

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 2, 2008

I. Johansson
M. Westerlund
Ericsson AB
Oct 30, 2007

Support for non-compound RTCP, opportunities and consequences
draft-ietf-avt-rtcp-non-compound

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 2, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This memo discusses benefits and issues that arise when allowing RTCP packets to be transmitted as non-compound packets, i.e. not follow the rules of [RFC 3550](#). Based on that analysis this memo proposes changes to the rules to allow feedback messages to be sent as non-compound RTCP packets when using the RTP AVPF profile ([RFC 4585](#)) under certain conditions.

Requirements 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

Table of Contents

1.	Introduction	3
2.	RTCP Compound Packets	3
3.	Benefits from non-compound packets	4
4.	Issues with non-compound RTCP packets	6
4.1.	Middle boxes	6
4.2.	Packet Validation	6
4.2.1.	Old RTCP Receivers	6
4.2.2.	Weakened Packet Validation	7
4.2.3.	Bandwidth consideration	7
4.2.4.	Computation of avg_rtcp_size	7
4.3.	Header compression	7
4.4.	RTP and RTCP multiplex on the same port	7
5.	Allowing non-compound RTCP packets	8
5.1.	Usage of non-compound packets in AVPF	8
5.1.1.	Verifying the delivery of non-compound packets	9
5.2.	SDP Signalling Attribute	9
6.	IANA Considerations	9
7.	Security Considerations	10
8.	Acknowledgements	10
9.	References	10
9.1.	Normative References	10
9.2.	Informative References	10
	Authors' Addresses	11
	Intellectual Property and Copyright Statements	13

1. Introduction

In RTP [[RFC3550](#)] it is currently mandatory to always use RTCP compound packets containing at least Sender Reports or Receiver reports, and a SDES packet containing at least the CNAME item. There are good reasons for this as discussed below (see [Section 2](#)). However this do result in that the minimal RTCP packets are quite large. The RTP profile AVPF [[RFC4585](#)] specifies new RTCP packet types for feedback messages. Some of these feedback messages would benefit from being possible to transmit with minimal delay and AVPF do provide some mechanism to enable this. However for environments with low-bitrate links this still consumes quite large amount of resources and introduce extra delay in the time it takes to completely send the compound packet in the network. There are also other benefits as discussed in [Section 3](#).

At least one application disregards the requirement in RTP and uses non-compound packets when transmitting certain events, namely Open Mobile Alliance (OMA) Push-to-talk over Cellular (PoC) [[OMA-PoC](#)]. The OMA POC service is primarily used over cellular links capable of IP transport, such as the GSM GPRS. The use of non-compound RTCP is also specified in [[3GPP-MTSI](#)].

The use of non-compound packets is, however, not without issues. This is discussed in [Section 4](#). These issues needs to be considered and is one motivation for this document.

In addition this document proposes how AVPF could be updated to allow the transmission of non-compound packets in a way that would not substantially affect the mechanisms that compound packets provide. The connection to AVPF is motivated by the fact that non-compound RTCP is mainly intended for event driven feedback purposes and that the AVPF early and immediate modes makes this possible.

2. RTCP Compound Packets

[Section 6.1 in RFC3550](#) [[RFC3550](#)] specifies that an RTCP packet must be sent in a compound packet consisting of at least two individual packets, first an Sender Report (SR) or Receiver Report (RR), followed by an SDES packet containing at least a CNAME Item for the transmitting source identifier (SSRC). Lets examine what these RTCP packet types are used for.

1. The sender and receiver reports (see [Section 6.4 of RFC 3550](#) [[RFC3550](#)]) provides the RTP session participant with the Sender Source Identifier (SSRC) of all RTCP senders. Having all participants send these packets periodically allows everyone to

determine the current number of participants. This information is used in the transmission scheduling algorithm. Thus this is particularly important for new participants so that they quickly can establish a good estimate of the group size. Failure to do this would result in RTCP senders consuming too much bandwidth.

2. The sender and receiver reports contain some basic statistics usable for monitoring of the transport and thus enable adaptation. These reports become more useful if sent regularly as the receiver of a report can perform analysis to find trends between the individual reports. When used for media transmission adaptation the information becomes more useful the more frequently it is received, at least until one report per round-trip time (RTT) is achieved. Therefore there are most cases no reason to not include the sender or receiver report in all RTCP packets.
3. The CNAME SDES item (See [Section 6.5.1 of RFC 3550](#) [[RFC3550](#)]) exist to allow receivers to determine which media flows that should be synchronized with each other between different RTP sessions carrying different media types. Thus it is important to quickly receive this for each media sender in the session when joining an RTP session.
4. Sender Reports (SR) is used in combination with the above SDES CNAME mechanism to establish inter media synchronization. After having determined which media streams should be synchronized using the CNAME field, the receiver uses the Sender Report's NTP and RTP timestamp fields to establish synchronization.

