

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 16, 2011

Q. Wu
Huawei
G. Zorn
Network Zen
R. Schott
Deutsche Telekom Laboratories
October 13, 2010

RTP Control Protocol Extended Reports (RTCP XR) Report Blocks for Real-time Video Quality Monitoring
[draft-wu-avt-rtcp-xr-quality-monitoring-04](#)

Abstract

This document defines a set of RTP Control Protocol Extended Reports (RTCP XR) Report Blocks and associated SDP parameters allowing the report of video quality metrics, primarily for video applications of RTP.

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 April 16, 2011.

Copyright Notice

Copyright (c) 2010 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	3
2.	Terminology	4
2.1.	Standards Language	4
2.2.	Acronyms	5
3.	Applicability	5
4.	Transport Layer Metrics	5
4.1.	RTP Flows Initial Synchronization Delay Report Block	6
4.2.	RTP Flow General Synchronization Offset Metrics Block	7
4.3.	Layered Streams Statistics Metrics Block	8
5.	Application Layer Metrics	9
5.1.	RTP Streams Statistics Summary Report Block	9
5.2.	Video Stream Loss and Discard Metrics Block	11
5.3.	Video Stream Burst Metrics Block	13
5.4.	Synthetical Multimedia Quality Metrics Block	15
6.	SDP Signaling	17
7.	IANA Considerations	19
8.	Security Considerations	20
9.	Acknowledgements	20
10.	References	20
10.1.	Normative References	20
10.2.	Informative References	21
	Authors' Addresses	22

1. Introduction

Along with the wide deployment of broadband access and the development of new IPTV services (e.g., broadcast video, video on demand), there is increasing interest in monitoring and managing networks and applications that deliver real-time applications over IP, to ensure that all end users obtain acceptable video/audio quality. The main drives come from operators, since offering performance monitoring capability can help diagnose network impairments, facilitate in root cause analysis and aid in verifying compliance with service level agreements (SLAs) between Internet Service Providers (ISPs) and content providers.

The factors that affect real-time application quality can be split into two categories. The first category consists of transport-dependent factors such as packet loss, delay and jitter (which also translates into losses in the playback buffer). The factors in the second category are application-specific factors that affect video quality and are sensitivity to network errors. These factors can be but not limited to video codec and loss recovery technique, coding bit rate, packetization scheme, and content characteristics.

Compared with application-specific factors, the transport-dependent factors sometimes are not sufficient to measure video quality, since the ability to analyze the video in the application layer provides quantifiable measurements for subscriber Quality of Experience (QoE) that may not be captured in the transmission layers or from the RTP layer down. In a typical scenario, monitoring of the transmission layers can produce statistics suggesting that quality is not an issue, such as the fact that network jitter is not excessive. However, problems may occur in the service layers leading to poor subscriber QoE. Therefore monitoring using only network-level measurements may be insufficient when application layer video quality is required.

In order to provide accurate measures of video quality for operators

when transporting video across a network, the video quality Metrics is highly required which can be conveyed in the RTCP XR packets[RFC3611] and may have the following three benefits:

- * Tuning the video encoder algorithm to satisfy video quality requirements
- * Determining which system techniques to use in a given situation and when to switch from one technique to another as system parameters change
- * Verifying the continued correct operation of an existing system

[RFC 3611](#) [RFC3611] defines seven report block formats for network

management and quality monitoring. However, there are no block types specifically designed for conveying video quality metrics. This document focuses on specifying new report block types used to convey video-specific quality metrics.

The report block types defined in this document fall into two categories. The first category consists of general information regarding transmission quality, to be generated and processed by the RTP transport. The report blocks in the second category convey metrics above transport that affect video quality and are sensitivity to network errors.

