Internet Engineering Task Force (IETF)
Request for Comments: 7005
Category: Standards Track
ISSN: 2070-1721

A. Clark
Telchemy
V. Singh
Aalto University
Q. Wu
Huawei

September 2013

RTP Control Protocol (RTCP) Extended Report (XR) Block for De-Jitter Buffer Metric Reporting

Abstract

This document defines an RTP Control Protocol (RTCP) Extended Report (XR) block that allows the reporting of de-jitter buffer metrics for 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/rfc7005.

Table of Contents

1. Introduction

1.1. De-Jitter Buffer Metrics Block

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

The new block type provides information on de-jitter buffer configuration and performance.

The metric belongs to the class of transport-related end-system metrics defined in [RFC6792].

Instances of this metrics block refer by synchronization source (SSRC) to the separate auxiliary Measurement Information Block [RFC6776], which contains information such as the SSRC of the measured stream, and RTP sequence numbers and time intervals indicating the span of the report.

1.2. RTCP and RTCP Extended Reports

The use of RTCP for reporting is defined in [RFC3550]. [RFC3611] defines 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

"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 on the reporting block format using RTCP XR. Metrics described in this document are in accordance with the guidelines in [RFC6390]and [RFC6792].

1.4. Applicability

Real-time applications employ a de-jitter buffer [RFC5481] to absorb jitter introduced on the path from source to destination. These metrics are used to report how the de-jitter buffer at the receiving end of the RTP stream behaves as a result of jitter in the network; they are applicable to a range of RTP applications.

These metrics correspond to terminal-related factors that affect real-time application quality and are useful for providing a better end-user quality of experience (QoE) when these terminal-related factors are used as inputs to calculate QoE metrics [QMB].

2. 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].

3. De-Jitter Buffer Operation

A de-jitter buffer is required to absorb delay variation in the network delivery of media packets. A de-jitter buffer works by holding media data for a period of time after it is received and before it is played out. Packets that arrive early are held in the de-jitter buffer longer. If packets arrive too early, they may be discarded if there is no available de-jitter buffer space. If packets are delayed excessively by the network, they may be discarded if they miss their playout time.

The de-jitter buffer can be considered a time window with the early edge aligned with the delay corresponding to the earliest arriving packet and the late edge representing the maximum permissible delay before a late arriving packet would be discarded. The delay applied to packets that arrive on time or at their expected arrival time is known as the nominal delay, and this is equivalent to the time difference/buffer size difference between the insertion point of the on-time packets and the point at which the packets are read out.

The reference for the expected arrival time may be, for example, the first packet in the session or the running average delay. If all packets arrived at their expected arrival time, then every packet would be held in the de-jitter buffer exactly the nominal delay.

The de-jitter buffer maximum delay is the delay that is applied to the earliest arriving packet that is not discarded and corresponds to the early edge of the de-jitter buffer time window.

3.1. Idealized De-Jitter Buffer

In practice, de-jitter buffer implementations vary considerably; however, they should behave in a manner conceptually consistent with an idealized de-jitter buffer, which is described as follows:

(i) Receive the first packet and delay playout by D ms. Keep the RTP timestamp (TS) and receive time as a reference.

RTP TS[1]

receive time[1]

Assume that both are normalized in ticks (there are 10,000 ticks in a millisecond).

(ii) Receive the next packet.

(iii) Calculate r = RTP TS[n] - RTP TS[1] and t = receive time[n] - receive time[1]. If r == t, then the packet arrived on time. If r < t, then the packet arrived late, and if r > t, then the packet arrived early.

(iv) Delay playout of packet by D + (r-t).

(v) Go back to (ii).

Note that this idealized implementation assumes that the sender's RTP clock is synchronized to the clock in the receiver, which is used to timestamp packet arrivals. If there is no such inherent synchronization, the system may need to use an adaptive de-jitter buffer or other techniques to ensure reliable reception.

