AVTCORE WG
Internet-Draft
Updates: 3550 (if approved)
Intended status: Standards Track
Expires: January 12, 2014

J. Lennox
Vidyo
M. Westerlund
Ericsson
Q. Wu
Huawei
C. Perkins
University of Glasgow
July 11, 2013

Sending Multiple Media Streams in a Single RTP Session: Grouping RTCP Reception Statistics and Other Feedback
draft-ietf-avtcore-rtp-multi-stream-optimisation-00

Abstract

RTP allows multiple media streams to be sent in a single session, but requires each Synchronisation Source (SSRC) to send RTCP reception quality reports for every other SSRC visible in the session. This causes the number of RTCP reception reports to grow with the number of SSRCs, rather than the number of endpoints. In many cases most of these RTCP reception reports are unnecessary, since all SSRCs of an endpoint are co-located and see the same reception quality. This memo defines a Reporting Group extension to RTCP to reduce the reporting overhead in such scenarios.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Real-time Transport Protocol (RTP) [RFC3550] is a protocol for group communication, supporting multiparty multimedia sessions. A single RTP session can support multiple participants sending at once, and can also support participants sending multiple simultaneous media streams. Examples of the latter might include a participant with multiple cameras who chooses to send multiple views of a scene, or a participant that sends audio and video flows multiplexed in a single RTP session. Rules for handling RTP sessions containing multiple media streams are described in [RFC3550] with some clarifications in [I-D.ietf-avtcore-rtp-multi-stream].

An RTP endpoint will have one or more synchronisation sources (SSRCs) that send and receive media streams (it will have one SSRC for each media stream it sends). Each SSRC has to send RTCP sender reports corresponding to the RTP packets it sends, and receiver reports for traffic it receives. That is, every SSRC will send RTCP packets to report on every other SSRC. This rule is simple, but can be quite inefficient for endpoints that send large numbers of media streams in a single RTP session. Consider a session comprising ten participants, each sending three media streams with their own SSRC. There will be 30 SSRCs in such an RTP session, and 30 RTCP reception reports will be sent per reporting interval as each SSRC reports on all the others. However, the three SSRCs comprising each participant will almost certainly see identical reception quality, since they are co-located. If there was a way to indicate that several SSRCs are co-located, and see the same reception quality, then two-thirds of those RTCP reports could be suppressed.

This memo defines such an RTCP extension, Reporting Groups. This extension is used to indicate the SSRCs that originate from the same endpoint, and therefore have identical reception quality, allowing the endpoint to suppress unnecessary RTCP reception reports.

2. Terminology

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

3. Grouping of RTCP Reception Statistics and Other Feedback

3.1. Semantics and Behavior of Reporting Groups

An RTCP Reporting Group indicates that a set of sources (SSRCs) that originate from a single entity (endpoint or middlebox) in an RTP session, and therefore all the sources in the group have an identical view of the network. If reporting groups are in use, two sources SHOULD be put into the same reporting group if their view of the network is identical; i.e., if they report on traffic received at the same interface of an RTP endpoint. Sources with different views of the network MUST NOT be put into the same reporting group.

If reporting groups are in use, an endpoint MUST NOT send reception reports from one source in a reporting group about another one in the same group ("self-reports"). Similarly, an endpoint MUST NOT send reception reports about a remote media source from more than one source in a reporting group ("cross-reports"). Instead, it MUST pick one of its local media sources as the "reporting" source for each remote media source, and use it to send reception reports about that remote source; all the other media sources in the reporting group MUST NOT send any reception reports about that remote media source.

This reporting source MUST also be the source for any RTP/AVPF Feedback [RFC4585] or Extended Report (XR) [RFC3611] packets about the corresponding remote sources as well. If a reporting source leaves the session (i.e., if it sends a BYE, or leaves the group without sending BYE under the rules of [RFC3550] section 6.3.7), another reporting source MUST be chosen if any members of the group still exist.

An endpoint or middlebox MAY use multiple sources as reporting sources; however, each reporting source MUST have non-overlapping sets of remote SSRCs it reports on. This is primarily to be done when the reporting source's number of reception report blocks is so large that it would be forced to round-robin around the sources. Thus, by splitting the reports among several reporting SSRCs, more consistent reporting can be achieved.

