

A/V Transport Payloads Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 18, 2018

T. Edwards
FOX
December 15, 2017

RTP Payload for SMPTE ST 291-1 Ancillary Data
draft-ietf-payload-rtp-ancillary-13

Abstract

This memo describes a real-time transport protocol (RTP) payload format for the Society of Motion Picture and Television Engineers (SMPTE) Ancillary data, as defined by SMPTE ST 291-1. SMPTE Ancillary data is generally used along with professional video formats to carry a range of ancillary data types, including time code, Closed Captioning, and the Active Format Description (AFD).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 18, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	RTP Payload Format for SMPTE ST 291 Ancillary Data	3
2.1.	Payload Header Definitions	5
3.	Payload Format Parameters	11
3.1.	Media Type Definition	11
4.	SDP Considerations	13
4.1.	Grouping ANC Streams with other Media Streams	14
5.	Offer/Answer Model and Declarative Considerations	15
5.1.	Offer/Answer Model	15
5.2.	Declarative SDP Considerations	15
6.	IANA Considerations	15
7.	Security Considerations	16
8.	References	16
8.1.	Normative References	16
8.2.	Informative References	18
	Author's Address	19

1. Introduction

This memo describes a real-time transport protocol (RTP) payload format for the Society of Motion Picture and Television Engineers (SMPTE) Ancillary data (ANC), as defined by SMPTE ST 291-1 [ST291]. ANC data is transmitted in the ancillary space of serial digital video interfaces, the space outside of the active video region of images intended for users to view. Ancillary space roughly corresponds to vertical and horizontal blanking periods of cathode ray tube type displays. ANC can carry a range of data types, including time code, Closed Captioning, and the Active Format Description (AFD).

ANC is generally associated with the carriage of metadata within the bit stream of a Serial Digital Interface (SDI), such as the standard definition (SD) Serial Digital Interface, the 1.5 Gb/s Serial Digital Interface for high definition (HD) television applications, or the 3 Gb/s Signal/Data Serial Interface (SMPTE ST 259 [ST259], SMPTE ST 292-1 [ST292], SMPTE ST 424 [ST424]).

ANC data packet payload definitions for a specific application are specified by a SMPTE Standard, Recommended Practice, Registered Disclosure Document, or by a document generated by another organization, a company, or an individual (an Entity). When a

payload format is registered with SMPTE, it is identified by a registered data identification word.

This memo describes an RTP payload that supports carriage of ANC data packets with origin from any location within any SMPTE defined SDI signal, or even if the ANC packets did not originate in an SDI signal. Sufficient information is provided to enable the ANC data packets at the output of the decoder to be restored to their original locations in the serial digital video signal raster (if that is desired). An optional Media Type parameter allows for signaling of carriage of one or more types of ANC data as specified by Data Identification (DID) and Secondary Data Identification (SDID) words. Another optional Media Type parameter allows for the identification of the Video Payload ID (VPID) Code of the source interface of ANC data packets.

Note that the ancillary data flag (ADF) word is not specifically carried in this RTP payload. The ADF might be specified in a document defining an interconnecting digital video interface, otherwise a default ADF is specified by SMPTE ST 291-1 [ST291].

This ANC payload can be used by itself, or used along with a range of RTP video formats. In particular, it has been designed so that it could be used along with RFC 4175 [RFC4175] "RTP Payload Format for Uncompressed Video" or RFC 5371 [RFC5371] "RTP Payload Format for JPEG 2000 Video Streams."

The data model in this document for the ANC data RTP payload is based on the data model of SMPTE ST 2038 [ST2038], which standardizes the carriage of ANC data packets in an MPEG-2 Transport Stream.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. RTP Payload Format for SMPTE ST 291 Ancillary Data

An example of the format of an RTP packet containing SMPTE ST 291 Ancillary Data is shown below:




Figure 1: SMPTE Ancillary Data RTP Packet Format

In this example, two ANC data packets are present. The first has four 10-bit User Data Words, and the second has five 10-bit User Data Words (note that few ANC data packets are this small, thus this example does not reflect actual defined ANC data packets, and does not specifically call out the DIDs and SDIDs). The ANC data packets are located on lines 9 and 10 of the SDI raster.

The term "network byte order" in the payload format SHALL refer to the Data Transmission Order as defined in Appendix B of RFC 791 [RFC0791].

