
Stream: Internet Engineering Task Force (IETF)
RFC: [8849](#)
Category: Standards Track
Published: December 2020
ISSN: 2070-1721
Authors: R. Even J. Lennox
8x8/Jitsi

RFC 8849

Mapping RTP Streams to Controlling Multiple Streams for Telepresence (CLUE) Media Captures

Abstract

This document describes how the Real-time Transport Protocol (RTP) is used in the context of the Controlling Multiple Streams for Telepresence (CLUE) protocol. It also describes the mechanisms and recommended practice for mapping RTP media streams, as defined in the Session Description Protocol (SDP), to CLUE Media Captures and defines a new RTP header extension (CaptureID).

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/rfc8849>.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Terminology
3. RTP Topologies for CLUE
4. Mapping CLUE Capture Encodings to RTP Streams
5. MCC Constituent CaptureID Definition
 - 5.1. RTCP CaptureID SDES Item
 - 5.2. RTP Header Extension
6. Examples
7. Communication Security
8. IANA Considerations
9. Security Considerations
10. References
 - 10.1. Normative References
 - 10.2. Informative References

Acknowledgments

Authors' Addresses

1. Introduction

Telepresence systems can send and receive multiple media streams. The CLUE Framework [RFC8845] defines Media Captures (MCs) as a source of Media, from one or more Capture Devices. A Media Capture may also be constructed from other Media streams. A middlebox can express conceptual Media Captures that it constructs from Media streams it receives. A Multiple Content Capture (MCC) is a special Media Capture composed of multiple Media Captures.

SIP Offer/Answer [RFC3264] uses SDP [RFC4566] to describe the RTP media streams [RFC3550]. Each RTP stream has a unique Synchronization Source (SSRC) within its RTP session. The content of the RTP stream is created by an encoder in the endpoint. This may be an original content from a camera or a content created by an intermediary device like a Multipoint Control Unit (MCU).

This document makes recommendations for the CLUE architecture about how RTP and RTP Control Protocol (RTCP) streams should be encoded and transmitted and how their relation to CLUE Media Captures should be communicated. The proposed solution supports multiple RTP topologies [[RFC7667](#)].

With regards to the media (audio, video, and timed text), systems that support CLUE use RTP for the media, SDP for codec and media transport negotiation (CLUE individual encodings), and the CLUE protocol for Media Capture description and selection. In order to associate the media in the different protocols, there are three mappings that need to be specified:

1. CLUE individual encodings to SDP
2. RTP streams to SDP (this is not a CLUE-specific mapping)
3. RTP streams to MC to map the received RTP stream to the current MC in the MCC.

2. 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.

Definitions from the CLUE Framework (see [Section 3](#) of [[RFC8845](#)]) are used by this document as well.

3. RTP Topologies for CLUE

The typical RTP topologies used by CLUE telepresence systems specify different behaviors for RTP and RTCP distribution. A number of RTP topologies are described in [[RFC7667](#)]. For CLUE telepresence, the relevant topologies include Point-to-Point, as well as Media-Mixing Mixers, Media-Switching Mixers, and Selective Forwarding Middleboxes.

In the Point-to-Point topology, one peer communicates directly with a single peer over unicast. There can be one or more RTP sessions, each sent on a separate 5-tuple, that have a separate SSRC space, with each RTP session carrying multiple RTP streams identified by their SSRC. All SSRCs are recognized by the peers based on the information in the RTCP Source description (SDS) report that includes the Canonical Name (CNAME) and SSRC of the sent RTP streams. There are different Point-to-Point use cases as specified in the CLUE use case [[RFC7205](#)]. In some cases, a CLUE session that, at a high level, is Point-to-Point may nonetheless have an RTP stream that is best described by one of the mixer topologies. For example, a CLUE endpoint can produce composite or switched captures for use by a receiving system with fewer displays than the sender has cameras. The Media Capture may be described using an MCC.

For the media mixer topology [RFC7667], the peers communicate only with the mixer. The mixer provides mixed or composited media streams, using its own SSRC for the sent streams. If needed by the CLUE endpoint, the conference roster information including conference participants, endpoints, media, and media-id (SSRC) can be determined using the conference event package [RFC4575] element.

