

Internet Engineering Task Force (IETF)
Request for Comments: 7244
Category: Standards Track
ISSN: 2070-1721

H. Asaeda
NICT
Q. Wu
R. Huang
Huawei
May 2014

RTP Control Protocol (RTCP) Extended Report (XR) Blocks
for Synchronization Delay and Offset Metrics Reporting

Abstract

This document defines two RTP Control Protocol (RTCP) Extended Report (XR) blocks that allow the reporting of initial synchronization delay and synchronization offset metrics for use in a range of RTP applications.

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 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7244>.

Copyright Notice

Copyright (c) 2014 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
1.1. Synchronization Delay and Offset Metrics Reporting Blocks ..	2
1.2. RTCP and RTCP XR Reports	3
1.3. Performance Metrics Framework	3
1.4. Applicability	3
2. Terminology	4
2.1. Standards Language	4
3. RTP Flow Initial Synchronization Delay Report Block	4
3.1. Metric Block Structure	5
3.2. Definition of Fields in RTP Flow Initial Synchronization Delay Metrics Block	5
4. RTP Flow Synchronization Offset Metrics Block	6
4.1. Metric Block Structure	7
4.2. Definition of Fields in RTP Flow General Synchronization Offset Metrics Block	7
5. SDP Signaling	9
5.1. SDP rtcp-xr-attrib Attribute Extension	9
5.2. Offer/Answer Usage	9
6. IANA Considerations	9
7. Security Considerations	10
8. Acknowledgements	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Appendix A. Metrics Represented Using the Template from RFC 6390 ..	12

1. Introduction

1.1. Synchronization Delay and Offset Metrics Reporting Blocks

This document defines two new block types to augment those defined in [RFC3611], for use in a range of RTP applications.

The first new block type supports reporting of the Initial Synchronization Delay to establish a multimedia session. Information is recorded about the time difference between the start of RTP sessions and the time the RTP receiver acquires all components of RTP sessions in the multimedia session [RFC6051].

The second new block type supports reporting of the relative synchronization offset time of two arbitrary streams (e.g., between audio and video streams), with the same RTCP CNAME included in RTCP Source description items (SDES) packets [RFC3550].

These metrics belong to the class of transport-level metrics defined in [RFC6792].

1.2. RTCP and RTCP XR Reports

The use of RTCP for reporting is defined in [RFC3550]. [RFC3611] defined an extensible structure for reporting -- the RTCP Extended Report (XR). This document defines a new Extended Report block for use with [RFC3550] and [RFC3611].

1.3. Performance Metrics Framework

"Guidelines for Considering New Performance Metric Development" [RFC6390] provides guidance on the definition and specification of performance metrics. "Guidelines for Use of the RTP Monitoring Framework" [RFC6792] provides guidance for reporting block format using RTCP XR. The metrics block described in this document is in accordance with the guidelines in [RFC6390] and [RFC6792].

1.4. Applicability

When joining each session in layered video sessions [RFC6190] or the multimedia session, a receiver may not synchronize playout across the multimedia session or layered video session until RTCP Sender Report (SR) packets have been received on all components of RTP sessions. The components of RTP sessions are per-media-type RTP sessions for the multimedia sessions or per-layer RTP sessions for the layered video sessions. For multicast sessions, the Initial Synchronization Delay metric varies with the session bandwidth, the number of members, and the number of senders in the session. The RTP Flow Initial Synchronization Delay Metrics Block defined in this document can be used to report such a metric, i.e., the Initial Synchronization Delay to receive all the RTP streams belonging to the same multimedia session or layered video session. In the absence of packet loss, the Initial Synchronization Delay is equal to the average time taken to receive the first RTCP packet in the RTP session with the longest RTCP reporting interval. In the presence of packet loss, the media synchronization should rely on the in-band mapping of RTP and NTP-format timestamps [RFC6051] or wait until the reporting interval has passed, and the next RTCP SR packet is sent.

Receivers of the RTP Flow Initial Synchronization Delay Metrics Block could use this metric to compare with targets (i.e., Service Level Agreement or thresholds of the system) to help ensure the quality of real-time application performance.

In an RTP multimedia session, there can be an arbitrary number of streams carried in different RTP sessions, with the same RTCP CNAME. These streams may be not synchronized with each other. For example, one audio stream and one video stream belong to the same session, and the audio stream is transmitted lagging behind the video stream for

multiple tens of milliseconds [TR-126]. The RTP Flow Synchronization Offset block can be used to report such synchronization offset between video and audio streams. This block is also applied to the case where an RTP session can contain media streams with media from multiple media types. The metrics defined in the RTP Flow Synchronization Offset Metrics Block can be used by the network manager for troubleshooting and dealing with user-experience issues.