RTP packet header fields SHALL be interpreted as per RFC 3550 [RFC3550], with the following specifics:

Timestamp: 32 bits

The timestamp field is interpreted in a similar fashion to RFC 4175 [RFC4175]:

For progressive scan video, the timestamp denotes the sampling instant of the frame to which the ancillary data in the RTP packet belongs. RTP packets MUST NOT include ANC data from multiple frames, and all RTP packets with ANC data belonging to the same frame MUST have the same timestamp.

For interlaced video, the timestamp denotes the sampling instant of the field to which the ancillary data in the RTP packet belongs. RTP packets MUST NOT include ANC data from multiple fields, and all RTP packets belonging to the same field MUST have the same timestamp.

If the sampling instant does not correspond to an integer value of the clock, the value SHALL be truncated to the next lowest integer, with no ambiguity. Section 3.1 describes timestamp clock rates.

Marker bit (M): 1 bit

The marker bit set to "1" indicates the last ANC RTP packet for a frame (for progressive scan video) or the last ANC RTP packet for a field (for interlaced video).

2.1. Payload Header Definitions

The ANC RTP payload header fields are defined as:

Extended Sequence Number: 16 bits

The high order bits of the extended 32-bit sequence number, in network byte order. This is the same as the Extended Sequence Number field in RFC 4175 [RFC4175].

Length: 16 bits

Number of octets of the ANC RTP payload, beginning with the "C" bit of the first ANC packet data, as an unsigned integer in network byte order. Note that all word_align fields contribute to the calculation of the length field.

ANC_Count: 8 bits

This field is the count of the total number of ANC data packets carried in the RTP payload, as an unsigned integer. A single ANC RTP packet payload cannot carry more than 255 ANC data packets.

If more than 255 ANC data packets need to be carried in a field or frame, additional RTP packets carrying ANC data MAY

be sent with the same RTP timestamp but with different sequence numbers. ANC_Count of 0 indicates that there are no ANC data packets in the payload (for example, for an RTP packet with the marker bit set indicating the last ANC RTP packet in a field/frame, even if that RTP packet carries no actual ANC data packets.) If ANC_Count is 0, Length will also be 0.

F: 2 bits

These two bits relate to signaling the field specified by the RTP timestamp in an interlaced SDI raster. A value of 0b00 indicates that either the video format is progressive or that no field is specified. A value of 0b10 indicates that the timestamp refers to the first field of an interlaced video signal. A value of 0b11 indicates that the timestamp refers to the second field of an interlaced video signal. The value 0b01 is not valid. Receivers SHOULD ignore an ANC data packet with an F field value of 0b01, and SHOULD process any other ANC data packets with valid F field values that are present in the RTP payload.

Note that some multi-stream SDI interfaces might use multiple interlaced signal streams to transmit progressive images, in which case the "F" bits would refer to the field of the interlaced stream used for transport of the ANC data packet.

reserved: 22 bits

The 22 reserved bits of value "0" follow the F field to ensure that the first ANC data packet header field in the payload begins 32-bit word-aligned to ease implementation.

For each ANC data packet in the payload, the following ANC data packet header fields MUST be present:

C: 1 bit

This flag, when set to "1", indicates that the ANC data corresponds to the color-difference data channel (C). When set to "0", this flag indicates either that the ANC data corresponds to the luma (Y) data channel, that the ANC data source is from an SD signal, or that the ANC data source has no specific luma or color-difference data channels. For ANC data from a multi-stream interface source, the C flag SHALL refer to the channel of the stream used to transport the ANC packet. For situations where there is no SDI source, if the ANC data type definition specifically requires the use of the C or Y data channel, the C flag SHALL reflect that requirement.

Line_Number: 11 bits

This field contains the line number (as defined in ITU-R BT.1700 [BT1700] for SD video or ITU-R BT.1120 [BT1120] for HD video) that corresponds to the location of the ANC data packet in an SDI raster as an unsigned integer in network byte order.

