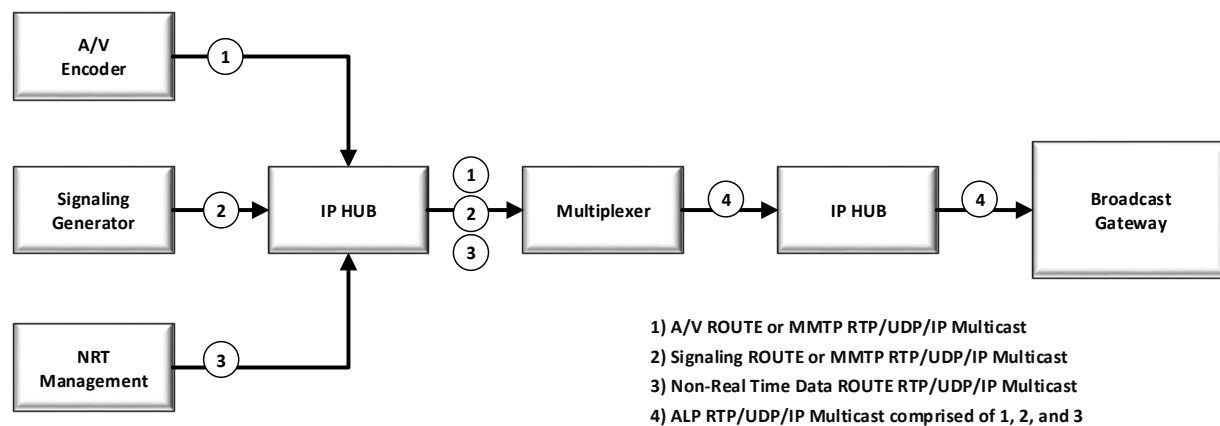


Figure B.1 Simple ALP encapsulation.

Figure B.2 shows a case in which separate A/V encoding systems produce audio and video for two PLPs. In this example, a signaling system would produce signaling information that would be included in both PLP Streams. Signaling could be included in either Stream individually or in a dedicated PLP in the same manner. The non-real-time (NRT) data generation management system also would produce ROUTE Streams that could be included in either or both PLPs as desired. In Figure B.2, NRT data is only included in PLP 1. Note that in this example, all Streams are on separate multicast addresses on the input to the multiplexer. The output of the multiplexer (5 & 6) would be on the same multicast address with different ports. Since there is only one source for each Stream in this example, there is no need for IGMPv3 Source Specific Multicasting (SSM) and all Network connections can be accomplished with relatively inexpensive hubs or switches.

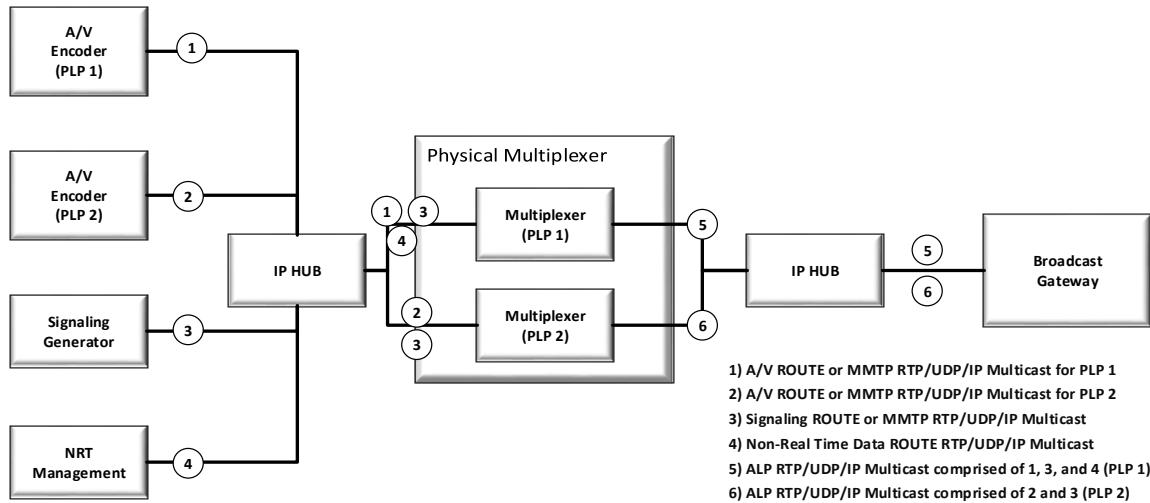


Figure B.2 Multiple PLP example.

A more typical production situation is diagrammed in Figure B.3. This example shows a redundant encoding and signaling environment, redundant multiplexer and redundant output systems feeding into separate STLs (not shown). The multiplexers combine various ROUTE and/or MMTP multicasts into individual ALP Streams, each destined for a unique PLP, removing RTP in the process. A single physical multiplexer likely would create multiple ALP Streams for different PLPs but, functionally, each PLP Stream can be considered separately.

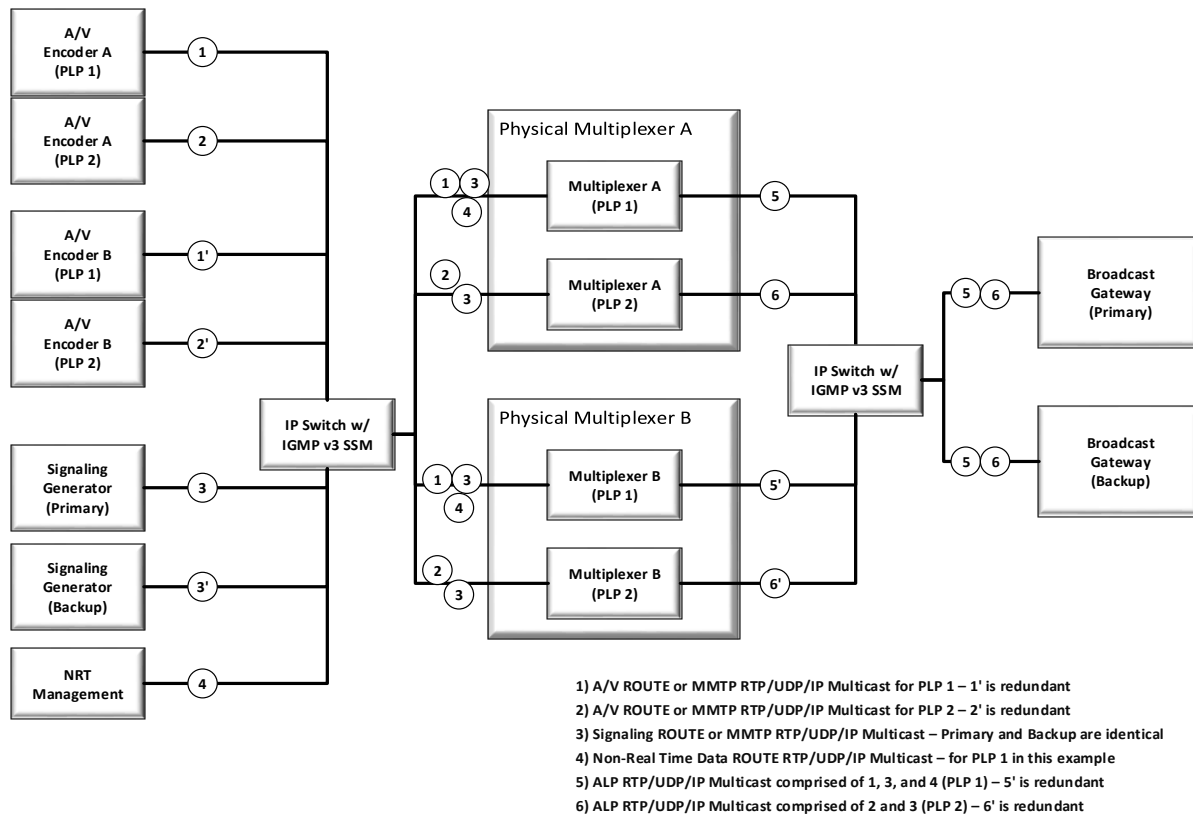


Figure B.3 Fully redundant routing example.

