

AVT Working Group
Internet-Draft
Intended status: Informational
Expires: January 13, 2012

V. Singh
J. Devadoss
J. Ott
Aalto University
I. Curcio
Nokia
July 12, 2011

Real-time Transport Control Protocol (RTCP) in Overlay Multicast
draft-ott-avtcore-rtcp-overlay-multicast-01.txt

Abstract

The Real-time Transport Control Protocol (RTCP) is designed to operate along with Real-time Transport Protocol (RTP) in unicast, single-source multicast and any-source multicast environments. With the availability of overlay multicast and Application Layer Multicast (ALM), the suitability of RTCP in such environments need to be analyzed. The applicability of the existing RTCP reporting architectures in overlay multicast and ALM environments are investigated and the new features that may be required are discussed in this document.

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 January 13, 2012.

Copyright Notice

Copyright (c) 2011 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
3. Overlay Multicast	4
4. Classifying feedback reporting mechanism	5
5. Classifying overlay multicast topologies and media/overlay operations	6
6. RTP Entities in an overlay multicast network	6
7. Applicability of RTCP reporting as defined in RFC 3550	8
8. Applicability of RTCP with unicast feedback target	9
9. Desirable features of RTCP in overlay multicast	10
10. An example explaining the issues	10
11. Endpoint Reporting	11
12. Hop-by-hop RTCP reporting	12
13. Security Considerations	12
14. IANA Considerations	12
15. Normative References	12
Authors' Addresses	13

1. Introduction

RTP [2] provides transport mechanisms suitable for unicast, any-source multicast and single-source multicast. RTP is used to carry audio and video streams together with the RTP control protocol to provide periodic feedback about the media streams received in a specific duration. The RTCP reporting interval is defined as part of the RTCP and depends on the number of participants, type of participants (sender or receiver) and the operating environment of the session (point to point or multicast).

The RFC3550 [2] specifies the maximum RTCP bandwidth to be 5% of the media bit rate. In point to point sessions, each participant gets a equal share of 2.5% of the media bit rate to be used for RTCP. In multicast (including any-source multicast and source-specific multicast) sessions, the senders usually share 25% of the allocated RTCP bandwidth and the receivers share the remaining 75% of the RTCP bandwidth. The RTCP bandwidth share may be modified using RFC 3556 [5], but will generally be kept small so as to have a smooth media stream flow.

In a multicast session, the RTCP reports are multicast so that each participant can receive the RTCP reports from every other participant and thus obtain a "global" view of the multicast session. This, however, requires multicast connectivity between all peers and thus cannot be applied to Source-specific Multicast (SSM) [7]. Specific RTCP extensions were developed for SSM [8] introducing a mechanism of RTCP reporting where the RTCP reports are unicast to a feedback target. This mechanism is specifically applicable in scenarios where many-to-many group communication is not available or not desired or a sender-controlled summarized reporting is desired. RTCP reports are unicast to a feedback target specified during initialization or inside RTCP reports. The RTCP reporting interval calculation specified in RTCP Extension for SSM uses the same mechanisms as specified in RFC3550 [2]; the distribution source may send RTCP RSI reports to control the rate at which RTCP reports are generated by the receivers. The aforementioned rules for bandwidth modifications apply as well.

Recent interest in overlay multicasting and ALM -- to substitute for the lack of globally available native IP any-source multicast -- motivates also carrying RTP-based media streams in such environments.

In ALM, the participating nodes form a distribution tree to forward the data. ALM involves the use of different mechanisms to construct and maintain the distribution tree. In overlay multicast, individual multicast enabled networks are connected by virtual unicast links. Overlay multicast involves mechanisms to construct optimal

interconnection of virtual links. An ALM can be viewed as a subclass of overlay multicast, if the individual multicast enabled networks have only one participating node. So, further in this document, when we refer to overlay multicast it also includes ALM. Depending on the abstraction chosen for the overlay multicast, the RTP/RTCP entities using it may or may not be aware of the hop-by-hop forwarding of the packets:

If they are not, regular any-source multicast operation of RTP and RTCP as per RFC 3550 is a workable, yet possibly not optimal solution.

