ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.280
(03/2013)

SERIES J: CABLE NETWORKS AND TRANSMISSION
OF TELEVISION, SOUND PROGRAMME AND OTHER
MULTIMEDIA SIGNALS

Digital transmission of television signals

Digital program insertion: Splicing application program interface

Recommendation ITU-T J.280

Recommendation ITU-T J.280

Digital program insertion: Splicing application program interface

Summary

The splicing application program interface (API) focuses on the single server, single splicer configuration. This version of the Recommendation increases the value of Revision_Num from 1 to 2.

History

Edition	Recommendation	Approval	Study Group	Unique ID[*]
1.0	ITU-T J.280	2004-03-15	9	11.1002/1000/7204-en
2.0	ITU-T J.280	2005-12-14	9	11.1002/1000/8668-en
3.0	ITU-T J.280	2013-03-01	9	11.1002/1000/11829-en

_____

[*] To access the Recommendation, type the URL http://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID. For example, http://handle.itu.int/11.1002/1000/11830-en.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Recommendation ITU-T J.280

Digital program insertion: Splicing application program interface

1 Scope

This application program interface (API) creates a standardized method of communication between servers and splicers for the insertion of content into any MPEG-2 output multiplex in the splicer. This API is flexible enough to support one or more servers attached to one or more splicers. Digital program insertion includes content such as spot advertisements of various lengths, program substitution, public service announcements or program material created by splicing portions of the program from a server.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T H.222.0]   Recommendation ITU-T H.222.0 (2012) | ISO/IEC 13818-1:2012, *Information technology – Generic coding of moving pictures and associated audio information: Systems*.

[ITU-T H.262]   Recommendation ITU-T H.262 (2012) | ISO/IEC 13818-2:2012, *Information technology – Generic coding of moving pictures and associated audio information: Video*.

[ANSI/SCTE 35]   ANSI/SCTE 35 2012 (2012), *Digital Program Insertion Cueing Message for Cable*.

[ANSI/SCTE 54]   ANSI/SCTE 54 2009 (2009), *Digital Video Service Multiplex and Transport System Standard for Cable Television*.

[ANSI/SCTE 128]   ANSI/SCTE 128 2010a (2010), *AVC Video Systems and Transport Constraints for Cable Television*.

3 Definitions

3.1 Terms defined elsewhere

None.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 API connection: A TCP/IP socket connection between a server and a splicer for transferring API messages.

3.2.2 back-to-back insertion: Insertion of two or more temporally contiguous sessions without returning to the primary channel between sessions.

3.2.3 channel: A channel is a synonym for a "Service" in DVB terminology, or a "Program" in MPEG terminology.

3.2.4   insertion channel: The insertion multiplex channel(s) that replaces the primary channel in whole or in part of the duration for a splice event.

3.2.5   insertion multiplex: This is the source of the insertion channel. A multiplex produced by a server may under some circumstances exclude program specific information (PSI), thus it is understood that this multiplex may be a non-compliant MPEG-2 transport stream.

3.2.6   multiplex: A multiplex is a collection of one or more channel(s) that may include the associated service information. A multiplex is an MPEG-2 transport stream with the possible exception of an insertion multiplex.

3.2.7   output channel: The channel that is produced at the output of the splicer.

3.2.8   output multiplex: The MPEG-2 transport stream produced by multiplexing one or more output channels.

3.2.9   primary channel: The primary multiplex channel that is replaced in whole or in part. A single primary channel may result in multiple output channels.

3.2.10   primary multiplex: This is the source of the primary channel(s).

3.2.11   server: The device that originates the insertion channel(s) to be spliced into the primary channel(s). This device communicates with the splicer about when and what to splice.

3.2.12   session: A session is the insertion of content (such as spot advertisements of various lengths, program substitution, public service announcements, or program material created by splicing portions of the program from a server). Each session is identified by a unique SessionID.

3.2.13   splice-in: The splice at the start of the insertion. This happens at the time specified in the Splice_Request message.

3.2.14   splice-out: The splice at the end of the insertion. The expected insertion end time is calculated by adding the start time and the duration specified in the Splice_Request message; however, this may occur earlier due to error conditions.

3.2.15   splicer: The device that splices the insertion channel(s) into the primary channel(s). It may receive ANSI/SCTE 35 cue messages. This device also communicates with the server about when and what to splice.

4   Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

API   Application Program Interface

DVB-ASI   Digital Video Broadcast – Asynchronous Serial Interface

ID   Identifier

MLD   Multicast Listener Discovery

MPEG   Moving Picture Experts Group

MPTS   Multi-Program Transport Stream

NTP   Network Time Protocol

PAT   Program Association Table

PCR   Program Clock Reference

PID   Packet Identifier

PMT   Program Map Table

PSI   Program Specific Information

SPTS	Single Program Transport Stream
tcimsbf	twos complement integer, most significant bit first
TCP/IP	Transport Control Protocol/Internet Protocol
uimsbf	unsigned integer, most significant bit first
UTC	Coordinated Universal Time

5   Conventions
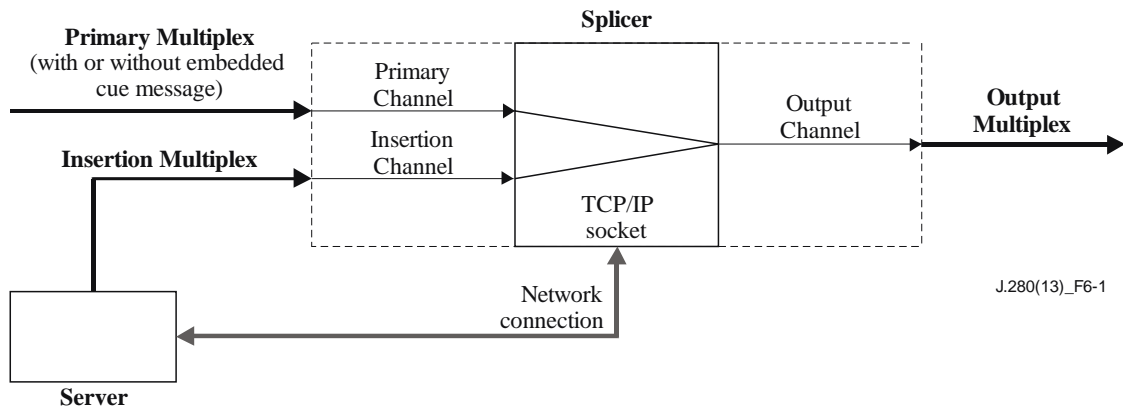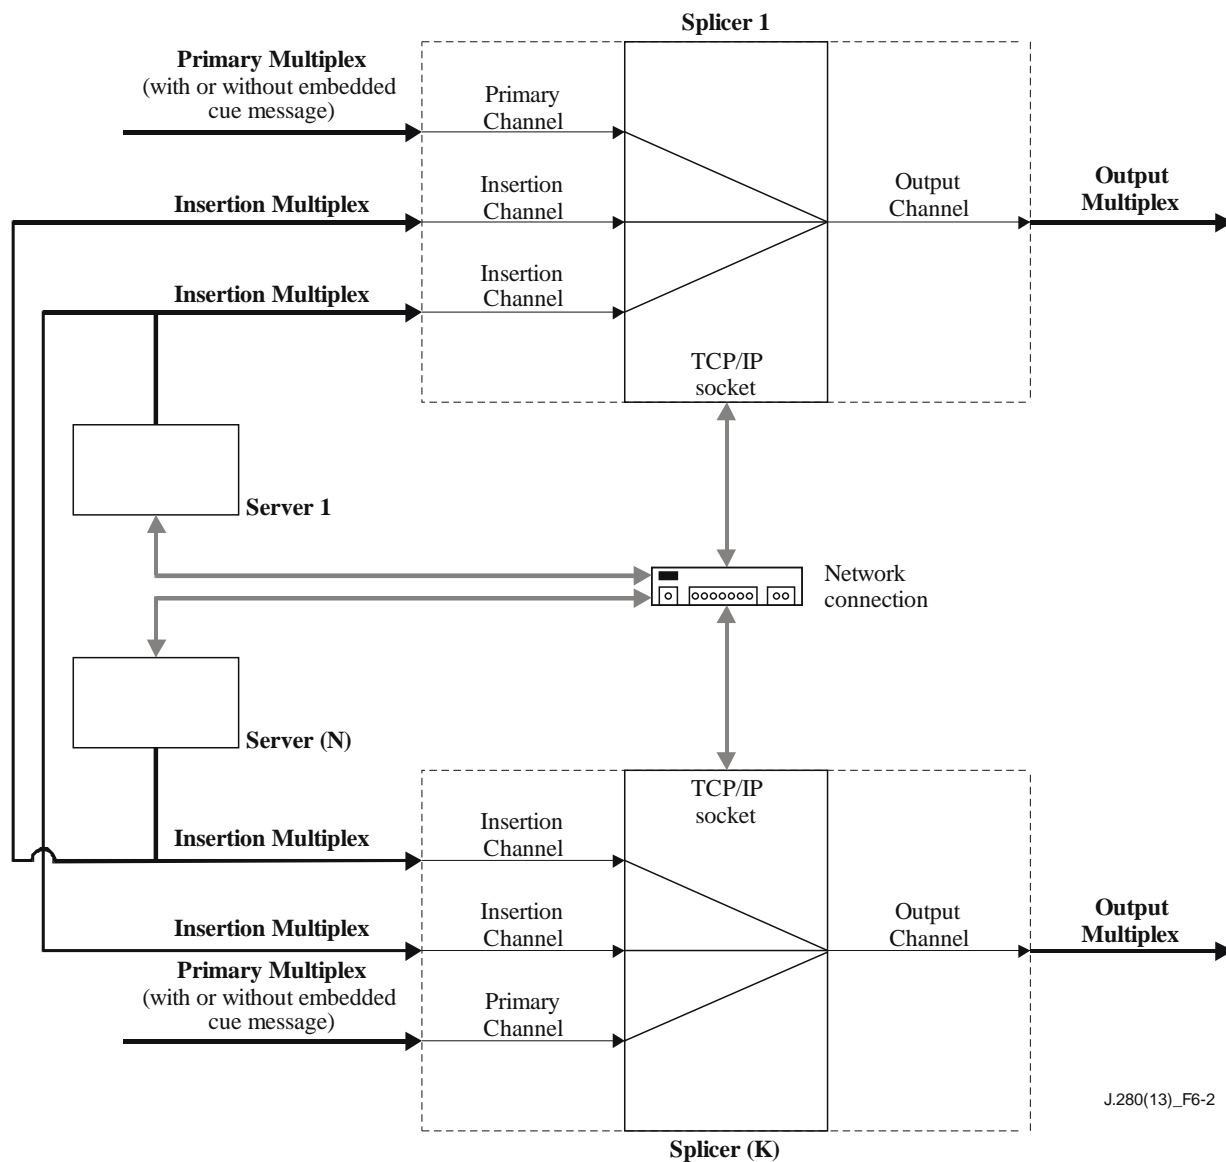
None.

6   Requirements

6.1   System block diagram

This API may be used with many different configurations of server(s) and splicer(s). This API focuses on the single server, single splicer configuration shown in Figure 6-1. However, this can be expanded to the multiple servers, multiple splicers configuration as shown in Figure 6-2.

