Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 5, 2013

H. Asaeda
NICT
R. Huang
Q. Wu
Huawei

February 1, 2013

RTP Control Protocol (RTCP) Extended Report (XR) Blocks for Synchronization Delay and Offset Metrics Reporting
draft-ietf-xrblock-rtcp-xr-synchronization-02

Abstract

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

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

1.1. Synchronization Delay and Offset Metrics Reporting Blocks

This draft 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 Initial Synchronization Delay to establish multimedia session. Information is recorded about 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 SDES packets [RFC3550]. Information is recorded about the synchronization offset time of each RTP stream relative to the reference RTP stream with the same CNAME and General Synchronization Offset of zero.

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 using an RTCP Extended Report (XR). This document defines a new Extended Report block for use with [RFC3550] and [RFC3611].

1.3. Performance Metrics Framework

The RTP Monitoring Architectures [RFC6792] provides guideline for reporting block format using RTCP XR. The new report block described in this memo is in compliance with the monitoring architecture specified in [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 SR packets have been received on all components of RTP sessions. The component RTP session are referred to as each RTP session for each media type in multimedia session or separate RTP session for each layer in the layered video session. For multicast session, 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 block defined in this document can be used to report such 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 equals 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 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 video stream for multiple tens of milliseconds [TR-126]. The RTP Flows Synchronization Offset block can be used to report such synchronization offset between video stream and audio stream. The metrics defined in the RTP flows synchronization Offset block can be used by network manager for trouble shooting 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 receiver to synchronize all components of a multimedia session [RFC6051].

Synchronization Offset:

Synchronization between two media streams must be maintained to ensure satisfactory QoE. Two media streams can be of the same media type belonging to one RTP session or in 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 Flows Initial Synchronization Delay Report Block

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

3.1. Metric Block Structure

The RTP Flows Initial Synchronization Delay Report Block has the following format:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    BT=RFISD   |   Reserved    |         Block length=2        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       SSRC of Source                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Initial Synchronization Delay                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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

Block type (BT): 8 bits

The RTP Flows Initial Synchronization Delay Report Block is identified by the constant <RFISD>.

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 MUST be 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 multimedia session [RFC6051] to the time when RTCP packets are received on all of the components RTP sessions. It is recommended that the beginning of 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 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 the average time taken to receive the first RTP header extension containing 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 Flows Synchronization Offset Metrics Block

In the RTP multimedia sessions, 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. The receiver associates RTP streams to be synchronized by means of RTCP CNAME contained in the RTCP Source Description (SDES) packets [RFC3550].

This block is sent by RTP receivers and reports synchronization offset of the arbitrary two RTP streams that needs to be synchronized in the RTP multimedia session. Information is recorded about the relative average time difference between the reporting stream and the reference stream with the same CNAME. For multimedia session with multiple media types (e.g., audio and video), it is recommended to choose the stream with the lower bandwidth as the reference stream. For layered video sessions, it is recommended to use the base layer

stream as the reference stream.

4.1. Metric Block Structure

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

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    BT=RFSO    |   Reserved    |         Block length=4        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        SSRC of source                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        SSRC of reference                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Synchronization Offset, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Synchronization Offset, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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

Block type (BT): 8 bits

The RTP Flow General Synchronization Offset Report Block is identified by the constant <RFSO>.

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 MUST be ignored by the receiver.

Block length: 16 bits

The constant 4, 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.

SSRC of Reference: 32 bits

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

Synchronization Offset: 64 bits

The synchronization offset of the reporting RTP stream relative to the reference RTP stream with the same CNAME. The calculation of Synchronization Offset is similar to Difference D calculation in the RFC3550. That is to say, if Si is the NTP timestamp from the reporting RTP packet i, and Ri is the time of arrival in NTP timestamp units for reporting RTP packet i, Sj is the NTP timestamp from the reference RTP packet j, and Rj 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

$$D(i,j) = (Rj - Ri) - (Sj - Si) = (Rj - Sj) - (Ri - Si)$$

If in-band delivery of NTP-format timestamps is supported [RFC6051], Si and Sj should be obtained directly from the RTP packets where NTP timestamps are available. If not, Si and Sj 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 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 value means that the reporting stream lags behind 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

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

[RFC3611]:

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

xr-rfisd-block = " init-syn-delay"
xr-rfso-block = " 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 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 Flows Initial Synchronization Delay
Value       <RFISD>
Reference:  Section 3

Name:       RFSO
Long Name:  RTP Flows Synchronization Offset Metrics Block
Value       <RFSO>
Reference:  Section 4
```

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

* "xr-rfisd "
* "xr-rfso"

The contact information for the registrations is:

Qin Wu
sunseawq@huawei.com
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012, China

7. Security Considerations

The new RTCP XR report blocks proposed in this document 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 and Yinliang Hu 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.

[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.

[TR-126] BBF Forum, "Triple-play Services Quality of Experience (QoE) Requirements", December 2006.

9.2. Informative References

[RFC6792] Wu, Q., "Guidelines for Use of the RTP Monitoring Framework", RFC 6792, November 2012.

[Y.1540] ITU-T, "ITU-T Rec. Y.1540, IP packet transfer and availability performance parameters", November 2007.

Appendix A. Change Log

Note to the RFC-Editor: please remove this section prior to publication as an RFC.

A.1. draft-ietf-xrblock-rtcp-xr-syncronization-02

The following are the major changes compared to previous version:

Editorial change based on comments raised on the list and in the IETF85 meeting

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

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

Email: Rachel@huawei.com

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

Email: sunseawq@huawei.com