Note that none of the redundant systems needs to be aware of its counterpart. For example, A/V Encoder A (PLP 1) can be configured identically to A/V Encoder B (PLP 1) except, of course, for its Network address. The emitted multicasts (in Figure B.3, the circled numbers {1} & {1'}) are identical in all aspects except that the UDP headers would contain different source addresses. Both Multiplexer A (PLP 1) and Multiplexer B (PLP 1) would be configured with the addresses of the PLP 1 A/V Encoders. Initially, they would both begin ALP encapsulation of A/V Stream {1} into ALP Streams {5} and {5'} by joining the multicast address and port using the address of the A/V Encoder A (PLP 1). If any problems were detected with {1} by either multiplexer, either would then switch to the second A/V Stream, {1'}. This would be accomplished by simply leaving the IGMP multicast group for A/V Encoder A (PLP 1) {1} and 'joining' the A/V Encoder B (PLP 1) multicast {1'}. No other communication is required. Similarly, the {5} or {5'} ALP PLP 1 multicasts would feed both Broadcast Gateway systems based on joins from each. Each Broadcast Gateway would necessarily need to know the source addresses of both Physical Multiplexer A and B so that it could switch between the two sources if the quality of any of the Streams degrades.

Note that a single router or switch with IGMPv3 SSM capability could be used to route all the multicast traffic between each system. There are many options for duplicating routers and even having ring and tree configurations. Be aware that all multicast addresses would need to be unique within the broadcast plant Network to enable a variety of routing options unless Virtual Local Area Networks were used.

Annex C Scheduler Functional Description

C.1 OVERVIEW

Section 5 above described the high-level functions and requirements for the Scheduler and its place within a typical ATSC 3.0 headend. This annex provides more detail on the function of a reference Scheduler.

C.2 SCHEDULER OPERATION

The operation of the Scheduler is constrained by a combination of dynamic, quasi-static, and static parameters. The exact definition of these constraints is left to implementation. This document defines a required function, but not the parameters of that function.

There are two categories of metadata associated with content to be delivered. There is metadata associated with the content which must transit the link as they comprise information needed to successfully decode and render the media. There is another class of metadata that is exclusive to the task of scheduling media in an ATSC 3.0 system. This scheduling metadata is referred to as delivery metadata. This section discusses the functions and application of this delivery metadata. There is some discussion with respect to the order in which metadata and content should be delivered in order to optimize channel change speed.

The cascade of functions involved in the process of scheduling is shown below in Figure C.1. This figure contains no requirements as it is merely an example of an architecture. References to time are with respect to Server Current Time (SCT) as discussed in [4].

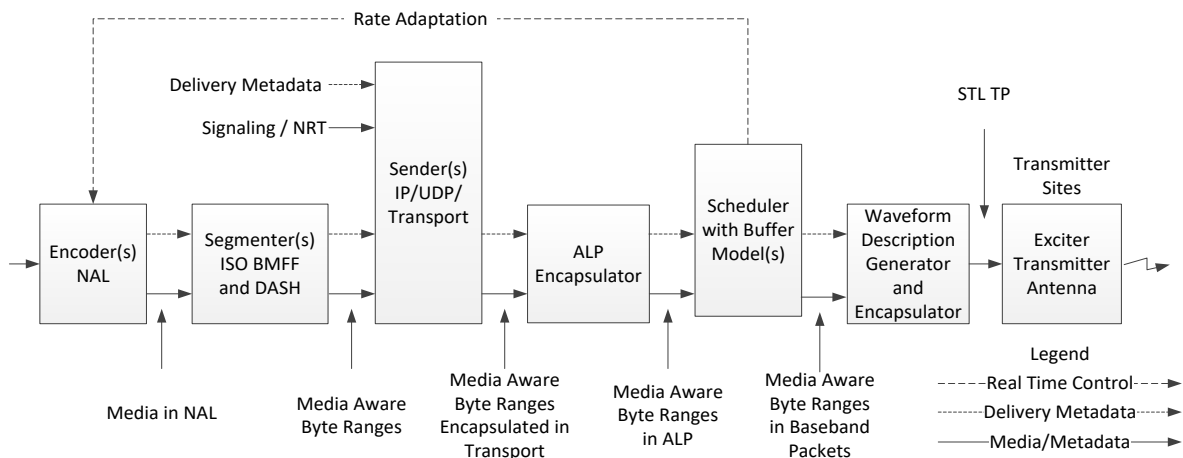


Figure C.1 Cascade of real-time functions involved with Scheduler.

Figure C.2 below shows a conceptual depiction of a Scheduler process flow. The central concept is that there are durations of time under construction with known target radiation times. The process generates an endless sequence of Media Segments, which are mapped into Physical Layer frames and transmitted. There are no requirements contained in this figure; it is an example.

The point of the figure is there is a process in time which results in Media Segments being constructed and radiated. The construction of said Media Segments and Physical Layer frames is a series of intermediate processes that each run on a deadline. For the purposes of this figure, there is one set of Media Segments (e.g., audio, video, and captions) per Service per Physical Layer frame. This is not required and may not be optimum. This figure is depicted with a one-to-one correspondence among Media Segments and Physical Layer frames as a convenience to illustrating the process flow in time. The relationship between an Analyzed Media Duration and Physical Layer frames is discussed in some detail in Section C.4 below.

