

ETSI TS 126 522 V18.2.0 (2025-07)



5G; 5G Real-time Media Transport Protocol Configurations (3GPP TS 26.522 version 18.2.0 Release 18)



Reference

RTS/TSGS-0426522vi20

Keywords

5G

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from the
[ETSI Search & Browse Standards application](#).

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on [ETSI deliver repository](#).

Users should be aware that the present document may be revised or have its status changed,
this information is available in the [Milestones listing](#).

If you find errors in the present document, please send your comments to
the relevant service listed under [Committee Support Staff](#).

If you find a security vulnerability in the present document, please report it through our
[Coordinated Vulnerability Disclosure \(CVD\)](#) program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.

The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2025.
All rights reserved.

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the [ETSI IPR online database](#).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™**, **LTE™** and **5G™** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM®** and the GSM logo are trademarks registered and owned by the GSM Association.

Legal Notice

This Technical Specification (TS) has been produced by ETSI 3rd Generation Partnership Project (3GPP).

The present document may refer to technical specifications or reports using their 3GPP identities. These shall be interpreted as being references to the corresponding ETSI deliverables.

The cross reference between 3GPP and ETSI identities can be found at [3GPP to ETSI numbering cross-referencing](#).

Modal verbs terminology

In the present document **"shall"**, **"shall not"**, **"should"**, **"should not"**, **"may"**, **"need not"**, **"will"**, **"will not"**, **"can"** and **"cannot"** are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"must" and **"must not"** are **NOT** allowed in ETSI deliverables except when used in direct citation.

Contents

Intellectual Property Rights	2
Legal Notice	2
Modal verbs terminology.....	2
Foreword.....	5
Introduction	6
1 Scope	7
2 References	7
3 Definitions of terms, symbols and abbreviations	8
3.1 Terms.....	8
3.2 Symbols.....	9
3.3 Abbreviations	9
4 RTP Functionalities	9
4.1 General	9
4.2 RTP Header Extension for PDU Set Marking.....	10
4.2.1 General.....	10
4.2.2 One-byte RTP Header Extension Format	10
4.2.3 Two-byte RTP Header Extension Format.....	10
4.2.4 Semantics.....	10
4.2.5 SDP Signaling.....	11
4.2.6 Guidelines for PDU Set Marking.....	12
4.2.6.1 End of Data Burst Field.....	12
4.2.6.2 PDU Set Importance Field	12
4.2.6.2.1 General	12
4.2.6.2.2 H.264/AVC Codec	13
4.2.6.2.3 H.265/HEVC Codec	13
4.2.6.2.4 PSI based on affected PDU Sets.....	15
4.2.6.2.5 PSI mapping based on PDU Set dependencies	15
4.2.6.3 PDU Set Size Field	16
4.2.7 Guidelines for AS	17
4.3 RTP Header Extension for XR Pose.....	17
4.3.1 General.....	17
4.3.2 SDP Signaling.....	17
4.3.3 Header Extension Format	18
4.4 RTP Header Extension for in-band end-to-end delay measurement	20
4.4.1 General.....	20
4.4.2 One-byte RTP Header Extension Format	21
4.4.3 Two-byte RTP Header Extension Format.....	21
4.4.4 Syntax	22
4.4.5 Semantics.....	22
4.4.6 SDP signaling	22
5 RTCP Feedback Reporting Procedures	24
5.1 General	24
5.2 Transmission of timing information data for QoE measurements.....	24
5.2.1 General.....	24
5.2.2 RTCP message-based transmission of timing information	24
5.2.2.1 General	24
5.2.2.2 Extended Report block for QoE timing information.....	25
5.2.3 SDP signaling and attributes	26
Annex A (informative): Guidelines for PDU Set identification.....	27
A.0 General	27

A.1	Leveraging RTP Header Extensions	27
A.2	Obtaining PDU Set information from RTP Header or Payload	27
A.2.0	General	27
A.2.1	RTP/SRTP header	27
A.2.2	RTP payload	28
A.2.2.1	General	28
A.2.2.2	RTP payload for H.264/AVC codec	28
A.2.2.3	RTP payload for H.265/HEVC codec	29
Annex B (informative): Examples of SDP offers and answers		31
B.1	SDP example for RTP header extension for XR pose	31
Annex C (informative): RTP Header Extension for Absolute Sender Time		32
Annex D (informative): IANA registration information for RTP Header Extensions		33
D.1	Introduction	33
D.2	urn:3gpp:pdu-set-marking:rel-18	33
D.3	urn:3gpp:xr-pose	33
D.4	urn:3gpp:delay-measurement-response:rel-18	33
Annex E (informative): IANA registration information for RTCP XR block types		35
E.1	Introduction	35
E.2	Timing Information for QoE Metrics Calculation	35
Annex F (informative): Change history		35
History		36

Foreword

This Technical Specification has been produced by the 3rd Generation Partnership Project (3GPP).

The contents of the present document are subject to continuing work within the TSG and may change following formal TSG approval. Should the TSG modify the contents of the present document, it will be re-released by the TSG with an identifying change of release date and an increase in version number as follows:

Version x.y.z

where:

- x the first digit:
 - 1 presented to TSG for information;
 - 2 presented to TSG for approval;
 - 3 or greater indicates TSG approved document under change control.
- y the second digit is incremented for all changes of substance, i.e. technical enhancements, corrections, updates, etc.
- z the third digit is incremented when editorial only changes have been incorporated in the document.

In the present document, modal verbs have the following meanings:

- shall** indicates a mandatory requirement to do something
- shall not** indicates an interdiction (prohibition) to do something

The constructions "shall" and "shall not" are confined to the context of normative provisions, and do not appear in Technical Reports.

The constructions "must" and "must not" are not used as substitutes for "shall" and "shall not". Their use is avoided insofar as possible, and they are not used in a normative context except in a direct citation from an external, referenced, non-3GPP document, or so as to maintain continuity of style when extending or modifying the provisions of such a referenced document.

- should** indicates a recommendation to do something
- should not** indicates a recommendation not to do something
- may** indicates permission to do something
- need not** indicates permission not to do something

The construction "may not" is ambiguous and is not used in normative elements. The unambiguous constructions "might not" or "shall not" are used instead, depending upon the meaning intended.

- can** indicates that something is possible
- cannot** indicates that something is impossible

The constructions "can" and "cannot" are not substitutes for "may" and "need not".

- will** indicates that something is certain or expected to happen as a result of action taken by an agency the behaviour of which is outside the scope of the present document
- will not** indicates that something is certain or expected not to happen as a result of action taken by an agency the behaviour of which is outside the scope of the present document
- might** indicates a likelihood that something will happen as a result of action taken by some agency the behaviour of which is outside the scope of the present document

might not indicates a likelihood that something will not happen as a result of action taken by some agency the behaviour of which is outside the scope of the present document

In addition:

is (or any other verb in the indicative mood) indicates a statement of fact

is not (or any other negative verb in the indicative mood) indicates a statement of fact

The constructions "is" and "is not" do not indicate requirements.

Introduction

TR 26.998 [8] (5G Glass-type AR/MR) identified multiple aspects of normative work to support "5G/AR Real-time Communication" (clause 8.4). TR 26.998 identified normative work needed to support delivery of immersive media via RTP for IMS-based and WebRTC-based conversational services. To support XR split rendering as described in clause 8.6 of TR 26.998, RTP is also needed to transport immersive media and metadata information between the edge and device.

To improve support for the above XR services and enablers, it is necessary to configure RTP with specific settings and features that enable immersive experiences. Further improvements in performance and QoE over the 5G system can be achieved by specifying RTP configurations that are integrated and optimized for the 5G system, and leverage cross-layer optimizations used by other 3GPP specifications.

As these RTP configurations will be specified for use by multiple services, service enablers, and potentially, application developers, it is very important that they do not introduce unnecessary complexities that would discourage commercial deployment of the configurations. Therefore, technologies specified here should be commercially relevant and not introduce implementation and interoperability complexity without clearly demonstrating performance gains or new relevant functionalities.

1 Scope

The present document focuses on RTP [4] over UDP [9] for eXtended Reality in 5G.

RTP Header Extensions and RTCP Feedback Reporting are introduced for real-time immersive media and associated metadata for use in 5G Systems.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication, edition number, version number, etc.) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies. In the case of a reference to a 3GPP document (including a GSM document), a non-specific reference implicitly refers to the latest version of that document *in the same Release as the present document*.

- [1] 3GPP TR 21.905: "Vocabulary for 3GPP Specifications".
- [2] ITU-T Rec H.264 (08/2021): "Advanced video coding for generic audiovisual services" | ISO/IEC 14496-10:2022: "Information technology – Coding of audio-visual objects – Part 10: Advanced Video Coding".
- [3] ITU-T Rec H.265 (08/2021): "High efficiency video coding" | ISO/IEC 23008-2:2023: "High Efficiency Coding and Media Delivery in Heterogeneous Environments – Part 2: High Efficiency Video Coding".
- [4] IETF RFC 3550 (2003): "RTP: A Transport Protocol for Real-Time Applications", H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson.
- [5] IETF RFC 6184 (2011): "RTP Payload Format for H.264 Video", Y.-K. Wang, R. Even, T. Kristensen, R. Jesup.
- [6] IETF RFC 7798 (2016): "RTP Payload Format for High Efficiency Video Coding (HEVC)", Y.-K. Wang, Y. Sanchez, T. Schierl, S. Wenger, M. M. Hannuksela.
- [7] 3GPP TR 26.928: "Extended Reality (XR) in 5G".
- [8] 3GPP TR 26.998: "Support of 5G glass-type Augmented Reality / Mixed Reality (AR/MR) devices".
- [9] IETF RFC 768 (1980): "User Datagram Protocol", J. Postel.
- [10] IETF RFC 5761 (2010): "Multiplexing RTP Data and Control Packets on a Single Port", C. Perkins, M. Westerlund.
- [11] IETF RFC 8285 (2017): "A General Mechanism for RTP Header Extensions", D. Singer, H. Desineni, R. Even.
- [12] 3GPP TS 23.501: "System architecture for the 5G System (5GS)".
- [13] IETF RFC 5905 (2010): "Network Time Protocol Version 4: Protocol and Algorithms Specification", D. Mills, J. Martin, J. Burbank, W. Kasch.
- [14] IEEE 1588-2019 – IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, June 2020.
- [15] IETF RFC 4574 (2006): "The Session Description Protocol (SDP) Label Attribute", O. Levin, G. Camarillo.