Reviewing the above it is obvious that both SR/RR and the CNAME is very important for new session participants to be able to utilize any received media and to avoid flooding the network with RTCP reports. In addition if not sent regular the dynamic nature of the information provided would make it less and less useful.

3. Benefits from non-compound packets

As mentioned in the introduction most advantages of using non-compound packets exists in cases when the available RTCP bit-rate is limited. This because non-compound packets will be substantially smaller than a compound packet. A compound packet is forced to contain both an RR or an SR and the CNAME SDES item. The RR containing a report block for a single source is 32 bytes, an SR is 52 bytes. Both may be larger if they contain report blocks for multiple sources. The SDES packet containing a CNAME item will be 10 bytes plus the CNAME string length. Here it is reasonable that the CNAME string is at least 10 bytes to get a decent collision

resistance. And if the recommended form of user@host is used, then most strings will be longer than 20 characters. Thus a non-compound packets can become at least 70-80 bytes smaller than the equivalent compound packet.

The following benefits exist for the smaller non-compound packets:

1. Shorter serialization time, i.e. the time it takes the link to transmit the packet. For slower links this time can be substantial. For example transmitting 120 bytes over an link interface capable of 30 kbps takes 32 milliseconds (ms) assuming uniform transmission rate.
2. For links where the packet loss rate grows with the packet size, smaller packets will be less likely to be dropped. Example of such links are radio links. In the cellular world there exist links that are optimized to handle RTP packets with speech and these packets common sizes, the rationale behind this is to be able to increase the capacity and coverage for voice services. Compound RTCP packets commonly are 2-3 times the size of a RTP packet with compressed speech. If the speech packet over such a bearer have a packet loss rate of p , then the RTCP packet will experience $1 - (1-p)^x$ where x is the number of fragments the compound packet will be split on the link layer, i.e. 2 or 3 commonly.
3. In fixed links there are also benefits with sending feedback in small non-compound RTCP. One such application is those that use RTCP AVPF in early or immediate mode to send frequent event driven feedback. Under these circumstances the use of non-compound RTCP minimizes the risk that the RTCP bandwidth becomes too high during periods of intense adaptation feedback signaling.
4. In cases when regular feedback is need, like in the profile under development for TCP friendly rate control (TFRC) for RTP [[I-D.ietf-avt-tfrc-profile](#)]. The size of compund RTCP can result in very high bandwidth requirements for the feedback in case the roundtrip time is short. For this particular application non-compound RTCP gives a very substantial improvement.

In cases when non-compound packets carry important and time sensitive feedback both shorter serialization time and the lower loss probability are important to enable the best possible functionality. Having a packet loss rate that is much higher for the feedback packets compared to media packets is not advantageous when for example trying to perform media adaptation to handle the e.g. changed performance present at the cell border in cellular system.

For high bit-rate applications there is usually no problem of supplying RTCP with sufficient bit-rates. When using AVPF one can use the "trr-int" parameter to restrict the regular reporting interval to approximately once per RTT or less often. As in most cases there are no reasons to provide regular reports with higher density than this. Any additional bandwidth can then be used for feedback messages. The benefits of non-compound packets in this case are limited, but exist, one typical example is video using generic NACK in cases where the RTT is low. Using non-compound packets would reduce the total amount of bits used for RTCP. This is primarily applicable if the number of non-compound packets is large. This would also result in lower processing delay and less complexity for the feedback packets as they do not need to query the RTCP database to construct the right messages.

4. Issues with non-compound RTCP packets

This section describes some of the known issues with non-compound RTCP packets

4.1. Middle boxes

Middle boxes in the network may discard RTCP packets that does not follow the rules outlined in [section 6.1 of RFC3550](#). The effect of this might for instance be that compound RTCP packets makes its way through while the non-compound feedback packets are lost.

4.2. Packet Validation

A non-compound packet will be discarded by the packet validation code in [Appendix A of RFC 3550](#) [[RFC3550](#)]. This has several impacts as described in the following sub sections.