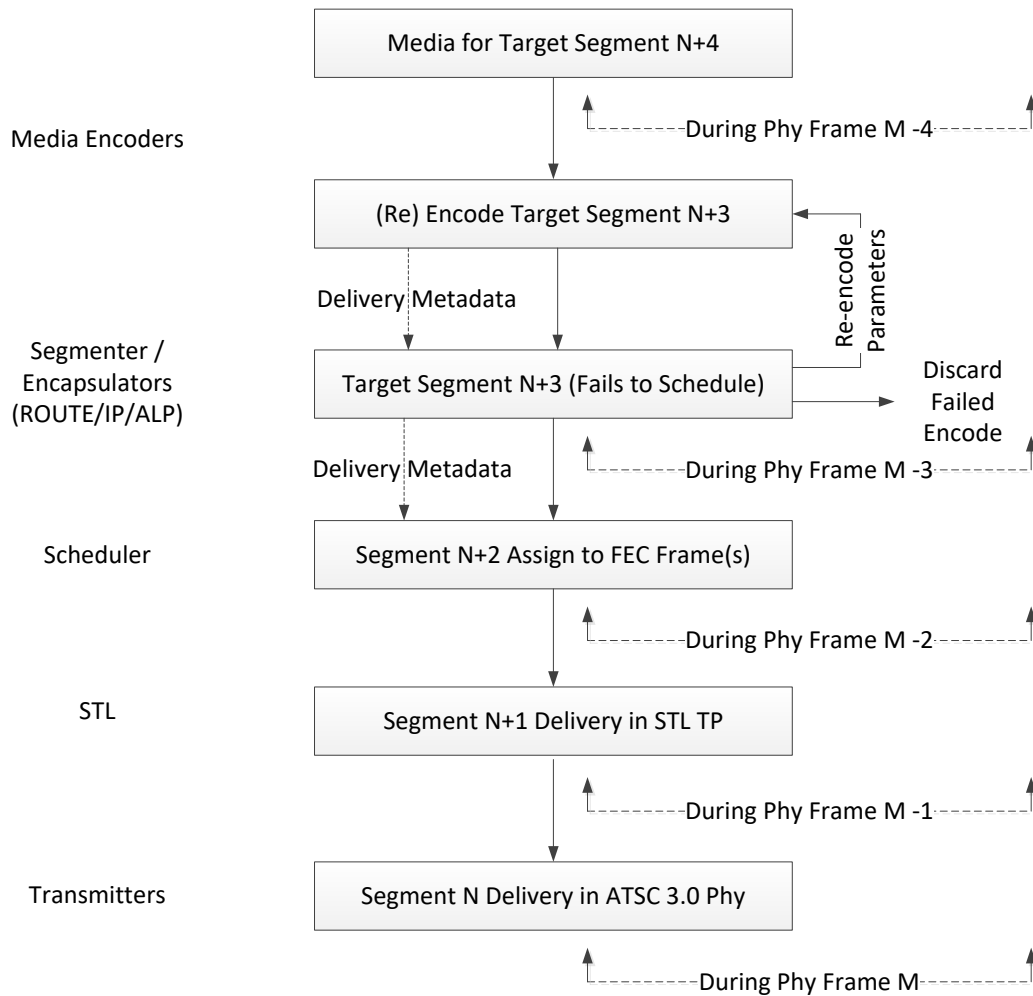


Figure C.2 Example depiction of a high-level Scheduler process flow.

C.3 KEY CONCEPTS OF SCHEDULER DELIVERY METADATA

The concept of earliest delivery at the Physical Layer is illustrated in Figure C.3 below. The data contained in this FEC Frame will radiate from the ATSC 3.0 Transmitter after this Earliest Time. As shown, this description is inclusive of all processes comprised in the generation of the transmitted FEC Frame. If the FEC Frame were being discussed in the context of the Sync and Delivery [4] the Physical Layer FEC Frame contains data that is a Data Delivery Event at the receiver.

Latest Time at the Physical Layer is illustrated in Figure C.3 below. Conceptually, this is constructed and constrained in a manner similar to Earliest Time above. The data contained in this FEC Frame will radiate from the ATSC 3.0 Transmitter before this Latest Time.

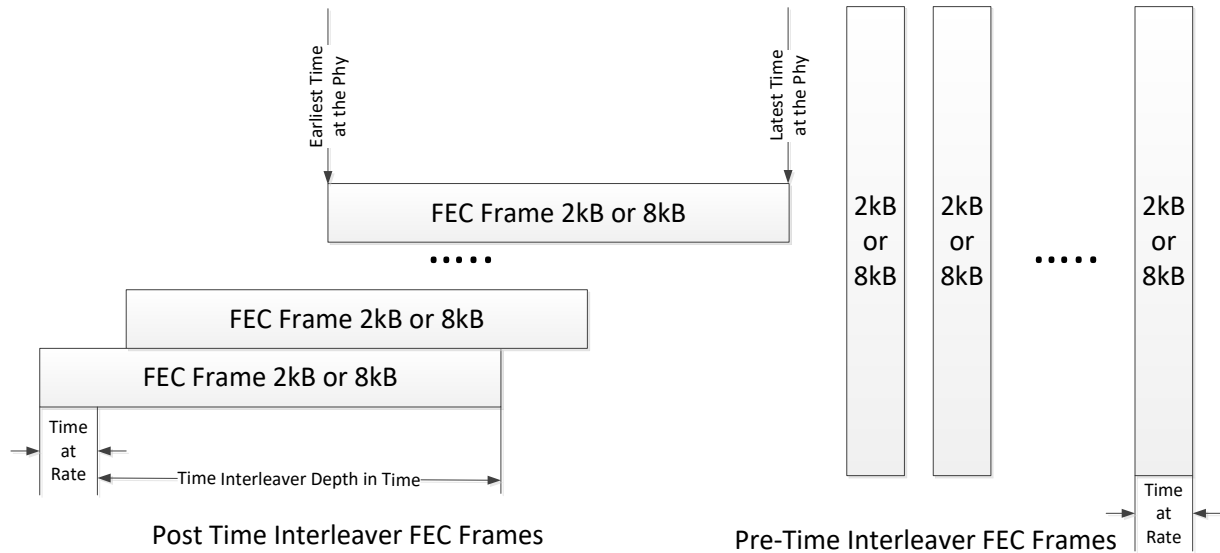


Figure C.3 Illustration of Earliest and Latest Time with a block interleaver.

Figure C.3 depicts the Earliest and Latest Times of a FEC Frame as realized in the Physical Layer. In practice the Scheduler function should receive a notably wider time range of delivery metadata, such that the Scheduler is not over-constrained.

Figure C.4 provides a conceptual view of Earliest and Latest Times in the context of a Media Segment playback duration. Figure C.4 depicts a broader notion of Earliest and Latest Times for multiple use cases.

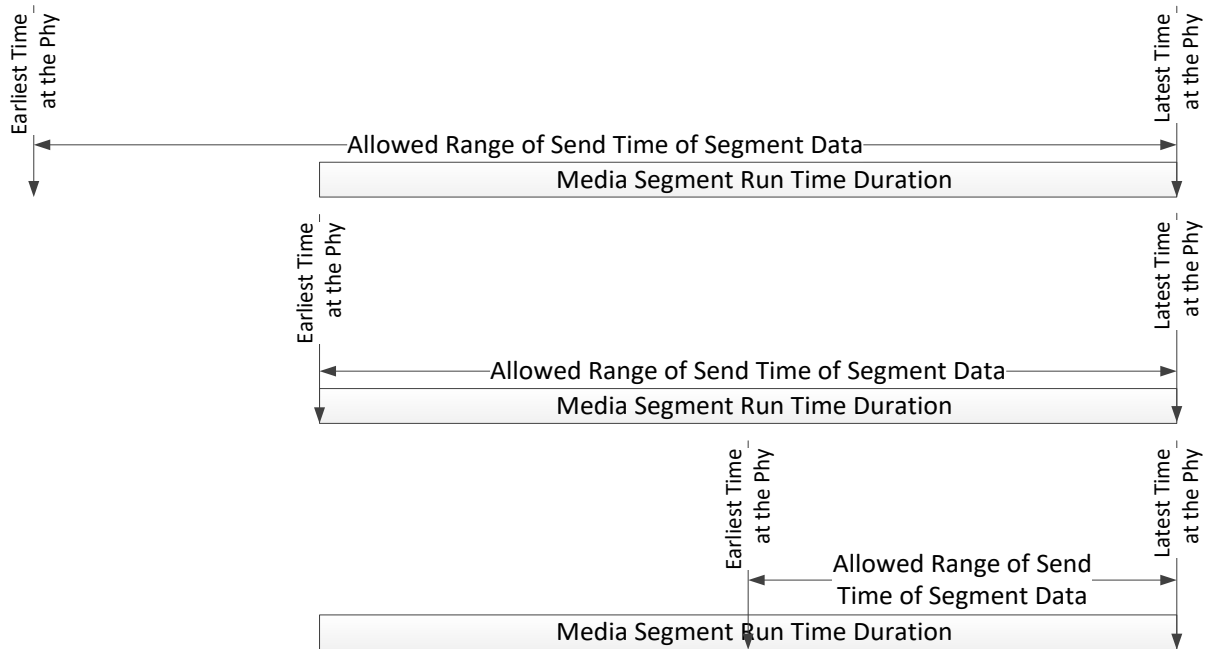


Figure C.4 Illustration of Earliest and Latest Time options relative to Media Segment play.