If RTP entities are aware of the forwarding process, additional RTCP reporting and aggregation mechanisms may be applied and the existing RTCP reporting mechanisms need to be investigated for their applicability in overlay multicast

In either case, in an overlay multicast environment, using RTP to transport media streams will be straightforward,

In this document, we take a very first stab at reviewing the RTCP reporting mechanism specified in RFC3550 [2] and the RTCP Extension for SSM to find its applicability in the overlay multicast environment. We also discuss on the specific characteristics of the overlay multicast and the type of reporting that is desired on such environments. This document also complements the RTP topologies overview [6].

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 BCP 14, RFC 2119 [1] and indicate requirement levels for compliant implementations.

The terminology defined in RTP [2], the RTP Profile for Audio and Video Conferences with Minimal Control [3], and the RTCP Extension for SSM [8], apply.

For the time being, in this document, by overlay multicast we refer to a multicast overlay offering media distribution from a single source, analogous to SSM.

3. Overlay Multicast

In overlay multicast, the media streams are multicast over an network

of virtual links that is constructed using mechanisms in line with the requirements of the application using the overlay. The construction and maintenance of the overlay involves mechanisms for bootstrapping the new participant to the overlay, routing, proactive/reactive repair, improving scalability and recovery from loss of a forwarding node or a virtual link. In this document, these mechanisms are further referred to as overlay operation. As the overlay network is formed by the participating nodes, the loss of a participating node brings change in the network structure. So, there is an inherent instability in the overlay multicast which are addressed by the overlay operation.

In overlay multicast, the participating nodes can use mechanisms for providing error resilience and congestion control that can proactively adapt or reactively repair the media streams depending on the receiver metrics reported by the nodes below its hierarchy. In this document, the error resilience mechanisms together with congestion control techniques are further referred to as media operation.

The media operation depends on the metrics (reported reception quality) reported by the participants. The overlay operation can depend on different types of metrics and may also include the media quality metrics like round trip time, observed packet loss, jitter etc. As the overlay operation depends on the application using the overlay, there can be significant overlap with the metrics used by media operation.

Using RTP in overlay multicast does not require any change to its specification. Every participating node that receives the RTP packet replicates the packet and sends down the distribution tree. RTCP SR follows the same path as RTP, so it does not bring in changes with its use in overlay multicast. In the case of RTCP RR, it reports the receiver's perceived media quality and it carries significant data based on which the algorithms of media and/or overlay operation can operate.

The multicast overlay maintenance mechanisms may operate entirely independent of RTCP reporting, they may choose to leverage parts of the RTCP reporting, or they may rely entirely on RTCP. In this document, we are concerned about the operation of RTCP reporting in such an environment and we do not take (at this point) any position on the way the overlay operates.

4. Classifying feedback reporting mechanism

Any feedback reporting mechanism in a session with n participants can

be classified into one of the following categories.

- o (i) 1 to 1 reporting
- o (ii) 1 to n reporting
- o (iii) 1 to m reporting ($m < n$, every node reports to m nodes)

An example of '1 to 1' reporting is RTCP in SSM with unicast feedback target (and, of course, in case of point-to-point sessions). An example of '1 to n' reporting is RTCP in multicast mode as specified in RFC 3550. In further sections, we shall be referring to these mechanisms while discussing the applicability of a RTCP reporting model.

5. Classifying overlay multicast topologies and media/overlay operations

The effectiveness of a chosen RTCP reporting model, directly depends on the type of overlay multicast topology, type of the media/overlay operation and the number of participants in the session. Therefore, when considering the applicability of a RTCP reporting model, they need to be evaluated against the every possible combinations of overlay multicast topology, media and overlay operation.

The overlay multicast topologies can be broadly classified into two categories.

- o (i) Overlay multicast, with no or very minimal overlay operations(virtual links are statically configured)
- o (ii) Overlay multicast, that relies on overlay operations to dynamically configure distribution tree.

The overlay and media operations can be classified into two categories.

- o (i) Centralized: A single entity decides on the type of media/overlay operations.
- o (ii) Distributed: At any instant more than one entity is attempting to perform media/overlay operation.