4.2.1. Old RTCP Receivers

Any RTCP receiver without updated packet validation code will discard the non-compound packets. Thus these receivers will not see the feedback contained in the these non-compound packets. The effect of this depends on the type of feedback message and the role of the receiver. For example this may cause complete function loss in the case of attempting to use a non-compound NACK message (see [Section 6.2.1 of RFC 4585](#) [[RFC4585](#)]) to non updated media sender in a session using the retransmission scheme defined by [RFC 4588](#) [[RFC4588](#)].

This type of discarding would also effect the feedback suppression defined in AVPF. The result would be a partitioning of the receivers within the session between old ones only seeing the compound RTCP

feedback messages and the newer ones seeing both. Where the old ones may send feedback messages for events already reported on in non-compound packets.

4.2.2. Weakened Packet Validation

The packet validation code needs to be rewritten to accept non-compound packets. One potential effect of this change is much weaker validation that received packets actually are RTCP packets, and not packets of some other type being wrongly delivered. Thus some consideration should be done to ensure the best possible validation is available. For example restricting non-compound packets to contain only some specific RTCP packet types, that is preferably signalled on a session basis.

4.2.3. Bandwidth consideration

The discarding of non-compound RTCP packets would effect the RTCP transmission calculation in the following way; the `avg_rtcp_size` value would become larger than for RTP receivers that exclude the non-compound in this calculation (assuming that non-compound packets are smaller than compound ones). Therefore these senders would under-utilize the available bit-rate and send with a longer interval than updated receivers. For most sessions this would should not be an issue. However for sessions with a large portion of non-compound packets may result in that the updated receivers time out non-updated senders prematurely.

4.2.4. Computation of `avg_rtcp_size`

Long intervals between compound RTCP packets and many non-compound RTCP packets inbetween may lead to a computation of the value `avg_rtcp_size` that varies greatly over time.

4.3. Header compression

The classifiers for header compression algorithms such as RoHC [[RFC3095](#)] and its profiles must be aware of the fact that, with the proposed non-compound RTCP packets, the first RTCP packet type might differ from 200 or 201. Otherwise they may wrongly classify the packets as something else than RTCP. This may have impact on the compression efficiency.

4.4. RTP and RTCP multiplex on the same port

In applications that multiplexes RTP and RTCP on the same port, as defined in [[I-D.ietf-avt-rtp-and-rtcp-mux](#)], care must be taken to ensure that the demultiplexing is done properly even though RTCP

packets are non-compound.

5. Allowing non-compound RTCP packets

Based on the above analysis it seems feasible to allow transmission in RTCP under some restrictions. First of all it is important that compound packets are regularly sent to ensure the feedback reporting works. The tracking of session size and number of participants is also important as this ensures that the RTCP bandwidth remain bounded independent of the number of session participants. As the compound packets also are used to establish the synchronization, any newly joining participant in a session would need to receive a compound packet from the media sender. In summary the regular usage of compound packets must be maintained throughout the complete session. Thus non-compound packets should be restricted to be used as extra feedback packets sent in cases when a regular compound packet would not have been sent.

If one is going to use non-compound packets it will be important to verify that they actually reaches the session participants. As outlined above in [Section 4.1](#) and [Section 4.2](#) packets may be discarded along the path or in the end-point. The end-points can be resolved by introducing signalling that inform if all session participants are capable of non-compound packets or not. The middlebox issue is more difficult and here one will be required to use heuristics to determine if the non-compound packets are delivered or not. However in many cases the feedback messages sent using non-compound packets will result in either explicit or implicit indications that they have been received. Example of such are the RTP retransmission [[RFC4588](#)] that result from a NACK message [[RFC4585](#)], the Temporary Maximum Media Bit-rate Notification message resulting from a Temporary Maximum Media Bit-rate Request [[I-D.ietf-avt-avpf-ccm](#)], or the presence of a Decoder Refresh Point [[I-D.ietf-avt-avpf-ccm](#)] in the video media stream resulting from the Full Intra Request sent.

5.1. Usage of non-compound packets in AVPF

The usage of non-compound RTCP packet SHALL only be done in RTP sessions operating in AVPF [[RFC4585](#)] Early RTCP or Immediate feedback mode. Non-compound packets SHALL NOT be sent until at least one compound packet has been sent. In Immediate feedback mode all feedback messages MAY be sent as non-compound packets. In early RTCP mode a feedback message scheduled for transmission as an Early RTCP packet, i.e. not a Regular RTCP packet, MAY be sent as a non-compound packet.