If RTP/AVPF feedback is in use, a reporting source MAY send immediate or early feedback at any point when any member of the reporting group could validly do so.

An endpoint SHOULD NOT create single-source reporting groups, unless it is anticipated that the group might have additional sources added to it in the future.

3.2. Determine the Report Group

A remote RTP entity, such as an endpoint or a middlebox needs to be able to determine the report group used by another RTP entity. To achieve this goal two RTCP extensions have been defined. For the SSRCs that are reporting on behalf of the reporting group, an SDES item RGRP has been defined for providing the report group with an identifier. For SSRCs that aren't reporting on any peer SSRC a new RTCP packet type is defined. This RTCP packet type "Reporting Sources" lists the SSRC that are reporting on this SSRC's behalf.

This divided approach has been selected for the following reasons:

1. To enable an explicit indication of who reports on this SSRC's behalf. Being explicit prevents the remote entity from detecting that is missing the reports if there issues with the reporting SSRC's RTCP packets.

2. To enable explicit identification of the SSRCs that are actively reporting as one entity.

3.3. RTCP Packet Reporting Group's Reporting Sources

This section defines a new RTCP packet type called "Reporting Group's Reporting Sources" (RGRS). It identifies the SSRC(s) that report on behalf of the SSRC that is the sender of the RGRS packet.

This packet consists of the fixed RTCP packet header which indicates the packet type, the number of reporting sources included and the SSRC which behalf the reporting SSRCs report on. This is followed by the list of reporting SSRCs.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|    SC   | PT=RGRS(TBA)  |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      SSRC of packet sender                    |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
:                      SSRC for Reporting Source                :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The RTCP Packets field has the following definition.

version (V): This field identifies the RTP version. The current version is 2.

padding (P): 1 bit If set, the padding bit indicates that the packet contains additional padding octets at the end that are not part of the control information but are included in the length field. See [RFC3550].

Source Count (SC): 5 bits Indicating the number of reporting SSRCs (1-31) that are included in this RTCP packet type.

Payload type (PT): 8 bits This is the RTCP packet type that identifies the packet as being an RTCP FB message. The RGRS RTCP packet has the value [TBA].

Length: 16 bits The length of this packet in 32-bit words minus one, including the header and any padding. This is in line with the definition of the length field used in RTCP sender and receiver reports [RFC3550].

SSRC of packet sender: 32 bits. The SSRC of the sender of this packet which indicates which SSRCs that reports on its behalf, instead of reporting itself.

SSRC for Reporting Source: A variable number (as indicated by Source Count) of 32-bit SSRC values. Each SSRC is an reporting SSRC belonging to the same Report Group.

Each RGRS packet MUST contain at least one reporting SSRC. In case the reporting SSRC field is insufficient to list all the SSRCs that are reporting in this report group, the SSRC SHALL round robin around the reporting sources.

Any RTP mixer or translator which forwards SR or RR packets from members of a reporting group MUST forward the corresponding RGRS RTCP packet as well.

3.4. RTCP Source Description (SDES) item for Reporting Groups

A new RTCP Source Description (SDES) item is defined for the purpose of identifying reporting groups.

The Source Description (SDES) item "RGRP" is sent by a reporting group's reporting SSRC. Syntactically, its format is the same as the RTCP SDES CNAME item [RFC6222], and MUST be chosen with the same global-uniqueness and privacy considerations as CNAME. This name MUST be stable across the lifetime of the reporting group, even if the SSRC of a reporting source changes.

Every source which belongs to a reporting group MUST either include an RGRP SDES item in an SDES packet (if it is a reporting source), or an RGRS packet (if it is not), in every compound RTCP packet in which it sends an RR or SR packet (i.e., in every RTCP packet it sends, unless Reduced-Size RTCP [RFC5506] is in use).

Any RTP mixer or translator which forwards SR or RR packets from members of a reporting group MUST forward the corresponding RGRP SDES items as well, even if it otherwise strips SDES items other than CNAME.

3.5. Middlebox Considerations

This section discusses middlebox considerations for Reporting groups.

To be expanded.

3.6. SDP signaling for Reporting Groups

TBD

3.7. Bandwidth Benefits of RTCP Reporting Groups

To understand the benefits of RTCP reporting groups, consider a scenario in which the two endpoints in a session each have a hundred sources, of which eight each are sending within any given reporting interval.