The following special Line_Number values code for the location of the ANC data packet to be in certain generic vertical regions of the SDI raster:

Line_Number	ANC data packet generic vertical location
0x7FF	Without specific line location within the field or frame
0x7FE	On any line in the range from the second line after the line specified for switching, as defined in SMPTE RP 168 [RP168], to the last line before active video, inclusive
0x7FD	On a line number larger than can be represented in 11 bits of this field (if needed for future formats)

Note that the lines that are available to convey ANC data are as defined in the applicable sample structure specification (e.g., SMPTE ST 274 [ST274], SMPTE ST 296 [ST296], ITU-R BT.656 [BT656]) and possibly further restricted per SMPTE RP 168 [RP168].

In multi-stream interfaces, this field refers to the line number that an ANC packet is carried on within a particular data stream in the interface.

Horizontal_Offset: 12 bits

This field defines the location of the ANC data packet in an SDI raster relative to the start of active video (SAV, a digital synchronizing signal present in SDI interfaces) as an unsigned integer in network byte order. A value of 0 means that the Ancillary Data Flag (ADF) of the ANC data packet begins immediately following SAV. For HD, this is in units of Y samples as defined in ITU-R BT.1120 [BT1120]. For SD, this is in units of (27MHz) multiplexed word numbers, as specified in SMPTE ST 125 [ST125].

The following special Horizontal_Offset values code for the location of the ANC data packet to be in certain generic horizontal regions of the SDI raster:

Horizontal_Offset	ANC data packet generic horizontal location
0xFFFF	Without specific horizontal location
0xFFFE	Within horizontal ancillary data space (HANC) as defined in SMPTE ST 291-1 [ST291]
0xFFFD	Within the ancillary data space located between SAV (Start of Active Video) and EAV (End of Active Video) markers of the serial digital interface
0xFFFC	Horizontal offset is larger than can be represented in the 12 bits of this field (if needed for future formats, or for certain low frame rate 720p formats)

In multi-stream interfaces, this field refers to the horizontal location where an ANC packet is placed on a line carried within a particular data stream in the interface.

Note that HANC data space will generally have higher luma sample numbers than any samples in the active digital line. Also note that SMPTE ST 296 [ST296] (1280 x 720 progressive active images) image sampling systems 7 and 8 (1280 x 720 progressive @ 24 fps and 1280 x 720 progressive @ 23.98 fps respectively) have a luma sample number maximum of 4124. It is unlikely that an actual implementation would have an ANC packet begin at a Horizontal_Offset beyond 4091 (0xFFB) in these formats, but should that occur, the Horizontal_Offset value 0xFFC can be used to signal a horizontal offset larger than can be represented in the field. Further note that the 12-bit field of Horizontal_Offset is kept that size in this memo to maintain easy conversion to/from SMPTE ST 2038 [ST2038] which also has a 12-bit Horizontal_Offset field.

S (Data Stream Flag): 1 bit

This field indicates whether the data stream number of a multi-stream data mapping used to transport the ANC data packet is specified. If the S bit is '0', then the StreamNum field provides no guidance regarding the source data stream number of the ANC data packet. If the S bit is '1', then the

StreamNum field carries information regarding the source data stream number of the ANC data packet.

StreamNum: 7 bits

If the S (Data Stream Flag) bit is '1', then the StreamNum field MUST carry identification of the source data stream number of the ANC data packet. If the data stream is numbered, then the StreamNum field SHALL carry the number of the source data stream minus one. If the source multi-stream interface does not have numbered data streams, the following numbers SHALL be used in this field: '0' for link A data stream, '1' for link B data stream. For stereoscopic multi-stream interface formats that do not have numbered streams, the following numbers SHALL be used in this field: '0' for left eye stream, '1' for right eye stream.

Note that in multi-link SDI connections, the physical link that a particular stream utilizes is typically specified by the interface standard. Also note that numbering of data streams is across the interface as a whole. For example, in the SMPTE ST 425-3 dual-link 3 Gb/s interface, the data streams are numbered 1-4 with data streams 1 and 2 on link 1 and data streams 3 and 4 on link 2.

An ANC data packet with the header fields Line_Number of 0x7FF and Horizontal_Offset of 0xFFF SHALL be considered to be carried without any specific location within the field or frame.

For each ANC data packet in the payload, immediately after the ANC data packet header fields, the following data fields MUST be present, with the fields DID, SDID, Data_Count, User_Data_Words, and Checksum_Word representing the 10-bit words carried in the ANC data packet, as per SMPTE ST 291-1 [ST291]:

DID: 10 bits

Data Identification Word

SDID: 10 bits

Secondary Data Identification Word. Used only for a "Type 2" ANC data packet. Note that in a "Type 1" ANC data packet, this word will actually carry the Data Block Number (DBN).