Media-Switching Mixers and Selective Forwarding Middleboxes behave as described in [RFC7667].

4. Mapping CLUE Capture Encodings to RTP Streams

The different topologies described in Section 3 create different SSRC distribution models and RTP stream multiplexing points.

Most video conferencing systems today can separate multiple RTP sources by placing them into RTP sessions using the SDP description; the video conferencing application can also have some knowledge about the purpose of each RTP session. For example, video conferencing applications that have a primary video source and a slides video source can send each media source in a separate RTP session with a content attribute [RFC4796], enabling different application behavior for each received RTP media source. Demultiplexing is straightforward because each Media Capture is sent as a single RTP stream, with each RTP stream being sent in a separate RTP session, on a distinct UDP 5-tuple. This will also be true for mapping the RTP streams to Capture Encodings, if each Capture Encoding uses a separate RTP session and the consumer can identify it based on the receiving RTP port. In this case, SDP only needs to label the RTP session with an identifier that can be used to identify the Media Capture in the CLUE description. The SDP label attribute serves as this identifier.

Each Capture Encoding **MUST** be sent as a separate RTP stream. CLUE endpoints **MUST** support sending each such RTP stream in a separate RTP session signaled by an SDP "m=" line. They **MAY** also support sending some or all of the RTP streams in a single RTP session, using the mechanism described in [RFC8843] to relate RTP streams to SDP "m=" lines.

MCCs bring another mapping issue, in that an MCC represents multiple Media Captures that can be sent as part of the MCC if configured by the consumer. When receiving an RTP stream that is mapped to the MCC, the consumer needs to know which original MC it is in order to get the MC parameters from the advertisement. If a consumer requested a MCC, the original MC does not have a Capture Encoding, so it cannot be associated with an "m=" line using a label as described in "CLUE Signaling" [RFC8848]. It is important, for example, to get correct scaling information for the original MC, which may be different for the various MCs that are contributing to the MCC.

5. MCC Constituent CaptureID Definition

For an MCC that can represent multiple switched MCs, there is a need to know which MC is represented in the current RTP stream at any given time. This requires a mapping from the SSRC of the RTP stream conveying a particular MCC to the constituent MC. In order to address this

5.2. RTP Header Extension

The CaptureID is also carried in an RTP header extension [RFC8285], using the mechanism defined in [RFC7941].

Support is negotiated within SDP using the URN "urn:ietf:params:rtp-hdext:sdes:CaptureID".

The CaptureID is sent in an RTP header extension because for switched captures, receivers need to know which original MC corresponds to the media being sent for an MCC, in order to correctly apply geometric adjustments to the received media.

As discussed in [RFC7941], there is no need to send the CaptId Header Extension with all RTP packets. Senders **MAY** choose to send it only when a new MC is sent. If such a mode is being used, the header extension **SHOULD** be sent in the first few RTP packets to reduce the risk of losing it due to packet loss. See [RFC7941] for further discussion.

6. Examples

In this partial advertisement, the media provider advertises a composed capture VC7 made of a big picture representing the current speaker (VC3) and two picture-in-picture boxes representing the previous speakers (the previous one -- VC5 -- and the oldest one -- VC6).

```
<ns2:mediaCapture
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:type="ns2:videoCaptureType" captureID="VC7"
  mediaType="video">
  <ns2:captureSceneIDREF>CS1</ns2:captureSceneIDREF>
  <ns2:nonSpatiallyDefinable>true</ns2:nonSpatiallyDefinable>
  <ns2:content>
    <ns2:captureIDREF>VC3</ns2:captureIDREF>
    <ns2:captureIDREF>VC5</ns2:captureIDREF>
    <ns2:captureIDREF>VC6</ns2:captureIDREF>
  </ns2:content>
  <ns2:maxCaptures>3</ns2:maxCaptures>
  <ns2:allowSubsetChoice>false</ns2:allowSubsetChoice>
  <ns2:description lang="en">big picture of the current
  speaker pips about previous speakers</ns2:description>
  <ns2:priority>1</ns2:priority>
  <ns2:lang>it</ns2:lang>
  <ns2:mobility>static</ns2:mobility>
  <ns2:view>individual</ns2:view>
</ns2:mediaCapture>
```

In this case, the media provider will send capture IDs VC3, VC5, or VC6 as an RTP header extension and RTCP SDES message for the RTP stream associated with the MC.