For ease of analysis, we can make the simplifying approximation that the duration of the RTCP reporting interval is equal to the total size of the RTCP packets sent during an RTCP interval, divided by the RTCP bandwidth. (This will be approximately true in scenarios where the bandwidth is not so high that the minimum RTCP interval is reached.) For further simplification, we can assume RTCP senders are following the recommendations regarding Compound RTCP Packets in [I-D.ietf-avtcore-rtp-multi-stream]; thus, the per-packet transport-layer overhead will be small relative to the RTCP data. Thus, only the actual RTCP data itself need be considered.

In a report interval in this scenario, there will, as a baseline, be 200 SDES packets, 184 RR packets, and 16 SR packets. This amounts to approximately 6.5 kB of RTCP per report interval, assuming 16-byte CNAMEs and no other SDES information.

Using the original [RFC3550] everyone-reports-on-every-sender feedback rules, each of the 184 receivers will send 16 report blocks, and each of the 16 senders will send 15. This amounts to approximately 76 kB of report block traffic per interval; 92% of RTCP traffic consists of report blocks.

If reporting groups are used, however, there is only 0.4 kB of reports per interval, with no loss of useful information. Additionally, there will be (assuming 16-byte RGRPs, and a single reporting source per reporting group) an additional 2.4 kB per cycle of RGRP SDES items and RGRS packets. Put another way, the unmodified [RFC3550] reporting interval is approximately 8.9 times longer than if reporting groups are in use.

3.8. Consequences of RTCP Reporting Groups

The RTCP traffic generated by receivers using RTCP Reporting Groups might appear, to observers unaware of these semantics, to be generated by receivers who are experiencing a network disconnection, as the non-reporting sources appear not to be receiving a given sender at all.

This could be a potentially critical problem for such a sender using RTCP for congestion control, as such a sender might think that it is sending so much traffic that it is causing complete congestion collapse.

However, such an interpretation of the session statistics would require a fairly sophisticated RTCP analysis. Any receiver of RTCP statistics which is just interested in information about itself needs to be prepared that any given reception report might not contain information about a specific media source, because reception reports in large conferences can be round-robined.

Thus, it is unclear to what extent such backward compatibility issues would actually cause trouble in practice.

4. Security Considerations

The security considerations of [RFC3550] and [I-D.ietf-avtcore-rtp-multi-stream] apply.

(tbd: any security considerations due to these extensions?)

5. IANA Considerations

(Note to the RFC-Editor: please replace "TBA" with the IANA-assigned value, and "XXXX" with the number of this document, prior to publication as an RFC.)

The IANA is requested to register one new RTCP SDES items in the "RTCP SDES Item Types" registry, as follows:

Value	Abbrev	Name	Reference
TBA	RGRP	Reporting Group	[RFCXXXX]

Figure 1: Item for the IANA Source Attribute Registry

The IANA is also requested to register one new RTCP packet type as follows:

Value	Abbrev	Name	Reference
TBA	RGRR	Reporting Group Reporting Sources	[RFCXXXX]

Figure 2: Item for the IANA RTCP Control Packet Types (PT) Registry

6. References

6.1. 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.

[RFC6222] Begen, A., Perkins, C., and D. Wing, "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)", RFC 6222, April 2011.

6.2. Informative References

[I-D.ietf-avtcore-rtp-multi-stream]
Lennox, J., Westerlund, M., Wu, W., and C. Perkins, "RTP Considerations for Endpoints Sending Multiple Media Streams", draft-ietf-avtcore-rtp-multi-stream-00 (work in progress), April 2013.

[RFC3611] Friedman, T., Caceres, R., and A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.

[RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July 2006.

[RFC5506] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", RFC 5506, April 2009.

Authors' Addresses

Jonathan Lennox
Vidyo, Inc.
433 Hackensack Avenue
Seventh Floor
Hackensack, NJ 07601
US

Email: jonathan@vidyo.com

Magnus Westerlund
Ericsson
Farogatan 6
SE-164 80 Kista
Sweden

Phone: +46 10 714 82 87
Email: magnus.westerlund@ericsson.com

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

Email: sunseawq@huawei.com

Colin Perkins
University of Glasgow
School of Computing Science
Glasgow G12 8QQ
United Kingdom

Email: csp@csperkins.org