3.2. Fixed De-Jitter Buffer

A fixed de-jitter buffer lacks provision to track the condition of the network and has a fixed size, and packets leaving the de-jitter buffer have a constant delay. For fixed de-jitter buffer implementation, the nominal delay is set to a constant value corresponding to the packets that arrive at their expected arrival time, while the maximum delay is set to a constant value corresponding to the fixed size of the de-jitter buffer.

3.3. Adaptive De-Jitter Buffer

An adaptive de-jitter buffer can adapt to the change in the network's delay and has variable size or variable delay. It allows the nominal delay to be set to a low value initially to minimize user perceived delay; however, it can automatically extend the late edge (and possibly also retract the early edge) of a buffer window if a significant proportion of the packets are arriving late (and hence being discarded).

4. De-Jitter Buffer Metrics Block

This block describes the configuration and operating parameters of the de-jitter buffer in the receiver of the RTP end system or RTP mixer that sends the report. Instances of this metrics block use the SSRC to refer to the separate auxiliary Measurement Information Block [RFC6776], which describes the measurement periods in use (see [RFC6776], Section 4.2). This metrics block relies on the measurement interval in the Measurement Information Block indicating the span of the report and MUST be sent in the same compound RTCP packet as the Measurement Information Block. If the measurement interval is not received in the same compound RTCP packet as this metrics block, this metrics block MUST be discarded.

4.1. Report Block Structure

De-Jitter Buffer (DJB) Metrics Block

```
 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=23     | I |C|  resv   |       Block Length=3          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          SSRC of Source                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          DJB nominal          |          DJB maximum          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     DJB high-water mark       |      DJB low-water mark       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Report Block Structure

4.2. Definition of Fields in De-Jitter Buffer Metrics Block

Block Type (BT): 8 bits

A De-Jitter Buffer Metrics Report Block is identified by the constant 23.

Interval Metric flag (I): 2 bits

This field is used to indicate whether the de-jitter buffer metrics are Sampled, Interval, or Cumulative metrics:

I=01: Sampled Value - the reported value is a sampled instantaneous value.

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.

In this document, de-jitter buffer metrics can only be sampled and cannot be measured over definite intervals. Also, the value I=00 is reserved for future use. Senders MUST NOT use the values I=00, I=10, or I=11. If a block is received with I=00, I=10, or I=11, the receiver MUST discard the block.

Jitter Buffer Configuration (C): 1 bit

This field is used to identify the de-jitter buffer method in use at the receiver, according to the following code:

0 = Fixed de-jitter buffer

1 = Adaptive de-jitter buffer

Reserved (resv): 5 bits

These bits are reserved. They MUST be set to zero by senders and ignored by receivers (see [RFC6709], Section 4.2).

Block Length: 16 bits

The length of this report block in 32-bit words, minus one, in accordance with the definition in [RFC3611]. This field MUST be set to 3 to match the fixed length of the report block.

SSRC of Source: 32 bits

As defined in Section 4.1 of [RFC3611].

De-jitter buffer nominal delay (DJB nominal): 16 bits

This is the current nominal de-jitter buffer delay (in milliseconds) that corresponds to the nominal de-jitter buffer delay for packets that arrive exactly on time. It is calculated based on the time spent in the de-jitter buffer for the packet that arrives exactly on time. This parameter MUST be provided for both fixed and adaptive de-jitter buffer implementations.

The measured value is an unsigned value. If the measured value exceeds 0xFFFD, the value 0xFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFF MUST be reported.

De-jitter buffer maximum delay (DJB maximum): 16 bits

This is the current maximum de-jitter buffer delay (in milliseconds) that corresponds to the earliest arriving packet that would not be discarded. It is calculated based on the time spent in the de-jitter buffer for the earliest arriving packet. In simple queue implementations, this may correspond to the size of the de-jitter buffer. In adaptive de-jitter buffer implementations, this value may vary dynamically. This parameter MUST be provided for both fixed and adaptive de-jitter buffer implementations.

The measured value is an unsigned value. If the measured value exceeds 0xFFFD, the value 0xFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFF MUST be reported.

De-jitter buffer high-water mark (DJB high-water mark): 16 bits

This is the highest value of the de-jitter buffer nominal delay (in milliseconds) that occurred at any time during the reporting interval. This parameter MUST be provided for adaptive de-jitter buffer implementations, and its value MUST be set to DJB maximum for fixed de-jitter buffer implementations.

The measured value is an unsigned value. If the measured value exceeds 0xFFFD, the value 0xFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFF MUST be reported.

De-jitter buffer low-water mark (DJB low-water mark): 16 bits

This is the lowest value of the de-jitter buffer nominal delay (in milliseconds) that occurred at any time during the reporting interval. This parameter MUST be provided for adaptive de-jitter buffer implementations, and its value MUST be set to DJB maximum for fixed de-jitter buffer implementations.

The measured value is an unsigned value. If the measured value exceeds 0xFFFD, the value 0xFFFE MUST be reported to indicate an over-range measurement. If the measurement is unavailable, the value 0xFFFF MUST be reported.

5. SDP Signaling

[RFC3611] defines the use of the Session Description Protocol (SDP) [RFC4566] for signaling the use of XR blocks. However, XR blocks MAY be used without prior signaling (see Section 5 of RFC 3611).

5.1. SDP rtcp-xr-attrib Attribute Extension

This section augments the SDP [RFC4566] attribute "rtcp-xr" defined in [RFC3611] by providing an additional value of "xr-format" to signal the use of the report block defined in this document.