Note that this is part of the full advertisement message example from the CLUE data model example [RFC8846] and is not a valid XML document.

7. Communication Security

CLUE endpoints **MUST** support RTP/SAVPF profiles and the Secure Real-time Transport Protocol (SRTP) [RFC3711]. CLUE endpoints **MUST** support DTLS [RFC6347] and DTLS-SRTP [RFC5763] [RFC5764] for SRTP keying.

All media channels **SHOULD** be secure via SRTP and the RTP/SAVPF profile unless the RTP media and its associated RTCP are secure by other means (see [RFC7201] and [RFC7202]).

All CLUE implementations **MUST** support DTLS 1.2 with the TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 cipher suite and the P-256 curve [FIPS186]. The DTLS-SRTP protection profile SRTP_AES128_CM_HMAC_SHA1_80 **MUST** be supported for SRTP. Implementations **MUST** favor cipher suites that support Perfect Forward Secrecy (PFS) over non-PFS cipher suites and **SHOULD** favor Authenticated Encryption with Associated Data (AEAD) over non-AEAD cipher suites. Encrypted SRTP Header extensions [RFC6904] **MUST** be supported.

Implementations **SHOULD** implement DTLS 1.2 with the TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 cipher suite. Implementations **MUST** favor cipher suites that support Perfect Forward Secrecy (PFS) over non-PFS cipher suites and **SHOULD** favor Authenticated Encryption with Associated Data (AEAD) over non-AEAD cipher suites.

NULL Protection profiles **MUST NOT** be used for RTP or RTCP.

CLUE endpoints **MUST** generate short-term persistent RTCP CNAMEs, as specified in [RFC7022], and thus can't be used for long-term tracking of the users.

8. IANA Considerations

This document defines a new extension URI in the "RTP SDES Compact Header Extensions" subregistry of the "Real-Time Transport Protocol (RTP) Parameters" registry, according to the following data:

Extension URI: urn:ietf:params:rtp-hdext:sdes:CaptId

Description: CLUE CaptId

Contact: Roni Even <ron.even.tlv@gmail.com>

Reference: RFC 8849

The IANA has registered one new RTCP SDES items in the "RTCP SDES Item Types" registry, as follows:

Value	Abbrev	Name	Reference
-------	--------	------	-----------

Value	Abbrev	Name	Reference
14	CCID	CLUE CaptId	RFC 8849

Table 1

9. Security Considerations

The security considerations of the RTP specification, the RTP/SAVPF profile, and the various RTP/RTCP extensions and RTP payload formats that form the complete protocol suite described in this memo apply. It is believed that there are no new security considerations resulting from the combination of these various protocol extensions.

The "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)" document [RFC5124] provides the handling of fundamental issues by offering confidentiality, integrity, and partial source authentication. A mandatory-to-implement and use media security solution is created by combining this secured RTP profile and DTLS-SRTP keying [RFC5764] as defined in the communication security section of this memo (Section 7).

RTCP packets convey a CNAME identifier that is used to associate RTP packet streams that need to be synchronized across related RTP sessions. Inappropriate choice of CNAME values can be a privacy concern, since long-term persistent CNAME identifiers can be used to track users across multiple calls. The communication security section of this memo (Section 7) mandates the generation of short-term persistent RTCP CNAMEs, as specified in [RFC7022], so they can't be used for long-term tracking of the users.

Some potential denial-of-service attacks exist if the RTCP reporting interval is configured to an inappropriate value. This could be done by configuring the RTCP bandwidth fraction to an excessively large or small value using the SDP "b=RR:" or "b=RS:" lines [RFC3556], or some similar mechanism, or by choosing an excessively large or small value for the RTP/AVPF minimal receiver report interval (if using SDP, this is the "a=rtcp-fb:... trr-int" parameter) [RFC4585]. The risks are as follows:

1. The RTCP bandwidth could be configured to make the regular reporting interval so large that effective congestion control cannot be maintained, potentially leading to denial of service due to congestion caused by the media traffic;
2. The RTCP interval could be configured to a very small value, causing endpoints to generate high-rate RTCP traffic, which potentially leads to denial of service due to the non-congestion-controlled RTCP traffic; and
3. RTCP parameters could be configured differently for each endpoint, with some of the endpoints using a large reporting interval and some using a smaller interval, leading to denial of service due to premature participant timeouts, which are due to mismatched timeout periods that are based on the reporting interval (this is a particular concern if endpoints use a small but non-zero value for the RTP/AVPF minimal receiver report interval (trr-int) [RFC4585], as discussed in [RFC8108]).