Data_Count: 10 bits

The lower 8 bits of Data_Count, corresponding to bits b7 (MSB) through b0 (LSB) of the 10-bit Data_Count word, contain the actual count of 10-bit words in User_Data_Words. Bit b8 is the even parity for bits b7 through b0, and bit b9 is the inverse (logical NOT) of bit b8.

User_Data_Words: integer number of 10 bit words

User_Data_Words (UDW) are used to convey information of a type as identified by the DID word or the DID and SDID words. The number of 10-bit words in the UDW is defined by the Data_Count field. The 10-bit words are carried in order starting with the most significant bit and ending with the least significant bit.

Checksum_Word: 10 bits

The Checksum_Word can be used to determine the validity of the ANC data packet from the DID word through the UDW. It consists of 10 bits, where bits b8 (MSB) through b0 (LSB) define the checksum value and bit b9 is the inverse (logical NOT) of bit b8. The checksum value is equal to the nine least significant bits of the sum of the nine least significant bits of the DID word, the SDID word, the Data_Count word, and all User_Data_Words in the ANC data packet. The checksum is initialized to zero before calculation, and any end carry resulting from the checksum calculation is ignored.

At the end of each ANC data packet in the payload:

word_align: bits as needed to complete 32-bit word

Word align contains enough "0" bits as needed to complete the last 32-bit word of the ANC packet's data in the RTP payload. If an ANC data packet in the RTP payload ends aligned with a word boundary, there is no need to add any word alignment bits. Word align SHALL be used even for the last ANC data packet in an RTP packet. Word align SHALL NOT be used if there are zero ANC data packets being carried in the RTP packet.

When reconstructing an SDI signal based on this payload, it is important to place ANC data packets into the locations indicated by the ANC payload header fields C, Line_Number and Horizontal_Offset, and also to follow the requirements of SMPTE ST 291-1 [ST291] Section 7 "Ancillary Data Space Formatting (Component or Composite Interface)", which include rules on the placement of initial ANC data into allowed spaces as well as the contiguity of ANC data packet sequences within those spaces in order to assure that the resulting ANC data packets in the SDI signal are valid. The optional Media Type parameter VPID_Code can inform receivers of the type of originating SDI interface. For multi-stream originating interfaces, the StreamNum field can provide information regarding which stream an ANC data packet can be placed in to match the ANC data location in the originating SDI interface.

Senders of this payload SHOULD transmit available ANC data packets as soon as practical to reduce end-to-end latency, especially if receivers will be embedding the received ANC data packet into an SDI signal emission. One millisecond is a reasonable upper bound for the amount of time between when an ANC data packet becomes available to a sender and the emission of an RTP payload containing that ANC data packet.

ANC data packets with headers that specify specific location within a field or frame SHOULD be sent in raster scan order, both in terms of packing position within an RTP packet and in terms of transmission time of RTP packets.

3. Payload Format Parameters

This RTP payload format is identified using the video/smpte291 media type, which is registered in accordance with RFC 4855 [RFC4855], and using the template of RFC 6838 [RFC6838].

Note that the Media Type Definition is in the "video" tree due to the expected use of SMPTE ST 291 Ancillary Data along with video formats.

3.1. Media Type Definition

Type name: video

Subtype name: smpte291

Required parameters:

Rate:

RTP timestamp clock rate.

When an ANC RTP stream is to be associated with an RTP video stream, the RTP timestamp rates SHOULD be the same to ensure that ANC data packets can be associated with the appropriate frame or field. Otherwise, a 90 kHz rate SHOULD be used.

Note that techniques described in RFC 7273 [RFC7273] can provide a common reference clock for multiple RTP streams intended for synchronized presentation.

Optional parameters:

DID_SDID:

Data identification and Secondary data identification words.

The presence of the DID_SDID parameters signals that all ancillary data packets of this stream are of a particular type or types, i.e., labeled with a particular DIDs and SDIDs. DID and SDID values of SMPTE Registered ANC packet types can be found on the at the SMPTE Registry for Data Identification Word Assignments [SMPTE-RA] web site.

"Type 1" ANC packets (which do not have SDIDs defined) SHALL be labeled with SDID=0x00.