2. Terminology

2.1. Standards Language

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

In addition, the following terms are defined:

Initial Synchronization Delay:

A multimedia session comprises a set of concurrent RTP sessions among a common group of participants, using one RTP session for each media type. The Initial Synchronization Delay is the average time for the receiver to synchronize all components of a multimedia session [RFC6051].

Synchronization Offset:

Synchronization between two media streams must be maintained to ensure satisfactory Quality of Experience (QoE). Two media streams can be of the same or different media types belonging to one RTP session, or of different media types belonging to one multimedia session. The Synchronization Offset is the relative time difference of the two media streams that need to be synchronized.

3. RTP Flow Initial Synchronization Delay Metrics Block

This block is sent by RTP receivers and reports the Initial Synchronization Delay beyond the information carried in the standard RTCP packet format. Information is recorded about the time difference between the start of the multimedia session and the time when the RTP receiver acquires all components of RTP sessions [RFC6051] measured at the receiving end of the RTP stream.

This block needs to be exchanged only occasionally, for example, sent once at the start of the RTP session.

3.1. Metric Block Structure

The RTP Flow Initial Synchronization Delay Metrics Block has the following format:

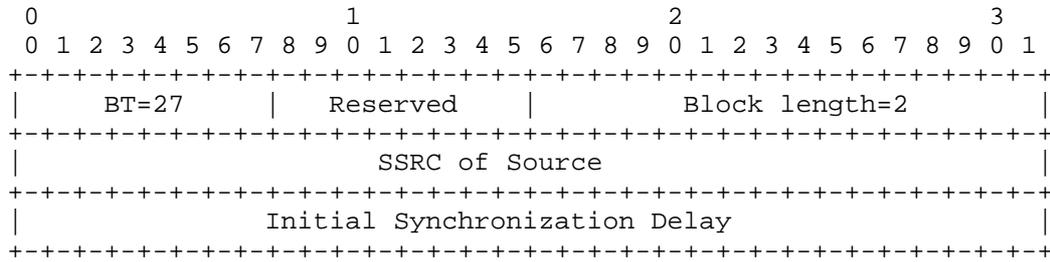


Figure 1: Report Block Structure

3.2. Definition of Fields in RTP Flow Initial Synchronization Delay Metrics Block

Block type (BT): 8 bits

The RTP Flow Initial Synchronization Delay Metrics Block is identified by the constant 27.

Reserved: 8 bits

This field is reserved for future definition. In the absence of such a definition, the bits in this field MUST be set to zero and ignored by the receiver.

Block length: 16 bits

The constant 2, in accordance with the definition of this field in Section 3 of RFC 3611 [RFC3611].

SSRC of source: 32 bits

The SSRC of the media source SHALL be set to the value of the SSRC identifier carried in any arbitrary component of RTP sessions belonging to the same multimedia session.

Initial Synchronization Delay: 32 bits

The average delay, expressed in units of 1/65536 seconds, from the beginning of the multimedia session [RFC6051] to the time when RTCP packets are received on all of the component RTP sessions. It is recommended that the beginning of the multimedia session is

chosen as the time when the receiver has joined the first RTP session of the multimedia session. The value of the Initial Synchronization Delay is calculated based on received RTCP SR packets or the RTP header extension containing the in-band mapping of RTP and NTP-format timestamps [RFC6051]. If there is no packet loss, the Initial Synchronization Delay is expected to be equal to the average time taken to receive the first RTCP packet in the RTP session with the longest RTCP reporting interval or to the average time taken to receive the first RTP header extension containing the in-band mapping of RTP and NTP-format timestamps.

If the measurement is unavailable, the value of this field with all bits set to 1 MUST be reported.

4. RTP Flow Synchronization Offset Metrics Block

In the RTP multimedia sessions or one RTP session, there can be an arbitrary number of media streams and each media stream (e.g., audio stream or video stream) is sent in a separate RTP stream. In case of one RTP session, each media stream or each medium uses a different SSRC. The receiver correlates these media streams that need to be synchronized by means of the RTCP CNAME contained in the RTCP Source Description (SDS) packets [RFC3550].

This block is sent by RTP receivers and reports the synchronization offset of two arbitrary RTP streams that need to be synchronized in the RTP multimedia session. Information is recorded about the relative average time difference between two arbitrary RTP streams (the reporting stream and the reference stream) with the same CNAME and measured at the receiving end of the RTP stream. In order to tell what the offset of the reporting stream is relative to, the block for the reference stream with synchronization offset of zero should be reported.

Instances of this block refer by synchronization source (SSRC) to the separate auxiliary Measurement Information block [RFC6776], which describes measurement periods in use (see Section 4.2 of [RFC6776]). This metrics block relies on the measurement period in the Measurement Information block indicating the span of the report and SHOULD be sent in the same compound RTCP packet as the Measurement Information Block. If the measurement period is not received in the same compound RTCP packet as this block, this block MUST be discarded.

