

---

Stream: Internet Engineering Task Force (IETF)  
RFC: [9627](#)  
Category: Standards Track  
Published: March 2025  
ISSN: 2070-1721  
Authors: J. Lennox   D. Hong   J. Uberti   S. Holmer   M. Flodman  
          8x8/Jitsi   Google   OpenAI   Google   Google

# RFC 9627

## The Layer Refresh Request (LRR) RTCP Feedback Message

---

### Abstract

This memo describes the RTCP Payload-Specific Feedback Message Layer Refresh Request (LRR), which can be used to request a state refresh of one or more substreams of a layered media stream. This document also defines its use with several RTP payloads for scalable media formats.

### Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9627>.

### Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction	2
2. Conventions and Terminology	3
2.1. Terminology	3
3. Layer Refresh Request	5
3.1. Message Format	5
3.2. Semantics	7
4. Usage with Specific Codecs	7
4.1. H.264 SVC	7
4.2. VP8	8
4.3. H.265	9
5. Usage with Different Scalability Transmission Mechanisms	10
6. SDP Definitions	10
7. Security Considerations	11
8. IANA Considerations	11
9. References	11
9.1. Normative References	11
9.2. Informative References	12
Authors' Addresses	13

## 1. Introduction

This memo describes an [RTCP \[RFC3550\] Payload-Specific Feedback Message \[RFC4585\]](#) "Layer Refresh Request" (LRR), which is designed to allow a receiver of a layered media stream to request that one or more of its substreams be refreshed. The stream can then be decoded by an endpoint that previously was not receiving those layers, without requiring that the entire stream be refreshed (as it would be if the receiver sent a [Full Intra Request \(FIR\) \[RFC5104\]](#); see also [\[RFC8082\]](#)).

The feedback message is applicable to both temporally and spatially scaled streams and to both single-stream and multi-stream scalability modes.

## 2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

### 2.1. Terminology

A "layer refresh point" is a point in a scalable stream after which a decoder, which previously had been able to decode only some (possibly none) of the available layers of stream, is able to decode a greater number of the layers.

For spatial (or quality) layers, in normal encoding, a subpicture can depend both on earlier pictures of that spatial layer and also on lower-layer pictures of the current picture. However, a layer refresh typically requires that a spatial-layer picture be encoded in a way that references only the lower-layer subpictures of the current picture, not any earlier pictures of that spatial layer. Additionally, the encoder must promise that no earlier pictures of that spatial layer will be used as reference in the future.

However, even in a layer refresh, layers other than the ones being refreshed may still maintain dependency on earlier content of the stream. This is the difference between a layer refresh and a FIR [RFC5104]. This minimizes the coding overhead of refresh to only those parts of the stream that actually need to be refreshed at any given time.

The spatial-layer refresh of an enhancement layer is shown below. The "<--" indicates a coding dependency.

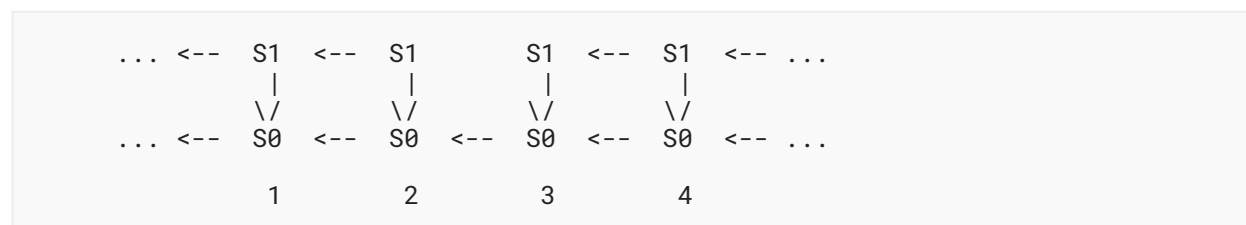


Figure 1: Refresh of a Spatial Enhancement Layer

In Figure 1, frame 3 is a layer refresh point for spatial layer S1; a decoder that had previously only been decoding spatial layer S0 would be able to decode layer S1 starting at frame 3.