- [16] IETF RFC 3611 (2003): "RTP Control Protocol Extended Reports (RTCP XR)", T. Friedman, R. Caceres, A. Clark.
- [17] 3GPP TS 26.119: "Media Capabilities for Augmented Reality".
- [18] IETF RFC 7656 (2015): "A Taxonomy of Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources ", J. Lennox, K. Gross, S. Nandakumar, G. Salgueiro, B. Burman.
- [19] IETF RFC 5888 "The Session Description Protocol (SDP) Grouping Framework", G. Camarillo et al.
- [20] ISO/IEC 60559:2020: "Floating-point arithmetic".

3 Definitions of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in TR 21.905 [1] and the following apply. A term defined in the present document takes precedence over the definition of the same term, if any, in TR 21.905 [1].

Age of content: The time duration between the moment the content is created and the time it is presented.

Estimated-at-time: Time when the pose was estimated.

Data Burst: A data burst is a set of multiple PDUs generated and sent by the application such that there is an idle period between two data bursts. A Data Burst can be composed of one or multiple PDU Sets.

Multimedia Session: An association among a group of participants engaged in the communication via one or more RTP sessions, as defined in section 2.2.4 of IETF RFC 7656 [18].

Orientation quaternion: Quaternion used to represent the orientation of an object.

PDU Set: One or more PDUs carrying the payload of one unit of information generated at the application level (e.g. frame(s), video slice(s), metadata, etc.).

PDU Set marking: Marking the PDUs carrying a payload with the PDU Set Information.

Rendered pose: An XR pose sent from a server to a client that was used for rendering at the server.

Roundtrip interaction delay: The sum of the *age of content* and the *user interaction delay*.

Scene Update Time: Time when the scene manager starts processing.

Split-render-output-time: Time of completing a rendering.

Split rendering server: Server to perform remote rendering.

Start-to-render-at-time: Time of starting a rendering.

User interaction delay: The time duration between the moment at which a user action is initiated and the time such an action is taken into account by the content creation engine.

XR Pose: A position and orientation in space relative to an XR Space.

XR Service: A service supporting XR use case as defined in clause 5 of [7].

XR Space: A frame of reference in which an application chooses to track the real world. An XR Space provides a relation of the user's physical environment with other tracked entities.

3.2 Symbols

For the purposes of the present document, the following symbols apply:

<i>Ih_p</i>	IP header overhead
<i>P</i>	Number of RTP packets
<i>R</i>	RTP payload data size
<i>Rh_p</i>	RTP header overhead, including any RTP header extensions
<i>S</i>	RTP packet maximum SDU size
<i>Uh_p</i>	UDP header overhead

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in TR 21.905 [1] and the following apply. An abbreviation defined in the present document takes precedence over the definition of the same abbreviation, if any, in TR 21.905 [1].

AP	Aggregation Packet
AVC	Advanced Video Coding
BLA	Broken Link Access
CRA	Clean Random Access
DoF	Degrees of Freedom
FU	Fragmentation Unit
HE	(RTP) Header Extension
HEVC	High Efficiency Video Coding
IDR	Instantaneous Decoder Refresh
IRAP	Intra Random Access Picture
NAL	Network Abstraction Layer
NRI	nal_ref_idc
NTP	Network Time Protocol
OS	Operating System
PACI	Payload Content Information
PPS	Picture Parameter Set
PSI	PDU Set Importance
PTP	Precision Time Protocol
RADL	Random Access Decodable Leading
RASL	Random Access Skipped Leading
RTCP	RTP Control Protocol
RTCP XR	RTCP eXtended Report
SPS	Sequence Parameter Set
SRS	Split Rendering Server
SRTP	Secure RTP
TID	Temporal Identifier
UPF	User Plane Function
UDP	User Datagram Protocol
VCL	Video Coding Layer
VPS	Video Parameter Set
XR	eXtended Reality

4 RTP Functionalities

4.1 General

This clause defines functional additions to RTP [4] for potential use in 5G Systems. The RTP protocol is designed to be generic and extensible, e.g., by using RTP Header Extensions (HE) [11] and by defining RTP payload formats for specific media codecs, e.g., [5][6].

- **End PDU of the PDU Set [E] (1 bit):** This field is a flag that shall be set to 1 for the last PDU of the PDU Set and set to 0 for all other PDUs of the PDU Set.
- **End of Data Burst [D] (1 bit):** This field is a flag that shall be set to 1 for the last PDU of a Data Burst. It shall be set to 0 for all other PDUs. A Data Burst may consist of one or more PDU Sets.

NOTE 1: The bit encodes the End of Data Burst indication as per the guidelines provided in clause 4.2.6.1.

- **Reserved [R] (2 bits):** This field is reserved for future usage (e.g., dynamic burst indication). It shall be set to 0 by the RTP sender and shall be ignored by the RTP receiver.
- **PDU Set Importance [PSI] (4 bits):** The PDU Set Importance field indicates the importance of this PDU Set compared to other PDU Sets within the same Multimedia Session. This information may help RAN to discard PDUs, when needed. Lower values shall indicate a higher importance PDU Set, with the highest importance PDU Set indicated by 1 and the lowest importance PDU Set indicated by 15. When the RTP sender cannot define an importance, it shall set the value to 0.

NOTE 2: A complete set of guidelines for setting the PSI field for the 3GPP audio/video codecs are provided in clause 4.2.6.2

- **PDU Set Sequence Number [PSSN] (10 bits):** The sequence number of the PDU Set to which the current PDU belongs, acting as a 10-bit numerical identifier for the PDU Set. The PSSN shall be incremented monotonically by 1 for each subsequent PDU Set.

NOTE 3: This value wraps around at 1023, however, using the 16-bit RTP packet sequence number and PSSN pair, a receiver may uniquely distinguish between any PDU Sets.

- **PDU Sequence Number within a PDU Set [PSN] (6 bits):** The sequence number of the current PDU within the PDU Set. The PSN shall be set to 0 for the first PDU in the PDU Set and incremented monotonically for every PDU in the PDU Set in the order of transmission from the sender.

NOTE 4: A receiver may use the RTP packet sequence number together with the PSN to distinguish between PDUs within a PDU Set that contains more than 64 PDUs.

- **PDU Set Size [PSSize] (24 bits):** The PDU Set Size indicates the total size of all PDUs of the PDU Set to which this PDU belongs. This field is optional and subject to an SDP signaling offer/answer negotiation, where the RTP sender shall indicate whether it will provide the size of the PDU Set for that RTP stream. If not enabled, the field shall not be present within the RTP HE. If enabled, but the RTP sender is not able to determine the PDU Set Size for a particular PDU Set, it shall set the value to 0 in all PDUs of that PDU Set. The PSSize shall indicate the size of a PDU Set including RTP/UDP/IP header encapsulation overhead of its corresponding PDUs. The PSSize shall be expressed in bytes. It is recommended to add the PDU Set Size field when the Number of PDUs in the PDU Set field is present.

NOTE 5: This field may be optionally present given the signaling of the "pdu-set-size" extension attribute in the SDP offer/answer negotiation as per clause 4.2.5.

- **Number of PDUs in the PDU Set [NPDS] (16 bits):** The number of PDUs within the PDU Set indicates the total number of PDUs belonging to the same PDU Set. This field is optional and subject to an SDP signaling offer/answer negotiation, where the RTP sender may indicate whether it will provide the number of PDUs within the PDU Set for that RTP stream. If enabled, but the RTP sender is not able to determine the Number of PDUs in the PDU Set, it shall set the value to 0 in all PDUs of that PDU Set. It is recommended to add the Number of PDUs in the PDU Set field when the PDU Set Size field is present.

NOTE 6: This field may be optionally present given the signaling of the "num-pdus-in-pdu-set" extension attribute in the SDP offer/answer negotiation as per clause 4.2.5.

NOTE 7: Guidelines to set the PDU Set Size in bytes by an RTP sender are provided in clause 4.2.6.3.

NOTE 8: The use of NPDS for the NAT46/NAT64 correction is FFS.

4.2.5 SDP Signaling

An RTP sender capable of sending RTP HE for PDU Set marking shall use the SDP extmap attribute for RTP HE for PDU Set marking in the media description of the RTP stream(s) carrying the RTP HE for PDU Set marking. An RTP

receiver that does not support RTP HE for PDU Set marking can ignore that RTP HE when included. The signaling of the PDU Set marking RTP HE shall follow the SDP signaling design and the syntax and semantics of the "extmap" attribute as outlined in RFC8285. The URN for the PDU Set marking shall be set to "**urn:3gpp:pdu-set-marking:rel-18**".

The ABNF syntax for the extmap attribute for the signaling of RTP HE for PDU Set marking is defined as follows, extending the ABNF in RFC 8285:

```
extensionname = "urn:3gpp:pdu-set-marking:rel-18"
extensionattributes = [pdu-set-attribute *2(SP pdu-set-attribute)]
pdu-set-attribute = format / "pdu-set-size" / "num-pdus-in-pdu-set"
format = "short" / "long"
```

The extension attributes have the following semantics:

- format: indicates if the RTP HE for PDU Set marking uses the 1-byte (short) or the 2-byte (long) format. This extension attribute shall not be included more than once.

NOTE: Regardless if this extension attribute is present or not, the use of long or short format is determined as described by section 4.1.2 of RFC 8285, i.e., based on what format other RTP HEs use in the same RTP session, unless both endpoints announced support for handling mixed format with "a=extmap-allow-mixed" as described by section 6 of RFC 8285.

- pdu-set-size: if present, this extension attribute indicates that the RTP sender will provide the PDU Set size in bytes in the RTP HE with every RTP packet. This results in an additional 3 bytes of length for the RTP HE. This extension attribute shall not be included more than once.
- num-pdus-in-pdu-set: if present, this extension attribute indicates that the RTP sender will provide the number of PDUs in the PDU Set (NPDS) in the RTP HE with every RTP packet. This results in an additional 2 bytes of length for the RTP HE. This extension attribute shall not be included more than once.