```
xr-format =/ xr-djb-block

xr-djb-block = "de-jitter-buffer"
```

5.2. Offer/Answer Usage

When SDP is used in Offer/Answer context [RFC3264], the SDP Offer/Answer usage defined in [RFC3611] for unilateral "rtcp-xr" attribute parameters applies. For detailed usage of Offer/Answer for unilateral parameters, refer to Section 5.2 of [RFC3611].

6. IANA Considerations

New block types for RTCP XR are subject to IANA registration. For general guidelines on IANA considerations for RTCP XR, refer to [RFC3611].

6.1. New RTCP XR Block Type Value

This document assigns the block type value 23 in the IANA "RTP Control Protocol Extended Reports (RTCP XR) Block Type Registry" to the "De-Jitter Buffer Metrics Block".

6.2. New RTCP XR SDP Parameter

This document also registers a new parameter "de-jitter-buffer" in the "RTP Control Protocol Extended Reports (RTCP XR) Session Description Protocol (SDP) Parameters Registry".

6.3. Contact Information for Registrations

The contact information for registrations is:

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

7. Security Considerations

It is believed that this RTCP XR block introduces no new security considerations beyond those described in [RFC3611]. This block does not provide per-packet statistics, so the risk to confidentiality documented in Section 7, paragraph 3 of [RFC3611] does not apply.

8. Contributors

Geoff Hunt wrote the initial draft of this document.

9. Acknowledgments

The authors gratefully acknowledge reviews and feedback provided by Bruce Adams, Philip Arden, Amit Arora, Claire Bi, Bob Biskner, Benoit Claise, Kevin Connor, Claus Dahm, Spencer Dawkins, Randy Ethier, Roni Even, Jim Frauenthal, Kevin Gross, Albert Higashi, Tom Hock, Shane Holthaus, Paul Jones, Rajesh Kumar, Keith Lantz, Mohamed Mostafa, Amy Pendleton, Colin Perkins, Mike Ramalho, Ravi Raviraj, Dan Romascanu, Albrecht Schwarz, Tom Taylor, Hideaki Yamada, and Glen Zorn.

10. References

10.1. Normative References

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

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

[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.

[RFC6776] Clark, A. and Q. Wu, "Measurement Identity and Information Reporting Using a Source Description (SDES) Item and an RTCP Extended Report (XR) Block", RFC 6776, October 2012.

10.2. Informative References

[QMB] Clark, A., "RTP Control Protocol (RTCP) Extended Report (XR) Blocks for QoE Metric Reporting", Work in Progress, May 2013.

[RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.

[RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.

[RFC6709] Carpenter, B., Aboba, B., and S. Cheshire, "Design Considerations for Protocol Extensions", RFC 6709, September 2012.

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

Appendix A. Metrics Represented Using the Template from RFC 6390

a. De-Jitter Buffer Nominal Delay Metric

* Metric Name: De-jitter buffer nominal delay in RTP