The use cases in Figure C.4 are intended to illustrate that the concept of Earliest and Latest Times is flexible. These use cases are not requirements. The three use cases depict that the mechanism can describe a variety of implementations. In the top use case, the specification of Earliest would allow the Scheduler to send Media Segment data significantly ahead of the actual demand time for the start of playback. The Latest Time in this example is bounded by the playback deadline for the Media Segment. The time shown is a duration – not the actual delivery or play time. The middle use case depicts a situation in which the media playback and media delivery are constrained to be within a Media Segment playback duration. This might be required, for instance, if there is any switching among PLPs, e.g., at a DASH Period boundary. The bottom use case illustrates a use case in which the delivery of a Media Segment is accomplished in less than a Media Segment run time duration. This sort of use case can result in faster channel change. These use cases are examples and are not intended to constrain the definition or usage of Earliest and Latest Times or actual implementation.




Figure C.5 Example of MDE and associated Earliest and Latest Times.

Figure C.5 illustrates the relationship of an individual MDE to its associated Earliest and Latest Times. The Latest Time in this case is related to 11/24ths of the media second as compared to a 0 Latest Time for the I frame MDE of the same “GoP”. The 0 Latest Time repeats in this example on 1-second boundaries. If in this example, the Physical Layer frame and Media Segment boundaries conform to whole seconds, Earliest Time could be N.0000... and Latest Time N.99999... Given that in-order delivery is enforced for streaming media, the functional requirement is within this whole second for this one-second-duration Media Segment delivery.

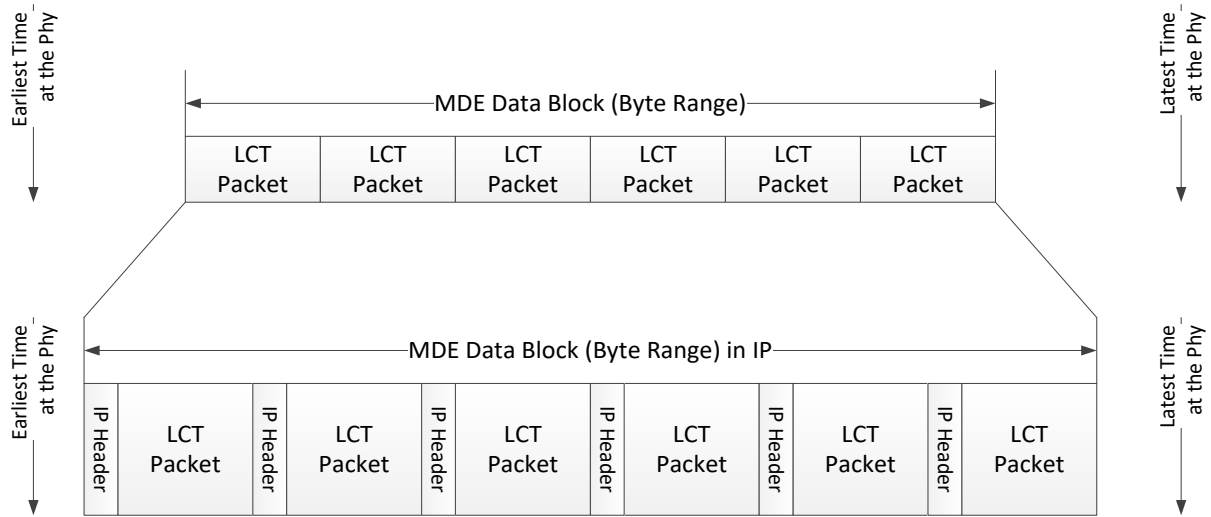


Figure C.6 Relationship of an MDE in LCT packets to IP encapsulation.

Figure C.6 is an example of an MDE with IP encapsulation shown. This figure contains no requirements beyond the Earliest and Latest Times being inherited from the MDE. There is no requirement that there is one LCT packet per IP header. There is no requirement that there is more or less than one LCT packet per IP Header. The only constraint is that IP packets containing the MDE are subject to the indicated Earliest and Latest Times of the source MDE data block.

C.4 HANDLING BOUNDARY CONDITIONS

To this point these concepts have been discussed in broad terms. Known metadata about the timeline of the encapsulated and to-be-encapsulated media is expressed down the stack to the Scheduler. This process may be summarized as the encapsulated media inherits the delivery requirements of contained media. There are assorted boundary conditions to consider. This section illustrates a few of these and describes how delivery metadata may be handled.

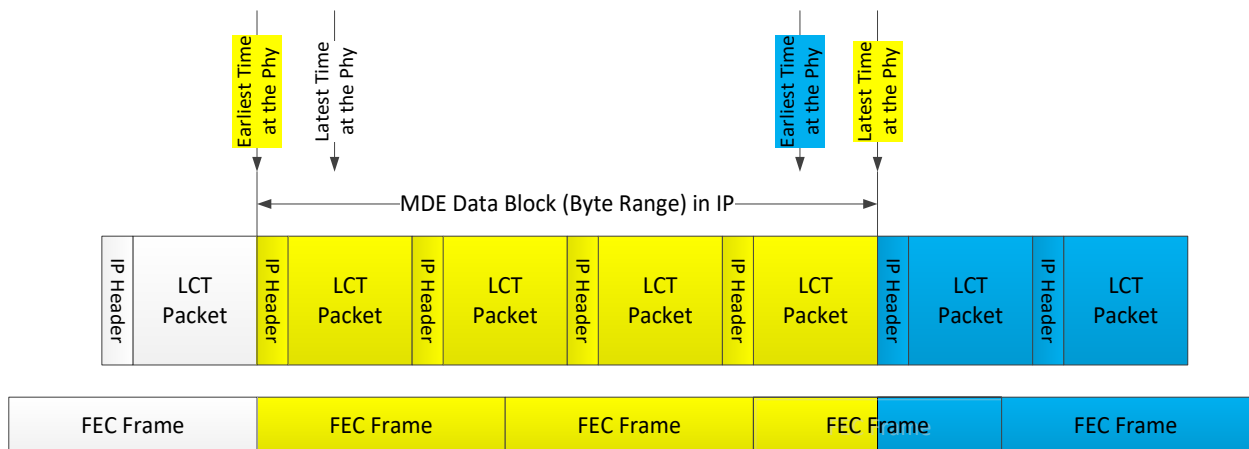


Figure C.7 Media is sent at early boundary; e.g., Period start.

Figure C.7 depicts a use case for which the leading data of the first FEC Frame of the center MDE has a forward-justified IP packet (and ALP) in the first as-assigned FEC Frame. This is caused by the leading LCT packet in for example MDE with RAP having been identified as “Start in new FEC Frame”. This could also occur because the Earliest Time of the center MDE in Figure C.7 is later than the Latest Time of the leading MDE, which is not shown and is more constraining for the Scheduler. The trailing FEC Frame after the center MDE contains a fragment of the trailing MDE. It should be noted that the sequence of FEC Frames may not be continuous in time and is only shown as such for the convenience of illustrating the mapping of IP packets to FEC Frames. The ALP encapsulation is not shown for convenience but is required. This sequence of encapsulated MDEs is assumed to be from a single content type with this media content having been declared as in-order delivery.