Below is an example:

```
a=extmap:7 urn:3gpp:pdu-set-marking:rel-18 short pdu-set-size
```

4.2.6 Guidelines for PDU Set Marking

4.2.6.1 End of Data Burst Field

NOTE: These detailed guidelines are FFS.

4.2.6.2 PDU Set Importance Field

4.2.6.2.1 General

In general, whenever the RAN needs to discard packets (e.g., under congestion situations), it is better to discard packets of lower importance rather than discarding packets randomly. If a discarded packet is critical for the media stream, the QoE may be severely degraded. For this reason, the PDU Set Importance (PSI) field can be used to mark PDU Sets with their importance level. The PSI field can then be used by the RAN to discard PDU sets. In case of congestion, PDU Sets with higher PSI values are more likely to be discarded.

PDU Sets that contain audio data should be assigned a lower PSI value (i.e., higher importance) compared with PDU Sets that contain other media types.

NOTE 1: PDU Sets that carry immersive audio data are not necessarily assigned a lower PSI value compared with the other media PDU Sets. The PSI value of immersive audio PDU Sets is FFS.

PDU Sets that contains the reference frames present in the video bitstream should be assigned a lower PSI value compared with PDU Sets that contain non-reference frames.

NOTE 2: It is assumed that the video bitstream uses referencing structures that have no coding delay caused by out-of-order output, as typically done for low-delay applications.

The following clauses provides the guidelines for the 3GPP video codecs on setting the PSI field in the RTP HE for PDU Set marking. For specific PSI value ranges, refer to clause 4.2.6.2.5.

4.2.6.2.2 H.264/AVC Codec

In an H.264/AVC bitstream, NAL units with the `nal_unit_type` field assigned the value 5 (refer to Table 7.1 in the H.264/AVC specification [2]) are Instantaneous Decoding Refresh (IDR) pictures. When the `Type` field value in the NAL Unit header of an RTP packet is 5, then the corresponding PDUs in that PDU Set should be set with higher importance.

The parameter set NAL units such as Sequence Parameter Set (SPS) and Picture Parameter Set (PPS) are important for decoding the bitstream. Therefore, PDU sets with a `Type` field value equal to 7, 8, 13 or 15 (refer to Table 7.1 in the H.264/AVC specification[2]) in the NAL Unit header of the RTP packet should be assigned a higher importance (lower PSI value) relative to PDU Sets with other `Type` field values.




Figure 4.2.6-1: Format of the H.264/AVC NAL unit header

The NAL unit type octet contains the NRI (`nal_ref_idc`) field highlighted in Figure 4.2.6-1. The NRI field indicate the relative transport priority. A value of b00 indicates that the content of the NAL unit is not used to reconstruct reference pictures for inter picture prediction. Such NAL units can be discarded by the RAN (in case of congestion) without risking the integrity of the reference pictures. Values greater than b00 indicate that the decoding of the NAL unit is required to maintain the integrity of the reference pictures. The highest transport priority is b11, followed by b10, and then by b01; finally, b00 is the lowest. PDU Sets with an NRI value b00 should be set with lower importance relative to the PDU Sets with other NRI values. PDU Sets with an NRI value b11 should be set with higher importance relative to the PDU Sets with other NRI field values.

The `Type` and NRI fields can be used to set the PSI. The PSI value assignment based on the `Type` and NRI field values is for further study.

4.2.6.2.3 H.265/HEVC Codec

Different from H.264/AVC, H.265/HEVC NAL unit header (shown in Figure 4.2.6-2) is two bytes, contains a 6-bit `Type` field, a 5-bit `LayerID` field, a 3-bit `TID` field, and no NRI field. The `Type` and `TID` fields in the NAL unit header indicate the relative transport priority and can be used to set the PSI.

NAL unit types 0–31 indicate Video Coding Layer (VCL) NAL unit types; 32–40 indicate non-VCL NAL unit types. NAL unit types 41–47 are reserved, and NAL unit types 48–63 are unspecified.




Figure 4.2.6-2: Format of the H.265/HEVC NAL unit header

All VCL NAL units of the same access unit have the same NAL unit type, which defines the type of the access unit and its coded picture. There are three basic classes of pictures in H.265/HEVC: intra random access point (IRAP) pictures, leading pictures, and trailing pictures.

In an H.265/HEVC bitstream, NAL units with the `nal_unit_type` field assigned a value in the range 16 to 23 (inclusive) (refer to Table 7.1 in the H.265/HEVC specification [3]) are Intra Random Access Pictures (IRAP) pictures. This includes IDR, CRA, and BLA picture types as well as types 22 and 23, which currently are reserved for future use.

When the `Type` field value in the NAL Unit header of RTP packet is in the range 16 to 23 (inclusive), then the corresponding PDUs in that PDU Set should be assigned higher importance (i.e., lower PSI value).

The parameter set NAL units such as Sequence Parameter Set (SPS), Picture Parameter Set (PPS), Video Parameter Set (VPS) are important for decoding the bitstream. Therefore, PDU Sets with payload `Type` field value in the NAL Unit header of RTP packet in the range 32 to 34 (inclusive) should be assigned higher importance (lower PSI value).

RFC 7798 [6] specifies Aggregation Packets (APs) to enable the reduction of packetization overhead for small NAL units, such as most of the non-VCL NAL units, which are often only a few octets in size. An AP aggregates NAL units within one access unit. Each NAL unit to be carried in an AP is encapsulated in an aggregation unit. An AP consists of a payload header (denoted as `PayloadHdr`) followed by two or more aggregation units. In an AP, the `Type` field in the `PayloadHdr` is equal to 48. APs are typically used to aggregate parameters sets (VPS, SPS, PPS) into a single packet. When APs are used, the sender should consider the NAL unit types of the aggregation units while assigning the PSI value. For example, if the aggregation unit contains parameter sets, PDU Sets containing those should be assigned higher importance (lower PSI value).

It could be that there are PDUs with different NAL unit types in a PDU Set. For example, if the first PDU in PDU Set is a prefix SEI message or Access Unit Delimiter (AUD), it would be misleading if the sender looked only at the first PDU of the PDU Set to determine the PSI value. The sender should ignore the NAL units with non-VCL NAL unit types 35 and 39 and instead consider NAL unit types of the subsequent VCL NAL units while determining the PSI value for such PDU Sets.

A leading picture is a picture that follows a particular IRAP picture in decoding order and precedes it in output order. There are two types of leading pictures in H.265/HEVC: Random access decodable leading (RADL) pictures and Random access skipped leading (RASL) pictures. A RADL picture is a leading picture that is guaranteed to be decodable when random access is performed at the associated IRAP picture. Therefore, RADL pictures are only allowed to reference the associated IRAP picture and other RADL pictures of the same IRAP picture. A RASL picture is a leading picture that may not be decodable when random access is performed from the associated IRAP picture. Only other RASL pictures are allowed to be dependent on a RASL picture. Hence, in H.265/HEVC bitstreams, RASL pictures can be discarded during random access. H.265/HEVC provides mechanisms to enable specifying the conformance of a bitstream wherein the originally present RASL pictures have been discarded. Consequently, system components can discard RASL pictures, when needed, without worrying about causing the bitstream to become non-compliant.

PDU Sets with `Type` field value equal to 6 or 7 (refer to Table 7.1 in H.265/HEVC specification [3]) in the NAL Unit header of RTP packet are RADL pictures. PDU Sets with `Type` field value equal to 8 or 9 (refer to Table 7.1 in H.265/HEVC specification [3]) in the NAL Unit header of RTP packet are RASL pictures. PDU Sets that contain RADL pictures should be assigned lower importance (higher PSI value) relative to the IRAP pictures and higher importance (lower PSI value) relative to the RASL pictures in the bitstream.

In video coding, temporal scalability is the option to decode only some of the frames in a video stream instead of the whole stream. This enables a media server to reduce the bitrate sent towards viewers that don't have enough bitrate or CPU to handle the whole stream. In H.265/HEVC, pictures with lowest temporal identifier value (TID) are used as reference pictures in the bitstream and are important for decoding the dependent frames. PDU Sets with TID value 1 (lowest possible value) should be set with higher importance (lower PSI value) relative to PDU Sets that have a higher TID value. The PSI value for such pictures should be lower for IRAP pictures and slightly higher for non-IRAP pictures compared to the pictures with higher TID values. Pictures with highest TID value cannot be used as reference pictures and can be discarded at the network level when the throughput is not good, or network conditions are unstable. PDU Sets with higher TID values should be set with lower importance (higher PSI value) compared with the PDU Sets with lower TID values.

In H.265/HEVC, each leading picture and trailing picture type has two type values. The even picture type numbers indicate sub-layer non-reference pictures and odd picture type numbers indicate sub-layer reference pictures. An encoder can use the sub-layer non-reference picture types for pictures that are not used for reference for prediction of any picture in the same temporal sub-layer. Note that a sub-layer non-reference picture may still be used as a reference picture for prediction of a picture in a higher temporal sub-layer.

PDU Sets that contain sub-layer reference picture types should be assigned a lower PSI value compared with the PDU sets with the corresponding sub-layer non-reference picture types.

4.2.6.2.4 PSI based on affected PDU Sets

When the transport layer is forced to perform immediate dropping/discarding of a PDU Set but has a freedom of selection among the PDU Sets, the PDU Set with smaller degrees of artifact would be the better choice in most cases. Dropping of a PDU Set may corrupt the decoded output of itself and the other PDU Sets though they may already be transmitted perfectly to the receiving end or yet in a queue waiting to be transmitted. The degrees of artifact can be explicitly transferred as the number of affected frames which precedes/follows the PDU Set, or can be implicitly transferred as the importance value where the lower value means the higher PDU Sets are affected while higher values proportionally mean less number of PDU Sets are affected, for example. By considering such a quantization of various affected PDU Sets can be translated into importance field, using 4 bits to represent 16 possible size ranges is recommended.

The information on the size of propagation error which caused by the dropping of each PDU Set may be provided by the application layer. The information may present the size of error propagation implicitly with a proportional mapping of error propagation size to an index such as the importance of the PDU Set in the media stream.

The importance value of a PDU Set in PDU Set Information (PSI) RTP HE is set as follows:

- The error propagation size is mapped to importance field value. The higher the error propagation size of a PDU Set, that PDU Set is more important, and it shall be assigned the lower PSI value. PDU Sets with low error propagation are of less importance and the PSI value for such PDU Sets shall be higher compared to PDU sets with higher error propagation size.