Seven report block formats are defined by this document. Of these, three are transport layer metrics:

- * RTP Flows Initial Synchronization Delay Report Block
- * Audio-Video Playout Offset Report Block
- * Layered Streams Statistics Metrics Block

The other four are application layer metrics:

- * Video Statistics Summary Report Block
- * Video Stream Loss and Discard Metrics Block
- * Video Stream Burst Metrics Block
- * Synthetical Multimedia Quality Metrics Block

[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:

Layered Component Packet

a RTP packet using layered codecs containing the specified layered component, e.g., encoded stream at the base layer or at the enhancement layer.

Picture Type

Picture types used in the different video algorithms compose of the key-frame and the Derivation frame. Key-frame is also called as reference frame and used as a reference for predicting other

pictures. It is coded without prediction from other pictures. The Derivation frame is derived from Key-frame using prediction from the reference frame.

[2.2.](#) Acronyms

SSRC

Synchronization Source [[RFC3550](#)]

TS

Transport Stream [[ISO-IEC.13818-1.2007](#)]

[3.](#) Applicability

All the report blocks defined in this document could be used by dedicated network monitoring applications. As specified in [RFC 3611](#) [[RFC3611](#)], for such an application it might be appropriate to allow more than 5% of RTP data bandwidth to be used for RTCP packets, thus allowing proportionately larger and more detailed report blocks.

The Flows General Synchronization Offset Block [Section 4.2](#) has been defined for various multimedia applications. Such applications can use this report block to monitor offset between two RTP streams synchronization to ensure satisfactory QoE. Tighter tolerances than typically used have been recommended for such applications.

The Flows Synchronization Delay Report Block has been defined primarily for layered or multi-description video coding applications. When joining a layered video session in such an application, a receiver may not synchronize playout across the multimedia session until RTCP SR packets have been received on all of the component RTP sessions. This report block can be used to ensure synchronization between different media layers for the same multimedia session.

The Video Stream Loss and Discard Metrics Report Block, Video Stream Burst Metrics report Block, Video Statistics Summary Report Block and Layered Video Statistics Metrics Block can be applied to any real time video application, while Synthetical Multimedia Quality Metrics Report Block can be used in any real-time AV application .

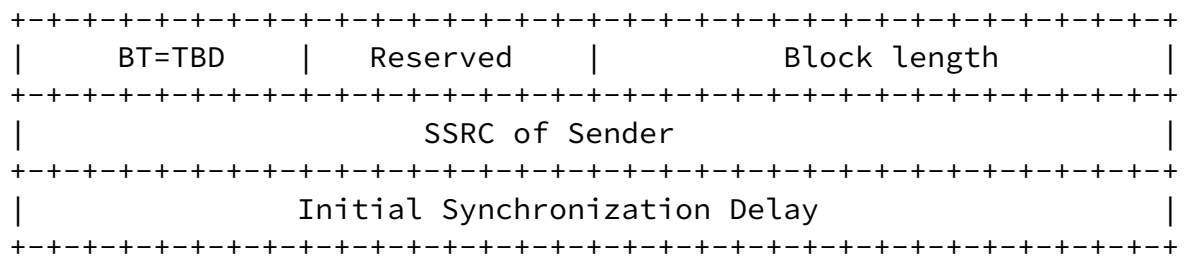
[4.](#) Transport Layer Metrics

[4.1.](#) RTP Flows Initial Synchronization Delay Report Block

This block reports the initial synchronization delay between RTP sessions of the same media stream sent using Multi-Session Transmission [[I-D.ietf-avt-rtp-svc](#)] or the initial synchronization delay between RTP session of the different media types [[I-D.ietf-avt-rapid-rtp-sync](#)], which is beyond the information carried in the standard RTCP packet format. Information is recorded about session bandwidth and synchronization delay.

The RTP Flows Intial 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 2 3 4 5 6 7 8 9	0 1		



Block type (BT): 8 bits

The Statistics Summary 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 3, in accordance with the definition of this field in [Section 3 of RFC 3611](#) [[RFC3611](#)].

SSRC of Sender: 32 bits

The SSRC of the RTP data packet source being reported upon by this report block. ([Section 4.1 of \[RFC3611\]](#)).