6. RTP Entities in an overlay multicast network

The below figure represents a sample overlay multicast network. In this section, we see how the RTP entities can be used to describe a overlay multicast network.

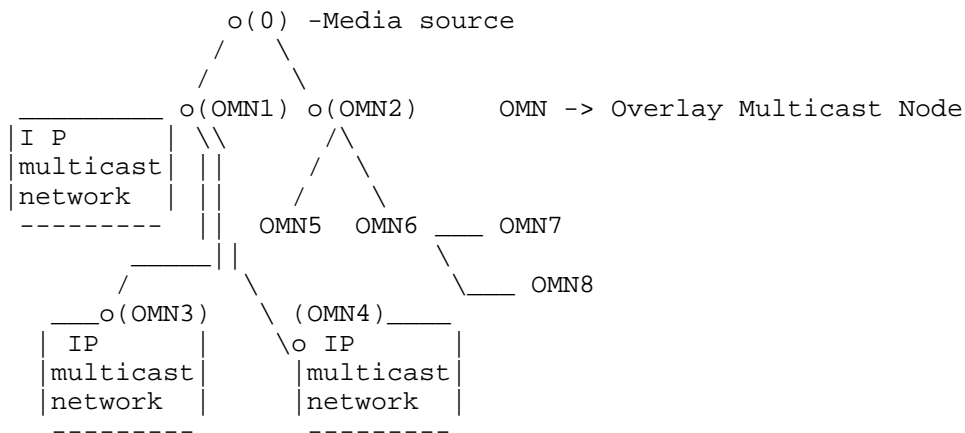


Figure 1: A sample overlay multicast topology

Describing an Overlay Multicast Node (OMN) using RTP entities: Each OMN can be considered as a combination of a RTP media handler and a RTP translator. The purpose of the translator is to replicate and forward the streams to different network destinations and to handle RTCP reports. The media handler receives RTP streams, processes received RTCP reports and generates RTCP receiver reports.

Role of RTP Translator in forwarding RTP streams: The RTP translator in each OMN forwards the received RTP packets(from its parent node) to all virtual links it is connected with. For example, RTP translator in OMN1 forwards RTP packets to its RTP media handler, OMN3, OMN4 and the IP multicast network it is connected with.

Role of RTP Translator in handling RTCP reports: When there is no change to the media encoding, the RTP translator uses the straight case of Replicate and Forward mechanism.

Replicate and Forward: The RTP translator replicates and forwards the RTCP RR received from one virtual link to all other virtual links. Below, we explain using OMN1 as example. The RTP translator in OMN1 receives RTCP RR from,

- o the media handler(which is within the same node)
- o the native multicast network
- o OMN3
- o OMN4
- o parent OMN (which is forwarding the RTCP RRs received from the other virtual links, it is connected with)

To explain the RTCP forwarding rules, we take an sample event where the translator in OMN1 receives a RTCP RR from IP multicast network.

The event response shall be that, it is forwarded to the media handler(within OMN1), OMN3, OMN4 and to the parent node(in the case of OMN1, it is media source). In this way, all participating nodes in the session shall receive RTCP RR from all other participating nodes.

In the overlay multicast scenario, the RTP translator in an OMN can be extended to do other operations, such as, filtering and aggregating RTCP reports, etc. But to realize it, the scope and definitions of an RTP translator needs to be analyzed. (TBD)

7. Applicability of RTCP reporting as defined in RFC 3550

If this reporting mechanism is to be used in an overlay multicast, then the RTP translator in each node replicates and forwards the RTCP reports on all (virtual) links except the incoming (virtual) link. It is a form of '1 to n' reporting, where all participants get a copy of the RTCP report sent by every other participant. Now, let us look at the effectiveness of this reporting model in the overlay multicast environment.

