

SIPREC Working Group	C. Eckel
Internet-Draft	Cisco
Intended status: Informational	June 16, 2011
Expires: December 18, 2011	

Real-time Transport Protocol (RTP) Recommendations for SIPREC draft-eckel-siprec-rtp-rec-01

Abstract

This document provides recommendations and guidelines for RTP and RTCP in the context of SIPREC. This document exists as a standalone document to facilitate discussion of the RTP recommendations, and it is anticipated that portions of this document will be incorporated into **[I-D.portman-siprec-protocol]** rather than this document itself being adopted as a working group document.

Status of this Memo

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

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

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

This Internet-Draft will expire on December 18, 2011.

Copyright Notice

Copyright (c) 2011 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 (<http://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. Definitions**
- 4. Roles**
 - 4.1. SRC acting as an RTP Translator**
 - 4.2. SRC acting as an RTP Mixer**
 - 4.3. SRC acting as an RTP Endpoint**
- 5. RTCP**
- 6. RTP Profile**
- 7. SSRC**
- 8. CSRC**

- [9. SDES](#)
- [10. CNAME](#)
- [11. Keepalive](#)
- [12. RTCP Feedback Messages](#)
 - [12.1. Full Intra Request](#)
 - [12.1.1. SIP INFO for FIR](#)
 - [12.2. Picture Loss Indicator](#)
 - [12.3. Temporary Maximum Media Stream Bit Rate Request](#)
 - [12.3.1. Renegotiation of SDP bandwidth attribute](#)
- [13. Symmetric RTP/RTCP for Sending and Receiving](#)
- [14. Security Considerations](#)
- [15. IANA Considerations](#)
- [16. Acknowledgements](#)
- [17. References](#)
 - [17.1. Normative References](#)
 - [17.2. Informative References](#)
- [§ Author's Address](#)

1. Introduction

TOC

This document provides recommendations and guidelines for RTP and RTCP in the context of SIPREC. In order to communicate most effectively, the Session Recording Client (SRC) and the Session Recording (SRS) SHOULD utilize the mechanisms provided by RTP in a well defined and predictable manner. It is the goal of this document to make the reader aware of these mechanisms and provide recommendations and guidelines. This document is completely informational. It includes no requirements and no normative language. This document exists as a standalone document to facilitate discussion of the RTP recommendations, and it is anticipated that portions of this document will be incorporated into [\[I-D.portman-siprec-protocol\]](#) rather than this document itself being adopted as a working group document.

2. Terminology

TOC

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

3. Definitions

TOC

This document uses the definitions provided in [\[I-D.ietf-siprec-architecture\]](#) for all SIPREC modules. A Session Recording Client (SRC), as defined within SIPREC, is expected to implement one or more of a variety of RTP roles within the context of various SIPREC use cases. These roles include, but are not limited to, acting as an end system, a mixer, or a translator. This document uses the definitions provided in "RTP: A Transport Protocol for Real-Time Application" [\[RFC3550\]](#).

4. Roles

TOC

An SRC has the task of gathering media from the various UAs in a CS and forwarding the information to the SRS in the context of the RS. For a given media stream, the SRC has a number of ways of doing this and may choose different ways for different media streams. The following subsections define a set of roles an SRC may choose to play from the perspective of its position between a UA within a CS, and an SRS within an RS.

Figure 1: SRC Positioning

4.1. SRC acting as an RTP Translator

TOC

The SRC may act as a translator, as defined in **[RFC3550]**. A defining characteristic of a translator is that it forwards RTP packets with their SSRC identifier intact. Optionally, an SRC acting as a translator can perform transcoding (e.g., from one codec to another), and this can result in a different rate of packets. With this approach, RTP received as separate streams from different sources (e.g., from different UAs with different SSRCs) cannot be mixed by the SRC and SHOULD be sent separately to the SRS. More specifically, the SRC cannot multiple RTP streams with different SSRCs into a single RTP stream multiplexed by SSRC and/or CSRC. All RTCP reports are passed by the SRC between the UAs and the SRS, such that they are able to detect any SSRC collisions.

RTCP Sender Reports generated by a UA sending a stream SHOULD be forwarded to the SRS. RTCP Receiver Reports generated by the SRS need to be forwarded to the relevant UA. This implies that the SRC cannot take multiplex RTP packets with different SSRC values into a single SSRC stream unless it was The SRC in the case may need to manipulate the RTCP Receiver Reports to take account of any transcoding that has taken place.