Figure 6-1 – Single server/Single splicer

Figure 6-2 – Multiple servers/Multiple splicers

The model referenced in this API has one or more splicers with one or more multiplex inputs. The splicer logically separates the channel(s) in the multiplex(es) and presents these channel(s) to a switch. This switch is capable of mapping any input to any output channel. The initial configuration maps primary channel(s) to the output channel(s). The server may then direct the splicer to switch from a primary channel to an insertion channel for a specified duration. It may then direct the splicer to switch to another insertion channel following the initial switch. The splicer may then create an MPEG-2 transport stream produced by multiplexing one or more output channels that shall be compliant with [ANSI/SCTE 54].

Logically, a splice involves two input channels and one output channel. The splicer is responsible for joining the various elementary streams (audio, video and data) together. The optimal splice point may occur at slightly different times for each elementary stream, so the splicer should perform the splice that will supply the best quality output. Splicing may not always be performed from the primary channel, i.e., programming network, to the insertion channel, i.e., spot advertisement, and back to the primary channel. The splicer may splice content that is stored solely on the server and arrives over a single input multiplex. It is possible to use this API in a situation where a server has one multi-program transport stream (MPTS) output that contains program and interstitial material and uses the splicer to create proper splices between the content.

This API supports all combinations of single and multiple servers communicating with single and multiple splicers. A separate API connection is associated with each output channel.

In some configurations, there can be either multiple servers or multiple channels within the insertion multiplex connected to a splicer. In these cases, the splicer will have multiple API connections associated with an output channel. When an ANSI/SCTE 35 cueing message is received in a primary channel, the **Cue_Request** message must be sent to servers over all of the API connections that were made for the associated output channel(s). It is also possible that more than one API connection will transport a **Splice_Request** message for the same insertion at the same time for an output channel.

6.2 Arbitration priorities

Different levels of access are used to ensure that the correct insertion channel is utilized. There are ten different levels of access, 0 to 9, with 9 being the highest priority which may override any lower priority connection. The **OverridePlaying** flag in the **Splice_Request** message specifies whether an insertion request is honoured when the splicer is currently queueing or performing an insertion. If the flag is set to 1, then the higher priority insertion can interrupt the same or lower priority currently playing insertion. If the flag is set to 0, the splicer will not replace the insertion currently playing, even if the new request is of a higher priority.

The **Splice_Request** message should be sent at least three seconds before the splice time( ) in order to be valid. If the three-second minimum is not met, the outcome of the **Splice_Request** message is not determined by this API. If multiple servers initiate splice requests for the same time with the same priority, the splicer will prioritize the requests on a first-come, first-served basis. All other requests will be denied and a collision error will be sent in the **Splice_Response** message (unless the **OverridePlaying** flag is set).

For example, during the period of time immediately preceding the initiation of an insertion, the following is true: if a priority 5 **Splice_Request** is received for the same splice time as a priority 3 **Splice_Request**, a collision error is returned for the priority 3 request. If a priority 7 **Splice_Request** is later received for the same time, a collision error is returned for the priority 5 request and the priority 7 request is queued. If a second priority 7 request is received with the **OverridePlaying** flag set to 0, then the second priority 7 request would receive a collision error. However, if the **OverridePlaying** flag is set to 1 on the second priority 7 request, the original priority 7 request would receive a collision error and be overridden.

In Figure 6-3, three splicer inputs are shown. The shaded areas indicate which input source will be directed to the output channel at any given moment.


Figure 6-3 – OverridePlaying flag operation

t1 – Server 1 issues a **Splice_Request** message and begins its stream to the splicer. The splicer switches this insertion channel stream to the output channel. The **Splice_Request** has requested an insertion duration from time t1 to time t5. The splicer shall send Server 1 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_in and a result code set to 100, "Successful Response".

t2 – Server 2 issues a **Splice_Request**, with the **OverridePlaying** flag set to 1 and has an equal or higher priority. At the time specified by the **Splice_Request**, Insertion Channel 2's stream is switched to the output channel (replacing the ongoing stream from Server 1). Server 2's **Splice_Request** requests a duration from time t2 to time t3. The splicer shall send Server 1 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_out and a result code set to 125, "Channel Override". The splicer shall send Server 2 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_in and a result code set to 100, "Successful Response". If Server 1 determines that the channel override is an error, it may send an **Abort_Request** and terminate its stream at this time. This behaviour is not shown in Figure 6-3.

t3 – The insertion duration is completed and the splicer returns to the material from Server 1 to direct to the output channel. Note that the splicer did not return to the primary channel for direction to the output channel. The splicer shall send Server 1 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_in and a result code set to 125, "Channel Override". The splicer shall send Server 2 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_out and a result code set to 100, "Successful Response".

t4 – Server 2 issues another **Splice_Request**, with the **OverridePlaying** flag set to 1. At the time specified by the **Splice_Request**, Insertion Channel 2's stream is switched to the output channel (replacing the still ongoing stream from Server 1). Server 2's **Splice_Request** requests a duration from time t4 to time t6. The splicer shall send Server 1 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_out and a result code set to 125, "Channel Override". The splicer shall send Server 2 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_in and a result code set to 100, "Successful Response".

t5 – Server 1's insertion stream ends with two portions of its duration having been played and two portions having been overridden by Server 2's stream.

t6 – The final insertion duration is completed and the splicer returns to the material from the primary channel for direction to the output channel. The splicer shall send Server 2 a **SpliceComplete_Response** message with the **SpliceType** flag set to Splice_out and a result code set to 100, "Successful Response".

It is also possible that multiple servers will need to split a **Cue_Request** message; a 60-second duration splicing opportunity where one server will use the first 30 seconds and the second server will use the last 30 seconds, for example. Depending on the priorities and when the **Splice_Request** messages are received, the splicer shall indicate a result code 109 (Splice Collision) if one exists. This API does not coordinate the ability of the two servers to be able to perform this functionality. This can be done by mutual agreement between the servers or by a server-to-server API.

6.3   Abnormal terminations

It is possible that an insertion will be overridden at some time during playback by a higher priority insertion. In this case, the splicer shall return to the overridden insertion at the end of the higher priority insertion. If the higher priority insertion is aborted by an **Abort_Request** message, the splicer shall return to the overridden insertion. If the initial insertion channel is no longer available, then the splicer shall return to the primary channel if possible.

If the server requests a splice on a primary channel that currently has no valid input, the splicer shall perform the splice and report a result code 111 (No Primary Channel Found) in the **SpliceComplete_Response** to the server. Likewise, a splice from an insertion channel back to a primary channel that has no valid input shall complete with the result code 111 (No Primary Channel Found).

The splicer vendor may consider adding software to ensure that the splicer always returns to the primary channel. It is highly desirable to have the splicer fail-safe to the primary channel on any error condition that would cause the output channel to stop transmitting.

6.4   Splicing requirements

The splicer requires information about the insertion channel before it can be spliced into the primary channel. Some of this information shall be sent in the API connection and some of it may be sent in the MPEG multiplex. All of the information is required before the splice.

ChannelName is used for output channel identification. This is a unique name assigned to each output channel (e.g., CNN, Cable News Network) in the splicer set-up and is needed by the server to determine which primary channel shall be replaced by each insertion channel.

The splicer needs to know which insertion channel to splice into the primary channel. This includes the insertion multiplex location and which channel in the insertion multiplex to use. This information is available in the **Splice_Request** message.

6.5   Communication

The communication between the server and the splicer is conducted over one TCP/IP socket connection per output channel. Once this API connection is established it remains established until one of the devices terminates it, at which time re-initialization is required to splice again.

All messages exchanged between the splicer and server share a common general format detailed in clause 7.1. Only messages adhering to this format shall be used for communication between the splicer and server. The format allows for a class of "User Defined" type messages that can be used as a template for private data messages between the server and splicer, but is beyond the scope of this Recommendation.

All request messages require a response from either the splicer or the server, depending on which device is making the request. Most of the response messages only indicate a result and do not contain any other data, but are needed to ensure that the message was received and interpreted correctly by the requester. If there are errors, the message can be resent.

7   API syntax

7.1   Splicing_API_Message syntax

Messages in this API all contain a general message structure that wraps the data for the specific message being sent. This is done so that when the message is received a common parsing routine can store the message, determine what the structure of the data is and ensure that the message is received correctly.

Table 7-1 – Splicing_API_Message

Syntax	Bytes	Type
Splicing_API_Message {		
MessageID	2	uimsbf
MessageSize	2	uimsbf
Result	2	uimsbf
Result_Extension	2	uimsbf
data( )	*	*
}		

MessageID – An integer value that indicates what message is being sent. See Table 7-2.

MessageSize – The size of the data( ) field being sent in bytes.

Result – The results to the requested message. See Appendix I for details on the result codes. On request messages, this is set to 0xFFFF.

Result_Extension – This shall be set to 0xFFFF unless used to send additional result information in a response message.

data( ) – Specific data structure for the message being sent. Details on each of the messages containing data are described below. The size of this field is equal to the **MessageSize** and is determined by the size of the data being added to the message. Not all messages utilize the data( ) field.

Table 7-2 – MessageID values

MessageID	Message name	Sent by	Description
0x0000	General_Response	Splicer or Server	Used to convey asynchronous information between the devices. There is no data( ) associated with this message.
0x0001	Init_Request	Server	Initial message to the splicer on port 5168.
0x0002	Init_Response	Splicer	Initial response to the server on the established connection.
0x0003	ExtendedData_Request	Server	Request for detailed playback information from the splicer.
0x0004	ExtendedData_Response	Splicer	Vendor unique response of extended playback data from the requested playback event.
0x0005	Alive_Request	Server	Sends an alive message to acquire current status.
0x0006	Alive_Response	Splicer	Response to the alive message indicating current status.
0x0007	Splice_Request	Server	Request to splice at a specific time.
0x0008	Splice_Response	Splicer	Response to indicate that the **Splice_Request** was received and that the splicer is preparing to splice.
0x0009	SpliceComplete_ Response	Splicer	Response at the splice in and splice out.
0x000A	GetConfig_Request	Server	Request to get the current splice configuration for this API connection.