4.2.6.2.5 PSI mapping based on PDU Set dependencies

RTP senders should consider that multiplexed RTP streams are treated as a single Multimedia Session and set the PSI field accordingly, i.e., the PSI field for one bitstream that depends on other RTP stream(s) in the same Multimedia Session may need to be set taking the PSI field for PDU Sets in other multiplexed RTP streams into account. In some cases, dependencies can exist across bitstreams even when they are not multiplexed, particularly for XR services.

In case of such dependencies, it may not be sufficient to set the PSI values based on codecs and media types alone. PSI values shall be set in this case based on the following, which are listed in an increasing order of importance, i.e., decreasing order of PSI values.

- The PDU Set is considered not necessary for the processing of any other PDU Set. Such PDU Sets should be assigned the highest PSI values 14-15. When multiplexing, if a PDU Set is assigned PSI value of 15, similar PDU Sets of all streams should be assigned the PSI value 15 to prevent unfair treatment. If interdependency is known, e.g., in stereo streams (left eye is more important than right eye), then the more important stream can be assigned the PSI value 14.
- In H.264/AVC, these include the PDU Sets with an NRI value equal to b00 in the NAL unit header.
- In H.265/HEVC, the NAL unit header does not contain a field like NRI that indicates the relative transport priority. Hence, it is up to the application to identify such PDU Sets.
- The PDU Set is necessary for the processing of some PDU Sets of the stream to which it belongs. Such PDU Sets should be assigned a PSI value in the range 9-13 (inclusive). The lower end of the range should be used for IDR/IRAP pictures since they are more important for decoding of the bitstream.
- In H.264/AVC, these include:
 - IDR pictures with nal_unit_type equal to 5
 - Non-IDR pictures with nal_unit_type in the range 1 to 4 (inclusive)
- In H.265/HEVC, these include:
 - IRAP pictures with nal_unit_type field assigned a value in the range 16 to 23 (inclusive)
 - RADL or RASL pictures with nal_unit_type in the range 6 to 9 (inclusive)
- The PDU Set is necessary for the processing of all the other PDU Sets of the stream to which it belongs. Such PDU Sets should be assigned a PSI value in the range 6-8 (inclusive).
- In H.264/AVC, these include:

- SPS, PPS, i.e., NAL units with the `nal_unit_type` field equal to 7, 8, 13 or 15
- In H.265/HEVC, these include:
 - SPS, PPS, VPS, i.e., NAL units with the `nal_unit_type` field in the range 32 to 34 (inclusive)
- The PDU Set is necessary for the processing of some PDU Sets of the stream to which it belongs and also necessary for the processing of some PDU Sets of some other streams to which it does not belong. Such PDU Sets should be assigned a PSI value in the range 4-5 (inclusive).

NOTE 1: Values in this and lower range shall be used for assigning PSI values to PDU Sets in multiplexed streams or if dependencies exist across non-multiplexed bitstreams. Use cases for those cases are FFS. In case only a single RTP stream is present, the ranges provided by the previous bullet points shall be used.

NOTE 2: Considerations for multiplexed audio streams are FFS.

- The PDU Set is necessary for the processing of all PDU Sets of the stream to which it belongs and also of some other streams to which it does not belong. Such PDU Sets should be assigned a PSI value in the range 2-3 (inclusive).
- The PDU Set is necessary for the processing of all PDU Sets of all streams. Such PDU Sets should be assigned the lowest PSI value 1.

4.2.6.3 PDU Set Size Field

The PDU Set Size field may be present in the RTP HE for PDU Set marking if appropriately enabled for an RTP sender as per clause 4.2.5. In case the PDU Set Size is enabled the application shall express the PDU Set Size in bytes as per the PSize semantics defined in clause 4.2.4.

The PDU Set Size value of a PDU Set should be determined by the RTP sender based on the RTP payload corresponding to the PDU Set, transmission path MTU Size, or alternatively, maximum RTP SDU size, and network IP transport configuration.

The RTP sender should follow the corresponding steps in determining the PDU Set Size.

1. The RTP sender should receive from a media encoder (e.g., a H.264/AVC encoder, a H.265/HEVC encoder) payload data corresponding to a PDU Set. It is recommended that all Non-VCL NAL units (e.g. SPS NAL unit) are handled together with the associated VCL NAL units within the same PDU Set. The size of the received payload data (*R*) should be determined in bytes.
2. The RTP sender should perform next RTP fragmentation and packetization of the payload data (*R*). The maximum size of an RTP packet SDU (*S*) should be determined given a transmission path MTU size, or alternatively, a preconfigured maximum RTP SDU payload size less than the path MTU size. The RTP sender should determine the number of RTP packets (*P*) post-fragmentation given *S* and a packetization configuration of the RTP payload. The RTP payload should implement the payload formatting according to the corresponding payload type of the PDU Set (e.g., RFC 6184 [5] for H.264/AVC, RFC 7798 [6] for H.265/HEVC) and the packetization configuration to yield the *P* RTP packets' SDUs. *P* corresponds to the number of PDUs of the PDU Set.

NOTE 1: Some WebRTC implementations in commercial user agents configure a maximum RTP SDU size of 1200 bytes compliant also with the recommendations of RFC 8200 and further corresponding to an MTU Size of 1280 bytes. Other valid configurations exploiting larger MTU Size based on path MTU discovery protocols, RFC 1191, or RFC 8201, may apply up to the RTP stack implementation capabilities.

NOTE 2: It is generally assumed that the configuration of the RTP payload ensures RTP packets resulting from packetization do not violate the MTU Size. In addition, the RTP payload may be configured by applications to favor low-latency delivery. For example, in some cases of RTP H.264/AVC payload types, the RTP payload may be configured to operate in packetization-mode 1 (i.e., "non-interleaved mode" as per RFC 6184 clause 6.3) to allow for RTP packets to contain NAL units in decoding order and to map an RTP packet to a single NAL unit packet (as per RFC 6184, clause 5.6), a STAP-A packet (as per RFC 6184, clause 5.7.1) or a FU-A packet (as per RFC 6184, clause 5.8). In other cases, applications may select other RTP payload configuration up to implementation and application requirements.

3. The RTP sender should determine for each one of the P RTP packets the size of the RTP header overhead including any RTP HE overhead (Rh_p) as configured based on the SDP offer-answer negotiation.

NOTE 3: It may be possible for different PDUs in a PDU Set to contain distinct RTP HE besides the common RTP HE for PDU Set marking such that Rh_p may differ among different PDUs of a PDU Set.

4. The RTP sender should further determine per RTP packet the size of the UDP/IP headers overhead associated with an OS UDP socket sending out the RTP packets. This may be done by the RTP sender using UDP socket options available programmatically over OS network stack API calls or based on SDP-configured IP endpoints and corresponding transmission IP addresses. The RTP sender should determine the type of the underlying IP version used for transport, i.e., IPv4 or IPv6, and determine accordingly the IP header overhead (Ih_p) for each encapsulated RTP packet. If IPv4 options are configured for the UDP socket, or alternatively, if IPv6 header extensions are sent over the UDP socket, the RTP sender should consider the additional incurred size these have to the IP header overhead (Ih_p) of each RTP packet. The RTP sender should consider a fixed size UDP header overhead (Uh) of 8 bytes for each RTP packet.

NOTE 4: In case no IPv4 header options are used, the RTP sender should consider Ih_p corresponding to 20 bytes per RTP packet for IPv4. Whereas, in case no IPv6 extension headers are used, the RTP sender should consider Ih_p corresponding to 40 bytes per RTP packet for IPv6.

NOTE 5: For example, in case of Linux-based open-source OSs, any additional IPv4 options up to 40 bytes may be set and accessed programmatically based on socket API calls, the RTP sender implementation is expected to determine additional optional overheads to the IP header overhead, Ih_p .

5. The RTP sender should determine the PDU Set Size as the sum in bytes of all RTP/UDP/IP headers overhead of each one of the P packets and the received RTP payload corresponding to the PDUs of the PDU Set, e.g., $PSSize = R + \sum_{p=1}^P (Ih_p + Uh_p + Rh_p)$. The value should be indicated in the PSSize field of the RTP HE for PDU Set marking for all PDUs of the PDU Set before the corresponding RTP PDUs are sent over the UDP socket.

In case any of the above steps fails to determine for a PDU Set any of the Ih_p , Uh_p , Rh_p , P , or R , the RTP sender should set the PSSize to 0 for the PDU Set.

NOTE 6: The PDU Set Size guidelines above are generally applicable to video and audio media payload types.

4.2.7 Guidelines for AS

This clause describes guidelines for an AS that is on the media path between two or more UEs, e.g., an MRF, MCU etc. Such an AS may receive media over RTP with RTP HE for PDU Set marking added by the sender UE.

NOTE: These detailed guidelines are FFS.

4.3 RTP Header Extension for XR Pose

4.3.1 General

An RTP sender that uses RTP to deliver rendered video streams to an RTP receiver should include an RTP HE for XR pose to indicate the XR pose used for rendering the media (rendered pose). The RTP HE for XR pose may be used for signaling either a 6DoF XR pose or a 3DoF XR pose. The RTP HE for XR pose may also be used with audio streams.

The RTP HE for XR pose may also be used by an RTP sender to indicate the XR pose to be rendered to an RTP receiver.

The IANA registration information for the RTP HE for XR pose is provided in Annex D.3.

4.3.2 SDP Signaling

An RTP client that supports the RTP HE for XR pose shall negotiate the use of the extension using SDP. The signaling of the RTP HE for XR pose shall follow the SDP signaling design, the syntax, and semantics of the "extmap" attribute as outlined in RFC8285 [11].

For IANA registration, the "reference" field in the registry is 3GPP TS 26.522.

The ABNF syntax for this RTP HE extends the "extmap" attribute as follows:

```
extensionname = "urn:3gpp:xr-pose"

extensionattributes = xr-pose-dof [SP xr-pose-media]

xr-pose-dof = "3DOF" / "6DOF"

xr-pose-media = "media=" xr-pose-media-id *(", " xr-pose-media-id)

xr-pose-media-id = token
```

The extension attribute "3DOF" indicates that the sender uses the RTP HE to signal a 3DoF XR pose, i.e., an XR pose that does not include the position fields x, y, z.