C.5 DELIVERY ORDER WITHIN AND ACROSS MULTIPLE SESSIONS

To this point in the discussion, there has been no statement with respect to whether a sequence of IP packets containing LCT packets represents one or more sessions. For the purpose of convenience, it has been assumed that examples have represented a single session, but possibly with multiple content types within the delivered Media Segment files. This section provides an example of a Presentation containing multiple sessions. Perhaps there is one content type per session.

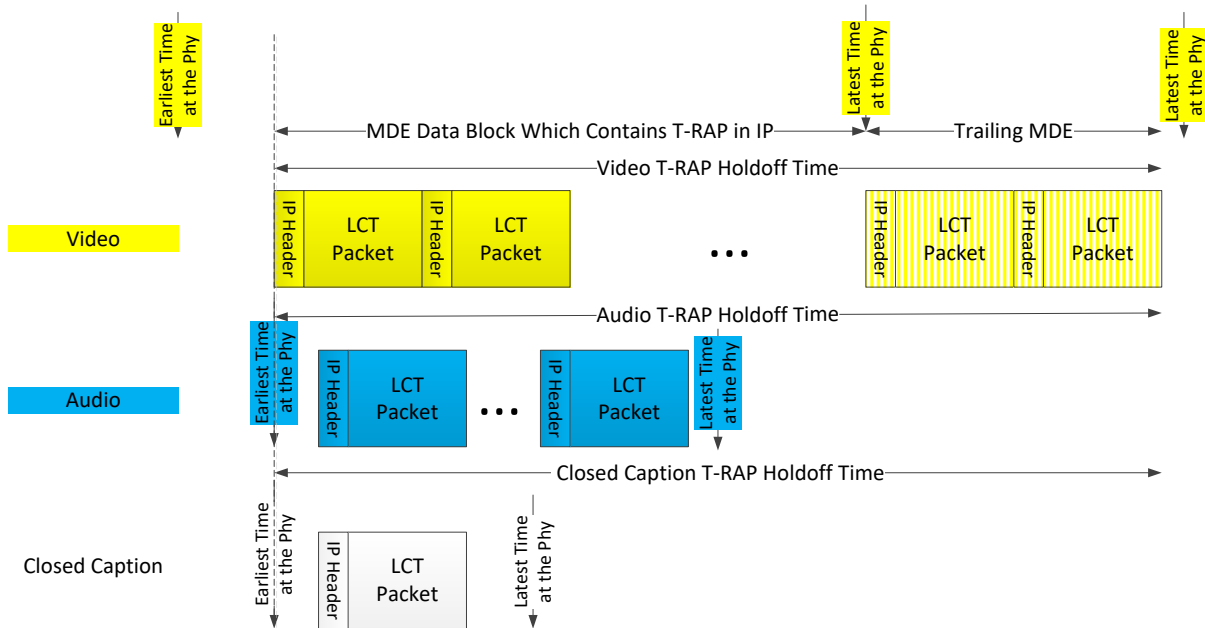


Figure C.8 Example of ordered delivery across multiple sessions.

In Figure C.8 an example is shown of delivery of multiple sessions across a single PLP or multiple PLPs. In this example the MPD is delivered in the T-RAP of video and the DASH client will see and request byte range delivery ahead of Media Segment availability start time. The video typically has the longest RAP cadence, so it is best to align MPD delivery to the video RAP. A key requirement for MDE is the request of byte range delivery ahead of Media Segment

availability start time. The MDE hold-off time due to `EXT_ROUTE_PRESENTATION_TIME - SCT` will expire ahead of Media Segment availability start time.

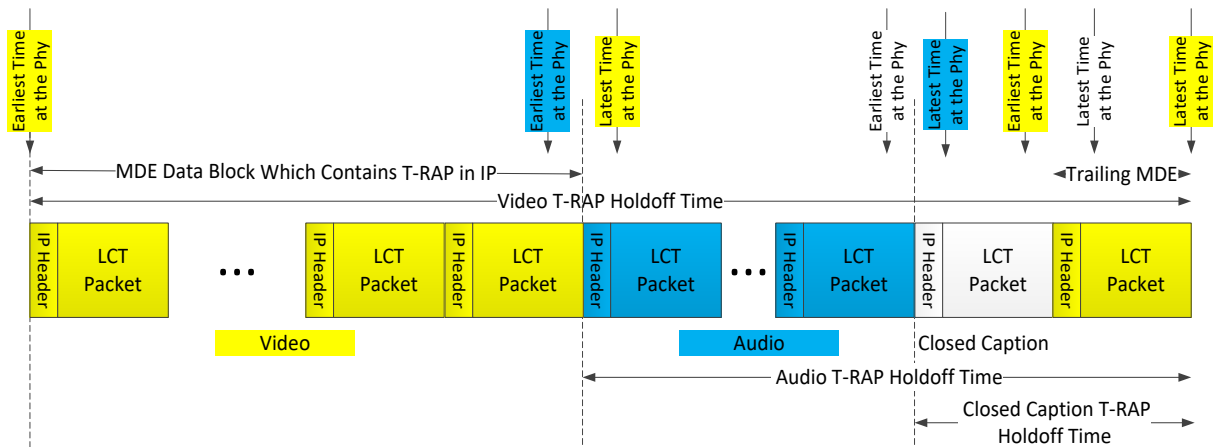


Figure C.9 Explicit delivery order on PLP with multiple sessions.

The order of the to-be-delivered IP packets may be preserved in a PLP. This is shown in Figure C.9. This is not relevant for a service spanning multiple PLPs. This requires that the sequence of LCT packets in IP or compressed IP is preserved in the emission. None of the timing description changes, but the PLP mapper is configured to observe in-order delivery of the IP Stream as delivered to a specific PLP at the Scheduler.

C.6 TIMELINES AND DEADLINES

The delivery of media as ISO BMFF files is an essentially endlessly repeating loop. Each Media Segment of media in Live Streaming is expressed in an ISO BMFF container that has an internal time line that starts at 0 and repeats on a, for example, N-second interval. A DASH MPD has a defined presentation timeline for the ISO BMFF files and the relationship of the send times at the Physical Layer has a nominally static relationship. The MDE method is adaptive to the stack delay of the receiver i.e., it starts as soon as possible without a stall. In each case, the correct timing must be defined by the Segmenter and encapsulator respectively.

If all the video “GoPs” run on the same time cycle, i.e., the IDR frames across Services align in time, most of the benefit of MDE mode will likely accrue to no time guard banding being required for stack delay variability. For Media Segment playback, the MPD must work for the slowest stack. MDE playback starts as soon after reception of a T-RAP as allowed. In order to achieve faster start-up, the use of staggered Media Segment start times likely is required as well as non-uniform bandwidth assignment priority vs. time. These are implementation details, which do not change the fundamental mechanisms, i.e., each MDE and Media Segment have a time deadline for delivery at the Physical Layer expressed by a Latest Time.

C.7 CONCEPT AND PRACTICE OF ANALYZED MEDIA DURATION

The discussion up to this point has assumed that media is encoded based on a one-second Media Segment duration; while this is conceptually convenient, it is a bit simplistic as compared to the capabilities of ATSC 3.0. The generalized construct for creating Physical Layer frames is an Analyzed Media Duration. As discussed above the delivery on the Physical Layer must meet a set of requirements driven by media time line(s). The Physical Layer frames defined by the Scheduler

do not have to correspond precisely to the Media Segment durations, as long as the time delivery requirements are met. The boundaries of Physical Layer frames may not or do not directly correspond to Media Segments, but they are related. The next paragraphs describe how the general method may be applied.