Table 7-2 – MessageID values

MessageID	Message name	Sent by	Description
0x000B	GetConfig_Response	Splicer	Contains all of the splice information for the API connection.
0x000C	Cue_Request	Splicer	Splicer sending the cue info section to the server.
0x000D	Cue_Response	Server	Acknowledgment that the cue info section was received.
0x000E	Abort_Request	Server	Request to immediately return to the primary channel or overridden insertion channel.
0x000F	Abort_Response	Splicer	Acknowledgment that the **Abort_Request** message was received. A SpliceComplete_Response shall also be generated if necessary.
0x0010	TearDownFeed_Request	Server	Request to delete an output channel created using the create_feed_descriptor( ).
0x0011	TearDownFeed_Response	Splicer	Response to indicate that the output channel has been deleted.
0x0012-0x7FFF 0xFFFF	Reserved		Range reserved for future standardization.
0x8000-0xFFFE	User defined		Range available for user-defined functions.

7.2 Conventions and requirements

1. Each message that contains data is outlined with its data fields and types below. Additional structures are indicated as functions and are described in clause 8.

2. All string lengths have space reserved for a null terminator character and must use null terminated strings. For example, a string that is defined as 16 characters can be at most 15 characters of data followed by a null (0x00) character immediately after the last data character. Once a null is encountered in scanning a string, the rest of the characters in the string are undefined. The size defined for the string is constant and will not vary depending on the length of the string. This specification uses 8-bit ASCII characters for strings.

3. All time values are UTC.

4. This specification uses all 1s for a DON'T CARE condition. For a 4-byte field, this value would be 0xFFFFFFFF.

5. Response messages shall be sent out without unnecessary delay. The device expecting a response should consider no response within 5 seconds to indicate a timeout. When a server suspects a timeout, it shall send an Alive_Request message. If the splicer does not answer as specified in this Recommendation, the connection for this channel shall be dropped and re-established.

6. A server receiving a response message indicating a failure to parse a message (error code 123) shall send an Alive_Request message. If it does not receive the appropriate Alive_Response message, the connection for this channel shall be dropped and re-established.

7. The Result field in the Splicing_API_Message is used to return a result code. Multiple response codes can be returned by sending multiple **General_Response** messages at any time.

8.   If the splicer or server cannot parse the Request message, it shall return a **General_Response** with the result code 123.

7.3   Initialization

The initial communication begins with the splicer listening on the predefined port 5168 and a server opening an API connection to the splicer. The server sends an **Init_Request** message to the splicer. The server then listens for the response from the splicer on the established API connection. All further communication is done on this API connection. Either the splicer or server may terminate communications by closing this API connection. Each device is responsible for detecting and properly handling a closed API connection. When the splicer initializes the TCP listener on port 5168, it should allow for at least three times the number of insertion channels for API connections to the splicer. For example, if the splicer controls 70 channels of which 40 are "spliceable", then it should allow 120, (40×3), simultaneous API connections.

7.3.1   Init_Request message

The data( ) field for this message contains the Init_Request_Data structure outlined below.

Table 7-3 – Init_Request_Data

Syntax	Bytes	Type
Init_Request_Data {		
Version( )		
ChannelName	32	String
SplicerName	32	String
Hardware_Config( )		
for (i=0; i<N; i++)		
splice_API_descriptor( )		
}		

Version( ) – See clause 8.1.

ChannelName – Logical name given to the output channel of this connection. This is also used to verify the correct API connection when the splicer responds to the server.

SplicerName – Name of the splicing device if the server uses the API to communicate to a device that controls multiple splicers.

Hardware_Config( ) – See clause 8.2.

splice_API_descriptor( ) – A descriptor that must follow the syntax defined in clause 8.5. The missing_Primary_Channel_action_descriptor( ) is a suitable descriptor for this request.

7.3.2   Init_Response message

After the **Init_Request** is sent, the splicer sends an **Init_Response** message on the opened API connection. The server verifies that the version sent by the splicer is supported and that it has an API connection to the correct primary channel.

The data( ) field for this message contains the Init_Response_Data structure outlined below.

Table 7-4 – Init_Response_Data

Syntax	Bytes	Type

Init_Response_Data {
     Version( )
     ChannelName
 }
	
	32
		
		String

Version( ) – See clause 8.1. The splicer shall respond with the highest version number of the API that it is capable of supporting.

ChannelName – Returned to the server to indicate the correct connection was made.

7.4 Embedded cueing messages

Splicers may have the ability to receive embedded cue messages based upon [ANSI/SCTE 35]. Once these cue messages are received by the splicer, they need to be passed to the server. The **Cue_Request** message is used to pass these cue messages to the server from the splicer. When a splicer receives a cue message it sends the entire splice_info_section( ) along with the splice time to the server. The server will acknowledge the message with a **Cue_Response** message. The **Cue_Response** message consists of just the Splicing_API_Message and has no associated data( ) but it may have a return code. The splicer will decrypt the splice_info_section( ) before sending it to the server if it is encrypted.

If the splicer receives a cue message that has an invalid cyclic redundancy check (CRC), it shall send a **General_Message** to the server with a result code of 117 (Invalid Cue Message). The splicer shall not send the **Cue_Request** message in this case.

It is suggested that the splicer be configurable as to which ANSI/SCTE 35 messages generate a Cue_Request message. Configurations can include the:

- ability to pass messages from newer versions of [ANSI/SCTE 35] that the splicer revision does not understand;
- ability to not pass bandwidth reservation messages; this should be the default setting;
- ability to not pass Splice_Null messages unless they have descriptors attached;
- ability to not pass messages that cannot be decrypted.

7.4.1 Cue_Request message

The data( ) field for this message contains the Cue_Request_Data structure outlined below.

Table 7-5 – Cue_Request_Data

Syntax	Bytes	Type

Cue_Request_Data {
     time( )
     splice_info_section( )
 }		

time( ) – This time is derived from the splice_time( ) in the splice_info_section( ) of the ANSI/SCTE 35 cueing message by the splicer. If component splice mode is used in the ANSI/SCTE 35 splice_info_section, the time( ) will refer to the default splice time detailed in Section 7.5.2.1 of [ANSI/SCTE 35]. In the case where the splice_info_section( ) does not contain a pts_time( ) that requires translation as in the splice_schedule( ) command, then the time structure shall be filled with all 1 s to denote no time specified. It is up to the splicer to determine how to map

the PTS time to UTC for communication with the server. This may vary for different splicers in order for them to properly manage their internal buffers. See clause 8.4 for the time( ) structure syntax.

splice_info_section( ) – The details of the structure can be found in [ANSI/SCTE 35].

7.5 Splice messages

After initializing and configuring the splicer, the server can issue the **Splice_Request** message to initiate a session. The two messages that are returned from the **Splice_Request** message are the **Splice_Response** message and the **SpliceComplete_Response** message. The server shall send a **Splice_Request** message at least 3 seconds prior to the time( ) in the **Splice_Request** message. This allows the splicer to set up its configuration and prepare for the splice. The insertion channel stream for the session must start between 300 and 600 milliseconds before time( ) as measured at the splicer input. A program clock reference (PCR) must be sent on or before the first video access unit of the insertion channel stream. The video stream of the insertion content shall start with a sequence header and an I-Frame. The splicer shall allow a minimum of 10 queued **Splice_Request** messages on a given API connection. If the splicer's message queue is full it will respond with the result code 114 (Splice Queue Full).

The details of the physical connection are supplied in the **Init_Request** message. There are two ways to indicate which channel in the insertion multiplex and which PIDs to use:

- If the **ServiceID** is not 0xFFFF in the **Splice_Request** message, the ServiceID field specifies the program number in the PAT which points to an associated PMT. The PAT and PMT must be stable in the insertion channel at least 200 ms before the **Splice_Request** message is sent and must remain stable for the duration of the session. These must be legal MPEG tables with revision increments as appropriate.

- If the **ServiceID** is 0xFFFF, use the splice_elementary_stream( ) structure (PCR, video, audio and data PIDs) in the **Splice_Request** message.

NOTE – If this method is used then the ServiceID shall be set to 0xFFFF. The splicer shall supply an MPEG-2 compliant transport stream to the output multiplex although the insertion multiplex need not include program specific information (PSI).

The order in which splice messages are sent is important. The first message sent for a given sequence of back-to-back insertions shall utilize time( ), while all of the other **Splice_Request** messages may utilize **PriorSession**. The **PriorSession** number must reference an existing session that has not yet completed. In all other cases, an error code 123 is returned pointing to the **PriorSession** or time( ) field.

The server chooses the PIDs of the elementary streams within an insertion multiplex. The PIDs may not be common between adjacent sessions from the same server via the same insertion multiplex. This is because the streams of adjacent sessions will occasionally overlap slightly in time, due to requirements in this API.

7.5.1 Splice_Request message

The data( ) field for this message contains the Splice_Request_Data structure outlined below.

Table 7-6 – Splice_Request_Data

Syntax	Bytes	Type
Splice_Request_Data {		
SessionID	4	uimsbf
PriorSession	4	uimsbf
time( )		
ServiceID	2	uimsbf
if (**ServiceID** = 0xFFFF)		
{		
PcrPID	2	uimsbf
PIDCount	4	uimsbf
for (j=0; j<**PidCount**; j++)		
splice_elementary_stream( )		
}		
Duration	4	uimsbf
SpliceEventID	4	uimsbf
PostBlack	4	uimsbf
AccessType	1	uimsbf
OverridePlaying	1	uimsbf
ReturnToPriorChannel	1	uimsbf
for (i=0; i<N; i++)		
splice_API_descriptor( )		
}		

SessionID – Identifier for the session. Used to distinguish this request from other requests that have been or are going to be issued. Multiple concurrent **Splice_Request** messages with the same **SessionID** are not permitted. If the **ExtendedData_Request** is used, an **ExtendedData_Response** must be received for that **SessionID** before that **SessionID** is reused. This field shall not have the value of 0xFFFFFFFF. Early versions of this standard (with Revision_Num = 0 or 1) allow the value 0xFFFFFFFF.