If SRTP is used on both the CS and the RS, decryption and/or re-encryption may occur. For example, if different keys are used, it must occur. If the same keys are used, it need not occur.

4.2. SRC acting as an RTP Mixer

TOC

In the case of the SRC acting as a RTP mixer, as defined in **[RFC3550]**, the SRC combines RTP streams from different UA and sends them towards the SRS using its own SSRC. The SSRCs from the contributing UA SHOULD be conveyed as CSRCs identifiers within this stream. The SRC may make timing adjustments among the received streams and generate its own timing on the stream sent to the SRS. Optionally an SRC acting as a mixer can perform transcoding, and can even cope with different codings received from different UAs. RTCP Sender Reports and Receiver Reports are not forwarded by an SRC acting as mixer, but there are requirements for forwarding RTCP Source Description (SDS) packets. The SRC generates its own RTCP Sender and Receiver reports toward the associated UAs and SRS. The use of SRTP between the SRC and the SRS for the RS is independent of the use of SRTP between the UAs and SRC for the CS.

4.3. SRC acting as an RTP Endpoint

TOC

The case of the SRC acting as an RTP endpoint, as defined in **[RFC3550]**, is similar to the mixer case, except that the RTP session between the SRC and the SRS is considered completely independent from the RTP session that is part of the CS. The SRC can, but need not, mix RTP streams from different participants prior to sending to the SRS. RTCP between the SRC and the SRS is completely independent of RTCP on the CS. The use of SRTP between the SRC and the SRS is independent of the use of SRTP on the CS.

5. RTCP

TOC

The RTP data transport is augmented by a control protocol (RTCP) to allow monitoring of the data delivery. RTCP, as defined in **[RFC3550]**, is based on the periodic transmission of control packets to all participants in the RTP session, using the same distribution mechanism

as the data packets. Support for RTCP is required by **[RFC3550]**, and it provides, among other things, the following important functionality in relation to SIPREC:

1) Feedback on the quality of the data distribution. This feedback is important for flow and congestion control functions, and to get feedback from the receivers to diagnose faults in the distribution. As such, RTCP is a well defined and efficient mechanism for the SRS to inform the SRC or any issues that arise in the reception of media that is to be recorded.

2) Carries a persistent transport-level identifier for an RTP source called the canonical name or CNAME. The SSRC identifier may change if a conflict is discovered or a program is restarted; in which case receivers can use the CNAME to keep track of each participant. Receivers may also use the CNAME to associate multiple data streams from a given participant in a set of related RTP sessions, for example to synchronize audio and video. Synchronization of media streams is also facilitated by the NTP and RTP timestamps included in RTCP packets by data senders.

6. RTP Profile

TOC

The RECOMMENDED RTP profiles for both the SRC and SRS are "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", **[RFC5124]** when using encrypted RTP streams, and "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", **[RFC4585]** when using non encrypted media streams. However, as this is not a requirement, some implementations may use "The Secure Real-time Transport Protocol (SRTP)", **[RFC3711]** and "RTP Profile for Audio and Video Conferences with Minimal Control", AVP **[RFC3551]**. Therefore, it is RECOMMENDED that the SRC and SRS not rely entirely on SAVPF or AVPF for core functionality that may be at least partially achievable using SAVP and AVP.

AVPF and SAVPF provide an improved RTCP timer model that allows more flexible transmission of RTCP packets as response to events, rather than strictly according to bandwidth. AVPF based codec control messages provide efficient mechanisms for an SRC and SRS to handle events such as scene changes, error recovery, and dynamic bandwidth adjustments. These messages are discussed in more detail later in this document.

SAVP and SAVPF provide media encryption, integrity protection, replay protection, and a limited form of source authentication. They do not contain or require a specific keying mechanism.

7. SSRC

TOC

The synchronization source (SSRC), as defined in **[RFC3550]**, is carried in the RTP header and in various fields of RTCP packets. It is a random 32-bit number that is required to be globally unique within an RTP session. It is crucial that the number be chosen with care in order that participants on the same network or starting at the same time are not likely to choose the same number. Guidelines regarding SSRC value selection and conflict resolution are provided in **[RFC3550]**.