An RTP client that supports the RTP HE for XR pose and receives an SDP offer with "a=extmap" attribute with the URN "urn:3gpp:xr-pose" and the extension attribute "3DOF", shall include the extension attribute "3DOF" in the SDP answer, if the SDP offer is accepted.

The extension attribute "6DOF" indicates that the sender uses the RTP HE to signal a 6DoF XR pose, i.e., an XR pose that includes both the position fields x, y, z and the orientation fields rx, ry, rz, rw.

An RTP client that supports the RTP HE for XR pose and receives an SDP offer with "a=extmap" attribute with the URN "urn:3gpp:xr-pose" and the extension attribute "6DOF", shall include the extension attribute "6DOF" in the SDP answer, if the SDP offer is accepted.

The extension attribute "media" is followed by a list of tokens for "mid" (as defined in RFC 5888 [19]) for media streams that can reuse the pose included in the RTP HE. Further details on reuse are provided later in the section.


An RTP endpoint that supports the RTP HE for XR pose and receives an SDP offer with "a=extmap" attribute with the URN "urn:3gpp:xr-pose" shall remove the attribute from the answer for any media that will not use the extension, and retain it for any media that will use it.


4.3.3 Header Extension Format

If the RTP HE for XR pose is for rendered pose, the RTP sender should use the RTP HE for XR pose to associate the selected pose with the rendered frame. The RTP sender delivers the rendered frames using one or more video streams, depending on the view and projection configuration that is selected by the UE.


If negotiated successfully, an RTP sender should add the RTP HE for XR pose to the RTP stream. The frequency of RTP HE for XR pose shall be at least once in a frame. It may be sent more often but not necessarily in every RTP packet.

The 2-byte (RFC 8285) RTP HE format shall be used for signalling the RTP HE. Format of the HE for 6DoF XR pose is shown below.





If the RTP HE for XR pose is used for signaling a 3DoF XR pose, the fields x, y, z shall be omitted. Format of the HE for 3DoF XR pose is shown below.



The fields rx, ry, rz, rw, x, y, z are defined in single-precision floating-point format (binary32 as per ISO/IEC 60559 [20]).

rx (32 bits): x coordinate of the orientation quaternion of the XR pose.

ry (32 bits): y coordinate of the orientation quaternion of the XR pose.

rz (32 bits): z coordinate of the orientation quaternion of the XR pose.

rw (32 bits): w coordinate of the orientation quaternion of the XR pose.

x (32 bits): x coordinate of the position of the XR pose in meters.

Y (32 bits): y coordinate of the position of the XR pose in meters.

z (32 bits): z coordinate of the position of the XR pose in meters.

XR timestamp (64 bits): Timestamp for the XR pose. If the RTP HE is used for rendered pose, this timestamp indicates the display time predicted by the XR runtime for the rendered image. Otherwise, this timestamp indicates the associated XR runtime display time for the predicted XR pose. XR timestamp uses the XR system clock and is represented in nanoseconds. The timestamp is passed to the XR runtime together with the rendered swapchain images (e.g. as part of the `xrEndFrame` call in OpenXR). A receiver may use the XR timestamp together with the RTP timestamp to determine the playout time of the media. XR timestamp shall not be used for media synchronization.

NOTE 1: It is left to the discretion of the receiver application how to use the XR timestamp. It is not specified how the receiver application determines the playout time using the XR timestamp together with the RTP timestamp. The receiver application may take both the media transport aspects and XR application aspects (e.g., reducing motion judder) into account while determining the playout time.

action_id (32 bits): A list of actions corresponding to the pose x, y, z, rx, ry, rz, rw coordinates. An action_id uniquely identifies an action and it may be an action identifier as defined in the action format of TS 26.119 [17] clause 6.2.3. The number of action identifiers in one RTP HE for XR pose shall be no more than 10. Hence, the size of the RTP HE is $36+2*n$, if a 6DoF XR pose is used, or $24+2*n$, if a 3DoF XR pose is used, where n is the number of action identifiers in the HE.

If the RTP HE for XR pose is for rendered pose, the RTP sender should contain an action_id field as defined above, with the list of action identifiers identifying the processed actions for the rendering of the frame.

If the RTP HE for XR pose is for pose to be rendered, the RTP sender should contain an `action_id` field as defined above, with the list of action identifiers identifying the action for which the pose coordinates apply.

NOTE 2: An XR server should be aware of an XR client's actions configuration in an action space. Signalling aspects for the actions configuration are defined in other specifications such as TS 26.119 and TS 26.565.

NOTE 3: An XR server should be aware of the XR space used by the XR client for the XR pose fields defined above. Signalling aspects for this XR space are defined in other specifications such as TS 26.119 and TS 26.565.

NOTE 4: When a receiver receives an RTP HE for XR pose, it is up to the application to determine whether the HE is for rendered pose or for pose to be rendered.

When both video and audio are delivered to an RTP receiver, or when either audio or video is delivered using multiple real-time streams (e.g., left eye + right eye), multiple RTP streams may be associated with the same RTP HE data, e.g., the same pose may have been used for generating multiple streams. This may lead to sending the same RTP HE data multiple times in different streams.

A sender may reuse the XR pose RTP HE of one stream for multiple RTP streams. For example, only the video stream carries the pose RTP HE, but the pose is applicable also for the audio bitstream. In this case, the sender shall include the extension attribute `media` followed by a space separated list of media ID (MID) values in the `"a=extmap"` attribute. The MID values indicate all media streams for which the pose RTP HE is applicable to. If the extension attribute `media` is present, then the media description of all bitstreams that reuse the RTP HE shall include the attribute `"mid"`, as defined in RFC 5888 [19].

NOTE 4: In case there is a mismatch between the frame rates of the streams, the receiver may use the few most recent samples from the source RTP stream to obtain a synchronized sample in the dependent stream via interpolation. Alternatively, the receiver may choose to not perform any interpolation and simply use the last available sample from the source RTP stream for the dependent stream. It is left to the discretion of the receiver application to select an appropriate synchronization method.

4.4 RTP Header Extension for in-band end-to-end delay measurement

4.4.1 General

An RTP HE that allows an RTP packet to carry timestamp(s) may help obtain measured delays that are representative of the end-to-end instantaneous delays experienced by the media in the user plane.

NOTE 1: End-to-end connections may imply in some cases a multi-hop link including non-3GPP network paths, such as a data network link and a tethering link.

Figure 4.4.1-1 shows how the RTP HEs are used to measure the delays, where T_1 , T_2 , T_3 and T_4 are the Originate Timestamp, the Receive Timestamp, the Transmit Timestamp, and the Destination Timestamp, respectively. The one-way delay from the Requester to the Responder is calculated as $T_2 - T_1$, the one-way delay in the opposite direction is calculated as $T_4 - T_3$, the RTT is calculated as $(T_4 - T_1) - (T_3 - T_2)$, and the processing delay on the Responder is calculated as $T_3 - T_2$.

NOTE 2: Time synchronization between the Requester and the Responder for example via PTP [14] is a prerequisite for computation of one-way delays in any direction.

NOTE 3: The Requester may use T_1 to group T_1 , T_2 , T_3 , T_4 measurements and index them to compute all the above delay measurements and any corresponding statistics. Specific means to achieve this are left to RTP implementers.

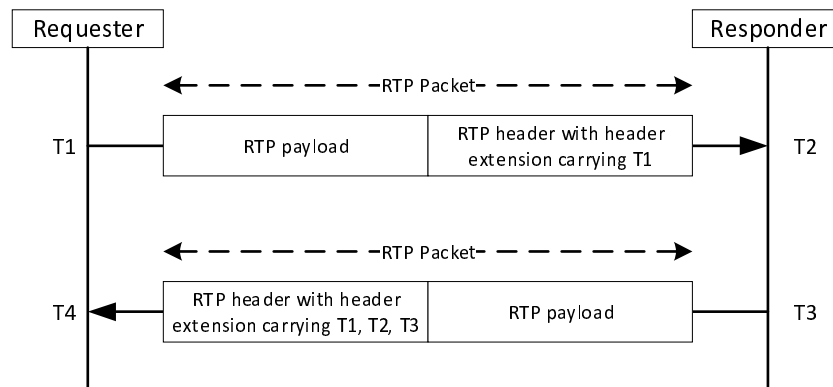



Figure 4.4.1-1: The RTP HEs for in-band end-to-end delay measurement.

The RTP HEs defined below follow RFC 8285 [11].


The IANA registration information for the RTP HE for in-band end-to-end delay measurement is provided in Annex D.4.

4.4.2 One-byte RTP Header Extension Format

The RTP HE element for the RTP packet that carries only one timestamp T1 is shown below. This is the same as the "RTP Header Extension for Absolute Sender Time" in Annex C.




The RTP HE element for the RTP packet that carries three timestamps T1, T2 and T3 is shown below.




4.4.3 Two-byte RTP Header Extension Format

The RTP HE element for the RTP packet that carries one timestamp T1 is shown below.



+--+--+--+--+--+--+--+--+

The RTP HE element for the RTP packet that carries three timestamps T1, T2 and T3 is shown below.



4.4.4 Syntax

T1: consists of 24 bits, taken from the 6 LSB bits of the integer part and the 18 MSB bits of the fractional part of the NTP timestamp format defined in RFC 5905 [13].

T2: follows the syntax of T1.

T3: follows the syntax of T1.

NOTE: The timestamps are expressed in seconds according to the above syntax, with a 64 second wraparound and a 3.8 microsecond resolution.

4.4.5 Semantics

T1: Originate Timestamp. It represents the time when the Requester transmits the RTP packet toward the Responder.

T2: Receive Timestamp. It represents the time when the Responder receives the RTP packet that carries the Originate Timestamp T1.

T3: Transmit Timestamp. It represents the time when the Responder transmits the RTP packet that carries the Originate Timestamp T1, the Receive Timestamp T2, and the Transmit Timestamp T3.

4.4.6 SDP signaling

The signaling of the delay measurement RTP HE shall follow the SDP signaling design and the syntax and semantics of the "extmap" attribute as outlined in RFC 8285 [11].