PriorSession – This field allows a simplified method of performing back-to-back insertions. The value of this field contains the **SessionID** of the session that immediately precedes it. When the value of this field is 0xFFFFFFFF, it indicates that this session uses time( ) to initiate its insertion, rather than the **SessionID** of the preceding session. This field shall have a valid **SessionID** only when the immediately preceding session originated from the same server. The time( ) field- rather than the **PriorSession** field- must be used when creating back-to-back insertions from multiple servers.

time( ) – The splice time for the event. This field will typically be the time( ) field from the **Cue_Request** message being echoed back to the splicer. If the event was not triggered by a **Cue_Request**, then this is the time that the server forces a splice event. This field is ignored if the **PriorSession** is not equal to 0xFFFFFFFF. If this value is not related to an ANSI/SCTE 35 cue message, there may be variation between splicers, depending on the buffer and splicing models of each, as to when the actual splice occurs. See clause 8.4 for the time( ) structure syntax.

ServiceID – The program number of the channel in the insertion multiplex which will be spliced in place of the primary channel. If this is set to 0xFFFF the splice_elementary_stream( ) and **PIDCount** are required.

PCR – Indicates the PCR PID.

PIDCount – The number of PIDs in the insertion channel (not including the PCR PID).

Duration – The number of 90 kHz clock ticks the server is requesting the splicer to insert. This field may override the ANSI/SCTE 35 duration value. This can be set to 0 to indicate that the splicer shall switch to the insertion channel until a new **Splice_Request** or **Abort_Request** arrives.

SpliceEventID – This is used to relate this insertion event back to the ANSI/SCTE 35 cue message that may have caused this splice to happen. This shall be equivalent to the splice_event_id from the splice_insert command of the associated ANSI/SCTE 35 cue message. This should be the same for all **Splice_Request** messages pertaining to the same ANSI/SCTE 35 cue message. For an event that was not initiated by an ANSI/SCTE 35 cue message, this field will be set to 0xFFFFFFFF.

PostBlack – Number of 90 kHz clock ticks of black video and muted audio to be played at the end of the insertion content playback. The **PostBlack** interval follows and is not included in the length of time specified by the **Duration**. If no **PostBlack** is requested then this field will be set to 0. **PostBlack** shall not be considered part of the currently playing insertion for the purposes of the **OverridePlaying** flag.

AccessType – Indicates the type of access this connection has. This is an integer from 0 to 9 with 0 being low priority and 9 being the highest priority.

OverridePlaying – When this flag is equal to 0, this **Splice_Request** cannot override a currently playing insertion. If this flag is set to 1, then this **Splice_Request** shall override any equal or lower priority currently playing insertion. A currently playing insertion occurs between the splice-in and the splice-out points.

ReturnToPriorChannel – When this flag is equal to 0, the splicer shall not return to the primary channel or the overridden insertion channel at the completion of this **Splice_Request**. It is expected that a new **Splice_Request** will be issued before this insertion completes. If a new **Splice_Request** is not received, then the splicer shall stop transmitting on this output channel. When this flag is equal to 1, it shall return to the prior channel unless a subsequent **Splice_Request** is received to indicate otherwise.

splice_API_descriptor( ) – A descriptor that must follow the syntax defined in clause 8.5. The playback_descriptor( ) and muxpriority_descriptor( ) are appropriate descriptors for this clause.

7.5.2   Splice_Response message

The data() field for the Splice_Response message contains the **Splice_Response_Data** structure outlined below. The **Splice_Response_Message** may contain an error code if appropriate. The **Splice_Offset** is used by the splicer to inform the server of a time offset for the delivery of the content for this message. This does not affect the point in the primary channel where the splice will occur.

Table 7-7 – Splice_Response_Data

Syntax	Bytes	Type
Splice_Response_Data {		
Splice_Offset	2	tcimsbf
}		

Splice_Offset – This shall be set to zero unless used to send delay information. The offset information is in milliseconds. A negative value is a request for the insertion channel content to be delivered earlier; a positive value is a request for the insertion channel content to be delivered later.

The **Splice_Offset** field is expected to be used by splicing devices which achieve seamless operation by altering the primary channel propagation delay through the splicer. When such a device is commanded to splice in the absence of ANSI/SCTE 35 cue messages, the splicer does not have the opportunity to advance or retard the ad server's ad timing via alteration of the equation in converting pts_time to UTC time (because there are no ANSI/SCTE 35 cue messages and hence no conversion and no cue request messages). In such a case, a splicer may utilize the new **Splice_Offset** field to advance or retard the ad server to match the splicer's output service timing following each Splice_Request message.

7.5.3   SpliceComplete_Response message

The **SpliceComplete_Response** message is sent when the insertion starts and finishes. This is true for back-to-back insertions as well. For example, if two pieces of content play, four **SpliceComplete_Response** messages are returned, one at the start of the first piece of content, one upon completion of the first piece of content, one upon the start of the second piece of content and one upon completion of the second piece of content. The result code in the header shall properly indicate the failure reason if the splice failed so that the server can take appropriate action. The splice-in and splice-out are separate events and shall be treated as such. If a splice between two pieces of content fails, the splice-out should indicate good status if the current piece of content is played in its entirety. The **SpliceComplete_Response** message shall be sent immediately upon failure of any splice event and shall not wait until the expected duration of the inserted content.

The data( ) field for this message contains the SpliceComplete_Response_Data structure outlined below.

Table 7-8 – SpliceComplete_Response_Data

Syntax	Bytes	Type
SpliceComplete_Response_Data {		
SessionID	4	uimsbf
SpliceTypeFlag	1	uimsbf
if (**SpliceTypeFlag** = 0)		
{		
time( )		
} else		
{		
Bitrate	4	uimsbf
PlayedDuration	4	uimsbf
}		
}		

SessionID – The session ID that the **Splice_Request** message used.

SpliceTypeFlag – This field shall be a 0 to indicate a splice-in (start) and a 1 to indicate a splice-out (end).

time( ) – The time that the splicer detected the first byte of the insertion stream from the server. The server may utilize this time( ) to adjust the arrival time at the splicer, of subsequent insertion channel content when the delivery mechanism is known to have a time varying latency.

Bitrate – This is the average bit rate for the session. This field is in bits-per-second (bps) including transport packet overhead for this channel.

PlayedDuration – This is the number of 90 kHz clock ticks actually played. This is exclusive of any post-black or transition frames.

7.6 Alive messages

Once the initialization is complete, the server can send **Alive_Request** messages to ensure that the splicer is still up and running. Each **Alive_Response** message contains a status from the splicer to the server. This status indicates the state of the device. If there has been no activity on the TCP/IP connection in the preceding 60 seconds, then an Alive_Request message shall be sent.

7.6.1 Alive_Request message

The data( ) field for the **Alive_Request** message contains the Alive_Request_Data structure outlined below.

Table 7-9 – Alive_Request_Data

Syntax	Bytes	Type
Alive_Request_Data{		
time( )		
}		

time( ) – The current UTC time clock of the sending device checked as close as possible to the sending of the message. This is designed to be used by the splicer and the server to check on how well the two systems are time synchronized. It is not expected that this will allow the systems to synchronize well enough to allow reliable splicing to occur, but the implementers may use this as they wish. See clause 8.4 for the time( ) structure syntax.

7.6.2 Alive_Response message

The data( ) field for the **Alive_Response** message contains the Alive_Response_Data structure outlined below.

Table 7-10 – Alive_Response_Data

Syntax	Bytes	Type
Alive_Response_Data{		
State	4	uimsbf
SessionID	4	uimsbf
time( )		
}		

State – This describes the state of the output channel.

Table 7-11 – Alive_Response message states

State	Description
0x00	No output
0x01	On primary channel
0x02	On insertion channel

SessionID – The **SessionID** of the currently playing insertion. Valid only for **State** = 0x02.

time( ) – The current UTC time clock of the sending device checked as close as possible to the sending of the message. This is designed to be used by the splicer and the server to check on how well the two systems are time synchronized. It is not expected that this will allow the systems to synchronize well enough to allow reliable splicing to occur, but the implementers may use this as they wish. See clause 8.4 for the time( ) structure syntax.

7.7 Extended data messages

This is a splicer defined structure to send detailed data about the playback to the server. After the **SpliceComplete_Response** has been received, then the extended data can be retrieved using the **ExtendedData_Request**. The **SessionID** used in this message is the same as the **SessionID** used in setting up this session and in the **SpliceComplete_Response**.

7.7.1 ExtendedData_Request message

The data( ) field for this message contains the ExtendedData_Request_Data structure outlined below.

Table 7-12 – ExtendedData_Request_Data

Syntax	Bytes	Type
ExtendedData_Request_Data {		
SessionID	4	uimsbf
ExtendedDataType	4	uimsbf
}		

SessionID – The **SessionID** of the completed session.

ExtendedDataType – The requested response data type from the splicer to the **ExtendedData_Response** message. This value may be set to 0xFFFFFFFF to indicate that the default data type is to be returned. This standard reserves 0x00000000 to 0x7FFFFFFF for future standardization. The range 0x80000000 to 0xFFFFFFFE is for vendor unique usage.

7.7.2 ExtendedData_Response message

The server shall use the **MessageSize** field to determine the amount of data it is required to read via the **ExtendedData_Response** message.

The data( ) field for this message contains the ExtendedData_Response_Data structure outlined below.

Table 7-13 – ExtendedData_Response_Data

Syntax	Bytes	Type
ExtendedData_Response_Data {		
SessionID	4	uimsbf
for(i=0;i<n;i++)		
splice_API_descriptor( )		
}		

SessionID – The **SessionID** that these data are valid for.

splice_API_descriptor( ) – A descriptor of the format defined in clause 8.5 that is splicer defined.

7.8 Abort messages

The server can send an **Abort_Request** at any time which will cause the splicer to immediately revert to the overridden insertion channel or primary channel. The splicer shall send an **Abort_Response** message to acknowledge the receipt of the **Abort_Request**. A **SpliceComplete_Response** with a result code 116 (Insertion Aborted) is sent if the **Abort_Request** caused a splice-out of the insertion. If no splice-out was needed then no **SpliceComplete_Response** message shall be reported.

All pending back-to-back insertions linked via the **PriorSession** field of the **Splice_Request** message to the **SessionID** of an **Abort_Request** message shall also be aborted. An error message shall be returned for each aborted **SessionID**. Consider the following example: three insertions are cued to run sequentially within a block of time – the first event is time-based; the second event is linked to the first **SessionID** using the **PriorSession**; the third event is linked to the second **SessionID** using that **PriorSession**. In this example, if the first insertion event is aborted, the two subsequently cued insertion events will also be aborted. The abort message does not abort any insertions that use a different API connection from a server to a splicer. The next splice that occurs for the primary channel requires the **PriorSession** in the splice message to be 0xFFFFFFFF.