The SSRC may also be used to separate different sources of media within a single RTP session. For this reason as well as for conflict resolution, it is important that the SRC and SRS handle changes in SSRC values and properly identify the reason of the change. The CNAME values carried in RTCP facilitate this identification.

8. CSRC

TOC

The contributing source (CSRC), as defined in **[RFC3550]**, identifies the source of a stream of RTP packets that has contributed to the combined stream produced by an RTP mixer. The mixer inserts a list of the SSRC identifiers of the sources that contributed to the generation of a particular packet into the RTP header of that packet. This list is called the CSRC list. It is RECOMMENDED that a SRC, when acting a mixer, sets the CSRC list accordingly, and that the SRS interprets the CSRC list appropriately when received.

9. SDES

TOC

The Source Description (SDES), as defined in [\[RFC3550\]](#), contains an SSRC/CSRC identifier followed by a list of zero or more items, which carry information about the SSRC/CSRC. End systems send one SDES packet containing their own source identifier (the same as the SSRC in the fixed RTP header). A mixer sends one SDES packet containing a chunk for each contributing source from which it is receiving SDES information, or multiple complete SDES packets.

10. CNAME

TOC

The Canonical End-Point Identifier (CNAME), as defined in [\[RFC3550\]](#), provides the binding from the SSRC identifier to an identifier for the source (sender or receiver) that remains constant. It is important the an SRC and SRS generate CNAMEs appropriately and use them for this purpose. Guidelines for generating CNAME values are provided in "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)" [\[RFC6222\]](#).

11. Keepalive

TOC

It is anticipated that media streams in SIPREC may exist in inactive states for extended periods of times for an of a number of valid reasons. In order to the bindings and any pinholes in NATs/firewalls to remain active during such intervals, it is RECOMMENDED to follow the keep-alive procedures in "Application Mechanism for Keeping Alive the NAT Mappings Associated to RTP/RTP Control Protocol (RTCP) Flows" [\[RFC6263\]](#) for all RTP media streams.

12. RTCP Feedback Messages

TOC

"Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)" [\[RFC5104\]](#) specifies extensions to the messages defined in AVPF [\[RFC4585\]](#). Support for and proper usage of these messages is important to SRC and SRS implementations. Note that these messages are applicable only when using the AVFP or SAVPF RTP profiles.

12.1. Full Intra Request

TOC

A Full Intra Request (FIR) Command, when received by the designate media sender, requires that the media sender sends a Decoder Refresh Point at the earliest opportunity. Using a decoder refresh point implies refraining from using any picture sent prior to that point as a reference for the encoding process of any subsequent picture sent in the stream.

Decoder refresh points, especially Intra or IDR pictures for H.264 video codecs, are in general several times larger in size than predicted pictures. Thus, in scenarios in which the available bit rate is small, the use of a decoder refresh point implies a delay that is significantly longer than the typical picture duration.

12.1.1. SIP INFO for FIR

TOC

"XML Schema for Media Control" [\[RFC5168\]](#) defines an Extensible Markup Language (XML) Schema for video fast update. Implementations are discouraged from using the method described except for backward compatibility purposes. Implementations SHOULD use FIR messages instead.

12.2. Picture Loss Indicator

TOC

Using the FIR command to recover from errors is explicitly disallowed, and instead the PLI message defined in AVPF **[RFC4585]** SHOULD be used. The PLI message reports lost pictures and has been included in AVPF for precisely that purpose.

12.3. Temporary Maximum Media Stream Bit Rate Request

TOC

A receiver, translator, or mixer uses the Temporary Maximum Media Stream Bit Rate Request (TMMBR) to request a sender to limit the maximum bit rate for a media stream to the provided value. Appropriate use of TMMBR facilitates rapid adaptation to changes in available bandwidth.

12.3.1. Renegotiation of SDP bandwidth attribute

TOC

If it is likely that the new value indicated by TMMBR will be valid for the remainder of the session, the TMMBR sender is expected to perform a renegotiation of the session upper limit using the session signaling protocol. Therefore for SIPREC, implementations are RECOMMENDED to use TMMBR for temporary changes, and renegotiation of bandwidth via SDP offer/answer of more permanent changes.