The spatial-layer refresh of a base layer is shown below. The "<--" indicates a coding dependency.

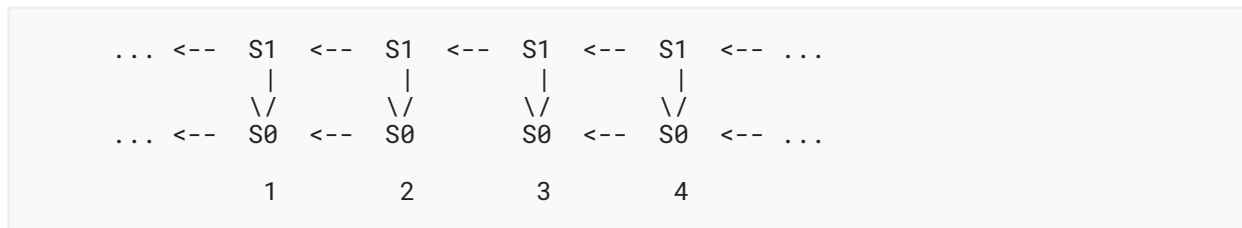


Figure 2: Refresh of a Spatial Base Layer

In [Figure 2](#), frame 3 is a layer refresh point for spatial layer S0; a decoder that had previously not been decoding the stream at all could decode layer S0 starting at frame 3.

For temporal layers, while normal encoding allows frames to depend on earlier frames of the same temporal layer, layer refresh requires that the layer be "temporally nested", i.e., use as reference only earlier frames of a lower temporal layer, not any earlier frames of this temporal layer and promise that no future frames of this temporal layer will reference frames of this temporal layer before the refresh point. In many cases, the temporal structure of the stream will mean that all frames are temporally nested; in this case, decoders will have no need to send LRR messages for the stream.

The temporal layer refresh is shown below. The "<--" indicates a coding dependency.




Figure 3: Refresh of a Temporal Layer

In [Figure 3](#), frame 6 is a layer refresh point for temporal layer T1; a decoder that had previously only been decoding temporal layer T0 would be able to decode layer T1 starting at frame 6.

An inherently temporally nested stream is shown below. The "<--" indicates a coding dependency.

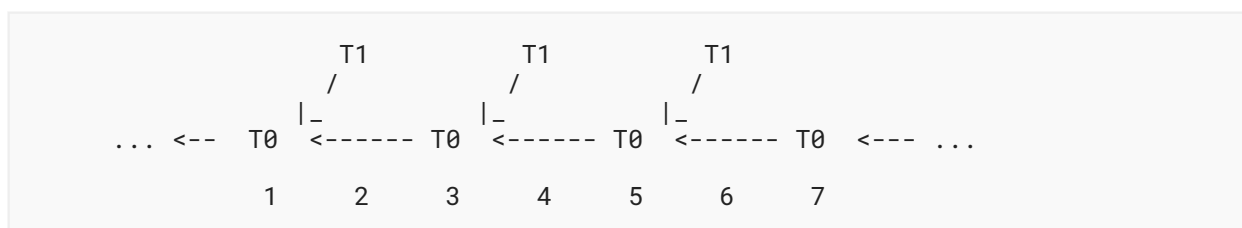


Figure 4: An Inherently Temporally Nested Stream

In [Figure 4](#), the stream is temporally nested in its ordinary structure; a decoder receiving layer T0 can begin decoding layer T1 at any point.

A "layer index" is a numeric label for a specific spatial and temporal layer of a scalable stream. It consists of both a "temporal-layer ID" identifying the temporal layer and a "layer ID" identifying the spatial or quality layer. The details of how layers of a scalable stream are labeled are codec specific. Details for several codecs are defined in [Section 4](#).