For each ALP Stream input, there is a Data Source. Video Data Sources can be the most demanding in terms of ISO BMFF Media Segment file size and therefore can drive requirements for ALP buffer size and Analyzed Media Duration. Large encoded Media Segment files may or may not be represented as MDE. An Analyzed Media Duration is a Period of time that is sufficient across all ALP Streams provided to the Scheduler such that an analysis Period is not dependent on other scheduling Periods. See Figure C.10 below as an example illustration of IDR positions and how the ALP buffer may go to zero (empty) at the end of each Media Segment. If it does not go empty, it is because a portion of the next Media Segment is already present. This description is valid for Media Segment delivery, irrespective of the ISO BMFF file delivery method for those Media Segments. If the Physical Layer has sufficient capacity and the scheduler has access to later data, some data from the Period of a later Analyzed Media Duration may be sent during the Physical Layer frame(s) corresponding to previous Analyzed Media Duration(s).

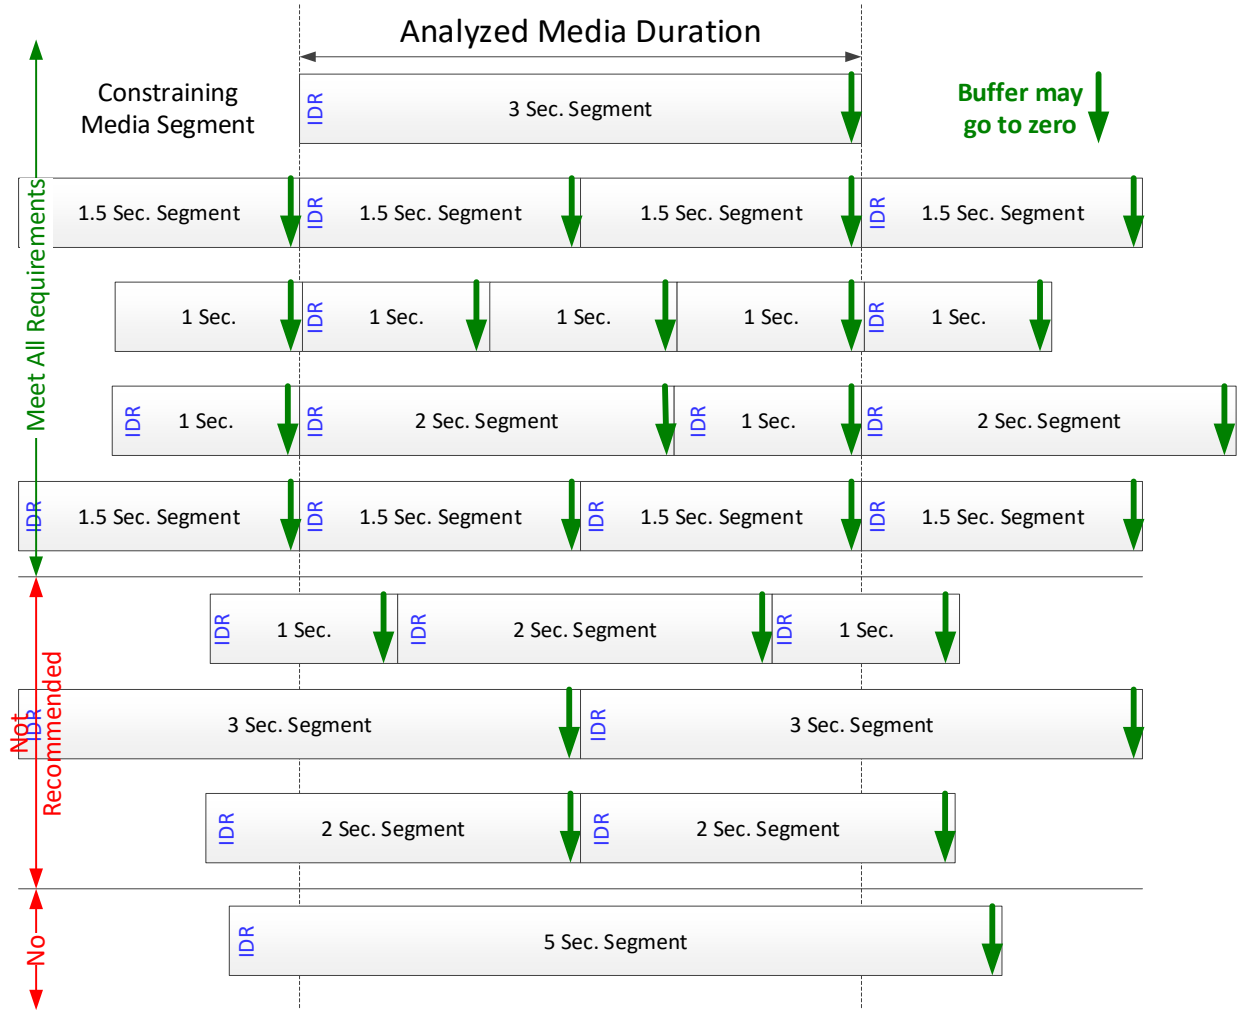


Figure C.10 Analyzed Media Duration.

Figure C.10 above has three main sections. The top section shows examples of Media Segment lengths providing an optimum ALP buffer length of one Media Segment (longest Media Segment among input ALP Streams). This Analyzed Media Duration starts with a Media Segment containing a SAP/RAP and ends with the ALP buffer going to zero. This is a sufficient set of conditions to allow single-Period scheduling. There is no need or requirement to have more than one Analyzed Media Duration in the ALP buffer. Such a configuration allows combinations of smaller Media Segments to be analyzed as well.

The middle ‘Not Recommended’ section of the figure shows examples of ALP buffer lengths that span more than one Media Segment. For the Scheduler to work properly, two Analyzed Media Durations need to be stored in this case.

The bottom section shows an example ALP buffer length that is less than the longest Media Segment present. This could require three Media Segment durations to schedule, which is not notionally supported. In this use case, the best approach would be to increase the Analyzed Media Duration and possibly align the Media Segment time boundary and the Analyzed Media Duration.

In general terms, the possibility of staggered Media Segment start times may be handled by a longer Analyzed Media Duration and use of two corresponding buffers.

There is an absolute time at which negotiations among Data Sources and the Scheduler must be completed and the data delivered. This time is a Data Delivery Deadline and the Scheduler needs to analyze entire Media Segments to meet this Data Delivery Deadline. Current Media Segments must be sent before next Media Segments are loaded to meet each Media Segment Data Delivery Deadline. The ALP buffer should send all Media Segment contents at the end of every Media Segment, independent of whether there is an IDR frame at the beginning. Every Media Segment may be used to ensure ALP buffer flushing. The time deadline for delivery of the complete Media Segment can result in the buffer going to zero. There is no requirement that the buffer go to zero, as it is possible for the Scheduler to “pull forward” media from a future Analyzed Media Duration. Consider that the earliest delivery time can allow the Scheduler to deliver media in a Physical Layer frame to be delivered ahead of the one that could meet the latest requirement.