For the RTP HE carrying only T1, the ABNF syntax for the "extmap" attribute is as follows:

extensionname = "http://www.webrtc.org/experiments/rtp-hdext/abs-send-time"

extensionattributes = ["short"/"long"]

If the *extensionattributes* is absent, the RTP HE follows the one-byte format, i.e., the "short" format. If *extensionattributes* is "short", the RTP HE follows the one-byte format. If *extensionattributes* is "long", the RTP HE follows the two-byte format.

NOTE 1: <http://www.webrtc.org/experiments/rtp-hdext/abs-send-time> is the extension URI of the RTP HE, and is currently implemented in WebRTC. This extension URI, instead of URN-based ones, allows for support from WebRTC without any change to the WebRTC implementation.

NOTE 2: This allows to reuse the "Absolute Sender Time" RTP HE in WebRTC without changes to the SDP syntax implemented in WebRTC.

NOTE 3: Regardless if this extension attribute is present or not, the use of long or short format is determined as described by section 4.1.2 of RFC 8285, i.e., based on what format other RTP HEs use in the same RTP session, unless both endpoints announced support for handling mixed format with "a=extmap-allow-mixed" as described by section 6 of RFC 8285.

Below is an example (Example 1):

```
a=extmap:4 http://www.webrtc.org/experiments/rtp-hdext/abs-send-time
```

For the RTP HE carrying T1, T2 and T3, the ABNF syntax for the "extmap" attribute is as follows:

```
extensionname = "urn:3gpp:delay-measurement-response:rel-18"
extensionattributes = [format SP] binding-info
format = "short"/"long"
binding-info = dependent-extmap-ID [";"m-line-label] [";"processing-ID]
dependent-extmap-ID = "dependent-extmap-ID="1*5DIGIT
m-line-label = "dependent-rtp-he-m-line-label="token
processing-ID = "processing-ID="token
; token as defined by RFC 4566
```

The extension attributes have the following semantics:

- dependent-extmap-ID: identifies the Requester sent RTP HE (i.e., carrying T1) on which the Responder sent RTP HE (i.e., carrying T1, T2 and T3) depends. Timestamps T1 and T2 included in the Responder sent RTP HE are the time the Requester sent RTP HE is transmitted at and the time the Requester sent RTP HE is received at the Responder, respectively.
- processing-ID: identifies a processing module on the Responder which takes data carried in RTP packets with the RTP HE identified by dependent-extmap-ID, processes them and produces data that are then carried in RTP packets with this RTP HE.

NOTE 3: The details of processing-ID are left to implementation at the application level.

- m-line-label: is the SDP "label" attribute defined in RFC 4574 [15], and it identifies a media stream from the Requester to the Responder and associates the RTP HE in that media stream to this RTP HE.

NOTE 4: There may be multiple media streams that carry RTP packets whose RTP HEs may be used for the binding.

Below is an example (Example 2):

```
a=extmap:5 urn:3gpp:delay-measurement-response:rel-18 short dependent-extmap-ID=4;dependent-rtp-he-m-line-label=2;processing-ID=7
```

In the example,

- 5 is the RTP HE ID
- 4 is the value of the attribute dependent-extmap-ID, which is the RTP HE ID of the RTP HE in Example 1. This establishes a binding between the two RTP HEs.
- 7 is the processing-ID.
- 2 is the SDP "label" attribute that identifies the media stream corresponding to "a=label:2" in the SDP signaling, and the RTP packets from the media stream are used for the binding.

5 RTCP Feedback Reporting Procedures

5.1 General

This clause defines the RTCP feedback reporting messages to transmit control information. There are a number of possible ways to carry a variety of control information using RTCP packets. This includes:

- profile-specific extensions to the sender (PT=200) and receiver report (PT=201),
- application-defined RTCP packet with payload type equal to 204 (PT=204),
- Generic RTP Feedback reports with payload type equal to 205 (RTPFB; PT=205), and
- payload-specific RTCP feedback messages with payload type equal to 206 (PSFB; PT= 206),
- extended reports (RTCP XR) with payload type equal to 207 (PT=207).

5.2 Transmission of timing information data for QoE measurements

5.2.1 General

In use cases for shared interactive immersive services, the user interaction information is sent from a UE to a server. The server handles the user's request to the immersive media scene (e.g., changing the context such as translation, rotation, and scaling or adding a new object in the scene). In the case of the edge-assisted UE type, the UE offloads the scene rendering.

In the context of interactive immersive services, one important parameter to estimate the user quality of experience is the *roundtrip interaction delay*. The *roundtrip interaction delay* is defined as the sum of the *age of content* and the *user interaction delay*.

The *age of content* is defined as the time duration between the moment the content is created and the time it is presented to the user. It is impacted by the downlink latency of the wireless network.

The *user interaction delay* is defined as the time duration between the moment at which a user action is initiated and the time such an action is taken into account by the content creation engine. It is impacted by the uplink latency of the wireless network.

The *estimated-at-time* (T1) and *start-to-render-at-time* (T3) provide the times when the pose was estimated and when the SRS started to render the rendered frame, respectively. The *split-renderer-output-time* (T5) provides the time when the output of the SRS for a rendered frame is available. This T5 information can be used to measure the server processing delay and the overall application delay excluding the server processing delay. The SRS processes the interaction according to the actions in the action message from the UE and updates the scene. The Scene Manager records the *sceneUpdateTime* (T6) timestamp when it starts to process the actions. The *sceneUpdateTime* is used to measure the user interaction delay, age of content and the roundtrip interaction delay. The details of *sceneUpdateTime*, measurement of *User-interaction-delay*, *Age-of-content* and *Roundtrip-interaction-delay* QoE interaction metrics.

The *user interaction delay*, *age of content*, and *round-trip interaction delay* measurements are described as quality of experience metrics for XR content. These delay measurement metrics need to be calculated at the UE for providing better QoE to the user. Also, the server processing delay measurement helps the UE in the adaptation process with the split rendering server for achieving better QoE.

5.2.2 RTCP message-based transmission of timing information

5.2.2.1 General

The timing information data recorded at the SRS or at the RTP sender can be transmitted to the UE by enhancing the RTCP XR packets, which are specified in IETF RFC 3611 [16]. The RTCP XR report is identified by payload type (PT)

equal to 207, which refers to an extended report block message. For transmission of timing information data using RTCP XR messages, the block type (BT) defined in RFC 3611 can be extended with a value TBD.

NOTE: The block type value for the QoE timing information RTCP XR message will be assigned by IANA and the specification will be updated with that block type value later.

5.2.2.2 Extended Report block for QoE timing information

This extended report block type permits detailed reporting of timing information recorded at the SRS. These reports can be used, for example, for calculating the QoE metrics such as *round-trip delay*, *server processing delay*, *user interaction delay*, *age of content* and the *round-trip interaction delay* at the UE.

The timing information required for measuring QoE metrics may be expressed in the same units as in the RTP timestamps of RTP data packets. This is so that, for each packet, one can establish the relation between the media data flowing and the corresponding QoE timing information recorded at the SRS for a specific media frame.

For optimum use of the RTCP bandwidth, the RTCP XR block payload may contain the whole or part of the timing information required to calculate the QoE metrics. `t_info` field present in Figure 1 represents the timing information present in an RTCP XR block report. When a bit is set to 'ONE' in `t_info` field the respective timing information shall be present in the payload. When a bit is set to 'ZERO' in `t_info` field, the respective time information shall not be present in the payload. E.g., when the sender like to transmit only T1 and T3 information, the `t_info` field is set to b0011 and only T1 and T3 information is present in the message payload.

The identifiers of all actions that were processed for the rendering of a frame at a specific time are reported in the RTP HE for XR pose defined in clause 4.3. The synchronization between the various timing information present in the below XR report and the action identifiers present in the RTP HE for XR pose is performed using the RTP timestamp information present in the RTP header of the packet containing the RTP HE for XR pose and the RTP timestamp field present in the below RTCP XR report block.

The QoE timing information Report Block has the following format:

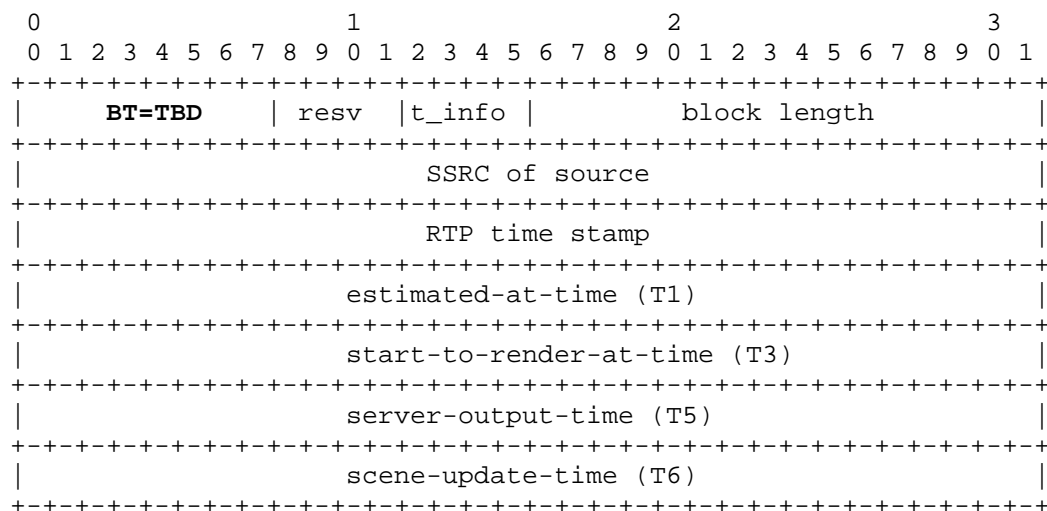


Figure 5.2.2.2-1: RTCP XR block format for QoE timing information data

The semantics of the fields in QoE time information Extended Report (RTCP XR) block are as follows:

- `block type (BT)` [8 bits]: A QoE time information Report Block is identified by a constant value.
- `block length` [16 bits]: The length of this report block, including the header, in 32-bit words minus one.
- `resv` [4 bits]: This field is reserved for future definition. In the absence of such definition, the bits in this field shall be set to zero by the sender and shall be ignored by the receiver.
- `t_info` [4 bits]: This field bits represent the timestamps that are present in the RTCP XR report block. When T1 is present in the RTCP XR report, first bit (least significant bit) is set to 1. When the LSB is set to 0, T1 information shall not be present. When T3, T5 and T6 are present in the RTCP XR block data, bits 2, 3 and 4 are

set to 1 respectively. When T1, T3, T5 and T6 are present in an RTCP XR block data, the `t_info` field value shall be `b1111`. The timing information when present shall follow the order T1, T3, T5 followed by T6. For example, when the `t_info` field value is `b0101`, the RTCP XR block carries the T1 information followed by T5. T3 and T6 timing information will not be present in that RTCP XR block content.