* Metric Description: The "expected arrival time" is the time that an RTP packet would arrive if there was no delay variation. The delay applied to packets that arrive at their expected time is known as the Nominal Delay.

* Method of Measurement or Calculation: See Section 4.2, de-jitter buffer nominal delay definition.

* Units of Measurement: See Section 4.2, de-jitter buffer nominal delay definition.

* Measurement Point(s) with Potential Measurement Domain: See Section 4.

* Measurement Timing: See Section 4 for measurement timing and Section 4.2 for Interval Metric flag.

* Use and Applications: See Section 1.4.

* Reporting Model: See RFC 3611.

b. De-Jitter Buffer Maximum Delay Metric

* Metric Name: De-jitter buffer maximum delay in RTP.

* Metric Description: It is the current maximum de-jitter buffer delay for RTP traffic that corresponds to the earliest arriving packet that would not be discarded.

* Method of Measurement or Calculation: See Section 4.2, de-jitter buffer maximum delay definition and Section 3, the last paragraph.

* Units of Measurement: See Section 4.2, de-jitter buffer maximum delay definition.

* Measurement Point(s) with Potential Measurement Domain: See Section 4.

* Measurement Timing: See Section 4 for measurement timing and Section 4.2 for Interval Metric flag.

* Use and Applications: See Section 1.4.

* Reporting Model: See RFC 3611.

c. De-Jitter Buffer High-Water Mark Metric

* Metric Name: De-jitter buffer high-water mark in RTP.

* Metric Description: It is the highest value of the de-jitter buffer nominal delay for RTP traffic which occurred at any time during the reporting interval.

* Method of Measurement or Calculation: See Section 4.2, de-jitter buffer high-water mark definition.

* Units of Measurement: See Section 4.2, de-jitter buffer nominal delay definition.

* Measurement Point(s) with Potential Measurement Domain: See Section 4.

* Measurement Timing: See Section 4 for measurement timing and Section 4.2 for Interval Metric flag.

* Use and Applications: See Section 1.4.

* Reporting Model: See RFC 3611.

d. De-Jitter Buffer Low-Water Mark Metric

* Metric Name: De-jitter buffer low-water mark in RTP.

* Metric Description: It is the lowest value of the de-jitter buffer nominal delay (for RTP traffic) that occurred at any time during the reporting interval.

* Method of Measurement or Calculation: See Section 4.2, de-jitter buffer low-water mark definition.

* Units of Measurement: See Section 4.2, de-jitter buffer low water mark definition.

* Measurement Point(s) with Potential Measurement Domain: See Section 4, 1st paragraph.

* Measurement Timing: See Section 4 for measurement timing and Section 4.2 for Interval Metric flag.

* Use and Applications: See Section 1.4.

* Reporting Model: See RFC 3611.

Authors' Addresses

Alan Clark
Telchemy Incorporated
2905 Premiere Parkway, Suite 280
Duluth, GA 30097
USA

EMail: alan.d.clark@telchemy.com

Varun Singh
Aalto University
School of Electrical Engineering
Otakaari 5 A
Espoo, FIN 02150
Finland

EMail: varun@comnet.tkk.fi

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

EMail: sunseawq@huawei.com