On each input ALP Stream, Media Segment durations can vary (for example with layered coding), but they usually have an integer relationship in duration. For example, there can be 0.5-second Media Segments with 2-second enhancement Media Segments, or 1.5-second Media Segments with 3-second enhancement Media Segments, but preferably not 0.5-second Media Segments with 0.75-second enhancement Media Segments. Note that 0.5-second and 0.75-second segments could run in an Analyzed Media Duration of 1.5 seconds or 3 seconds. The constraining duration is the longest duration segment when all segment relationships are integer. The smallest Analyzed Media Duration allows all segment durations present to be present in integer numbers. Smaller Media Segments can combine to form the length of the largest Media Segment as shown in Figure C.10.

C.8 SEQUENCE OF REQUIRED DATA AND MEDIA EVENTS FOR ACQUISITION

Figure C.11 shows required data and associated events to achieve streaming media service with or without an application. There are two sequences of events. The first grouping is related to the Physical Layer. The Scheduler needs to understand that packets containing, for example, the SLT and SLS need to occur in tight time proximity to the Bootstrap and Preamble. The cyclic temporal location of the Bootstrap and Preamble likely is aligned to media T-RAP timeline, so as to minimize wait states. Multiple staggered media start times and T-RAPs may require that multiple Bootstraps and the associated signaling are required to minimize channel change time. If ROHC-U [16] header compression is being utilized, then there is a need to synchronize the context refresh to the T-RAP. This should be supported optionally as shown below in Figure C.11. The first Baseband Packets after the Preamble, as depicted below, may be delivered in a dedicated signaling PLP.

The delivery of an application may be incremental, which is to say that each of the utilized or utilizable data objects for the application may not be delivered entirely within the RAP. It is reasonable and good practice to define the T-RAP delivered portion of the app such that service acquisition and start may occur. Such details are out of scope for this document.

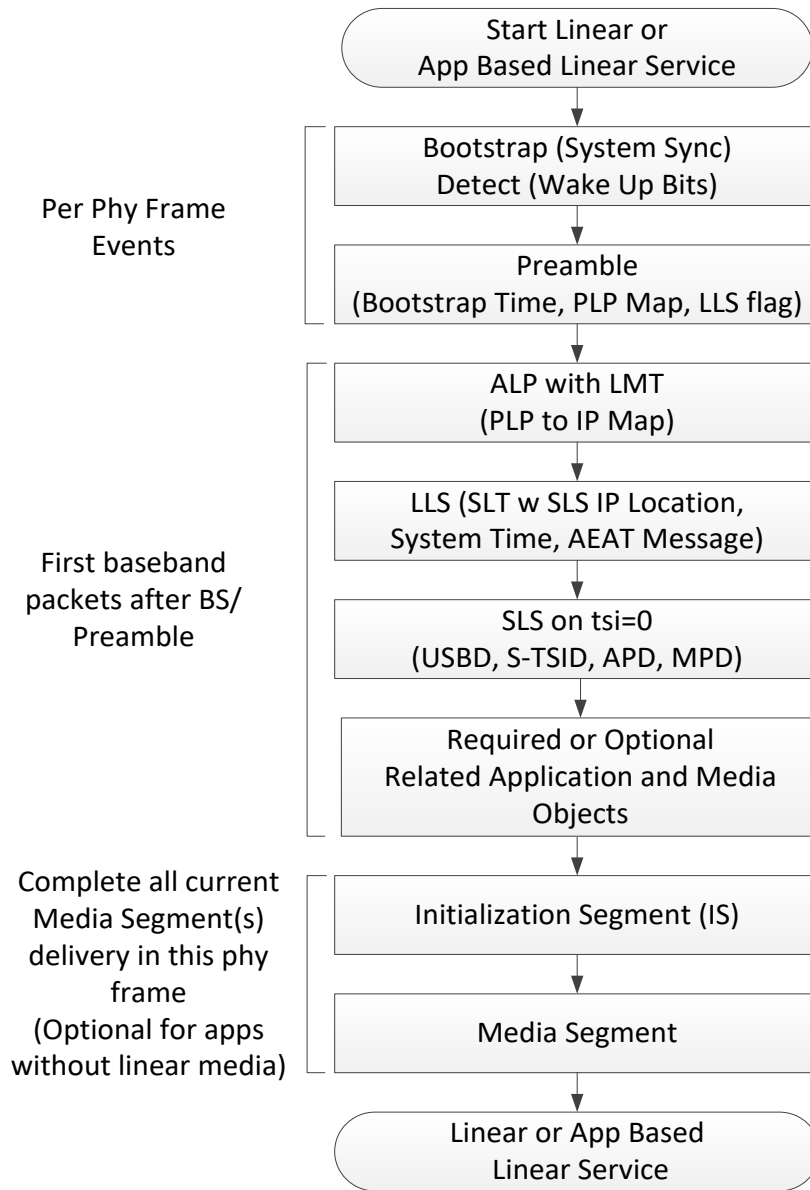


Figure C.11 Order of events to optimize linear service acquisition w/wo an Application.

Annex D Unicast for Outer Tunneling Packets

As described in Section 9, the outer “Tunneling” packets (STLTP) are expected to be delivered through an RTP/UDP/IP Multicast protocol stack that uses a dedicated network link for reliable packet delivery. When such multicast capability is not available over the STL (e.g., when using a public network), the outer “Tunneling” packets can be delivered through an RTP/UDP/IP Unicast protocol stack. Since a public network may not be reliable, which can cause packet drops over the STL, redundant Streams of the outer “Tunneling” packets may be employed by enabling the **redundancy** field in Table 6.1, as shown in the example in Figure D.1.

In an ATSC 3.0 Transmitter, the redundant outer “Tunneling” Streams can be aggregated to further process in the ATSC 3.0 Physical Layer chain (BICM, Framing and Interleaving, and Waveform Generation). The ATSC 3.0 Transmitter may use the sequence numbers carried identically on each of the RTP/UDP/IP Unicast Stream to determine which packets have been received or lost along each of the multiple unicast paths. Note that since the intention of employing unicast is to use unidirectional delivery from a Broadcast Gateway to ATSC 3.0 Transmitter(s), the Broadcast Gateway will have to configure the unicast destination address(es) of the ATSC 3.0 Transmitter(s), which will be assigned by the Internet Service Provider provisioning each delivery service.




Figure D.1 RTP/UDP/IP Unicast for outer “Tunneling” packets using STLTP redundancy.

Annex E Bonded-Channel Distribution

Channel Bonding allows spreading of content over multiple channels. In the current version of this standard, the number of Bonded Channels is limited to two. See Figure E.1. Data is spread over the same PLP_IDs among these channels (per A/322 [3], Annex K).




Figure E.1 Channel Bond over PLP_ID 0 between Channels A and B.

As discussed in Section 9.6, the inner Tunneled Packet Streams for Channels A and B can be carried on two separate outer Tunnel Data Streams, one for each channel, as depicted below in Figure E.2.