In order to get as much as possible out of the non-compound RTCP concept, it is necessary to relax the sending rules in [RFC 4585](#) [[RFC4585](#)] in such a way that it is possible to transmit more than one early feedback RTCP between regular RTCP, presently this is controlled by means of the `allow_early` flag which has the effect that only one early feedback message is allowed between the regularly scheduled RTCP. The proposed relaxation is reasonable as the size of the non-compound RTCP is expected to be small compared to the regular RTCP. The relaxation however needs to be done in a way that it is both bandwidth conservative and robust/scalable for an increased number of users. A possible outline of an algorithm that does this overrides the `allow_early` flag to the value TRUE if the bandwidth consumed by non-compound RTCP is within some predefined limits, one such limit may be that non-compound RTCP should not be allowed to consume more than a predefined fraction of the total RTCP bandwidth. Another option may be to use the SDP attribute "trr-int" as limiting factor.

The details of the algorithm that handles the relaxation of the sending rules is TBD.

[5.1.1.](#) Verifying the delivery of non-compound packets

A proposed algorithm to detect consistent failure of delivery of non-compound packets needs to be written. The details of this algorithm is application dependent and therefore outside the scope of this document.

If the verification fails it is strongly recommended that only compound RTCP according to the rules outlined in [RFC3550](#) is transmitted.

[5.2.](#) SDP Signalling Attribute

We request to define the a "a=ncp" [[RFC4566](#)] attribute to indicate if the session participant is capable of supporting non-compound packets. It is required that a participant that proposes the use of non-compound RTCP itself supports the reception of non-compound RTCP.

[6.](#) IANA Considerations

IANA will be required to register the SDP signalling attribute defined in [Section 5.2](#).

7. Security Considerations

The security considerations of RTP [[RFC3550](#)] and AVPF [[RFC4585](#)] will apply also to non-compound packets. The reduction in validation strength for received packets on the RTCP port may result in a higher degree of acceptance of spurious data as real RTCP packets. This vulnerability can mostly be addressed by usage of an security mechanism that provide authentication, e.g. SRTP[RFC3711].

8. Acknowledgements

The authors would like to thank all the people who gave feedback on this document.

This document also contain some text copied from [RFC 4585](#) [[RFC4585](#)] and we take the opportunity to thank the author of said document.

9. References

9.1. Normative References

- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 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.

9.2. Informative References

- [3GPP-MTSI]
3GPP, "Specification : 3GPP TS 26.114 (v7.1.0 preliminary), http://www.3gpp.org/ftp/tsg_sa/WG4_CODECS/Specs_update_after_SA36/26114-710.zip", March 2007.
- [I-D.ietf-avt-avpf-ccm]
Wenger, S., Chandra, U., Westerlund, M., and B. Burman, "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)", [draft-ietf-avt-avpf-ccm-10](#) (work in progress), October 2007.
- [I-D.ietf-avt-rtp-and-rtcp-mux]
Perkins, C. and M. Westerlund, "Multiplexing RTP Data and Control Packets on a Single Port",

[draft-ietf-avt-rtp-and-rtcp-mux-07](#) (work in progress),
August 2007.

[I-D.ietf-avt-tfrc-profile]

Gharai, L., "RTP with TCP Friendly Rate Control",
[draft-ietf-avt-tfrc-profile-10](#) (work in progress),
July 2007.

[OMA-PoC] Open Mobile Alliance, "Specification : Push to talk Over
Cellular User Plane, [http://www.openmobilealliance.org/
release_program/docs/PoC/V1_0_1-20061128-A/
OMA-TS-PoC-UserPlane-V1_0_1-20061128-A.pdf](http://www.openmobilealliance.org/release_program/docs/PoC/V1_0_1-20061128-A/OMA-TS-PoC-UserPlane-V1_0_1-20061128-A.pdf)",
November 2006.

[RFC3095] Bormann, C., Burmeister, C., Degermark, M., Fukushima, H.,
Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le,
K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K.,
Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header
Compression (ROHC): Framework and four profiles: RTP, UDP,
ESP, and uncompressed", [RFC 3095](#), July 2001.

[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.

[RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R.
Hakenberg, "RTP Retransmission Payload Format", [RFC 4588](#),
July 2006.

Authors' Addresses

Ingemar Johansson
Ericsson AB
Laboratoriegatan 11
SE-971 28 Lulea
SWEDEN

Phone: +46 73 0783289
Email: [ingemar.s.johansson \(AT\) ericsson.com](mailto:ingemar.s.johansson@ericsson.com)

Magnus Westerlund
Ericsson AB
Torshamnsgatan 21-23
SE-164 83 Stockholm
SWEDEN

Phone: +46 8 7190000

Email: magnus.westerlund (AT) ericsson.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