Initial Synchronization Delay: 32 bits

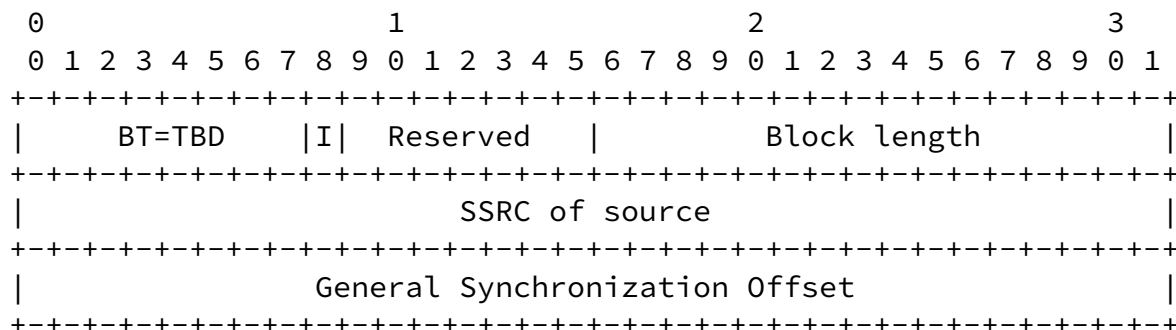
The average delay, expressed in units of 1/65536 seconds, between the RTCP packets received on all of the components RTP sessions and the beginning of session [[I-D.ietf-avt-rapid-rtp-sync](#)]. The value is calculated as follows:

The average time, expressed in units of 1/65536 seconds, taken to receive the first RTCP packet in the RTP session with the longest RTCP reporting interval [[I-D.ietf-avt-rapid-rtp-sync](#)]

[4.2.](#) RTP Flow General Synchronization Offset Metrics Block

In an RTP multimedia session, there can be an arbitrary number of streams, with the same RTCP CNAME. This block reports the general

Synchronization offset requirements of these RTP streams beyond the information carried in the standard RTCP packet format. Information is recorded about the synchronization offset time of each RTP stream relative to the reference RTP stream with the same CNAME and General Synchronisation Offset of zero.. The RTP Flow General Synchronization Offset Report Block has the following format:



Block type (BT): 8 bits

The Statistics Summary Report Block is identified by the constant <AVPO>.

Interval Metric flag (I): 1 bit

This field is used to indicate whether the Audio-Video synchronization metrics are Interval or Cumulative metrics, that is, whether the reported values applies to the most recent measurement interval duration between successive metrics reports (I=1) (the Interval Duration) or to the accumulation period characteristic of cumulative measurements (I=0) (the Cumulative Duration).

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.

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

SSRC of source: 32 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

General synchronization offset: 32 bits

This field indicates the synchronization offset time of one RTP stream in milliseconds relative to the reference RTP stream with the same CNAME and General Synchronisation Offset of zero [I-D.ietf-avt-rapid-rtp-sync] This value is calculated based on the interarrival time between arbitray RTP packet and the reference RTP packet with the same CNAME , and timestamps of this arbitray RTP packet and the reference RTP packet with the same CNAME.

4.3. Layered Streams Statistics Metrics Block

This block reports layered streams statistics beyond the information carried in the Statistics Summary Report Block RTCP packet specified in the [section 4.6 of RFC 3611](#) [RFC3611]. Information is recorded about lost layered component packets, duplicated layered component packets. Such information can be useful for network management and video quality monitoring.

The report block contents are dependent upon a series of flag bits carried in the first part of the header. Not all parameters need to be reported in each block. Flags indicate which parameters are reported and which are not. The fields corresponding to unreported parameters MUST be present, but are set to zero. The receiver MUST ignore any Layered Streams Statistics Metrics Block with a non-zero value in any field flagged as unreported.