4.1. Metric Block Structure

The RTP Flow General Synchronization Offset Metrics Block has the following format:

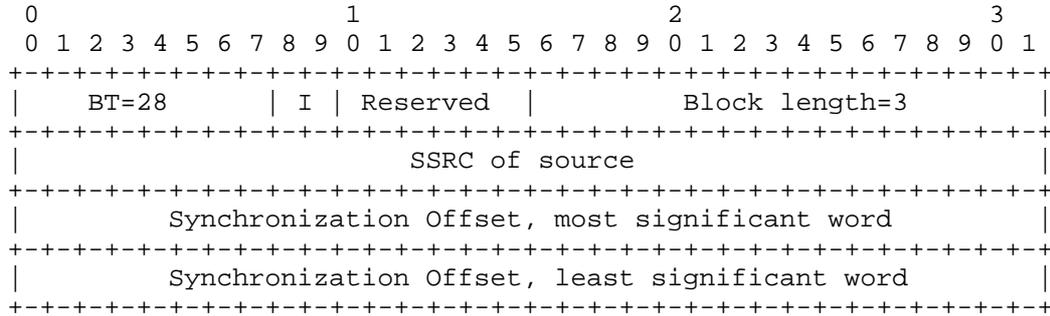


Figure 2: Report Block Structure

4.2. Definition of Fields in RTP Flow General Synchronization Offset Metrics Block

Block type (BT): 8 bits

The RTP Flow General Synchronization Offset Metrics Block is identified by the constant 28.

Interval Metric Flag (I): 2 bits

This field is used to indicate whether the Burst/Gap Discard Summary Statistics metrics are Sampled, Interval, or Cumulative metrics:

- I=10: Interval Duration - the reported value applies to the most recent measurement interval duration between successive metrics reports.
- I=11: Cumulative Duration - the reported value applies to the accumulation period characteristic of cumulative measurements.
- I=01: Sampled Value - the reported value is a sampled instantaneous value.

In this document, the value I=00 is the reserved value and MUST NOT be used. If the value I=00 is received, then the XR block MUST be ignored by the receiver.

Reserved: 6 bits

This field is reserved for future definition. In the absence of such a definition, the bits in this field MUST be set to zero and MUST be ignored by the receiver.

Block length: 16 bits

The constant 3, in accordance with the definition of this field in Section 3 of RFC 3611 [RFC3611].

SSRC of Source: 32 bits

The SSRC of the media source SHALL be set to the value of the SSRC identifier of the reporting RTP stream to which the XR relates.

Synchronization Offset: 64 bits

The synchronization offset of the reporting RTP stream relative to the reference stream with the same CNAME. The calculation of Synchronization Offset is similar to the Difference D calculation in the RFC 3550. That is to say, if S_i is the NTP timestamp from the reporting RTP packet i , R_i is the time of arrival in NTP timestamp units for reporting RTP packet i , S_j is the NTP timestamp from the reference RTP packet j , and R_j is the time of arrival in NTP timestamp units for reference RTP packet j , then the value of the Synchronization Offset D may be expressed as

$$D(i,j) = (R_j - R_i) - (S_j - S_i) = (R_j - S_j) - (R_i - S_i)$$

If in-band delivery of NTP-format timestamps is supported [RFC6051], S_i and S_j should be obtained directly from the RTP packets where NTP timestamps are available. If not, S_i and S_j should be calculated from their corresponding RTP timestamps. The value of the Synchronization Offset is represented using a 64-bit signed NTP-format timestamp as defined in [RFC5905], which is a 64-bit signed fixed-point number with the integer part in the first 32 bits and the fractional part in the last 32 bits. A positive value of the Synchronization Offset means that the reporting stream leads before the reference stream, while a negative one means the reporting stream lags behind the reference stream. The Synchronization Offset of zero means the stream is the reference stream.

If the measurement is unavailable, the value of this field with all bits set to 1 MUST be reported.

5. SDP Signaling

[RFC3611] defines the use of SDP (Session Description Protocol) [RFC4566] for signaling the use of XR blocks. XR blocks MAY be used without prior signaling.

5.1. SDP rtcp-xr-attrib Attribute Extension

Using the Augmented Backus-Naur Form (ABNF) [RFC5234], two new parameters are defined for the two report blocks defined in this document to be used with SDP [RFC4566]. They have the following syntax within the "rtcp-xr" attribute [RFC3611]:

```
xr-format =/ xr-rfisd-block
           / xr-rfso-block
```

```
xr-rfisd-block = "rtp-flow-init-syn-delay"
xr-rfso-block = "rtp-flow-syn-offset"
```

Refer to Section 5.1 of RFC 3611 [RFC3611] for a detailed description and the full syntax of the "rtcp-xr" attribute.