DID and SDID values can be registered with SMPTE as per SMPTE ST 291-1 [ST291].

The absence of the DID_SDID parameter signals that determination of the DID and SDID of ANC packets in the payload can only be achieved through direct inspection of the ANC data packet fields.

The ABNF description of the DID_SDID parameter is described in Section 4 of [this RFC].

VPID_Code:

This integer parameter specifies the Video Payload ID (VPID) Code of the source interface of ANC data packets using the value from byte 1 of the VPID as defined in SMPTE ST 352 [ST352]. The integer SHALL be made with bit 7 of VPID byte 1 being the most significant bit, and bit 0 of VPID byte 1 being the least significant bit. For example, 132 refers to SMPTE ST 292-1, 720-line video payloads on a 1.5 Gbps (nominal) serial digital interface.

Encoding considerations: This media type is framed and binary; see Section 4.8 of RFC 6838 [RFC6838].

Security considerations: See Section 7 of [this RFC]

Interoperability considerations: Data items in smpte291 can be very diverse. Receivers might only be capable of interpreting a subset of the possible data items. Some implementations might care about the location of the ANC data packets in the SDI raster, but other implementations might not care.

Published specification: [this RFC]

Applications that use this media type: Devices that stream real-time professional video, especially those that interoperate with legacy serial digital interfaces (SDI).

Additional Information:

Deprecated alias names for this type: N/A

Magic number(s): N/A

File extension(s): N/A

Macintosh file type code(s): N/A

Person & email address to contact for further information: T. Edwards <thomas.edwards@fox.com>, IETF Payload Working Group <payload@ietf.org>

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP RFC 3550 [RFC3550]. Transport within other framing protocols is not defined at this time.

Author: T. Edwards <thomas.edwards@fox.com>

Change controller: The IETF PAYLOAD working group, or other party as designated by the IESG.

4. SDP Considerations

The mapping of the above defined payload format media type and its parameters SHALL be done according to Section 3 of RFC 4855 [RFC4855].

- o The type name ("video") goes in SDP "m=" as the media name.
- o The subtype name ("smpte291") goes in SDP "a=rtpmap" as the encoding name, followed by a slash ("/") and the rate parameter.
- o The optional parameters VPID_Code and DID_SDID, when present, are included in the "a=fmtp" attribute line of SDP as a semicolon-separated list of parameter=value pairs.

DID and SDID values SHALL be specified in hexadecimal with a "0x" prefix (such as "0x61"). The ABNF as per RFC 5234 [RFC5234] of the DID_SDID optional parameter SHALL be:

```
TwoHex = "0x" 1*2(HEXDIG)
DidSdid = "DID_SDID={" TwoHex " ," TwoHex "}"
```

For example, EIA 608 Closed Caption data would be signalled with the parameter `DID_SDID={0x61,0x02}`. If a `DID_SDID` parameter is not specified, then the ancillary data stream might potentially contain ancillary data packets of any type.

Multiple `DID_SDID` parameters can be specified (separated by semicolons) to signal the presence of multiple types of ANC data in the stream. `DID_SDID={0x61,0x02};DID_SDID={0x41,0x05}`, for example, signals the presence of EIA 608 Closed Captions as well as AFD/Bar Data. Multiple `DID_SDID` parameters do not imply any particular ordering of the different types of ANC packets in the stream.

If the optional parameter `VPID_Code` is present, it SHALL be present only once in the semicolon-separated list, taking a single integer value.

A sample SDP mapping for ancillary data is as follows:

```
m=video 30000 RTP/AVP 112
a=rtpmap:112 smpte291/90000
a=fmtp:112 DID_SDID={0x61,0x02};DID_SDID={0x41,0x05};VPID_Code=132
```

In this example, a dynamic payload type 112 is used for ancillary data. The 90 kHz RTP timestamp rate is specified in the "a=rtpmap" line after the subtype. In the "a=fmtp:" line, `DID 0x61` and `SDID 0x02` are specified (registered to EIA 608 Closed Caption Data by SMPTE), and also `DID 0x41` and `SDID 0x05` (registered to AFD/Bar Data). The `VPID_Code` is 132 (referring to SMPTE ST 292-1, 720-line video payloads on a 1.5 Gbps serial digital interface).