- The transmission frequency of T1, T3, T5 and T6 time information in RTCP XR report block can be negotiated during the configuration phase.
- `SSRC of source` [32 bits]: The SSRC of the RTP data packet source being reported upon by this report block.
- `RTP timestamp` [32 bits]: This field represents the RTP timestamp of the media frame at which the corresponding QoE timing information date was recorded at the SRS. This correspondence may be used for synchronization between the media data and the QoE timing information measurements recorded at the SRS for a specific media frame.
- `estimated-at-time (T1)` [32 bits]: This field represents the time when the pose estimation was made. This time information is expressed in the same units and with the same random offset as the RTP timestamps in data packets.
- `start-to-render-at-time (T3)` [32 bits]: This field represents the time when the renderer in the split rendering server started to render the associated media frame. This time information is expressed in the same units and with the same random offset as the RTP timestamps in data packets.
- `server-output-time (T5)` [32 bits]: This field represents the recorded time at the output of the split rendering server. This time information is expressed in the same units and with the same random offset as the RTP timestamps in data packets.
- `scene-update-time (T6)` [32 bits]: This field represents the time when the Scene manager processes the interaction task according to the actions in the action message from the UE and updates the scene. This time information is expressed in the same units and with the same random offset as the RTP timestamps in data packets.

5.2.3 SDP signaling and attributes

RFC 3611 [16] defines the SDP attribute "`rtcp-xr`" that can be used to signal to participants in a media session that they should use the specified RTCP XR blocks. This attribute is extendable with new parameters to cover any new type of XR report blocks.

The extended RTCP XR blocks with QoE time information SDP attribute extends the "`a=rtcp-xr`" attribute definition in RFC 3611 [16] as described below in Augmented Backus-Naur Form (ABNF).

```
xr-format =/ qoe-timing-info
```

```
qoe-timing-info = "qoe-timing-info" [ "=" max-size ] ; max-size from RFC 3611
```

The new parameter name and its corresponding RTCP XR format is:

Parameter name	XR block (block type and name)
-----	-----
qoe-timing-info	TBD Timing Information for QoE Metrics Calculation Block

The "`qoe-timing-info`", parameter may specify an integer value. This value indicates the largest size the whole report block should have in octets.

Annex A (informative): Guidelines for PDU Set identification

A.0 General

This informative annex provides guidelines for network functions like the UPF, which needs to determine PDU Set information, as described in TS 23.501 [12], clause 5.37.5. The network function is typically provisioned with at least the Service Data Flow Filter to identify the Service Data Flow, and optionally additional information about the presence of RTP HEs according to IETF RFC 8285 [11], the used RTP Payload Type, the used RTP Payload Format and other information.

When the RTP sender multiplexes RTP data and control packets onto the same Service Data Flow using a single port, the RTP Sender should implement the Payload Type separation according to IETF RFC 5761, clause 4 [10] and the network function should separate RTP data from RTCP data accordingly.

To avoid IP fragmentation, the RTP sender should select a sufficiently small RTP payload.

A.1 Leveraging RTP Header Extensions

When the PDU Set related RTP HEs are available within the RTP headers, the network function only needs to parse the RTP header and the RTP HEs. The RTP HE for PDU Set Marking is defined in clause 4.2.

An intermediate network function determines based on the RTP header X bit being set to 1, whether the optional HE fields are present in the RTP packet, after the SSRC and the (optional) CSRC fields in the RTP header. All information for the PDU Set identification is present within the RTP HE and the network function does not need to know the RTP Payload format. The RTP Payload may be encrypted (i.e. SRTP).

When multiple RTP HEs are present within the RTP header, the network function uses the RTP HE ID for finding the PDU Set related HE.

A.2 Obtaining PDU Set information from RTP Header or Payload

A.2.0 General

When the RTP HE for PDU Set marking is not available, some or all of PDU Set information can be derived from the RTP/SRTP header and/or payload, e.g., by a network function like the UPF. The following clauses describe how the PDU Set information can be derived from the RTP/SRTP header and/or payload.

A.2.1 RTP/SRTP header

When RFC 6184 [5] or RFC 7798 [6] are used as payload formats, a network function can obtain some of the PDU Set information from RTP headers by following these guidelines.

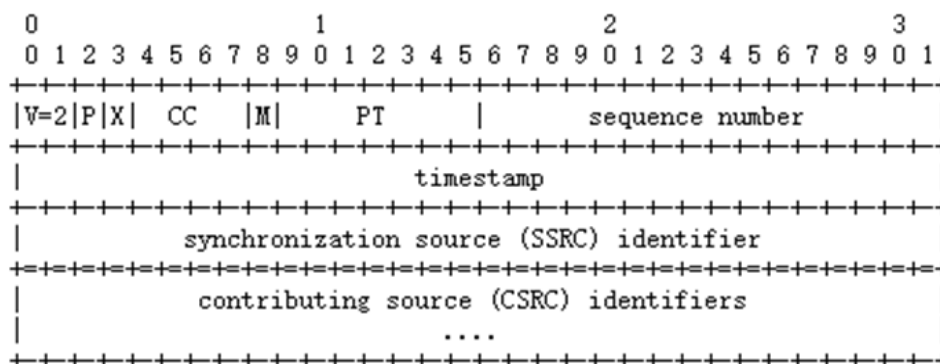


Figure A.2.1-1: RTP header fields as defined in RFC 3550 [4]

When the RTP/SRTP is used to convey the video content and when the PDU Set represents a video frame, the video frame may be identified based on the RTP header fields as following:

- The "marker (M)" bit is used with the video payload formats in clause A.2 to indicate the frame boundary, by setting the M bit on the last PDU of a frame. With the "M" bit and the sequence number in RTP header, the Indication of End PDU of a PDU Set and PDU SN within a PDU Set/frame can be derived. The network function should monitor the preceding packets to detect and compensate for potential packet reordering.
- The "timestamp" field indicates the sampling instant of the first octet in the RTP data packet and all RTP packets in the video frame is generally marked with the same timestamp. Therefore, with the "timestamp" field and the sequence number in RTP header, the Indication of End PDU of a PDU Set and PDU SN within a PDU Set/frame can be derived.

NOTE 1: When multiple RTP streams multiplexed over a single RTP session, the "M" bit, "timestamp" field, and sequence number information can be used together with the synchronization source (SSRC) in the RTP header to identify the boundary of video frame for each of the RTP streams that can be separated by their different SSRC values.

NOTE 2: For the timestamp-based solution, generally, the end PDU of the PDU Set can only be determined when a PDU with new RTP timestamp arrives, which may introduce additional latency.

- PDU Set Size can only be determined by a network function with reception of the last PDU belonging to the PDU Set, by summing up the individual PDU contributions to the PDU Set Size. PDU Set Importance cannot be derived using the RTP header fields.

A.2.2 RTP payload

A.2.2.1 General

When the RTP Payload is not encrypted, intermediate network functions may obtain additional information from the RTP payload.

The PDU Set information identification based on the RTP payload format is presented in this clause, including information on the RTP payload formats for H.264/AVC [5] and H.265/HEVC [6] codecs. The information about the used RTP Payload format for a service data flow is provided in advance to 5GC (e.g., UPF).

It is generally recommended that the network function considers non-VCL NAL units as part of the PDU Set of the associated VCL NALUs, e.g. identified by the same timestamp.

A.2.2.2 RTP payload for H.264/AVC codec

For a video content with H.264/AVC RTP payload, the PDU Set Information can be obtained by the following approach.

According to RFC 6184 [5], the first octet in the RTP payload indicates the content of the RTP payload, e.g. coded slice of an IDR frame, coded slice of a P frame, and also the possible structures of the RTP payload, e.g. single NAL unit packet, AP and fragmentation unit (FU). Depending on the indication of the first octet of the RTP payload, a second octet (the FU header) should also be processed.

- For single NAL unit packets and APs, it can be easily detected that each RTP packet can be treated as a complete PDU Set when the first Type field of Figure A.2.2-1 is less than 28.
- In case of APs, the network function may need to process all embedded NAL units.
- When the first Type field in Figure A.2.2-1 is 28 or 29, one NAL unit is carried over multiple RTP packets. In this case, the first byte of RTP payload is also named the FU indicator and the following byte is the FU header. The NAL unit type is contained in the Type field of the FU header (Figure A.2.2.-1). In the FU header, the "S" bit and "E" bit separately represents the start and end of the NAL unit. Therefore, based on the NAL unit type (also known as FU indicator for FU) and the FU header, the start/end of the PDU Set can be identified.

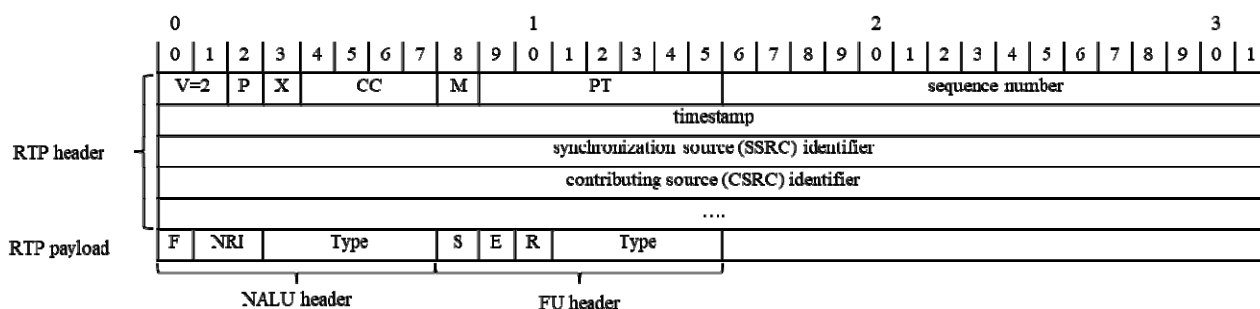


Figure A.2.2-1: RTP header [4] and NALU header format for H.264 [2]

With the RTP payload (i.e. NALU header and optionally FU header) and the sequence number in the RTP header, the indication of the End PDU of the PDU Set and the PDU SN within a PDU Set can be derived.