### 3. Layer Refresh Request

A layer refresh frame can be requested by sending a Layer Refresh Request (LRR), which is an [RTCP \[RFC3550\] payload-specific feedback message \[RFC4585\]](#) asking the encoder to encode a frame that makes it possible to upgrade to a higher layer. The LRR contains one or two tuples, indicating the temporal and spatial layer the decoder wants to upgrade to and (optionally) the currently highest temporal and spatial layer the decoder can decode.

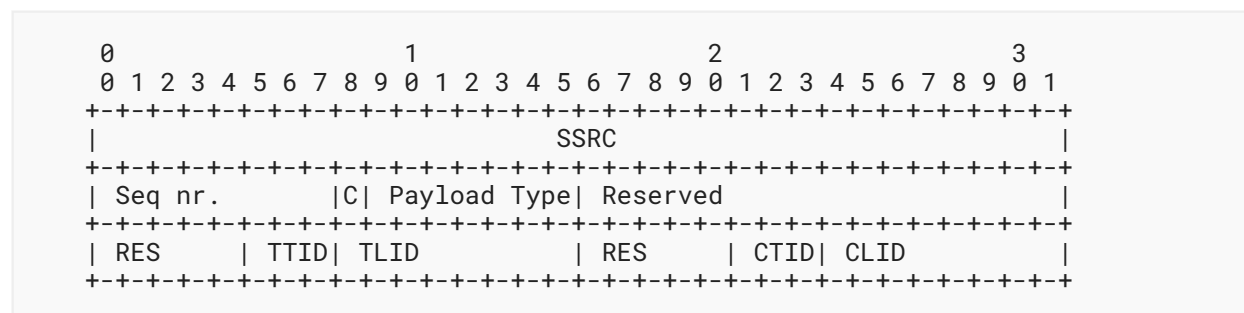
The specific format of the tuples, and the mechanism by which a receiver recognizes a refresh frame, is codec dependent. Usage for several codecs is discussed in [Section 4](#).

The design of LRR follows the FIR model ([Section 3.5.1 of \[RFC5104\]](#)) for its retransmission, reliability, and use in multipoint conferences.

The LRR message is identified by RTCP packet type value PT=PSFB and FMT=10. The Feedback Control Information (FCI) field **MUST** contain one or more LRR entries. Each entry applies to a different media sender, identified by its Synchronization Source (SSRC).

#### 3.1. Message Format

The FCI for the Layer Refresh Request consists of one or more FCI entries, the content of which is depicted in [Figure 5](#). The length of the LRR feedback message **MUST** be set to  $2+3*N$  32-bit words, where N is the number of FCI entries.



*Figure 5: Layer Refresh Request FCI Format*

Synchronization Source (SSRC) (32 bits):

The SSRC value of the media sender that is requested to send a layer refresh point.

**Seq nr. (8 bits):**

The command sequence number. The sequence number space is unique for each pairing of the SSRC of command source and the SSRC of the command target. The sequence number **SHALL** be increased by 1 for each new command (modulo 256, so the value after 255 is 0). A repetition **SHALL NOT** increase the sequence number. The initial value is arbitrary.

**C (1 bit):**

A flag bit indicating whether the Current Temporal-layer ID (CTID) and Current Layer ID (CLID) fields are present in the FCI. If this bit is 0, the sender of the LRR message is requesting refresh of all layers up to and including the target layer.

**Payload Type (7 bits):**

The RTP payload type for which the LRR is being requested. This gives the context in which the target layer index is to be interpreted.

**Reserved (RES) (three separate fields of 16 bits / 5 bits / 5 bits):**

All bits **SHALL** be set to zero by the sender and **SHALL** be ignored on reception.

**Target Temporal-layer ID (TTID) (3 bits):**

The temporal-layer ID of the target layer for which the receiver wishes a refresh point.