Premature participant timeout can be avoided by using the fixed (non-reduced) minimum interval when calculating the participant timeout [RFC8108]. To address the other concerns, endpoints **SHOULD** ignore parameters that configure the RTCP reporting interval to be significantly longer than the default five-second interval specified in [RFC3550] (unless the media data rate is so low that the longer reporting interval roughly corresponds to 5% of the media data rate) or that configure the RTCP reporting interval small enough that the RTCP bandwidth would exceed the media bandwidth.

The guidelines in [RFC6562] apply when using variable bit rate (VBR) audio codecs such as Opus.

Encryption of the header extensions is **RECOMMENDED**, unless there are known reasons, like RTP middleboxes performing voice-activity-based source selection or third-party monitoring that will greatly benefit from the information, and this has been expressed using API or signaling. If further evidence is produced to show that information leakage is significant from audio level indications, then the use of encryption needs to be mandated at that time.

In multi-party communication scenarios using RTP middleboxes, the middleboxes are **REQUIRED**, by this protocol, to not weaken the sessions' security. The middlebox **SHOULD** maintain confidentiality, maintain integrity, and perform source authentication. The middlebox **MAY** perform checks that prevent any endpoint participating in a conference to impersonate another. Some additional security considerations regarding multi-party topologies can be found in [RFC7667].

The CaptureID is created as part of the CLUE protocol. The CaptId SDES item is used to convey the same CaptureID value in the SDES item. When sending the SDES item, the security considerations specified in Section 6 of [RFC7941] and in the communication security section of this memo (see Section 7) are applicable. Note that since the CaptureID is also carried in CLUE protocol messages, it is **RECOMMENDED** that this SDES item use at least similar protection profiles as the CLUE protocol messages carried in the CLUE data channel.

10. References

10.1. 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>>.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <<https://www.rfc-editor.org/info/rfc3711>>.
- [RFC5763] Fischl, J., Tschofenig, H., and E. Rescorla, "Framework for Establishing a Secure Real-time Transport Protocol (SRTP) Security Context Using Datagram Transport Layer Security (DTLS)", RFC 5763, DOI 10.17487/RFC5763, May 2010, <<https://www.rfc-editor.org/info/rfc5763>>.

-
- [RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, DOI 10.17487/RFC5764, May 2010, <<https://www.rfc-editor.org/info/rfc5764>>.
 - [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
 - [RFC6904] Lennox, J., "Encryption of Header Extensions in the Secure Real-time Transport Protocol (SRTP)", RFC 6904, DOI 10.17487/RFC6904, April 2013, <<https://www.rfc-editor.org/info/rfc6904>>.
 - [RFC7941] Westerlund, M., Burman, B., Even, R., and M. Zanaty, "RTP Header Extension for the RTP Control Protocol (RTCP) Source Description Items", RFC 7941, DOI 10.17487/RFC7941, August 2016, <<https://www.rfc-editor.org/info/rfc7941>>.
 - [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>>.
 - [RFC8843] Holmberg, C., Alvestrand, H., and C. Jennings, "Negotiating Media Multiplexing Using the Session Description Protocol (SDP)", RFC 8843, DOI 10.17487/RFC8843, December 2020, <<https://www.rfc-editor.org/info/rfc8843>>.
 - [RFC8845] Duckworth, M., Ed., Pepperell, A., and S. Wenger, "Framework for Telepresence Multi-Streams", RFC 8845, DOI 10.17487/RFC8845, December 2020, <<https://www.rfc-editor.org/info/rfc8845>>.
 - [RFC8846] Presta, R. and S P. Romano, "An XML Schema for the Controlling Multiple Streams for Telepresence (CLUE) Data Model", RFC 8846, DOI 10.17487/RFC8846, December 2020, <<http://www.rfc-editor.org/info/rfc8846>>.