The Layered Stream Statistics metrics 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	2	3	4	5	6	7	8	9																								
BT=TBD																T																rsd.																block length															
SSRC of source																																																															
begin_seq																																end_seq																															
Lost_Layered Component Packets																																																															
Dup Layered Component_Packets																																																															

[illegible]

Block type (BT): 8 bits

The Layered stream Statistics Metrics Block is identified by the constant <LSSM>.

Layer Type flag (T): 1 bits

This field is used to indicate the Layer Type of layered video to be reported. LT is set to 0 if the loss_component_packet field and dup_component packet contain the base layer packet in layered codecs, e.g, SVC in [[I-D.ietf-avt-rtp-svc](#)], 1 if the loss_component packet field and dup_component packet contain enhancement layer packet in layered codec.

Rsd.: 3 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

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

begin_seq: 16 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

end seq: 16 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

Lost_Layered Component Packets: 32 bits

Number of lost_component packets in the above sequence number interval.

Dup_Layered Component Packets: 32 bits

Number of dup_component packets in the above sequence number interval.

5. Application Layer Metrics

5.1. RTP Streams Statistics Summary Report Block

This block reports statistics beyond the information carried in the Statistics Summary Report Block RTCP packet specified in the section

4.6 of [RFC 3611](#) [[RFC3611](#)]. Information is recorded about lost frame packets, duplicated frame packets, lost layered component packets, duplicated layered component packets. Such information can be useful for network management and video quality monitoring.

The report block contents are dependent upon a series of flag bits carried in the first part of the header. Not all parameters need to be reported in each block. Flags indicate which parameters are reported and which are not. The fields corresponding to unreported parameters MUST be present, but are set to zero. The receiver MUST ignore any Video Statistics Summary Report Block with a non-zero value in any field flagged as unreported.

The RTP Streams Statistics Summary 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=TBD      |T|P|      rsd.  |      block length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SSRC of source          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      begin_seq      |      end_seq      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     lost_frames              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     dup frames                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     lost_partial frame packets |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     dup partial frame_packets  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Block type (BT): 8 bits

The Video Statistics Summary Report Block is identified by the constant <VSS>.

Picture type indicator (T): 1 bits

Picture types used in the different video algorithms compose of key-frame and derivation frame. This field is used to indicate the frame type to be reported. Bits set to 0 if the lost_frames field or dup_frames field contain a key_frame report or reference frame report, 1 if the lost_frames field and dup_frames field contain other derivation frame report.

P: 1 bit

Bit set to 1 if the lost_partial frame packets field or the dup_partial_frame packets field contains a report, 0 otherwise.

Rsd.: 3 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 5, in accordance with the definition of this field in [Section 3 of RFC 3611](#) [RFC3611].

SSRC of source: 32 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

begin_seq: 16 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

end_seq: 16 bits

As defined in [Section 4.1 of RFC 3611](#) [RFC3611].

lost_frames: 32 bits

Number of lost_frames in the above sequence number interval.

dup_frames: 32 bits

Number of dup_frames in the above sequence number interval.

lost_partial frame packets: 32 bits

Number of lost_partial frame packets in the above sequence number interval.

the picture type to be reported. Bits set to 0 if the Loss rate field and discard rate field contain a Key_frame report or reference frame report, 1 if the Loss rate field and discard rate field contain other derivation frame reports.

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 1, in accordance with the definition of this field in [Section 3 of RFC 3611](#) [[RFC3611](#)].

SSRC of source: 32 bits

The SSRC of the RTP data packet source being reported upon by this report block. in accordance with the definition of this field in [Section 3 of RFC 3611](#) [[RFC3611](#)].