4.1. Grouping ANC Streams with other Media Streams

To indicate the association of an ancillary data stream with a particular video stream, implementers MAY group the "m" lines together using Flow Identificaiton ("FID") semantics as defined in RFC 5888 [RFC5888].

A sample SDP mapping for grouping ANC data with RFC 4175 video is as follows:

```
v=0
o=A1 123456 11 IN IP4 host.example.com
s=Professional Networked Media Test
i=A test of synchronized video and ANC data
t=0 0
a=group:FID V1 M1
m=video 50000 RTP/AVP 96
c=IN IP4 233.252.0.1/255
a=rtpmap:96 raw/90000
a=fmtp:96 sampling=YCbCr-4:2:2; width=1280; height=720; depth=10
a=mid:V1
m=video 50010 RTP/AVP 97
c=IN IP4 233.252.0.2/255
a=rtpmap:97 smpte291/90000
a=fmtp:97 DID_SDID={0x61,0x02};DID_SDID={0x41,0x05}
a=mid:M1
```

5. Offer/Answer Model and Declarative Considerations

5.1. Offer/Answer Model

Receivers might wish to receive ANC data streams with specific DID_SDID parameters. Thus when offering ANC data streams using the Session Description Protocol (SDP) in an Offer/Answer model [RFC3264], the offeror MAY provide a list of ANC streams available with specific DID_SDID parameters in the fmtp line. The answerer MAY respond with all or a subset of the streams offered along with fmtp lines with all or a subset of the DID_SDID parameters offered. Or the answerer MAY set the corresponding port number to 0 to decline the smpte291 stream if not in the same media section as a corresponding video stream, or MAY remove the corresponding payload type if the smpte291 stream is in the same media section as a corresponding video stream. There are no restrictions on updating DID_SDID parameters in a subsequent offer.

5.2. Declarative SDP Considerations

For declarative use of SDP, nothing specific is defined for this payload format. The configuration given by the SDP MUST be used when sending and/or receiving media in the session.

6. IANA Considerations

One media type (video/smpte291) has been defined and needs registration in the media types registry. See Section 3.1

7. Security Considerations

RTP packets using the payload format defined in this specification are subject to the security considerations discussed in the RTP specification [RFC3550], and in any applicable RTP profile such as RTP/AVP [RFC3551], RTP/AVPF [RFC4585] RTP/SAVP [RFC3711] or RTP/SAVPF [RFC5124]. However, as "Securing the RTP Protocol Framework: Why RTP Does Not Mandate a Single Media Security Solution" [RFC7202] discusses, it is not an RTP payload format's responsibility to discuss or mandate what solutions are used to meet the basic security goals like confidentiality, integrity and source authenticity for RTP in general. This responsibility lays on anyone using RTP in an application. They can find guidance on available security mechanisms and important considerations in Options for Securing RTP Sessions [RFC7201]. Applications SHOULD use one or more appropriate strong security mechanisms. The rest of this security consideration section discusses the security impacting properties of the payload format itself.

To avoid potential buffer overflow attacks, receivers SHOULD validate that the ANC data packets in the RTP payload are of the appropriate length (using the Data_Count field) for the ANC data type specified by DID & SDID. Also the Checksum_Word SHOULD be checked against the ANC data packet to ensure that its data has not been damaged in transit, but the Checksum_Word is unlikely to provide a payload integrity check in case of a directed attack.

Some receivers will simply move the ANC data packet bits from the RTP payload into a serial digital interface (SDI). It might still be a good idea for these "re-embedders" to perform the above mentioned validity tests to avoid downstream SDI systems from becoming confused by bad ANC data packets, which could be used for a denial of service attack.

"Re-embedders" into SDI SHOULD also double check that the Line_Number and Horizontal_Offset leads to the ANC data packet being inserted into a legal area to carry ancillary data in the SDI video bit stream of the output video format.

8. References

8.1. Normative References

- [BT1120] ITU-R, "BT.1120-8, Digital Interfaces for HDTV Studio Signals", January 2012.