7.9 Abort_Request message

The data( ) field for this message contains the Abort_Request_Data structure outlined below.

Table 7-14 – Abort_Request_Data

Syntax	Bytes	Type

Abort_Request_Data {
     SessionID
}
	4	uimsbf

SessionID – The **SessionID** and all subsequent sessions linked through the **PriorSession** field that are to be aborted.

7.10 Abort_Response message

The **Abort_Response** indicates that the **Abort_Request** message was received. This message may contain a result code if appropriate.

The data( ) field for this message contains the Abort_Response_Data structure outlined below.

Table 7-15 – Abort_Response_Data

Syntax	Bytes	Type

Abort_Response_Data {
     SessionID
}
	4	uimsbf

SessionID – The **SessionID** and all subsequent sessions linked through the **PriorSession** field that were aborted.

7.11 TearDownFeed_Request message

The **TearDownFeed_Request** message contains no data and is only valid for the connection this message is sent over. This message is used to tear down feeds created by the Init_Request message with the create_feed_descriptor(). This message is used to tear down only those feeds.

7.12 TearDownFeed_Response message

The **TearDownFeed_Response** message contains no data and indicates that the **TearDownFeed_Request** message was received and the action occurred. This message may contain a result code if appropriate.

7.13 Requesting configuration settings

The current configuration settings for the API connection can be returned. This includes some of the information in the **Init_Request**. The **GetConfig_Request** contains no additional data.

7.13.1 GetConfig_Request message

The **GetConfig_Request** message contains no data.

7.13.2 GetConfig_Response message

The data( ) field for this message contains the GetConfig_Response_Data structure outlined below.

Table 7-16 – GetConfig_Response_Data

Syntax	Bytes	Type
GetConfig_Response_Data {		
ChannelName	32	String
Hardware_Config( )		
TS_program_map_section( )		
}		

ChannelName – Logical name given to the output channel of this connection.

Hardware_Config( ) – See clause 8.2 for the syntax of the Hardware_Config( ) structure.

TS_program_map_section( ) – This is the entire PMT section for the output channel as defined in [ITU-T H.222.0]. If the splicer changes the PMT, it should signal this change to the server with a result code 128 in the **General_Response** message.

7.14 General_Response message

The General_Response message is used to convey asynchronous information between the server and the splicer. There is no data( ) associated with this message. Any result code may be sent in this message. This message will typically be used to indicate output channel PMT changes or invalid request messages.

8 Additional structures

8.1 Version

The "Version" structure is used to maintain the proper versioning within the API. It is expected that this API will evolve over time and to allow for this expansion, the version is specified in the **Init_Request** and **Init_Response** messages to ensure that the splicer supports the same version as the server.

<p style="text-align: center;">Table 8-1 – Version( )</p>

Syntax	Bytes	Type

Version {
    Revision_Num
 	
}
	2	uimsbf

Revision_Num – This field is two (2) in this version.

The server and splicer should set and check this field to ensure that both components are capable of operating at the appropriate revision.

8.2 Hardware_Config

This structure describes the hardware interface between the server and the splicer. It is important for the splicer to know exactly where the server is connected so that the splicer knows what multiplex is being referenced. An example of this link would be a DVB-ASI connection from the server to the splicer.

<p style="text-align: center;">Table 8-2 – Hardware_Config( )</p>

Syntax	Bytes	Type
Hardware_Config{		
Length	2	uimsbf
Chassis	2	uimsbf
Card	2	uimsbf
Port	2	uimsbf
Logical_Multiplex_Type	2	uimsbf
Logical_Multiplex( )		uimsbf
}		

Length – This gives the length in bytes of this structure following this field.

Chassis – An integer indicating to which splicer chassis the server's insertion multiplex is connected. In cases where the card is labelled alphabetically the translation is made to an integer value (i.e., A – 1; B – 2, etc.).

Card – An integer indicating the splicer card to which the server's insertion multiplex is connected. In cases where the card is labelled alphabetically the translation is made to an integer value (i.e., A – 1; B – 2, etc.).

Port – The hardware port number where the server's insertion multiplex is connected.

Logical_Multiplex_Type – A value from the following table:

<p style="text-align: center;">Table 8-3 – Logical multiplex type</p>

Type	Bytes	Name	Description
0x0000	0	Not used	The **Logical_Multiplex** field is not needed to identify the multiplex.
0x0001	Variable	User defined	The usage of the **Logical_Multiplex** field is not defined by this specification and must be agreed upon between the splicer and the server.

Table 8-3 – Logical multiplex type

Type	Bytes	Name	Description
0x0002	6	MAC address	The **Logical_Multiplex** field contains the IEEE Media Access Control address of the multiplex as a 6-byte address.
0x0003	6	IPv4 address	The most significant 4 bytes of the **Logical_Multiplex** field contain the Internet Protocol (IP) address of the multiplex, and the remaining 2 bytes contain the IP port number where the multiplex can be found.
0x0004	18	IPv6 address	The most significant 16 bytes of the **Logical_Multiplex** field contain the Internet Protocol (IPv6) address of the multiplex, and the remaining 2 bytes contain the IP port number where the multiplex can be found.
0x0005	5	ATM address	The **Logical_Multiplex** field contains the coordinates of the asynchronous transfer mode (ATM) circuit over which the multiplex is carried. The most significant two bytes of the logical multiplex field contain the virtual path identifier (VPI) and the next two bytes contain the virtual channel identifier (VCI) of the circuit. The least significant byte contains the ATM adaptation layer (AAL) number.
0x0006	Variable	IPv4 address with SPTS support	See description following this table.
0x0007	Variable	IPv6 address with SPTS support	See description following this table.
0x0008-0xFFFF	Variable	Reserved	Reserved for future standardization.

Type 0x0006 – IPv4 address with single program transport stream (SPTS) support

Type 0x0006 is utilized by VOD and ad servers where remapping of PIDs is impractical or undesirable. In these cases it is desirable to use an SPTS per UDP port.

Table 8-4 – Type 0x0006 structure

Syntax	Bytes	Type
Type 0x0006 structure {		
number_of_destination_ips	1	uimsbf
for (j=0; j< **number_of_destination_ips**; j++) {		
dest_ip_address	4	uimsbf
}		
number_of_source_ips	1	uimsbf
for (j=0; j< **number_of_source_ips**; j++) {		
source_ip_address	4	uimsbf
}		
base_port	2	uimsbf
number_of_ports	1	uimsbf
}		

number_of_destination_ips – Specifies how many **dest_ip_address**(es) follow. The valid range is 1 to 32.

dest_ip_address – The IPv4 address that the splicer shall use for the content associated with the splice.

number_of_source_ips – Specifies how many **source_ip_address**(es) follow. The valid range is 0-32.

source_ip_address – The source IPv4 address(es) that the splicer may use in an IGMP V3 join for the associated multicast **dest_ip_address**(es).

base_port – The initial UDP port that the splicer shall use for the content associated with a Splice_Request with time( ) specified. The base UDP port range shall be assigned by IANA.

number_of_ports – This byte contains the number of contiguous ports to reserve. The number_of_ports value may range from 1 to 4 and includes the base port. Allowed port numbers are determined, in order, by the base port, followed by base Port +1, followed by base Port +2, followed by base Port +3.

All Splice_Requests that use time( ) shall use the base IPv4 address:port unless the port_selection_descriptor( ) is used. The first Splice_Request of an avail shall use time( ). Subsequent sessions of the same avail that also use time( ) shall also use the base IPv4 address:port unless the port_selection_descriptor( ) is used. The next and subsequent splice_requests using PriorSession instead of time( ), shall use the base IP and port+1, then port+2 and so on until the requested number of ports is used and it shall then revert to the base port for the next splice_request.

The port_selection_descriptor( ) may be utilized in any Splice_Request command that has a hardware config with Logical_Multiplex type 0x0006 to alter the default operation of the ports.

The port may be any valid unicast or multicast IPv4 address:port combination. The splicer shall perform an IGMP join on a multicast IP.

Logical_Multiplex – If the **Port** carries multiple insertion multiplexes on a single input, then this field allows the splicer to determine which to use when splicing from this server. The meaning and format of this field is defined by the **Logical_Multiplex_Type** field. In the event that a non-standard definition for the **Logical_Multiplex** is required, the **Logical_Multiplex_Type** shall be set to 1 for user defined.

Type 0x0007 – IPv6 address with single program transport stream (SPTS) support

Type 0x0007 is utilized by VOD and ad servers where remapping of PIDs is impractical or undesirable. In these cases it is desirable to use an SPTS per UDP port.

Table 8-5 – Type 0x0007 structure

Syntax	Bytes	Type
Type 0x0007 structure {		
number_of_destination_ips	1	uimsbf
for (j=0; j< **number_of_destination_ips**; j++)		
{		
dest_ip_address	16	uimsbf
}		
number_of_source_ips	1	uimsbf
for (j=0; j< **number_of_source_ips**; j++) {		

Table 8-5 – Type 0x0007 structure

Syntax	Bytes	Type
source_ip_address	16	uimsbf
}		
base_port	2	uimsbf
number_of_ports	1	uimsbf
}		

number_of_destination_ips – Specifies how many **dest_ip_address**(es) follow. The valid range is 1 to 32.

dest_ip_address – The IPv6 address that the splicer shall use for the content associated with the splice.

number_of_source_ips – Specifies how many **source_ip_address**(es) follow. The valid range is 0-32.

source_ip_address – The source IPv6 address(es) that the splicer may use in an MLD V2 join for the associated multicast **dest_ip_address**(es).

base_port – The initial UDP port that the splicer shall use for the content associated with a Splice_Request with time( ) specified. The base UDP port range shall be assigned by IANA.

number_of_ports – This byte contains the number of contiguous ports to reserve. The number_of_ports value may range from 1 to 4 and includes the base port. Allowed port numbers are determined, in order, by the base port, followed by base Port +1, followed by base Port +2, followed by base Port +3.

All Splice_Requests that use time( ) shall use the base IPv6 address:port unless the port_selection_descriptor( ) is used. The first Splice_Request of an avail shall use time( ). Subsequent sessions of the same avail that also use time( ) shall also use the base IPv6 address:port unless the port_selection_descriptor( ) is used. The next and subsequent splice_requests using PriorSession instead of time( ), shall use the base IP and port+1, then port+2 and so on until the requested number of ports is used and it shall then revert to the base port for the next splice_request.