Loss rate: 8 bits

The fraction of RTP data packets from the source lost since the beginning of reception, expressed as a fixed point number with the binary point at the left edge of the field. This value is calculated by dividing the total number of lost packets containing

specified frame (e.g., Key frame) (after the effects of applying any error protection such as FEC) by the total number of packets expected, multiplying the result of the division by 256, limiting the maximum value to 255 (to avoid overflow), and taking the integer part. The numbers of duplicated packets and discarded packets do not enter into this calculation. Since receivers cannot be required to maintain unlimited buffers, a receiver MAY categorize late-arriving packets as lost. The degree of lateness that triggers a loss SHOULD be significantly greater than that which triggers a discard.

Discard rate: 8 bits

The fraction of RTP data packets from the source that have been discarded since the beginning of reception, due to late or early arrival, under-run or overflow at the receiving jitter buffer. This value is expressed as a fixed point number with the binary point at the left edge of the field. It is calculated by dividing the total number of discarded packets containing specified frame

(e.g., Key Frame) (excluding duplicate packet discards) by the total number of packets expected, multiplying the result of the division by 256, limiting the maximum value to 255 (to avoid overflow), and taking the integer part.

5.3. Video Stream Burst Metrics Block

This block reports Burst metrics statistics beyond the information carried in the standard RTCP packet format. It reports on the combined effect of losses and discards, as both have equal effect on video quality.

In order to properly assess the quality of a video stream, it is desirable to consider the degree of burstiness of packet loss [RFC 3357](#) [RFC3357]. Following the one-way loss pattern sample metrics discussed in [RFC3357], a measure of the spacing between consecutive network packet loss or error events, is a "loss distance". The loss distance metric captures the spacing between the loss periods. The duration of a loss or error event (e.g. and how many packets are lost in that duration) is a "loss period", the loss period metric captures the frequency and length (burstiness) of loss once it starts. Delay reports include the transit delay between RTP end points and the end system processing delays, both of which contribute to the user perceived delay.

Implementations **MUST** provide values for all the fields defined here. For certain metrics, if the value is undefined or unknown, then the specified default or unknown field value **MUST** be provided.

Max Loss Duration	Reserved.
-------------------	-----------

block type (BT): 8 bits

A Video Stream Metrics Report Block is identified by the constant <VSBM>.

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 RTP data packet source being reported upon by this report block. in accordance with the definition of this field in [Section 3 of RFC 3611](#) [RFC3611].

Loss Distance: 16 bits

The mean duration, expressed in milliseconds, of the loss intervals that have occurred since the beginning of reception [DSLF]. The duration of each loss distance is calculated based upon the frames that mark the beginning and end of that period. It is equal to the timestamp of the end frame, plus the duration of the end frame, minus the timestamp of the beginning frame. If the actual values are not available, estimated values MUST be used. If there have been no burst periods, the burst duration value MUST be zero.

Loss Period: 16 bits

The mean duration, expressed in milliseconds, of the burst loss periods that have occurred since the beginning of reception [DSLF]. The duration of each period is calculated based upon the frame that marks the end of the prior burst loss and the frame

that marks the beginning of the subsequent burst loss. It is equal to the timestamp of the subsequent burst frame, minus the timestamp of the prior burst packet, plus the duration of the prior burst packet. If the actual values are not available,

estimated values MUST be used. In the case of a gap that occurs at the beginning of reception, the sum of the timestamp of the prior burst packet and the duration of the prior burst packet are replaced by the reception start time. In the case of a gap that occurs at the end of reception, the timestamp of the subsequent burst packet is replaced by the reception end time. If there have been no gap periods, the gap duration value MUST be zero.

Max Loss Duration of a single error: 16 bits

The maximum loss duration, expressed in milliseconds, of the loss periods that have occurred since the beginning of reception. The recommended max loss duration is specified as less than 16 ms in [\[DSLIF\]](#), which provides a balance between interleaver depth protection from xDSL errors induced by impulse noise, delay added to other applications and video service QoE requirements to reduce visible impairments.