5.2. Offer/Answer Usage

When SDP is used in the offer/answer context, the SDP Offer/Answer usage defined in [RFC3611] applies.

6. IANA Considerations

New report block types for RTCP XR are subject to IANA registration. For general guidelines on IANA allocations for RTCP XR, refer to Section 6.2 of [RFC3611].

This document assigns two new block type values in the RTCP XR Block Type Registry:

```
Name:          RFISD
Long Name:     RTP Flow Initial Synchronization Delay
Value         27
Reference:    Section 3
```

```
Name:          RFSO
Long Name:     RTP Flow Synchronization Offset
Value         28
Reference:    Section 4
```

This document also registers two new SDP [RFC4566] parameters for the "rtcp-xr" attribute in the RTCP XR SDP Parameters Registry:

- * "rtp-flow-init-syn-delay "
- * "rtp-flow-syn-offset"

The contact information for the registrations is:
RAI Area Directors <rai-ads@tools.ietf.org>

7. Security Considerations

When using Secure RTP [RFC3711], or other media-layer security, reporting accurate synchronization offset information can expose some details about the timing of the cryptographic operations that are used to protect the media. There is a possibility that this timing information might enable a side-channel attack on the encryption. For environments where this attack is a concern, implementations need to take care to ensure cryptographic processing and media compression take the same amount of time irrespective of the media content, to avoid the potential attack.

Besides this, it is believed that this RTCP XR block introduces no new security considerations beyond those described in [RFC3611].

8. Acknowledgements

The authors would like to thank Bill Ver Steeg, David R. Oran, Ali Begen, Colin Perkins, Roni Even, Kevin Gross, Jing Zhao, Fernando Boronat Segui, Mario Montagud Climent, Youqing Yang, Wenxiao Yu, Yinliang Hu, Jonathan Lennox, and Stephen Farrel for their valuable comments and suggestions on this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3611] Friedman, T., Caceres, R., and A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.

- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.
- [RFC6051] Perkins, C. and T. Schierl, "Rapid Synchronisation of RTP Flows", RFC 6051, November 2010.
- [RFC6190] Wenger, S., Wang, Y., Schierl, T., and A. Eleftheriadis, "RTP Payload Format for Scalable Video Coding", RFC 6190, May 2011.
- [RFC6776] Wu, Q., "Measurement Identity and information Reporting using SDES item and XR Block", RFC 6776, August 2012.

9.2. Informative References

- [RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", RFC 6390, October 2011.
- [RFC6792] Wu, Q., "Guidelines for Use of the RTP Monitoring Framework", RFC 6792, November 2012.
- [TR-126] Broadband Forum, "Triple-play Services Quality of Experience (QoE) Requirements", Technical Report TR-126, December 2006.
- [Y.1540] ITU-T, "IP packet transfer and availability performance parameters", ITU-T Recommendation Y.1540, November 2007.

Appendix A. Metrics Represented Using the Template from RFC 6390

a. Initial Synchronization Delay Metric

- * Metric Name: RTP Initial Synchronization Delay
- * Metric Description: See the definition of "Initial Synchronization Delay" in Section 2.1.
- * Method of Measurement or Calculation: See the definition of the "Initial Synchronization Delay" field in Section 3.2.
- * Units of Measurement: See the definition of the "Initial Synchronization Delay" field in Section 3.2.
- * Measurement Point(s) with Potential Measurement Domain: See the first paragraph of Section 3.
- * Measurement Timing: See the second paragraph of Section 3.
- * Use and applications: See Section 1.4.
- * Reporting model: See RFC 3611.

b. Synchronization Offset Metric

- * Metric Name: RTP Synchronization Offset Delay
- * Metric Description: See the definition of "Synchronization Offset" in Section 1.2.
- * Method of Measurement or Calculation: See the definition of the "Synchronization Offset" field in Section 4.2.
- * Units of Measurement: See the definition of the "Synchronization Offset" field in Section 4.2.
- * Measurement Point(s) with Potential Measurement Domain: See the second paragraph of Section 4.
- * Measurement Timing: See the third paragraph of Section 4.2 for measurement timing and the Interval Metric flag.
- * Use and applications: See Section 1.4.
- * Reporting model: See RFC 3611.

Authors' Addresses

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

EMail: asaeda@nict.go.jp

Qin Wu
Huawei Technologies Co., Ltd.
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

EMail: bill.wu@huawei.com

Rachel Huang
Huawei Technologies Co., Ltd.
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

EMail: Rachel@huawei.com