The port_selection_descriptor( ) may be utilized in any Splice_Request command that has a hardware config with Logical_Multiplex type 0x0007 to alter the default operation of the ports.

The port may be any valid unicast or multicast IPv6 address:port combination. The splicer shall perform an MLD join on a multicast IP.

8.3 splice_elementary_stream( )

Packet identifiers (PIDs) are identifiers for parts of the transport stream, video, audio, data, etc. This structure is used to describe one of the elements in the program in the MPTS. The **Splice_Request** message may contain a splice_elementary_stream( ) structure for each of the transport stream components (except for the PCR PID). The **StreamType**s are based on the MPEG PMT table definitions.

This specification has not defined how to map multiple audio/video/data PIDs to output PIDs. It has also not defined splicer behaviour when multiple audio tracks may be either present or missing in the insertion channel compared with the primary channel.

Table 8-6 – splice_elementary_stream( )

Syntax	Bytes	Type
splice_elementary_stream {		
Length	1	uimsbf
PID	2	uimsbf
StreamType	2	uimsbf
AvgBitrate	4	uimsbf
MaxBitrate	4	uimsbf
MinBitrate	4	uimsbf
HResolution	2	uimsbf
VResolution	2	uimsbf
for(i=0;I<N;I++)		
descriptor( )		
}		

The PCR PID is required.

Length – Total length in bytes of the splice_elementary_stream( ) structure including this field.

PID – The PID number that is being used. This is a 2-byte field (16 bits) and shall contain the 13-bit PID right aligned as a 16-bit integer. (0x0000 to 0x1FFF)

StreamType – The type of PID (audio, video, etc.). This number corresponds with the PMT specification found in [ITU-T H.222.0].

AvgBitrate – The bit rate for this PID averaged over the entire piece of content in bits-per-second (bps). This is set to 0xFFFFFFFF if the bit rate is not known.

MaxBitrate – The maximum bit rate for this PID. This is set to 0xFFFFFFFF if the bit rate is not known.

MinBitrate – The minimum bit rate for this PID. This is set to 0xFFFFFFFF if the bit rate is not known.

HResolution – The width in number of pixels of the video pictures using this PID. If the PID does not contain video pictures or if the server cannot supply this value, it shall be set to 0xFFFF.

VResolution – The height in number of pixels of the video pictures using this PID. If the PID does not contain video pictures or if the server cannot supply this value, it shall be set to 0xFFFF.

descriptor( ) – Any valid descriptor used in a PMT. For multiple audio PIDs the language descriptors as defined in [ITU-T H.222.0] are required.

8.4 time( ) field definition

The time structure is used to define various times in this specification.

Table 8-7 – time( )

Syntax	Bytes	Type
time {		
Seconds	4	uimsbf
MicroSeconds	4	uimsbf
}		

Seconds – Elapsed seconds since 12:00 AM January 1, 1970 UTC.

MicroSeconds – Offset in microseconds of the **Seconds** field.

8.5 splice_API_descriptor( ) field definition

This is a template for adding descriptors in any message defined within this Recommendation. The **Splice_Request**, **ExtendedData_Response** and **Init_Request** messages may use descriptors. The use of descriptors in messages defined by this standard is optional. The following table is the general format for descriptors used in this standard.

Table 8-8 – splice_api_descriptor( )

Syntax	Bytes	Type
splice_API_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
for (i=0;i<n;i++)		
Private_Byte	1	uimsbf
}		

Splice_Descriptor_Tag – A value from 0x00 to 0xFF to denote the specific descriptor being used. Tag values 0x00 to 0xFF are reserved for use by this standard. The vendor may use a vendor unique **Splice_API_Identifier** to allow for a larger tag range and a more robust method of adding vendor unique descriptors.

Descriptor_Length – This gives the length in bytes of the descriptor following this field. Descriptors are limited to 256 bytes, so this value is limited to 254.

Splice_API_Identifier – An identifier of the organization that has defined this descriptor. For all descriptors within this Recommendation, the identifier is 0x53415049 (ASCII "SAPI"). This has been chosen to not conflict with descriptors of any other known identifier.

Private_Byte – The remainder of the descriptor is dedicated to data fields as required by the descriptor being defined.

8.5.1 playback_descriptor( ) field definitions

The playback_descriptor( ) is an implementation of the splice_API_descriptor( ) which is intended for use in the **Splice_Request** message.

The abort criteria examine the playback rate, defined as the output channel's bit rate averaged over a one second period. The sliding displacement of the averaging window is recommended to be one second or less.

Table 8-9 – playback_descriptor( )

Syntax	Bytes	Type
playback_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
BitrateRule	1	uimsbf
MinPlaybackRate	4	uimsbf
}		

Splice_Descriptor_Tag – 0x01.

Descriptor_Length – 0x09.

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

BitrateRule – Flag used to indicate the rules for **MinPlaybackRate**.

Table 8-10 – BitrateRule values

BitrateRule	Description
0x00	Ignore **MinPlaybackRate.**
0x01	Return a result code 127 immediately using the **General_Response** message if the playback rate falls below the **MinPlaybackRate** but do not abort.
0x02	Abort if the playback rate falls below the **MinPlaybackRate**.
0x03	Cancel the session prior to the splice-in if the splicer determines that the **MinPlaybackRate** will not be met. The splicer will send a **SpliceComplete_Response** or **General_Response** with a result code 127.

MinPlaybackRate – The minimum aggregate bit rate of the output channel averaged over one second for the duration of the splice that it can play at before the **BitrateRule** is triggered. Setting this value to 0 indicates there is no minimum rate.

8.5.2 muxpriority_descriptor( ) field definitions

The muxpriority_descriptor( ) is an implementation of the splice_API_descriptor( ) which is intended for use in the **Splice_Request** message.

Table 8-11 – muxpriority_descriptor( )

Syntax	Bytes	Type
muxpriority_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
MuxPriorityValue	1	uimsbf
}		

Splice_Descriptor_Tag – 0x02

Descriptor_Length – 0x05

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

MuxPriorityValue – This number ranges from 1 to 10 (1 being the lowest, 5 is the average and 10 being the highest). This number modifies the stored **MuxPriorityValue** of the primary channel in the splicer. **A MuxPriorityValue** of 5 will not modify the output channel's priority. A **MuxPriorityValue** of less than 5 will subtract from the output channel's priority level and a **MuxPriorityValue** greater than 5 will add to the output channel's priority.

Using the **MuxPriorityValue** will not ensure that the content is played with any specific level of quality. The actual effect of the **MuxPriorityValue** depends on the overall spliced multiplex configuration and how much the splicer needs to lower the total multiplex bit rate at any given time. This will also be dependent on how the splicer operates and as such will be a very splicer vendor dependent field.

8.5.3   missing_Primary_Channel_action_descriptor( ) field definitions

The missing_Primary_Channel_action_descriptor( ) is an implementation of the splice_API_descriptor( ) which is intended for use in the **Init_Request** message.

If the primary channel has terminated for any reason during an insertion, the result at the decoder may be to display a freeze-frame of the last inserted frame at the conclusion of the insertion. This descriptor allows the splicer to be directed to insert black video and silent audio in order to clear the decoder's buffer, if the primary channel is no longer present when it would normally become the output audio/video source.

Table 8-12 – missing_Primary_Channel_action_descriptor ( )

Syntax	Bytes	Type
missing_Primary_Channel_action_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_Api_Identifier	4	uimsbf
MissingPrimaryChannelAction	1	uimsbf
}		

Splice_Descriptor_Tag – 0x03

Descriptor_Length – 0x05

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

MissingPrimaryChannelAction – This parameter has three possible values, 0, 1 and 2. A value of 0 means do nothing. A value of 1 means insert one black I frame and one frame of audio silence. A value of 2 means continue to transmit black and silence until the primary signal returns.

8.5.4   port_selection_descriptor( ) field definitions

The port_selection_descriptor( ) is an implementation of the splice_API_descriptor( ) which shall only be used in the **Splice_Request** message when the logical multiplex type 0x0006 or 0x0007 is used in the hardware configuration. If during a sequence of insertions the server sends a port_selection_descriptor( ), the server shall continue to send the port_selection_descriptor( ) until the next Splice_Request based on time occurs.

The port_selection_descriptor( ) may be utilized to alter the default operation of the ports or select a new dynamically set up IPv4 or IPv6 Address:Port combination.

The splicer shall dynamically set up a destination port if the **ps_ip_address** was not defined in the hardware config. If the **ps_ip_address** is multicast, the splicer shall issue an IGMP join or an MLD join request within 400 milliseconds after the arrival of the Splice_Request message. Latency for setting up a multicast group shall be less than 2 seconds, which is derived as follows:

The 3 second arrival of the Splice_Request message (clause 7.5).

Less the 600 milliseconds stream start time (clause 7.5).

Less the 400 milliseconds for the splicer to issue the IGMP join or the MLD join request.

Table 8-13 – IPv4 port_selection_descriptor ( )

Syntax	Bytes	Type
port_selection_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
ps_ip_address	4	uimsbf
ps_port	2	uimsbf
ps_number_of_source_ip	1	uimsbf
for (j=0; j< **ps_number_of_source_ip**; j++) {		
ps_source_ip_address	4	uimsbf
}		
}		

Splice_Descriptor_Tag – 0x04

Descriptor_Length – Variable. The length in bytes of the descriptor following this field.

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

ps_ip_address – The IPv4 Internet protocol address that the splicer shall use for the content associated with the splice. If this address:port combination is different from the address in the Logical_Mux_Type 0x0006 table, it shall be considered a dynamic port set-up request.

ps_port – The UDP port that the splicer shall use for the content associated with the splice. This port number shall override the automatic port selection method of the Logical_Multiplex_Type 0x0006.

ps_number_of_source_ip – Specifies how many **ps_source_ip_address**(es) follow. The valid range is 0-32.

ps_source_ip_address – The source IPv4 address(es) that the splicer shall use in an IGMP V3 join for the associated multicast **ps_ip_address**.

Table 8-14 – IPv6 port_selection_descriptor ( )

Syntax	Bytes	Type
port_selection_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
ps_ip_address	16	uimsbf
ps_port	2	uimsbf
ps_number_of_source_ip	1	uimsbf
for (j=0; j< **ps_number_of_source_ip**;		
j++) {		
ps_source_ip_address	16	uimsbf
}		
}		

Splice_Descriptor_Tag – 0x05

Descriptor_Length – Variable. The length in bytes of the descriptor following this field.