**Target Layer ID (TLID) (8 bits):**

The layer ID of the target spatial or quality layer for which the receiver wishes a refresh point. Its format is dependent on the payload type field.

**Current Temporal-layer ID (CTID) (3 bits):**

If C is 1, the ID of the current temporal layer being decoded by the receiver. This message is not requesting refresh of layers at or below this layer. If C is 0, this field **SHALL** be set to zero by the sender and **SHALL** be ignored on reception.

**Current Layer ID (CLID) (8 bits):**

If C is 1, the layer ID of the current spatial or quality layer being decoded by the receiver. This message is not requesting refresh of layers at or below this layer. If C is 0, this field **SHALL** be set to zero by the sender and **SHALL** be ignored on reception.

When C is 1, TTID **MUST NOT** be less than CTID, and TLID **MUST NOT** be less than CLID; at least one of either TTID or TLID **MUST** be greater than CTID or CLID, respectively. That is to say, the target layer index <TTID, TLID> **MUST** be a layer upgrade from the current layer index <CTID, CLID>. A sender **MAY** request an upgrade in both temporal and spatial or quality layers simultaneously.

A receiver receiving an LRR feedback packet that does not satisfy the requirements of the previous paragraph, i.e., one where the C bit is present but the TTID is less than the CTID or the TLID is less than the CLID, **MUST** discard the request.

Note: The syntax of the TTID, TLID, CTID, and CLID fields match, by design, the TID and LID fields in [RFC9626].

### 3.2. Semantics

Within the common packet header for feedback messages (as defined in [Section 6.1](#) of [\[RFC4585\]](#)), the "SSRC of packet sender" field indicates the source of the request, and the "SSRC of media source" is not used and **SHALL** be set to zero. The SSRCs of the media senders to which the LRR command applies are in the corresponding FCI entries. An LRR message **MAY** contain requests to multiple media senders, using one FCI entry per target media sender.

Upon reception of an LRR, the encoder **MUST** send a decoder refresh point (see [Section 2.1](#)) as soon as possible.

The sender **MUST** respect bandwidth limits provided by the application of congestion control, as described in [Section 5](#) of [\[RFC5104\]](#). As layer refresh points will often be larger than non-refreshing frames, this may restrict a sender's ability to send a layer refresh point quickly.

An LRR **MUST NOT** be sent as a reaction to picture losses due to packet loss or corruption; it is **RECOMMENDED** to use a [PLI \(Picture Loss Indication\)](#) [\[RFC4585\]](#) instead. An LRR **SHOULD** be used only in situations where there is an explicit change in a decoder's behavior: for example, when a receiver will start decoding a layer that it previously had been discarding.

## 4. Usage with Specific Codecs

In order for an LRR to be used with a scalable codec, the format of the temporal and layer ID fields (for both the target and current layer indices) needs to be specified for that codec's RTP packetization. New RTP packetization specifications for scalable codecs **SHOULD** define how this is done. (The [VP9 payload](#) [\[RFC9628\]](#), for instance, has done so.) If the payload also specifies how it is used with the Video Frame Marking RTP Header Extension described in [\[RFC9626\]](#), the syntax **MUST** be defined in the same manner as the TID and LID fields in that header.

### 4.1. H.264 SVC

[H.264 SVC](#) [\[RFC6190\]](#) defines temporal, dependency (spatial), and quality scalability modes.

```

+-----+-----+
|0|1|2|3|4|5|6|7|0|1|2|3|4|5|6|7|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RES      | TID |R|  DID  | QID |
+-----+-----+

```

*Figure 6: H.264 SVC Layer Index Fields Format*

[Figure 6](#) shows the format of the layer index fields for H.264 SVC streams. The "R" and "RES" fields **MUST** be set to zero on transmission and ignored on reception. See [Section 1.1.3](#) of [\[RFC6190\]](#) for details on the dependency\_id (DID), quality\_id (QID), and temporal\_id (TID) fields.