With the increase in number of receivers, the RTCP reporting interval increases. The larger the reporting interval, fewer the options available for media operation. For example, with 'x' number of participants, it can be possible to do retransmission based on an RTCP report indicating a lost packet. But with the increasing number of participants (say n*x), the interval gets bigger by n times leading to fewer options of error repair mechanisms. In IP multicast, the error resilience mechanisms like adaptive FEC, retransmission can be performed only by the source or one designated participant/observer. But in the case of overlay multicast, there are more options available to deploy these mechanisms at intermediary nodes which become an inherent part of the forwarding tree. Furthermore, the importance of overlay operation increases with increasing number of participants: the larger the number of nodes, the greater the importance of overlay operation in maintaining stability i.e., the importance of the reports that influence the quality of the overlay operation grows.

The proportional relationship of number of nodes with the RTCP reporting interval was designed to limit the bandwidth consumed by the RTCP and provide scalability so as to evolve as a multicast session with many number of participants. But in the overlay multicast scenario, with the reducing number of reports per time unit the quality of the overlay is reduced due to reduced effectiveness of media and overlay operation.

In contrast to any-source multicast, however, the RTCP reports sent by a particular node are not automatically received by all other nodes. Instead, the RTCP reports travel along the virtual links formed between participating nodes. This may be used to limit their spread and may allow for mechanisms improving overall efficiency of RTCP reporting. Thus, while the total number of participants in an RTP session limits the RTCP bandwidth consumption within 5% of the media session, this limit may not need to be applied the total consumption across the entire overlay multicast, but be maintained in local regions only.

8. Applicability of RTCP with unicast feedback target

If RTCP/SSM reporting mechanism is to be used in an overlay multicast, then the RTCP RR reports are unicast to a designated node that is either within or outside the overlay multicast. It is a form of '1 to 1' reporting and here we discuss on its applicability in the overlay multicast.

RTCP/SSM defines the use of a single distribution source in conjunction with one or more feedback targets. If multiple feedback targets are to be used, their respective setup and coordination is outside the scope of the RTCP/SSM specification. Conceptually, however, the RTCP/SSM mechanisms support the idea of hierarchical report aggregation and forwarding, even though the RTP media as well as the RTCP sender reporting distribution paths are supposed to use SSM multicasting at the IP layer. This notion of RTCP aggregation would need to become more explicit, but could provide a first basis for realizing overlay multicast.

Overlay multicast also provides explicit support for using intermediary (participating nodes in the distribution tree) media error resilience mechanisms. Stepwise RTCP aggregation would also make the necessary feedback information available to the individual intermediate nodes and could thus provide the hook for invoking repair mechanisms.

The (kind of) centralized reporting and centralized decision making in RTCP/SSM would need to be expanded to allow for more flexible media and/or overlay operation and, in particular, would need to cover the assignment and maintenance of feedback targets as regular part of the overlay operation.

An interesting issue is whether this source-specific type of operation could be expanded later towards multi-source operation in order to support any-source multicast overlays as well. While RTCP/SSM seems to be a promising starting point, this latter step is left

for further study.

9. Desirable features of RTCP in overlay multicast

The following list is an initial set of desirable functions for media/overlay operation:

- o Need for a mechanism of RTCP reporting, where the reporting interval is decoupled from the number of participants in the media session. But still the original reason behind this linkage (limiting the RTCP bandwidth consumption) can be maintained through other suitable mechanisms. In essence, fine-granular RTCP reporting could be confined to subgroups while the global bandwidth limitation is still maintained.
- o Overlay multicast brings in the option to use media operation at intermediate nodes. With more frequent reporting, their effectiveness increases. So, to increase the reporting frequency and at the same time limit the bandwidth consumption, a RTCP RR reporting mechanism that provides the feature of '1 to m' ($n > m$, n - number of participants in the session) reporting is needed. By choosing a small m , the RTCP RR reporting frequency can be increased as it is directed only to a small subset of participating nodes. Such subset reporting could be carried out at a single hierarchy level inside an overlay multicast or across multiple such levels.
- o Media operation should be able to leverage the additionally available reporting information to optimize performance (for error control, possibly congestion control, etc.)
- o Overlay operation can be either centralized or distributed. For example, the decisions related to overlay operations can be made only by a single node for the whole network or can be distributed to many groups, with each groups making the decisions depending on the collected metrics. In the later form of overlay operation, the availability of '1 to m' reporting provides timely and more precise information for the algorithms used in the overlay operation. As each group can act independent from the other groups, there may not be a need for reporting across the groups but there may be a need for a higher reporting frequency within the group.