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

ps_ip_address – The IPv6 Internet protocol address that the splicer shall use for the content associated with the splice. If this address:port combination is different from the address in the Logical_Mux_Type 0x0007 table, it shall be considered a dynamic port set-up request.

ps_port – The UDP port that the splicer shall use for the content associated with the splice. This port number shall override the automatic port selection method of the Logical_Multiplex_Type 0x0007.

ps_number_of_source_ip – Specifies how many **ps_source_ip_address**(es) follow. The valid range is 0-32.

ps_source_ip_address – The source IPv6 address(es) that the splicer shall use in an IGMP V3 join or an MLD join for the associated multicast **ps_ip_address**.

8.5.5 asset_id_descriptor( ) field definitions

The asset_id_descriptor( ) is an implementation of the splice_API_descriptor( ) which shall only be used in the **Splice_Request** message. This descriptor is intended to be used when the asset playout is being performed by the splicer, but can be used to identify the asset being played also. It is suggested for advertising content that the spot id or other identifier be used rather than a fully qualified name. For long form VOD type content, a full asset name will probably be required. How the server and splicer manage the name format is not specified in this Recommendation.

Table 8-15 – asset_id_descriptor ( )

Syntax	Bytes	Type
asset_id_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_Api_Identifier	4	uimsbf
Asset_Upid_Type	1	uimsbf
Asset_Upid_Length	1	uimsbf
Asset_Upid( )		uimsbf
}		

Splice_Descriptor_Tag – 0x06

Descriptor_Length – Variable. The length in bytes of the descriptor following this field.

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

Asset_Upid_Type – A value from the segmentation_upid_type table (Table 8-6 of [ANSI/SCTE 35]). When the type value from Table 8-6 of [ANSI/SCTE 35] is 0x09 (ADI), the requirements regarding this ADI-related value in section 8.3.3.1 of [ANSI/SCTE 35] shall form a part of this specification.

Asset_Upid_Length – Length in bytes of the **Asset_Upid** structure. The maximum value of this field is 245.

Asset_Upid( ) – Length and identification from the segmentation_upid_type table (Table 8-6 of [ANSI/SCTE 35]. This structure's contents and length are determined by the Asset_Upid_Type and Asset_Upid_Length fields. An example is a type of 0x06 for V-ISAN and a length of 12 bytes. This field would then contain the V-ISAN identifier for the content to which this descriptor refers.

8.5.6 create_feed_descriptor( ) field definitions

The create_feed_descriptor( ) is an implementation of the splice_API_descriptor( ) which shall only be used in the **Init_Request** message. This descriptor will give the splicer enough information to create the output SPTS.

The usage of this descriptor shall be as follows:

1.   The ad server shall open a new TCP connection to the splicer of the to-be-created feed.

2.   The ad server shall send an **Init_Request** message with this descriptor on the newly-created TCP connection.

3.   The **ChannelName** field in the **Init_Request** message shall reflect the name of the newly-created feed.

4.   The splicer shall create the feed according to the **Init_Request** message. The newly-created feed shall be identical to the feed denoted by **OriginalChannelName** (see below). The identity shall encompass all PID values, types, descriptors, PCR PID and program number within the newly-created feed's PMT and PAT.

5.   The splicer shall send an **Init_Response** message only after the new feed is created.

Table 8-16 – create_feed_descriptor ( )

Syntax	Bytes	Type
create_feed_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
OriginalChannelName	32	String
Create_Feed_Descriptor_Type	1	uimsbf
if (**Create_Feed_Descriptor_Type** == 0) {		
IPV4_Dest_Address	4	uimsbf
Destination_Port	2	uimsbf
}		
else if (**Create_Feed_Descriptor_Type** == 1) {		
IPV6_Dest_Address	16	uimsbf
Destination_Port	2	uimsbf
}		
}		

Splice_Descriptor_Tag – 0x07

Descriptor_Length – Variable. The length in bytes of the descriptor following this field.

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

OriginalChannelName – The logical name of the output channel used as the template for the newly-created feed. This is a null-terminated string.

Create_Feed_Descriptor_Type – This field shall be a 0 to indicate that an IPv4 address is used and shall be a 1 to indicate that an IPv6 address is used.

IPV4_Dest_Address – The destination IP address for the created output service. This is in the IPv4 format.

IPV6_Dest_Address – The destination IP address for the created output service. This is in the IPv6 format.

Destination_Port – The destination UDP port for the created output service.

8.5.7   source_info_descriptor( ) field definitions

The source_info_descriptor( ) is an implementation of the splice_API_descriptor( ) which may be used in the **Cue_Request** message. This descriptor will give the server enough information to match an insertion feed to the same type, resolution and frame rate as the primary channel. This descriptor shall only be used if the primary channel has one and only one video present.

Table 8-17 – source_info_descriptor ( )

Syntax	Bytes	Type
source_info_descriptor {		
Splice_Descriptor_Tag	1	uimsbf
Descriptor_Length	1	uimsbf
Splice_API_Identifier	4	uimsbf
StreamType	1	uimsbf
HResolution	2	uimsbf
VResolution	2	uimsbf
frame_rate_code	1	uimsbf
progressive_sequence	1	uimsbf
}		

Splice_Descriptor_Tag – 0x08

Descriptor_Length – 0x0A

Splice_API_Identifier – 0x53415049, ASCII "SAPI".

StreamType – This number corresponds with the PMT specification found in [ITU-T H.222.0].

HResolution – The width in number of pixels of the video pictures.

VResolution – The height in number of pixels of the video pictures.

frame_rate_code – For MPEG-2 video, this parameter shall be coded as per the frame_rate_code from Table 6-4 of [ITU-T H.262]. For AVC video, this parameter shall be coded as per the frame_rate_code value from Table 6-4 of [ITU-T H.262] that matches the coded frame rate listed in Table 11 of [ANSI/SCTE 128]. Frame rate codes are shown in Appendix III.

progressive_sequence – For MPEG-2 video, this parameter shall be coded as per clause 6.3.5 of [ITU-T H.262]. For AVC video, this parameter shall be set to '1' when the letter 'P' in Table 11 of [ANSI/SCTE 128] is listed with the corresponding frame rate; otherwise, this parameter shall be set to '0'.

9 Time synchronization

Time synchronization is required due to the communication of time between the server and the splicer. The delay on a TCP/IP message is somewhat unpredictable and is affected by other machines on the network. By having the machines synchronized, time can be communicated between the two machines without concern for normal network delays keeping the splicing very accurate. One possible method is to use a network time protocol (NTP) to keep the server and the splicer in synchronization. It is likely that servers already keep some time synchronization, and thus could provide the NTP service and the splicer could be an NTP client. A network common host system NTP server could also be used since this also typically exists in a cable headend that has a network infrastructure.

The time synchronization system must be able to keep the splicer and server within ±15 ms of each other. The system may use the **Alive_Request/Alive_Response** messages to detect if the two devices have proper synchronization and to alert the operator if synchronization is lost.

The bit stream representing the primary channel is subject to various delays, which may include upstream splicing, satellite links and other transmission and conditioning processes. These delays may total from milliseconds to seconds. However, these delays do not affect the accuracy of a cue

message embedded in the primary channel. The cue message uses the PCR to indicate the correct time of insertion so that it retains its original accuracy relative to the content.

The server providing the insertion channel content knows only about clock time (UTC) and the insertion windows with which it has been programmed are relative to clock time. However, it depends on the splicer to tell it the exact moment that it is to begin streaming that content.

When the splicer receives the program bit stream, all delays to that stream have occurred. The splicer can take the PCR and relate it to the clock time, then send a message to the server that specifies the exact UTC at which it must begin streaming. The insertion channel from the server now arrives at the splicer exactly synchronized with the primary channel, and a perfect splice can be achieved. Any additional delays that occur within the splicer are irrelevant, since the input bit streams were synchronized.

If all parts of a facility do not use UTC time, that is, parts use GPS-based time, then all splicing system implementations should consider the conversion from GPS-based time of day to UTC and consistently include or ignore the GPS to UTC offset in computing times.

10   System timing

10.1   DPI splice signal flow

Figures 10-1 and 10-2 convey specific details regarding the usage and ordering of the various messages allowed by this API. The actual usage of API messages may not be limited to these examples.
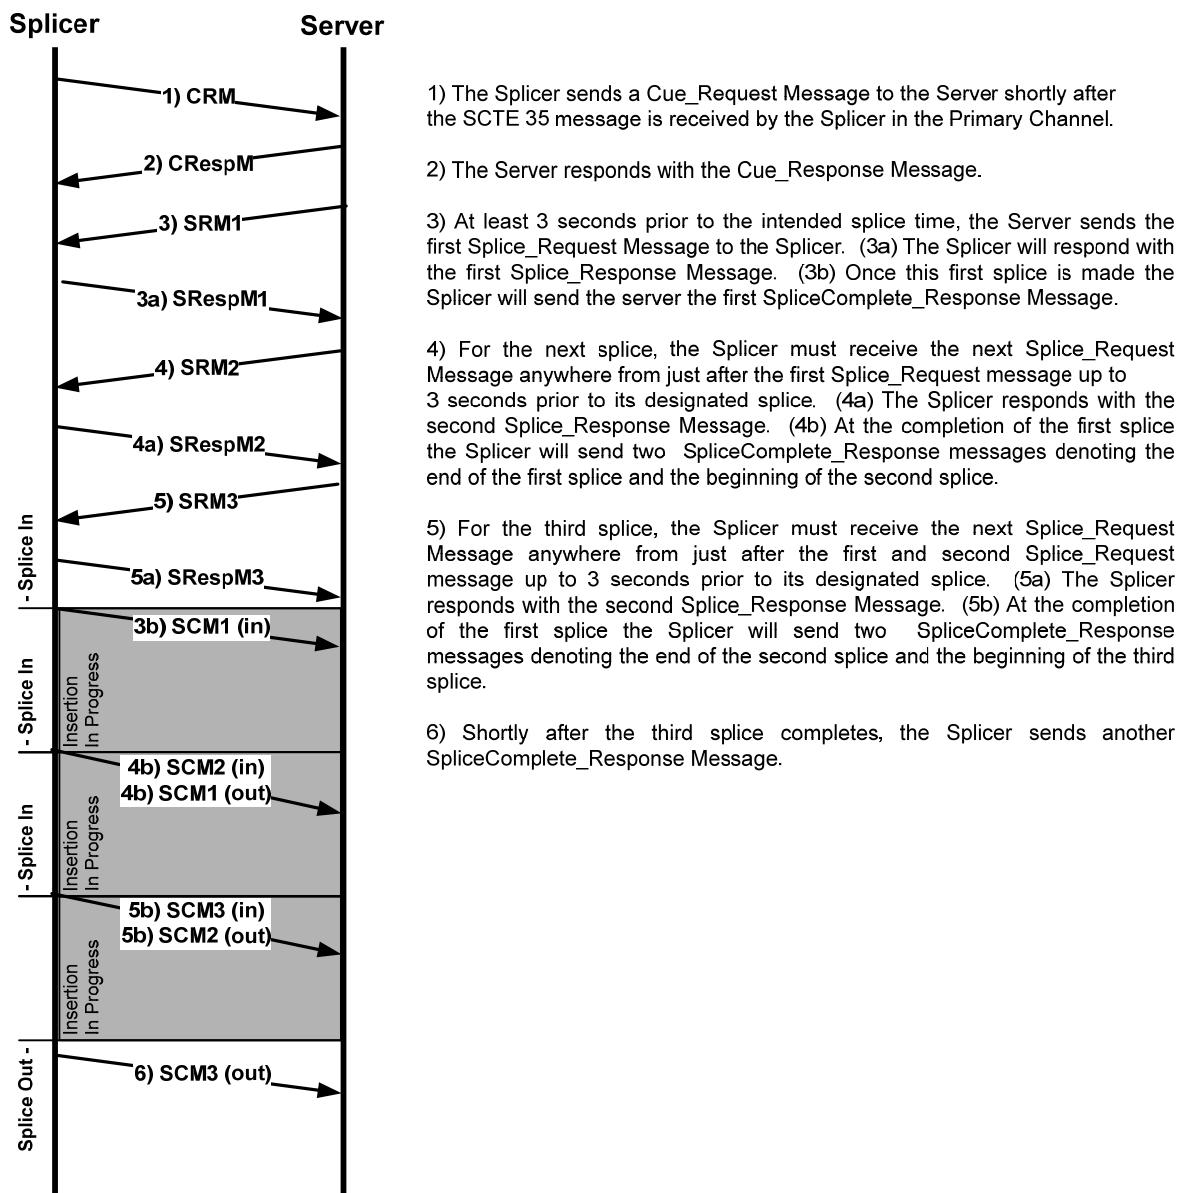