- [BT1700] ITU-R, "BT.1700, Characteristics of Composite Video Signals for Conventional Analogue Television Systems", February 2005.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <<https://www.rfc-editor.org/info/rfc3264>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<https://www.rfc-editor.org/info/rfc3550>>.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, DOI 10.17487/RFC3551, July 2003, <<https://www.rfc-editor.org/info/rfc3551>>.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <<https://www.rfc-editor.org/info/rfc3711>>.
- [RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, DOI 10.17487/RFC4585, July 2006, <<https://www.rfc-editor.org/info/rfc4585>>.
- [RFC4855] Casner, S., "Media Type Registration of RTP Payload Formats", RFC 4855, DOI 10.17487/RFC4855, February 2007, <<https://www.rfc-editor.org/info/rfc4855>>.
- [RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, DOI 10.17487/RFC5124, February 2008, <<https://www.rfc-editor.org/info/rfc5124>>.

- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC5888] Camarillo, G. and H. Schulzrinne, "The Session Description Protocol (SDP) Grouping Framework", RFC 5888, DOI 10.17487/RFC5888, June 2010, <<https://www.rfc-editor.org/info/rfc5888>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/info/rfc6838>>.
- [RP168] SMPTE, "RP 168:2009, Definition of Vertical Interval Switching Point for Synchronous Video Switching", 2009.
- [ST291] SMPTE, "ST 291-1:2011, Ancillary Data Packet and Space Formatting", 2011.
- [ST352] SMPTE, "ST 352:2013, Payload Identification Codes for Serial Digital Interfaces", 2013.
- [ST424] SMPTE, "ST 424:2012, 3 Gb/s Signal/Data Serial Interface", 2012.

8.2. Informative References

- [BT656] ITU-R, "BT.656-5, Interfaces for Digital Component Video Signals in 525-Line and 625-Line Television Systems Operating at the 4:2:2 Level of Recommendation ITU-R BT.601", December 2007.
- [RFC4175] Gharai, L. and C. Perkins, "RTP Payload Format for Uncompressed Video", RFC 4175, DOI 10.17487/RFC4175, September 2005, <<https://www.rfc-editor.org/info/rfc4175>>.
- [RFC5371] Futemma, S., Itakura, E., and A. Leung, "RTP Payload Format for JPEG 2000 Video Streams", RFC 5371, DOI 10.17487/RFC5371, October 2008, <<https://www.rfc-editor.org/info/rfc5371>>.
- [RFC7201] Westerlund, M. and C. Perkins, "Options for Securing RTP Sessions", RFC 7201, DOI 10.17487/RFC7201, April 2014, <<https://www.rfc-editor.org/info/rfc7201>>.

- [RFC7202] Perkins, C. and M. Westerlund, "Securing the RTP Framework: Why RTP Does Not Mandate a Single Media Security Solution", RFC 7202, DOI 10.17487/RFC7202, April 2014, <<https://www.rfc-editor.org/info/rfc7202>>.
- [RFC7273] Williams, A., Gross, K., van Brandenburg, R., and H. Stokking, "RTP Clock Source Signalling", RFC 7273, DOI 10.17487/RFC7273, June 2014, <<https://www.rfc-editor.org/info/rfc7273>>.
- [SMPTE-RA] SMPTE Registration Authority, LLC, "SMPTE ST 291 Ancillary Data Identification Word Assignments for Registered DIDs", 2011, <<http://www.smp-te-ra.org/smp-te-ancillary-data-smp-te-st-291>>.
- [ST125] SMPTE, "ST 125:2013, SDTV Component Video Signal Coding 4:4:4 and 4:2:2 for 13.5 MHz and 18 MHz Systems", 2013.
- [ST2038] SMPTE, "ST 2038:2008, Carriage of Ancillary Data Packets in an MPEG-2 Transport Stream", 2008.
- [ST259] SMPTE, "ST 259:2008, SDTV Digital Signal/Data - Serial Digital Interface", 2008.
- [ST274] SMPTE, "ST 274:2008, 1920 x 1080 Image Sample Structure, Digital Representation and Digital Timing Reference Sequences for Multiple Picture Rates", 2008.
- [ST292] SMPTE, "ST 292-1:2012, 1.5 Gb/s Signal/Data Serial Interface", 2012.
- [ST296] SMPTE, "ST 296:2012, 1280 x 720 Progressive Image 4:2:2 and 4:4:4 Sample Structure - Analog and Digital Representation and Analog Interface", 2012.

Author's Address

Thomas G. Edwards
FOX
10201 W. Pico Blvd.
Los Angeles, CA 90035
USA

Phone: +1 310 369 6696
Email: thomas.edwards@fox.com