Reserved: 16 bits

All bits SHALL be set to 0 by the sender and SHALL be ignored on reception.

block length: 16 bits

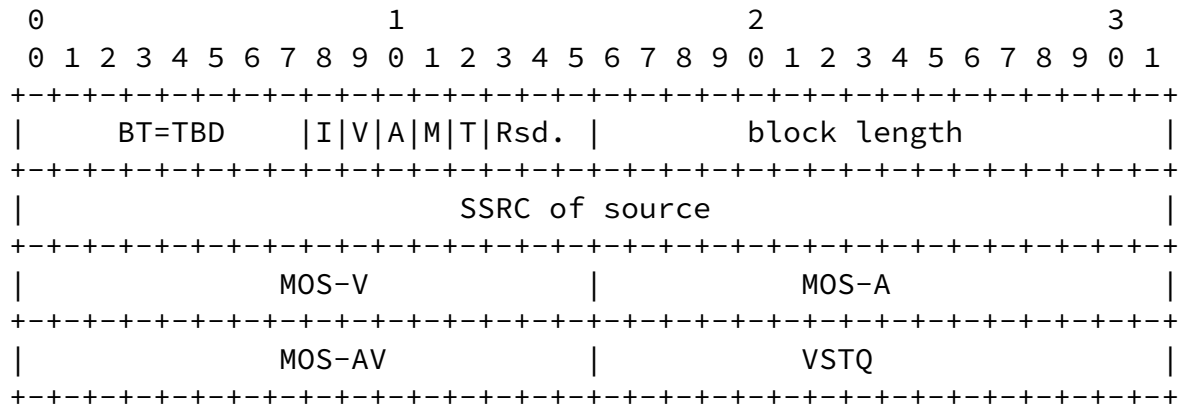
The constant 2, in accordance with the definition of this field in [Section 3 of RFC 3611](#) [\[RFC3611\]](#).

[5.4.](#) Synthetical Multimedia Quality Metrics Block

This block reports the multimedia quality metrics beyond the information carried in the standard RTCP packet format. Information is recorded about Video MOS, Audio MOS, Audio Video MOS, Video Service Transmission Quality [\[G.1082\]](#) [\[P.NAMS\]](#).

The report block contents are dependent upon a series of flag bits carried in the first part of the header. Not all parameters need to be reported in each block. Flags indicate which are and which are not reported. The fields corresponding to unreported parameters MUST be present, but are set to zero. The receiver MUST ignore any Perceptual Quality Metrics Block with a non-zero value in any field flagged as unreported.

The Synthetical Multimedia Quality Metrics Block has the following format:



Block type (BT): 8 bits

The Perceptual Quality Metrics Block is identified by the constant <SMQM>.

Interval Metric flag (I): 1 bit

This field is used to indicate whether the Basic Loss/Discard metrics are Interval or Cumulative metrics, that is, whether the reported values applies to the most recent measurement interval duration between successive metrics reports (I=1) (the Interval Duration) or to the accumulation period characteristic of cumulative measurements (I=0) (the Cumulative Duration).

MOS-V flag (V): 1 bit

Bit set to 1 if the MOS-V field and MOS-AV field contain a report, 0 otherwise.

MOS-A flag (A): 1 bit

Bit set to 1 if the MOS-A field contain a report, 0 otherwise.

MOS-AV flag (M): 1 bit

Bit set to 1 if the MOS-AV field contain a report, 0 otherwise.

Video Service Transmission Quality flag (T): 1 bit

Bit set to 1 if the VSTQ field contains a report, 0 otherwise.

Rsd.: 3 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.

SSRC of source: 32 bits

As defined in [Section 4.1 of \[RFC3611\]](#).

MOS-V: 16 bits

The estimated mean opinion score for video quality (MOS-V) is a video quality metric on a scale from 1 to 5, in which 5 represents excellent and 1 represents unacceptable [[G.1082](#)][P.NAMS]. This metric is defined as not including the effects of audio impairments and can be compared to MOS scores obtained from video quality tests. It is expressed as an integer in the range 10 to 50, corresponding to MOS x 10. For example, a value of 35 would correspond to an estimated MOS score of 3.5.