A dependency or quality layer refresh of a given layer in H.264 SVC can be identified by the I bit (`idr_flag`) in the extended Network Abstraction Layer (NAL) unit header, present in NAL unit types 14 (prefix NAL unit) and 20 (coded scalable slice). Layer refresh of the base layer can also be identified by its NAL unit type of its coded slices, which is "5" rather than "1". A dependency or quality layer refresh is complete once this bit has been seen on all the appropriate layers (in decoding order) above the current layer index (if any, or beginning from the base layer if not) through the target layer index.

Note that as the I bit in a Payload Content Scalability Information (PACSI) header is set if the corresponding bit is set in any of the aggregated NAL units it describes; thus, it is not sufficient to identify layer refresh when NAL units of multiple dependency or quality layers are aggregated.

In H.264 SVC, temporal layer refresh information can be determined from various Supplemental Encoding Information (SEI) messages in the bitstream.

Whether an H.264 SVC stream is scalably nested can be determined from the Scalability Information SEI message's `temporal_id_nesting` flag. If this flag is set in a stream's currently applicable Scalability Information SEI, receivers **SHOULD NOT** send temporal LRR messages for that stream, as every frame is implicitly a temporal layer refresh point. (The Scalability Information SEI message may also be available in the signaling negotiation of H.264 SVC as the `sprop-scalability-info` parameter.)

If a stream's `temporal_id_nesting` flag is not set, the Temporal Level Switching Point SEI message identifies temporal layer switching points. A temporal layer refresh is satisfied when this SEI message is present in a frame with the target layer index, if the message's `delta_frame_num` refers to a frame with the requested current layer index. (Alternately, temporal layer refresh can also be satisfied by a complete state refresh, such as an Instantaneous Decoding Refresh (IDR).) Senders that support receiving an LRR for streams that are not temporally nested **MUST** insert Temporal Level Switching Point SEI messages as appropriate.

## 4.2. VP8

The [VP8 RTP payload format \[RFC7741\]](#) defines temporal scalability modes. It does not support spatial scalability.

```

+-----+-----+
|0|1|2|3|4|5|6|7|0|1|2|3|4|5|6|7|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RES      | TID | RES      |
+-----+-----+

```

Figure 7: VP8 Layer Index Field Format

Figure 7 shows the format of the layer index field for VP8 streams. The "RES" fields **MUST** be set to zero on transmission and be ignored on reception. See [Section 4.2](#) of [\[RFC7741\]](#) for details on the TID field.



A VP8 layer refresh point can be identified by the presence of the Y bit (see [RFC7741]) in the VP8 payload header. When this bit is set, this and all subsequent frames depend only on the current base temporal layer. On receipt of an LRR for a VP8 stream, a sender that supports LRRs **MUST** encode the stream so it can set the Y bit in a packet whose temporal layer is at or below the target layer index.

Note that in VP8, not every layer switch point can be identified by the Y bit since the Y bit implies layer switch of all layers, not just the layer in which it is sent. Thus, the use of an LRR with VP8 can result in some inefficiency in transmission. However, this is not expected to be a major issue for temporal structures in normal use.

### 4.3. H.265

The initial version of the [H.265 payload format](#) [RFC7798] defines temporal scalability, with protocol elements reserved for spatial or other scalability modes (which are expected to be defined in a future version of the specification).

```

+-----+-----+
|0|1|2|3|4|5|6|7|0|1|2|3|4|5|6|7|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RES      | TID | RES | layer ID |
+-----+-----+

```

Figure 8: H.265 Layer Index Fields Format

Figure 8 shows the format of the layer index field for H.265 streams. The "RES" fields **MUST** be set to zero on transmission and ignored on reception. See Section 1.1.4 of [RFC7798] for details on the layer ID and TID fields.