When using FUs (Type equals 28 or 29), the size of the NALU can only be determined after reception of the last packet of this FU. Thus, a network function can only determine the PDU Set Size with the reception of the last PDU of this FU.

As described in clause 4.2.6.2.2, the Type and NRI value in the NAL unit header indicates the relative transport priority and can be used to set the PSI. Besides, different NRI values can also indicate different requirements, which can be used to provide different protections against transmission losses, e.g. reliabilities (tolerable frame/slice error rate), and priorities.

A.2.2.3 RTP payload for H.265/HEVC codec

For a video content with H.265/HEVC RTP payload, the identification of the PDU Set can be realized by following approach.

According to RFC 7798 [6], within the RTP packet, the first two octets of the RTP payload indicate the content of RTP packet. Besides, it also indicates the possible structures of the RTP payload, e.g. single NAL unit packet, APs, FUs, and Payload Content Information (PACI) carrying RTP packet.

- For single NAL unit packets and APs, it can be easily detected that each RTP packet can be treated as a single PDU Set when the NAL unit type is less than 49.
- When NAL unit type is 49, one NAL unit is carried over multiple RTP packets. In this case, the first two-byte of RTP payload is also named the payload header (denoted as NAL unit, NALU, header) and the following byte is the FU header. In the FU header, the "S" bit and "E" bit separately represents the start and end of the NAL unit. The FuType field contains the actual NAL unit type. Therefore, based on the Type field of the first two octets (also known as FU indicator for FU) and the FU header, the start/end of the PDU Set can be identified.
- When NAL unit type is 50, this is a PACI packet which may carry a single NAL unit packet or FU. In this case, the first two-byte of RTP payload is also named as the PACI header (denoted as NAL unit header). In the following two bytes, the "A" bit is the copy of "F" bit and cType field is the copy of Type field in the PACI payload NAL unit. Then the following is the PHES field, whose length is determined by the PHSize. Finally, the

following is the PACI payload NAL unit, during which the first byte is FU header when cType (within the PACI payload header) is 49. Therefore, based on the PACI header and PACI payload NAL unit, the start/end of the PDU Set can be identified.




Figure A.2.3-1: The Structure of the HEVC NAL Unit Header [6]




Figure A.2.3-2: The Structure of FU Header

With the RTP payload (i.e. NAL unit header and optionally FU header) and the sequence number in the RTP header, the indication of the End PDU of the PDU Set and the PDU SN within a PDU Set can be derived.

As described in clause 4.2.6.2.3, the Type field and the TID field in the NAL unit header indicates the relative transport priority and can be used to be mapped to the PSI. They can also indicate different requirements, which can be used to provide different protections against transmission losses, e.g. reliabilities (tolerable frame/slice error rate), and priorities.

When using FUs (Type equals 49, or 50 where cType is 49), the size of the NAL unit can only be determined after reception of the last packet of this FU. Thus, a network function can only determine the PDU Set Size with the reception of the last PDU of this FU.

Annex B (informative): Examples of SDP offers and answers

B.1 SDP example for RTP header extension for XR pose

An example SDP description using the RTP HE for XR pose (defined in clause 4.3) is presented below. Using the extension attribute `media`, the RTP HE for XR pose with URI `urn:3gpp:xr-pose` provided in the video stream with MID `m1` is also applicable to another video stream with MID `m2`.

```
v=0
o=alice 2890844526 2890844526 IN IP4 host.atlanta.example.com
s=SDP Session
c=IN IP4 host.atlanta.example.com
t=0 0
m=application 1001 UDP/DTLS/SCTP webrtc-datachannel
a=sendonly
m=video 23458 RTP/AVP 96
a=mid:m1
a=recvonly
a=rtpmap:96 H264/90000
a=extmap:1 urn:3gpp:xr-pose media: m2
m=audio 23468 RTP/AVP 97
a=mid:m2
a=recvonly
a=rtpmap:97 PCMU/8000
m=video 23478 RTP/AVP 97
a=mid:m3
a=recvonly
a=rtpmap:96 H264/90000
```

Annex C (informative): RTP Header Extension for Absolute Sender Time

The information below is about the "RTP Header Extension for Absolute Sender Time" and it was retrieved from <https://webrtc.googlesource.com/src/+refs/heads/main/docs/native-code/rtp-hdext/abs-send-time> on January 31, 2024.

Absolute Send Time

The Absolute Send Time extension is used to stamp RTP packets with a timestamp showing the departure time from the system that put this packet on the wire (or as close to this as we can manage). Contact solenberg@google.com for more info.

Name: "Absolute Sender Time" ; "RTP Header Extension for Absolute Sender Time"

Formal name: <http://www.webrtc.org/experiments/rtp-hdext/abs-send-time>

SDP "a=name": "abs-send-time" ; this is also used in client/cloud signaling.

Not unlike [RTP with TFRC](#)

Wire format: 1-byte extension, 3 bytes of data. total 4 bytes extra per packet (plus shared 4 bytes for all extensions present: 2 byte magic word 0xBEDE, 2 byte # of extensions). Will in practice replace the "toffset" extension so we should see no long term increase in traffic as a result.

Encoding: Timestamp is in seconds, 24 bit 6.18 fixed point, yielding 64s wraparound and 3.8us resolution (one increment for each 477 bytes going out on a 1Gbps interface).

Relation to NTP timestamps: $\text{abs_send_time_24} = (\text{ntp_timestamp_64} \gg 14) \& 0x00ffff$; NTP timestamp is 32 bits for whole seconds, 32 bits fraction of second.

Notes: Packets are time stamped when going out, preferably close to metal. Intermediate RTP relays (entities possibly altering the stream) should remove the extension or set its own timestamp.

Annex D (informative): IANA registration information for RTP Header Extensions

D.1 Introduction

This Annex provides the RTP HE registration information that is referenced from the IANA registry at <http://www.iana.org/>.

D.2 urn:3gpp:pdu-set-marking:rel-18

The desired extension naming URI:

urn:3gpp:pdu-set-marking:rel-18

A formal reference to the publicly available specification:

TS 26.522

A short phrase describing the function of the extension:

PDU Set marking and signalling of end of data burst, see clause 4.2

Contact information for the organization or person making the registration:

3GPP Specifications Manager

3gppContact@etsi.org

+33 (0)492944200

D.3 urn:3gpp:xr-pose

The desired extension naming URI:

urn:3gpp:xr-pose

A formal reference to the publicly available specification:

TS 26.522

A short phrase describing the function of the extension:

Signalling of a 6DoF or 3DoF XR pose, see clause 4.3

Contact information for the organization or person making the registration:

3GPP Specifications Manager

3gppContact@etsi.org

+33 (0)492944200

D.4 urn:3gpp:delay-measurement-response:rel-18

The desired extension naming URI:

urn:3gpp:delay-measurement-response:rel-18

A formal reference to the publicly available specification:

TS 26.522

A short phrase describing the function of the extension:

In-band end-to-end delay measurement, see clause 4.4

Contact information for the organization or person making the registration:

3GPP Specifications Manager

3gppContact@etsi.org

+33 (0)492944200

Annex E (informative): IANA registration information for RTCP XR block types

E.1 Introduction

This Annex provides the RTCP XR block type and SDP parameter registration information that is referenced from the IANA registry at <http://www.iana.org/>.

E.2 Timing Information for QoE Metrics Calculation

This document assigns the block type value TBD in the IANA "RTP Control Protocol Extended Reports (RTCP XR) Block Type Registry" to the "Timing Information for QoE Metrics Calculation Block".

NOTE: The above "TBD" value will be replaced by the IANA-assigned value and this note can then be removed.

The desired, corresponding SDP parameter name in the "RTP Control Protocol Extended Reports (RTCP XR) Session Description Protocol (SDP) Parameters Registry":

qoe-timing-info

A formal reference to the publicly available specification:

TS 26.522

Contact information for the organization or person making the registration:

3GPP Specifications Manager

3gppContact@etsi.org

+33 (0)492944200

Annex F (informative): Change history

Change history							
Date	Meeting	TDoc	CR	Rev	Cat	Subject/Comment	New version
2023-04	SA4#123-e	S4-230719				Initial version, with text from WID in SP-220613 and S4-230713	0.0.1
2023-05	SA4#124	S4-231044				Implementing S4-230848, S4-230965, S4-231026, S4-231028	0.0.2
2023-05	SA4#124	S4-231101				Agreed version	0.1.0
2023-08	SA4#125	S4-231544				Implementing S4-231440, S4-231524, S4-231533	0.1.1
2023-11	SA4#126	S4-231752				Implementing S4aR230101, S4aR230106	0.1.2
2023-11	SA4#126	S4-231983				Implementing S4-231756, S4-231758, S4-231925, S4-231927, S4-231928, S4-231929, S4-231930, S4-232028	0.2.0
2024-01	SA4#127	S4-240332				Implementing S4-240053, S4-240228, S4-240233, S4-240237, S4-240324, S4-240325, S4-240362	0.3.0
2024-02	SA4#127	S4-240493				Implementing S4aR240010	0.4.0
2024-03	SA#103	SP-240038				Version 1.0.0 created by MCC	1.0.0
2024-03						Version 18.0.0 created by MCC	18.0.0
2024-06	SA#104	SP-240687	0001	1	F	ABNF Corrections	18.1.0
2024-06	SA#104	SP-240687	0002	3	F	On the RTP header extension for the XR pose	18.1.0
2024-06	SA#104	SP-240687	0004	1	F	Clarification on PDU Sets combining	18.1.0
2024-06	SA#104	SP-240687	0005	2	F	Corrections of E2E delay measurements signaling	18.1.0
2025-07	SA#108	SP-250696	0018	1	F	[5G_RTP] Correction of Guidelines for PDU Set identification	18.2.0

History

Document history		
V18.0.0	May 2024	Publication
V18.1.0	July 2024	Publication
V18.2.0	July 2025	Publication