A value of 127 indicates that this parameter is unavailable. Values other than 127 and the valid range defined above MUST NOT be sent and MUST be ignored by the receiving system.

MOS-A: 16 bits

The estimated mean opinion score for Audio quality (MOS-A) is defined as including the effects of delay and other effects that would affect Audio-Video quality [[RFC3611](#)]. It is expressed as an integer in the range 10 to 50, corresponding to MOS x 10, as for MOS-A.

A value of 127 indicates that this parameter is unavailable. Values other than 127 and the valid range defined above MUST not be sent and MUST be ignored by the receiving system.

MOS-AV: 16 bits

The estimated mean opinion score for Audio-Video quality (MOS-AV) is defined as including the effects of delay and other effects that would affect Audio-Video quality [[G.1082](#)][P.NAMS]. It is expressed as an integer in the range 10 to 50, corresponding to MOS x 10, as for MOS-AV. A value of 127 indicates that this parameter is unavailable. Values other than 127 and the valid range defined above MUST NOT be sent and MUST be ignored by the receiving system.

VSTQ: 16 bits

Video Service Transmission Quality (TBC) .

6. SDP Signaling

Six new parameters are defined for the six 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]:

Wu, et al.

Expires April 16, 2011

[Page 17]

Internet-Draft

RTCP XR Video Quality Report Blocks

October 2010

```
rtcp-xr-attrib = "a=rtcp-xr:"
                  [xr-format *(SP xr-format)] CRLF

xr-format = RTP-flows-syn
           / audio-video-offset
           / multimedia-quality-metrics
           / video-stream-loss-metrics
           / video-stream-burst-metrics
           / video-stat-summary
           / layered-video-stat-metrics

RTP-flows-syn = "RTP-flows-syn"
               ["=" max-size]
               max-size = 1*DIGIT ; maximum block size in octets

audio-video-offset = "audio-video-offset"
                   ["=" max-size]
                   max-size = 1*DIGIT ; maximum block size in octets

video-stream-burst-metrics = "video-stream-burst-metrics"
                           ["=" max-size]
                           max-size = 1*DIGIT ; maximum block size in octets

video-stream-loss-metrics = "video-stream-loss-metrics"
                           ["=" stat-flag *(", " stat-flag)]
                           stat-flag = "key Frame loss and duplication"
                                      / "derivation Frame loss and duplication"

video-stat-summary = "video-stat-summary"
                   ["=" stat-flag *(", " stat-flag)]
                   stat-flag = "key Frame loss and duplication"
                              / "derivation Frame loss and duplication"
```

```

layered-stream-stat-metrics = "layered-stream-stat-metrics"
                                ["=" stat-flag *(", " stat-flag)]
stat-flag = "base layer packet"
            / "enhancement layer packet"

multimedia-quality-metrics = "multimedia-quality-metrics"
                             ["=" stat-flag *(", " stat-flag)]
stat-flag = "Interval Metric"
            / "MOS-V"
            / "MOS-A"
            / "MOS-AV"
            / "VSTQ"

```

Refer to [Section 5.1 of RFC 3611](#) [[RFC3611](#)] for a detailed description

and the full syntax of the "rtcp-xr" attribute.

[7.](#) 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 six new block type values in the RTCP XR Block Type Registry:

Name: RFISD
Long Name: RTP Flows Initial Synchronization Delay
Value: <RFISD>
Reference: [Section 4.1](#)

Name: AVPO
Long Name: Audio-Video Playout Offset
Value: <AVPO>
Reference: [Section 4.2](#)

Name: VSS
Long Name: Video Statistics Summary
Value: <VSS>
Reference: [Section 5.1](#)

Name: LSSM
Value <LSSM>
Long Name: Layered Stream Statistics Metrics
Reference: [Section 4.3](#)

Name: VSLDM
Long Name: Video Stream Loss and Discard Metrics
Value <VSLDM>
Reference: [Section 5.2](#)