10. An example explaining the issues

We take the below example for explaining the desired form of RTCP reporting in overlay multicast.




Figure 2: A sample overlay multicast topology

In the above figure, the nodes {1, 3, 5 and 7} and {2, 4, 6 and 8} are considered as a group with node 1 and 2 acting as their respective group leaders. The overlay operation involves changing the group leader based on the metrics reported by the participating nodes.

In the above example, the nodes {1, 3, 5 and 7} are more interested in RTCP reports from the nodes within their group rather than in the RTCP reports from {2 or 4 or 6 or 8}. If the RTCP reporting interval is supposed to be proportional to the number of participants in the session, then the individual groups see reduced reporting due to the increase in the number of participants. This reduces the effectiveness of both media and overlay operation.

The paper "Construction of an Efficient Overlay Multicast Infrastructure for Real-time Applications" is an overlay multicast solution, that dynamically re-arranges the distribution tree based on the metrics reported(or measured with) 'm' other nodes. (http://www.ieee-infocom.org/2003/papers/37_03.PDF)

11. Endpoint Reporting

For a Distribution Source (DS), RFC5760 [8] defines a feedback reflection mechanism and a feedback summary mechanism (RSI). In the case of overlay networks, each node forwards RTP packets and therefore, also generates and receives RTCP packets. Therefore, each node SHOULD also perform feedback reflection or summarization. However, unlike the RFC5760 [8], where the DS transmits the resulting feedback over the multicast RTCP channel, the intermediate node in an overlay networks needs to transmit the feedback packets both upstream and downstream. In the section, feedback information refers to Receiver Reports or RSI, as the case may be. In a given setup the feedback information between intermediate nodes and distribution source MUST be consistent and the leaf nodes MUST always send RRs.

In this section, we describe the behavior of an intermediate node.

- o Local Information: is the node's receiver report. Typically, this information will be used to create summary reports (RSI) or compounded with other reports (in case of feedback reflection). The local information is always sent upstream to its parent node.
- o Source Information: Feedback information received from an upstream node is sent downstream. The node MUST add or summarize its local information before transmitting it downstream.
- o Sub-tree Information: Feedback information received from the child nodes is summarized or reflected in both upstream and downstream directions. The node need not add or summarize its local information before transmitting the sub-tree information.

12. Hop-by-hop RTCP reporting

In an overlay network, the Sender Reports (SR) originate from the source and percolate through the distribution tree to the leaf nodes. Since the intermediate nodes just forward RTP packets, they do not generate SRs. Furthermore, they may either send RRs or RSIs downstream, this information is insufficient to measure the performance (of the links) between two consecutive intermediate nodes. To assist with better performance measurements, each intermediate node MUST generate an Intermediate Forwarder Report (IFR) and send it to its parent and child nodes -- the IFR carries loss and discard metrics, and NTP Timestamps for making RTT measurements.

13. Security Considerations

TBD.

14. IANA Considerations

There are no specific IANA action necessary for this document.

15. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

- [3] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.
- [4] 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.
- [5] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, July 2003.
- [6] Westerlund, M. and S. Wenger, "RTP Topologies", RFC 5117, January 2008.
- [7] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.
- [8] Ott, J., Chesterfield, J., and E. Schooler, "RTP Control Protocol (RTCP) Extensions for Single-Source Multicast Sessions with Unicast Feedback", RFC 5760, February 2010.

Authors' Addresses

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

Email: varun@comnet.tkk.fi

Jegadish Devadoss
Aalto University
School of Electrical Engineering
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: jegadish@netlab.tkk.fi

Joerg Ott
Aalto University
School of Electrical Engineering
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: jo@comnet.tkk.fi

Igor D.D. Curcio
Nokia
P.O. Box 88 (Tieteenkatu 1)
Tampere, FIN 33721
Finland

Email: igor.curcio@nokia.com