H.265 streams signal whether they are temporally nested by using the `vps_temporal_id_nesting_flag` in the Video Parameter Set (VPS) and the `sps_temporal_id_nesting_flag` in the Sequence Parameter Set (SPS). If this flag is set in a stream's currently applicable VPS or SPS, receivers **SHOULD NOT** send temporal LRR messages for that stream, as every frame is implicitly a temporal layer refresh point.

If a stream's `sps_temporal_id_nesting_flag` is not set, the NAL unit types 2 to 5 inclusively identify temporal layer switching points. A layer refresh to any higher target temporal layer is satisfied when a NAL unit type of 4 or 5 with TID equal to 1 more than current TID is seen. Alternatively, layer refresh to a target temporal layer can be incrementally satisfied with a NAL unit type of 2 or 3. In this case, given current TID = TO and target TID = TN, layer refresh to TN is satisfied when a NAL unit type of 2 or 3 is seen for TID = T1, then TID = T2, all the way up to TID = TN (note that TN and TO refer to nonce variables in this instance). During this incremental process, layer refresh to TN can be completely satisfied as soon as a NAL unit type of 2 or 3 is seen.

Of course, temporal layer refresh can also be satisfied whenever any Intra-Random Access Point (IRAP) NAL unit type (with values 16-23, inclusively) is seen. An IRAP picture is similar to an IDR picture in H.264 (NAL unit type of 5 in H.264) where decoding of the picture can start without any older pictures.

In the (future) H.265 payloads that support spatial scalability, a spatial-layer refresh of a specific layer can be identified by NAL units with the requested layer ID and NAL unit types between 16 and 21, inclusive. A dependency or quality layer refresh is complete once NAL units of this type have been seen on all the appropriate layers (in decoding order) above the current layer index (if any, or beginning from the base layer if not) through the target layer index.

## 5. Usage with Different Scalability Transmission Mechanisms

Several different mechanisms are defined for how scalable streams can be transmitted in RTP. [Section 3.7](#) of "A Taxonomy of Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources" [[RFC7656](#)] defines three mechanisms: Single RTP stream on a Single media Transport (SRST), Multiple RTP streams on a Single media Transport (MRST), and Multiple RTP streams on Multiple media Transports (MRMT).

The LRR message is applicable to all these mechanisms. For MRST and MRMT mechanisms, the "media source" field of the LRR FCI is set to the SSRC of the RTP stream containing the layer indicated by the Current Layer Index (if "C" is 1) or the stream containing the base encoded stream (if "C" is 0). For MRMT, the LRR message is sent on the RTP session on which this stream is sent. On receipt, the sender **MUST** refresh all the layers requested in the stream, simultaneously in decode order.

## 6. SDP Definitions

[Section 7](#) of [[RFC5104](#)] defines Session Description Protocol (SDP) procedures for indicating and negotiating support for Codec Control Messages (CCM) in SDP. This document extends this with a new codec control command, "lrr", which indicates support of the LRR.

[Figure 9](#) gives a formal [Augmented Backus-Naur Form \(ABNF\)](#) [[RFC5234](#)] showing this grammar extension, extending the grammar defined in [[RFC5104](#)].

```
rtcp-fb-ccm-param =/ SP "lrr" ; Layer Refresh Request
```

*Figure 9: Syntax of the "lrr" CCM*

The Offer-Answer considerations defined in [Section 7.2](#) of [[RFC5104](#)] apply.

## 7. Security Considerations

All the security considerations of [FIR feedback packets \[RFC5104\]](#) apply to LRR feedback packets as well. Additionally, media senders receiving LRR feedback packets **MUST** validate that the payload types and layer indices they are receiving are valid for the stream they are currently sending, and discard the requests if not.