13. Symmetric RTP/RTCP for Sending and Receiving

TOC

Within an SDP offer/answer exchange, RTP entities choose the RTP and RTCP transport addresses (i.e., IP addresses and port numbers) on which to receive packets. When sending packets; however, they may use a different source IP address or port number for RTP, RTCP, or both. When the IP address and port used to send and receive RTP is the same, it is termed "symmetric RTP" **[RFC4961]**. Likewise, when the IP address and port used to send and receive RTCP is the same, it is termed "symmetric RTCP" **[RFC4961]**.

If an SRC or SRS is sending RTP, it is RECOMMENDED to use symmetric RTP. If an SRC or SRS is sending RTCP, it is RECOMMENDED to use symmetric RTCP. Note, symmetric RTP and symmetric RTCP are different from RTP/RTCP multiplexing **[RFC5761]**.

14. Security Considerations

TOC

In many scenarios it will be critical that the media transported between the SRC and SRS using RTP/RTCP be protected. Media encryption is an important element of the overall SIPREC solution. The details regarding the encryption of the RTP/RTCP media will be addressed in **[I-D.portman-siprec-protocol]**.

15. IANA Considerations

TOC

None.

16. Acknowledgements

TOC

Thank you to Andrew Hutton, Ram Mohan, and Muthu Perumal, John Elwell, and Dan Wing for their valuable contributions.

17. References

TOC

17.1. Normative References

TOC

[RFC2119] [Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels,"](#) BCP 14, RFC 2119, March 1997 ([TXT](#), [HTML](#), [XML](#)).

17.2. Informative References

TOC

- [I-D.ietf-siprec-architecture] Hutton, A., Portman, L., Jain, R., and K. Rehor, "[An Architecture for Media Recording using the Session Initiation Protocol](#)," draft-ietf-siprec-architecture-02 (work in progress), April 2011 ([TXT](#)).
- [I-D.portman-siprec-protocol] Portman, L., Lum, H., Johnston, A., and A. Hutton, "[Session Recording Protocol](#)," draft-portman-siprec-protocol-04 (work in progress), May 2011 ([TXT](#)).
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "[RTP: A Transport Protocol for Real-Time Applications](#)," STD 64, RFC 3550, July 2003 ([TXT](#), [PS](#), [PDF](#)).
- [RFC3551] Schulzrinne, H. and S. Casner, "[RTP Profile for Audio and Video Conferences with Minimal Control](#)," STD 65, RFC 3551, July 2003 ([TXT](#), [PS](#), [PDF](#)).
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "[The Secure Real-time Transport Protocol \(SRTP\)](#)," RFC 3711, March 2004 ([TXT](#)).
- [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, July 2006 ([TXT](#)).
- [RFC4961] Wing, D., "[Symmetric RTP / RTP Control Protocol \(RTCP\)](#)," BCP 131, RFC 4961, July 2007 ([TXT](#)).
- [RFC5104] Wenger, S., Chandra, U., Westerlund, M., and B. Burman, "[Codec Control Messages in the RTP Audio-Visual Profile with Feedback \(AVPF\)](#)," RFC 5104, February 2008 ([TXT](#)).
- [RFC5124] Ott, J. and E. Carrara, "[Extended Secure RTP Profile for Real-time Transport Control Protocol \(RTCP\)-Based Feedback \(RTP/SAVPF\)](#)," RFC 5124, February 2008 ([TXT](#)).
- [RFC5168] Levin, O., Even, R., and P. Hagendorf, "[XML Schema for Media Control](#)," RFC 5168, March 2008 ([TXT](#)).
- [RFC5761] Perkins, C. and M. Westerlund, "[Multiplexing RTP Data and Control Packets on a Single Port](#)," RFC 5761, April 2010 ([TXT](#)).
- [RFC6222] Begen, A., Perkins, C., and D. Wing, "[Guidelines for Choosing RTP Control Protocol \(RTCP\) Canonical Names \(CNAMEs\)](#)," RFC 6222, April 2011 ([TXT](#)).
- [RFC6263] Marjou, X. and A. Sollaud, "[Application Mechanism for Keeping Alive the NAT Mappings Associated with RTP / RTP Control Protocol \(RTCP\) Flows](#)," RFC 6263, June 2011 ([TXT](#)).

Author's Address

TOC

Charles Eckel
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States
Email: eckelcu@cisco.com