10.2. Informative References

- [FIPS186] National Institute of Standards and Technology (NIST), "Digital Signature Standard (DSS)", FIPS, PUB 186-4, DOI 10.6028/NIST.FIPS.186-4, July 2013, <<https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf>>.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <<https://www.rfc-editor.org/info/rfc3264>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<https://www.rfc-editor.org/info/rfc3550>>.
- [RFC3556] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, DOI 10.17487/RFC3556, July 2003, <<https://www.rfc-editor.org/info/rfc3556>>.

-
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <<https://www.rfc-editor.org/info/rfc4566>>.
 - [RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, Ed., "A Session Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, DOI 10.17487/RFC4575, August 2006, <<https://www.rfc-editor.org/info/rfc4575>>.
 - [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>>.
 - [RFC4796] Hautakorpi, J. and G. Camarillo, "The Session Description Protocol (SDP) Content Attribute", RFC 4796, DOI 10.17487/RFC4796, February 2007, <<https://www.rfc-editor.org/info/rfc4796>>.
 - [RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, DOI 10.17487/RFC5124, February 2008, <<https://www.rfc-editor.org/info/rfc5124>>.
 - [RFC5506] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", RFC 5506, DOI 10.17487/RFC5506, April 2009, <<https://www.rfc-editor.org/info/rfc5506>>.
 - [RFC6562] Perkins, C. and JM. Valin, "Guidelines for the Use of Variable Bit Rate Audio with Secure RTP", RFC 6562, DOI 10.17487/RFC6562, March 2012, <<https://www.rfc-editor.org/info/rfc6562>>.
 - [RFC7022] Begen, A., Perkins, C., Wing, D., and E. Rescorla, "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)", RFC 7022, DOI 10.17487/RFC7022, September 2013, <<https://www.rfc-editor.org/info/rfc7022>>.
 - [RFC7201] Westerlund, M. and C. Perkins, "Options for Securing RTP Sessions", RFC 7201, DOI 10.17487/RFC7201, April 2014, <<https://www.rfc-editor.org/info/rfc7201>>.
 - [RFC7202] Perkins, C. and M. Westerlund, "Securing the RTP Framework: Why RTP Does Not Mandate a Single Media Security Solution", RFC 7202, DOI 10.17487/RFC7202, April 2014, <<https://www.rfc-editor.org/info/rfc7202>>.
 - [RFC7205] Romanow, A., Botzko, S., Duckworth, M., and R. Even, Ed., "Use Cases for Telepresence Multistreams", RFC 7205, DOI 10.17487/RFC7205, April 2014, <<https://www.rfc-editor.org/info/rfc7205>>.
 - [RFC7667] Westerlund, M. and S. Wenger, "RTP Topologies", RFC 7667, DOI 10.17487/RFC7667, November 2015, <<https://www.rfc-editor.org/info/rfc7667>>.
 - [RFC8108] Lennox, J., Westerlund, M., Wu, Q., and C. Perkins, "Sending Multiple RTP Streams in a Single RTP Session", RFC 8108, DOI 10.17487/RFC8108, March 2017, <<https://www.rfc-editor.org/info/rfc8108>>.

- [RFC8285] Singer, D., Desineni, H., and R. Even, Ed., "A General Mechanism for RTP Header Extensions", RFC 8285, DOI 10.17487/RFC8285, October 2017, <<https://www.rfc-editor.org/info/rfc8285>>.
- [RFC8848] Hanton, R., Kyzivat, P., Xiao, L., and C. Groves, "Session Signaling for Controlling Multiple Streams for Telepresence (CLUE)", RFC 8848, DOI 10.17487/RFC8848, December 2020, <<https://www.rfc-editor.org/info/rfc8848>>.

Acknowledgments

The authors would like to thank Allyn Romanow and Paul Witty for contributing text to this work. Magnus Westerlund helped draft the security section.

Authors' Addresses

Roni Even

Tel Aviv

Israel

Email: ron.even.tlv@gmail.com

Jonathan Lennox

8x8, Inc. / Jitsi

Jersey City, NJ 07302

United States of America

Email: jonathan.lennox@8x8.com