Name: VSBM
Long Name: Video Stream Burst Metrics
Value <VSBM>
Reference: [Section 5.3](#)

Name: SMQM
Long Name: Synthetical Multimedia Quality Metric

Value <SMQM>
Reference: [Section 5.4](#)

This document also registers seven SDP [[RFC4566](#)] parameters for the "rtcp-xr" attribute in the RTCP XR SDP Parameters Registry:

- * "RTP-flows-syn"
- * "audio-video-offset"
- * "multimedia-quality-metrics"
- * "video-stream-loss-metrics"
- * "video-stream-burst-metrics"
- * "video-stat-summary"
- * "layered-stream-stat-metrics"

The contact information for the registrations is:

Qin Wu
sunseawq@huawei.com
101 Software Avenue, Yuhua District

8. Security Considerations

TBC

9. Acknowledgements

The authors would like to thank Bill Ver Steeg, David R Oran, Ali Begen, Colin Perkins, Roni Even, Youqing Yang, Wenxiao Yu and Yinliang Hu for their valuable comments and suggestions on this document.

10. References

10.1. Normative References

[I-D.ietf-avt-rapid-rtp-sync]

Perkins, C. and T. Schierl, "Rapid Synchronisation of RTP Flows", [draft-ietf-avt-rapid-rtp-sync-12](#) (work in progress), July 2010.

[I-D.ietf-avt-rtp-svc]

Wenger, S., Wang, Y., Schierl, T., and A. Eleftheriadis, "RTP Payload Format for Scalable Video Coding", [draft-ietf-avt-rtp-svc-23](#) (work in progress), October 2010.

Wu, et al.

Expires April 16, 2011

[Page 20]

Internet-Draft

RTCP XR Video Quality Report Blocks

October 2010

[ISO-IEC.13818-1.2007]

International Organization for Standardization, "Information technology - Generic coding of moving pictures and associated audio information: Systems", ISO International Standard 13818-1, October 2007.

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

[RFC2250] Hoffman, D., Fernando, G., Goyal, V., and M. Civanlar, "RTP Payload Format for MPEG1/MPEG2 Video", [RFC 2250](#), January 1998.

- [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.

[10.2.](#) Informative References

- [DSLIF] Rahrer, T., Ed., Fiandra, Ed., and Wright, Ed., "Triple-play Services Quality of Experience (QoE) Requirements", DSL Forum Technical Report TR-126, December 2006.
- [G.1082] ITU-T, "Measurement-based methods for improving the robustness of IPTV performance", ITU-T Recommendation G.1082, April 2009.
- [I-D.ietf-fecframe-interleaved-fec-scheme] Begen, A., "RTP Payload Format for 1-D Interleaved Parity FEC", [draft-ietf-fecframe-interleaved-fec-scheme-09](#) (work in progress), January 2010.
- [I-D.ietf-fecframe-raptor] Waston, M., "Raptor FEC Schemes for FECFRAME", [draft-ietf-fecframe-raptor-02](#) (work in progress), March 2010.
- [IEEE] IEEE, "Human Perception of Jitter and Media

Synchronization", IEEE Journal on Selected Areas in Communications Vol. 14, No. 1, January 1996.

- [P.NAMS] ITU-T, "Non-intrusive parametric model for the Assessment of performance of Multimedia Streaming", ITU-T

Recommendation P.NAMS, November 2009.

[RFC3357] Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", [RFC 3357](#), August 2002.

Authors' Addresses

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

Email: sunseawq@huawei.com

Glen Zorn
Network Zen
77/440 Soi Phoomjit, Rama IV Road
Phra Khanong, Khlong Toie
Bangkok 10110
Thailand

Phone: +66 (0) 87 502 4274
Email: gwz@net-zen.net

Roland Schott
Deutsche Telekom Laboratories
Deutsche-Telekom-Allee 7
Darmstadt 64295
Germany

Email: Roland.Schott@telekom.de