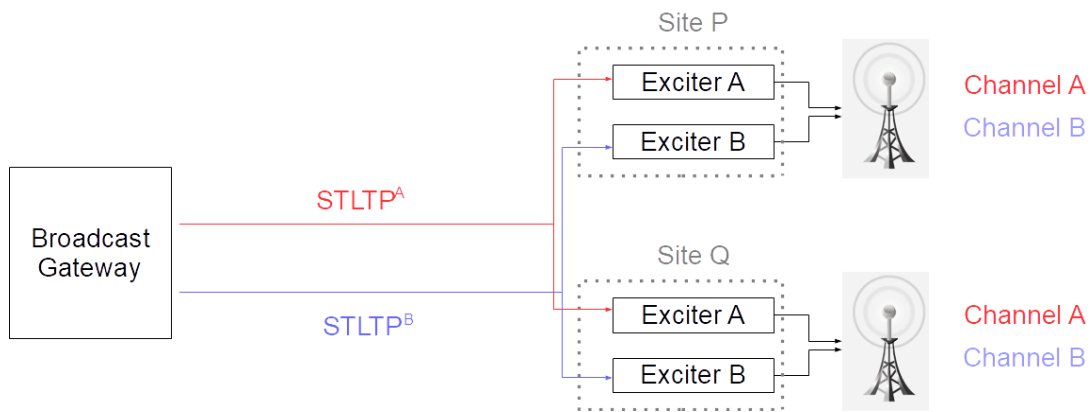


Figure E.2 Channel bonding with two separate STLTPs.

In the configuration of Figure E.2, the **number_of_channels** field is set to “0” in the RTP header of each Tunnel Data Stream. The Tunnel Data Streams each shall contain all the Tunneled Packet Streams required by Physical Layer processing.

Another strategy is to combine all the Tunneled Packet Streams inside one outer Tunnel Data Stream as depicted below in Figure E.3.

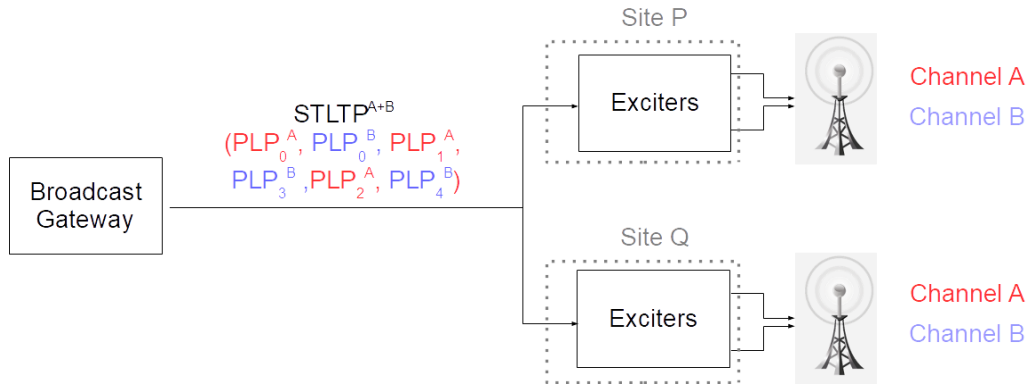


Figure E.3 Channel bonding with combined STLTP.

In this configuration, the `number_of_channels` field is set to “1” in the RTP header of the Tunnel Data Stream. The UDP port number ranges for the two groups of inner Tunneled Packet Streams are 30000 – 30066 and 30100 – 30166.

The Broadcast Gateway shall support both the separate and combined STLTP configurations. The choice between them will depend on the Channel Bonding methods employed plus the network and Transmitter-site topologies.

For example, in the case of Plain Channel Bonding, the Broadcast Gateway may output two separate STLTP Tunnel Data Streams, one for each channel, as shown in Figure E.4. Since there is no cell exchange required in the modulator processing for Plain Channel Bonding, the Physical Layer chain generating the RF signal for a particular channel only needs the STLTP inner Tunneled Packet Streams for its channel.

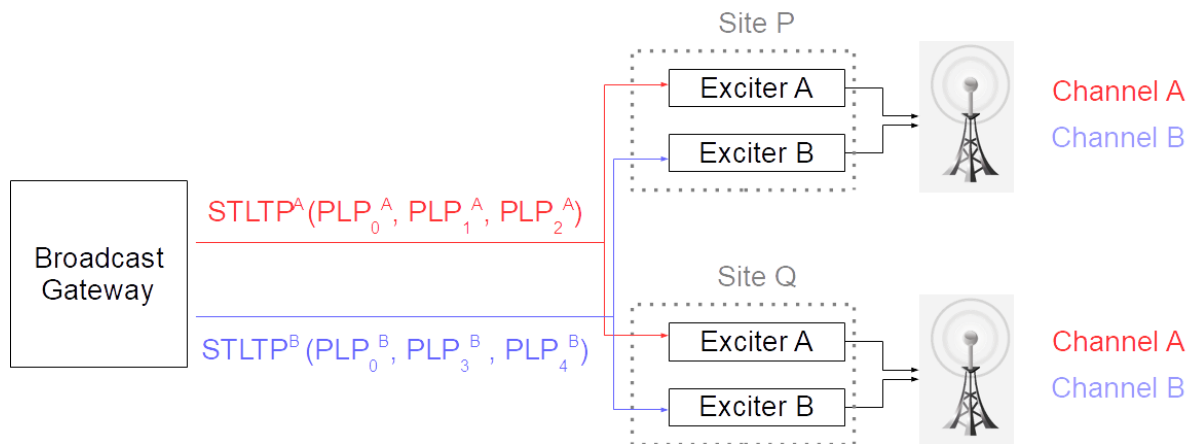


Figure E.4 Example of network topology for Plain Channel Bonding with co-located sites and separate STLTP Tunnel Data Streams.

In the case of SNR Averaging, cell-exchange processing between the two channels is required. Consequently, the modulator requires the inner Tunneled Packet Streams of the bonded PLPs of both channels. The Broadcast Gateway may output a single STLTP Tunnel Data Stream combining the inner Tunneled Packet Streams of both channels, as depicted below in Figure E.5.

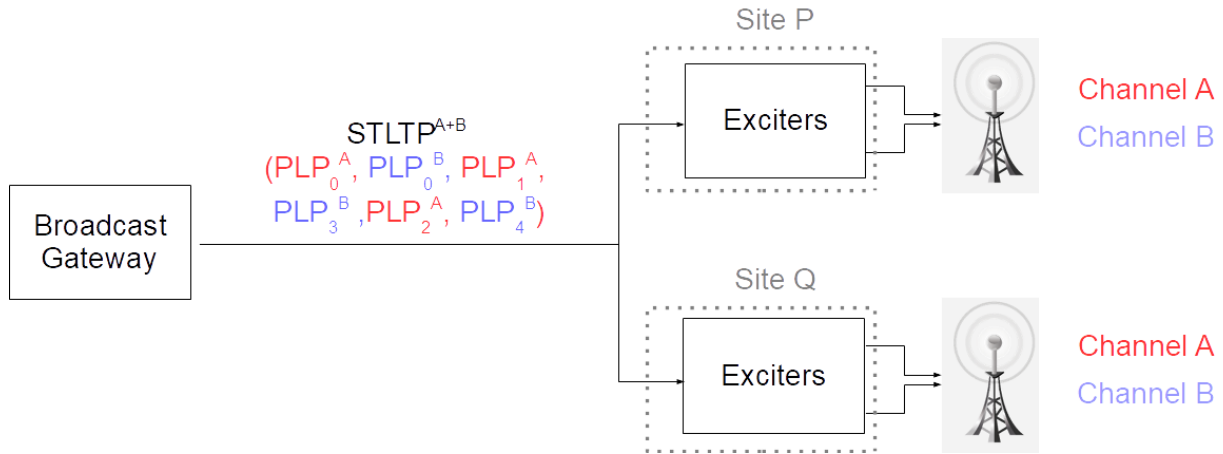


Figure E.5 Example of network topology for SNR Averaging with co-located sites and combined STLTP Tunnel Data Stream.

End of Document