## 8. IANA Considerations

This document defines a new entry to the "Codec Control Messages" registry of the "Session Description Protocol (SDP) Parameters" registry group, according to the following data:

Value Name: lrr  
Long Name: Layer Refresh Request Command  
Usable with: ccm  
Mux: IDENTICAL-PER-PT  
Reference: RFC 9627

This document also defines a new entry to the "FMT Values for PSFB Payload Types" registry of the "Real-Time Transport Protocol (RTP) Parameters" registry group, according to the following data:

Name: LRR  
Long Name: Layer Refresh Request Command  
Value: 10  
Reference: RFC 9627

## 9. References

### 9.1. Normative References

- [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>>.
- [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>>.
- [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>>.

- [RFC5104] Wenger, S., Chandra, U., Westerlund, M., and B. Burman, "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)", RFC 5104, DOI 10.17487/RFC5104, February 2008, <<https://www.rfc-editor.org/info/rfc5104>>.
- [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>>.
- [RFC6190] Wenger, S., Wang, Y.-K., Schierl, T., and A. Eleftheriadis, "RTP Payload Format for Scalable Video Coding", RFC 6190, DOI 10.17487/RFC6190, May 2011, <<https://www.rfc-editor.org/info/rfc6190>>.
- [RFC7741] Westin, P., Lundin, H., Glover, M., Uberti, J., and F. Galligan, "RTP Payload Format for VP8 Video", RFC 7741, DOI 10.17487/RFC7741, March 2016, <<https://www.rfc-editor.org/info/rfc7741>>.
- [RFC7798] Wang, Y.-K., Sanchez, Y., Schierl, T., Wenger, S., and M. M. Hannuksela, "RTP Payload Format for High Efficiency Video Coding (HEVC)", RFC 7798, DOI 10.17487/RFC7798, March 2016, <<https://www.rfc-editor.org/info/rfc7798>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC9626] Zanaty, M., Berger, E., and S. Nandakumar, "Video Frame Marking RTP Header Extension", RFC 9626, DOI 10.17487/RFC9626, March 2025, <<https://www.rfc-editor.org/info/rfc9626>>.

## 9.2. Informative References

- [RFC7656] Lennox, J., Gross, K., Nandakumar, S., Salgueiro, G., and B. Burman, Ed., "A Taxonomy of Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources", RFC 7656, DOI 10.17487/RFC7656, November 2015, <<https://www.rfc-editor.org/info/rfc7656>>.
- [RFC8082] Wenger, S., Lennox, J., Burman, B., and M. Westerlund, "Using Codec Control Messages in the RTP Audio-Visual Profile with Feedback with Layered Codecs", RFC 8082, DOI 10.17487/RFC8082, March 2017, <<https://www.rfc-editor.org/info/rfc8082>>.
- [RFC9628] Uberti, J., Holmer, S., Flodman, M., Hong, D., and J. Lennox, "RTP Payload Format for VP9 Video", RFC 9628, DOI 10.17487/RFC9628, March 2025, <<https://www.rfc-editor.org/info/rfc9628>>.

## Authors' Addresses

**Jonathan Lennox**

8x8, Inc. / Jitsi  
Jersey City, NJ 07302  
United States of America  
Email: [jonathan.lennox@8x8.com](mailto:jonathan.lennox@8x8.com)

**Danny Hong**

Google, Inc.  
315 Hudson St.  
New York, NY 10013  
United States of America  
Email: [dannyhong@google.com](mailto:dannyhong@google.com)

**Justin Uberti**

OpenAI  
1455 3rd St  
San Francisco, CA 94158  
United States of America  
Email: [justin@uberti.name](mailto:justin@uberti.name)

**Stefan Holmer**

Google, Inc.  
Kungsbron 2  
SE-111 22 Stockholm  
Sweden  
Email: [holmer@google.com](mailto:holmer@google.com)

**Magnus Flodman**

Google, Inc.  
Kungsbron 2  
SE-111 22 Stockholm  
Sweden  
Email: [mflodman@google.com](mailto:mflodman@google.com)