Figure 10-1 – Single event splice

Figure 10-2 – Multiple event splice

10.2 DPI splice initiation timeline

Figure 10-3 gives a timing example of the events leading up to the beginning of a program (or advertisement) insertion. Times in a real situation may vary from the timing shown in this figure. The interval of time shown is applicable to the discussion of priority arbitration as presented in clause 6.2. Operation in conjunction with ANSI/SCTE 35 cueing messages is also shown.

In the figure, bold black lines indicate the flow of MPEG information on the primary channel line and on the insertion channel line. Thin black lines indicate that MPEG information is either not flowing at that moment or is unimportant (i.e., not selected to appear at the output channel).

Figure 10-3 – DPI splice initiation timeline

Appendix I

Result codes

(This appendix does not form an integral part of this Recommendation.)

See the table on the following pages.

Result	Result Name	Description	Response Message
100	Successful Response		All
101	Unknown Failure		All
102	Invalid Version	Server and splicer are using different versions of this API.	**Init_Response**
103	Access Denied	Possible licence problem.	**Init_Response** **Splice_Response**
104	Invalid/Unknown **ChannelName**	Possible configuration error.	**Init_Response**
105	Invalid Physical Connection	Possible configuration error.	**Init_Response**
106	No Configuration Found	Splicer unable to determine the configuration for this connection.	**GetConfig_Response**
107	Invalid Configuration	One or more of the parameters in the configuration for this connection is invalid.	**GetConfig_Response**
108	Splice Failed – Unknown Failure		**SpliceComplete_Response**
109	Splice Collision	A higher or same priority is already set to splice.	**Splice_Response** **SpliceComplete_Response**
110	No Insertion Channel Found	This error shall be returned if the insertion channel is missing.	**SpliceComplete_Response**
111	No Primary Channel Found	This result shall be returned if the primary channel is missing at the splice-in or splice-out times.	**SpliceComplete_Response**
112	**Splice_Request** Was Too Late	The **Splice_Request** message was not received early enough (3 seconds) for the splicer to initiate the splice.	**Splice_Response**
113	No Splice Point Was Found	The splicer was unable to find a valid point to splice in to the Primary Channel.	**SpliceComplete_Response**
114	Splice Queue Full	Too many outstanding **Splice_Request** messages.	**Splice_Response**
115	Session Playback Suspect	Splicer has detected video or audio discrepancies that may have affected playback.	**SpliceComplete_Response**
116	Insertion Aborted	An **Abort_Request** message caused a splice-out.	**SpliceComplete_Response**
117	Invalid Cue Message	The splicer or server could not parse the cue message.	**General_Response** **Cue_Response**
118	Splicing Device Does Not Exist	**SplicerName** was not found.	**Init_Response**
119	**Init_Request** Refused	The splicer refuses to allow the server to connect.	**Init_Response**
120	Unknown **MessageID**	Use Splicing_API_Message to send response to requester. Echo back the unknown **MessageID**.	All
121	Invalid **SessionID**	The Splicer has no knowledge of the specified **SessionID**.	**Splice_Response** **Abort_Response** **ExtendedData_Response**
122	Session Did Not Complete	Splicer was not able to play the complete duration. This includes the case where the server did not supply the sufficient content.	**SpliceComplete_Response**
123	Invalid Request Message data( )	Splicer or server was not able to parse a field in the request message successfully. The invalid field position is returned in the **Result_Extension** field of	All

124	Descriptor Not Implemented	The splicer does not currently understand or implement the requested descriptor.	Responses to all messages that allow descriptors.
125	Channel Override	This result code is used to indicate to the currently playing insertion that it has been overridden with a splice-out status or has been re-entered with a Splice_in status.	**SpliceComplete_Response**
126	Insertion Channel Started Early	This error may be issued if the insertion channel started early and the splicer was not able to determine the correct start of the insertion stream.	**SpliceComplete_Response**
127	Playback Rate Below Threshold	See playback_descriptor( ) for details.	**SpliceComplete_Response**
128	PMT changed	This is used to indicate to the server that the PMT for this primary channel has changed.	**General_Response**
129	Invalid message size	The message was not the correct length as determined by this specification.	All
130	Invalid message syntax	Fields defined by this specification are not within the valid range.	All
131	Port Collision Error	The splicer was not able to utilize the specified IP:port combination requested. The combination is either in use or not valid on this splicer.	Init_Response
132	Splice Failed – EAS active	This error shall be returned if the emergency alert system is active.	**SpliceComplete_Response** **Splice_Response**
133	Insertion Component Not Found	One or more of the insertion stream components was not found. The result extension has the stream type of the first missing component.	**SpliceComplete_Response**
134	Resources Not Available	Returned by a splicer which was requested to provide a service for which it had no resources. i.e. the create_feed_descriptor() in an Init_Request message could not be honoured by the splicer due to the splicer's output B/W being fully occupied.	Init_Response

135	Component Mismatch	This code MAY be returned by a splicer when the Splice_Request fails to define an insertion channel component needed to match a primary channel component. Examples: 1. The primary channel contains an AC-3 audio component, but the Splice_Request defines only an MPEG2 audio component. 2. The primary channel contains an AVC video component, but the Splice_Request defines only an MPEG-2 video component.	Splice_Response
Note: All result codes may be used in the **General_Response** message.			

Appendix II

Example use of Logical_Multiplex type 0x0006 and the port_selection_descriptor( )

(This appendix does not form an integral part of this Recommendation.)

II.1 Informative example 1

The following example illustrates the use of multiple Splice_Requests in sequence and the incrementing of port numbers between the subsequent requests when port_selection_descriptors are not present. (See clause 8.2)

All ports are statically set up on the Init_Request.

Base IP:Port = 192.168.134.9:2000

Number of ports = 4

The following events occur sequentially in time during a single avail.

1) Splice_Request with time( ) set, server uses port 2000
2) Splice_Request with PriorSession, server uses port 2001
3) Splice_Request with PriorSession, server uses port 2002
4) Splice_Request with PriorSession, server uses port 2003
5) Splice_Request with PriorSession, server uses port 2000
6) Splice_Request with PriorSession, port_selection_descriptor port = 2000, server uses port 2000
7) Splice_Request with PriorSession, port_selection_descriptor port = 2003, server uses port 2003.

Next avail

1) Splice_Request with time( ) set, server uses port 2000.

II.2 Informative example 2

Use of port_selection_descriptor( ) to dynamically set up a port. Base port is established statically in the Init_Request message.

Base IP:Port = 192.168.134.9:3000

Number of ports = 1

The following events occur sequentially in time during a single avail.

1) Splice_Request with time( ) set, server uses port 3000
2) Splice_Request with PriorSession, port_selection_descriptor IP = 192.168.134.9 port = 2010, server sets up and uses 192.168.134.9:2010
3) Splice_Request with PriorSession, port_selection_descriptor IP = 239.192.0.2 port = 2010, server sets up and uses 239.192.0.2:2010.

Next avail

1) Splice_Request with time( ) set, server uses port 2000.

Appendix III

Frame rate codes

(This appendix does not form an integral part of this Recommendation.)

Table III.1 – Frame rate codes

Frame rate (Hz)	frame_rate_code	AVC time_scale	AVC num_units_in_tick
24/1.001 (23.976...)	'0001'	48000	1001
24	'0010'	48	1
30/1.001 (29.97...)	'0100'	60000	1001
30	'0101'	60	1
60/1.001 (59.94...)	'0111'	120000	1001
60	'1000'	120	1

Bibliography

[b-ANSI/SCTE 67]   ANSI/SCTE 67 (2010), *Recommended Practice for SCTE 35 Digital Program Insertion Cueing Message for Cable.*
                                    <http://www.scte.org/documents/pdf/Standards/ANSI_SCTE%2067%202010.pdf>

[b-IETF RFC 3810]   IETF RFC 3810 (2004), *Multicast Listener Discovery Version 2 (MLDv2) for IPv6.*
                                    <http://tools.ietf.org/html/rfc3810>

[Cable Television Laboratories]   Cable Television Laboratories (1997), *Cable Advertising*. White Paper. Louisville, CO, USA.

[Kar, Narasimhan, Prodan]   Kar, M. Narasimhan S., and Prodan R. (2000), *Local Commercial Insertion in the Digital Headend*, In: Proceedings of the NCTA 2000 Conference. New Orleans, USA.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	**Cable networks and transmission of television, sound programme and other multimedia signals**
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems