

AVTCORE
Internet-Draft
Intended status: Standards Track
Expires: November 27, 2011

T. VanCaenegem
Alcatel-Lucent
B. VerSteeg
A. Begen
Cisco
May 26, 2011

Retransmission for Source-Specific Multicast (SSM) Sessions
draft-vancaenegem-avtcore-retransmission-for-ssm-00

Abstract

This document describes RTP retransmission for source-specific multicast (SSM) architectures with unicast feedback. RTP payload format for retransmissions has been defined in RFC 4588, whereas the RTP profile an RTP receiver could use to issue RTP Control Protocol (RTCP) feedback messages and the format of these feedback messages have been defined in RFC 4585. RFC 5760 defines the operation for SSM architectures with unicast feedback.

First, we document potential issues that could arise when providing a retransmission service using RTP retransmission (RFC 4588 and RFC 4585) for RFC 5760 architectures based on rules described in the relevant RFCs. We then present solutions that allow to avoid unnecessary feedback suppression, provide enhanced retransmission services and address congestion control for retransmission in an SSM environment. These solutions specify certain rules that apply to the distribution source, feedback target(s), retransmission source(s) and RTP receivers.

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 November 27, 2011.

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	4
2. Requirements Notation	5
3. Definitions	5
4. Feedback Suppression and Retransmission for SSM: Applicable Specifications	6
4.1. RTCP Feedback Suppression in RFC 4585	6
4.2. RTCP Feedback Message Distribution in RFC 5760	7
4.3. RAMS	8
5. Pitfalls When Combining RTCP Feedback Handling and Retransmission Using RFC 5760 Rules	10
5.1. Case of DS Forwarding/Reflecting All RTCP NACKs	10
5.2. Case of DS Terminating All RTCP NACKs	10
5.3. Case of Multiple FTs	11
5.4. Takeaways	12
6. Rules for SSM with Retransmissions	13
6.1. Feedback Suppression Behavioral Requirements for DS, FT and RTP_Rx's	13
6.2. Informing RTP_Rx's with the SSRC Value of FT/RS	15
6.3. Unsolicited Retransmissions	17
6.4. Congestion Control Considerations	18
6.4.1. Congestion Control in RFC 4585 and in RAMS	18
6.4.2. Congestion Handling for SSM with Retransmissions	19
6.5. New RSI Message and RAMS-I TLV Extension	21
6.5.1. RSI Message 'FT SSRC List'	21
6.5.2. RAMS-I TLV Extension 'FT SSRC'	21
7. Security Considerations	22
8. IANA Considerations	22
9. RTP Retransmission Service in DVB IPTV	22
10. Contributors	23
11. Acknowledgments	23

12. References	23
12.1. Normative References	23
12.2. Informative References	24
Authors' Addresses	24

1. Introduction

The RTP toolkit is used to deliver a variety services over IP multicast. [RFC5760] defines a source-specific multicast (SSM) architecture with one or several Media Senders, a Distribution Source (DS) (sourcing the SSM), one or several Feedback Targets (FT) that may or may not be co-joint with the DS and RTP receivers (RTP_Rx) that send unicast feedback to a FT. This document explores the complications one might face when adding one or more Retransmission Sources (RSO) in this architecture for reliable multicast distribution.

In this specification, we assume that FT and RSO logical entities are combined in a single physical entity and they share state. In the remainder of the text, the term Retransmission Server (RS) is used whenever appropriate, to refer to this single physical entity. As in [I-D.ietf-avt-rapid-acquisition-for-rtp], the RS receives the RTP packets sent in an SSM session (as a normal RTP_Rx), and buffers these packets for a certain time for retransmission purposes.

The term RAMS (Rapid Acquisition of Multicast Sessions) will be used to refer to the protocol and architecture defined in [I-D.ietf-avt-rapid-acquisition-for-rtp]. The RAMS specification presents a method in which an RTP_Rx can rapidly acquire reference information transported in an SSM session. This is enabled by having the RTP_Rx retrieve a unicast burst of data from an RS, formatted as RTP retransmission packets prior to joining the SSM session. While the present document will include some requirements and observations expressed in [I-D.ietf-avt-rapid-acquisition-for-rtp], it will particularly focus on the impact of packet loss events experienced in an SSM distribution, and how the elements in the SSM architecture interact best to ensure impacted RTP_Rx's are provided with appropriate RTP retransmissions.

A packet loss event in an SSM session, will in general trigger one or multiple interactions between the RTP_Rx and the FT/RS/DS entities:

- o The RTP_Rx sending an RTCP NACK message to the FT ([RFC4585])
- o The RTP_Rx receiving RTP retransmission(s) from the RS
- o The RTP_Rx receiving an RTCP feedback message from the DS

This document describes the behavior for all entities defined in an SSM architecture with unicast feedback as described in [RFC5760] (DS, FT, RS and RTP_Rx) addressing feedback implosion (avoidance), retransmission efficacy and congestion concerns.

One application that can benefit from this document is IPTV. In many IPTV deployments, streams are transported over SSM. Transporting this data over SSM enables a packet loss recovery by means of RTP retransmissions and/or forward error correction (FEC) [I-D.ietf-fecframe-dvb-al-fec]. This document focuses on retransmission. Note that Digital Video Broadcasting (DVB) forum has specified an RTP retransmission solution for its managed Linear Media Broadcast (LMB) service, and which references most of the relevant RTP specifications. An overview of this solution is provided in Section 9.

2. Requirements Notation

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

3. Definitions

This document uses some acronyms and definitions which have been introduced [I-D.ietf-avt-rapid-acquisition-for-rtp] as the baseline. Some of these definitions have been slightly retailored to fit the restricted scope of this document.

(Primary) SSM (or Multicast) Session: The multicast session to which RTP_Rx's can join at a random point in time. A primary SSM session can carry multiple RTP streams.

Primary Multicast RTP Session: The multicast RTP session that is of interest to an RTP receiver. It comprises the primary multicast RTP stream(s) and an associated unicast RTCP feedback stream to a Feedback Target.

Primary Multicast (RTP) Streams: The RTP stream(s) carried in the primary multicast RTP session.

Source Filtering Group Management Protocol (SFGMP): Following the definition in [RFC4604], SFGMP refers to the Internet Group Management Protocol (IGMP) version 3 [RFC3376] and the Multicast Listener Discovery Protocol (MLD) version 2 [RFC3810] in the IPv4 and IPv6 networks, respectively. In the remainder of this document, SFGMP will refer to any group management protocol that has Join and Leave functionalities.

Feedback Target (FT): Unicast RTCP feedback target as defined in

[RFC5760]. FT_Ap denotes a specific feedback target running on a particular address and port.

Retransmission Packet: An RTP packet that is formatted as defined in Section 4 of [RFC4588]. The payload of a retransmission packet comprises the retransmission payload header followed by the payload of the original RTP packet.

(Unicast) Retransmission RTP Session: The unicast RTP session used to send one or more unicast retransmission RTP streams and their associated RTCP messages. This session is setup as described in [I-D.ietf-avtcore-ports-for-ucast-mcast-rtp].

Retransmission Server (RS): The RTP/RTCP endpoint that can generate the retransmission packets.

New definitions introduced in this document are:

Solicited Retransmission: Retransmission packet(s) that are received by an RTP_Rx which have been explicitly requested by this RTP_Rx by means of one or more RTCP NACKs.

Unsolicited Retransmission: Retransmission packet(s) that are received by an RTP_Rx which have not been explicitly requested by this RTP_Rx by means of an RTCP NACK.

Feedback Implosion/Storm: The situation where multiple RTP receivers start sending (quasi-) simultaneously (RTCP) feedback, often triggered by an upstream packet loss event in the SSM tree impacting a large set of RTP receivers.

4. Feedback Suppression and Retransmission for SSM: Applicable Specifications

In this section we first look at the various specifications in place that describe the architecture for SSM with unicast feedback and determine how feedback suppression and retransmission are accomplished.

4.1. RTCP Feedback Suppression in RFC 4585

[RFC4585] defines an RTP profile for an RTP receiver that wishes to provide fast feedback by means of RTCP feedback messages, such as is applicable to an RTP retransmission scenario. [RFC4585] defines the timing rules by which such RTCP feedback messages can be transmitted both in unicast and multicast/multi-party sessions, and it also defines the general format of RTCP feedback messages, as well as some

specific feedback messages. One of the transport-layer feedback messages that is particularly relevant here is the RTCP NACK message through which an RTP receiver can report missing packets. The missing packets are identified by means of an SSRC identifier (associated with the primary multicast stream for which the DS uses this SSRC as media sender) and an RTP sequence number (SN). The RS receiving such an RTCP NACK message can respond by sending retransmission packets whose format has been described in [RFC4588]. An important aspect covered by [RFC4585] is feedback suppression for RTP receivers involved in a multi-party session in order to avoid feedback implosion. To that extent, the RTP receiver waits for a (short) random dithering interval to check whether it sees a feedback message from any other RTP receiver reporting the same event. If such a feedback message is received, the RTP receiver refrains from sending the feedback message and continues to follow the regular RTCP transmission schedule. Note that a feedback message from one RTP receiver is typically not visible to other RTP receivers, so the FT might need to provide information to the RTP receivers. We will investigate the implications of this feedback suppression rule for SSM architectures with unicast feedback as defined in [RFC5760], where one or multiple RS's are deployed. To that extent we first provide an overview of how RTCP feedback messages transmitted by the RTP receivers are handled in the SSM architecture as specified by [RFC5760].

4.2. RTCP Feedback Message Distribution in RFC 5760

Two models in terms of DS behavior have been defined in [RFC5760]:

- o In Feedback Summary Model, the unicast RTCP receiver report messages from the RTP_Rx's are by default aggregated at the DS and their information is transmitted as Receiver Summary Information (RSI) messages in the SSM session. The RTCP feedback messages are by default terminated by the DS. However, the DS might also aggregate or forward RTCP feedback messages and transmit them in the SSM session, when this is explicitly signaled. Note that from the RTP perspective, the DS is an RTP_Rx generating its own RTCP receiver reports as well as other RTCP packets and sending them to the receiver group and media senders.
- o In Simple Feedback Model, the DS reflects all RTCP messages (including RTCP feedback) received in unicast via the FT from the RTP_Rx's. Note that many network topologies have a high degree of fanout, and also have a constrained link between the FT and the RTP_Rx's, so that reflecting all of the feedback messages is often not feasible.

The communication between DS and disjoint FT(s) occurs as follows:

[RFC5760] indicates that for the Simple Feedback Model where the FT(s) are disjoint from the DS, the FT must forward all RTCP packets to the DS.

[RFC5760] indicates the following for the Feedback Summary Model where the FT(s) are disjoint from the DS:

- o The FTs may simply forward all RTCP packets incoming from the RTP_Rx's to the DS.
- o The FTs may also perform aggregation of incoming RTCP packets and send only aggregated information to the DS.
- o If the FT performs summarization functions, it must also act as a receiver and choose a unique SSRC for its own reporting towards the DS.

4.3. RAMS

Assume now that the FT is combined with an RSo, constituting together the RS, which provides retransmissions in response to unicast RTCP NACK messages received from RTP_Rx's. Such architecture is described in [I-D.ietf-avt-rapid-acquisition-for-rtp] where the FT is co-joint with a Burst/Retransmission Source (BRS).

Figure 1 shows the main entities involved in SSM with retransmission enabled. They are

- o Distribution Source
- o Feedback Target (FT)
- o Retransmission Source (RSo)
- o RTP Receiver (RTP_Rx)

The figure also shows the various RTP/RTCP flows and associated sessions. The difference with RAMS is that the Burst/Retransmission Source (BRS) is replaced with a Retransmission Source (RSo). Another difference - not shown in the figure - is that next to the primary multicast session, there might also be a (source-specific) multicast retransmission session that can be sourced by the same DS entity (acting as an RS itself) or by the same RS that provides the unicast retransmissions. This is discussed in more detail in Section 6.3. The unicast retransmission session must be established as per [I-D.ietf-avtcore-ports-for-ucast-mcast-rtp]. Note that when RAMS and retransmissions are provided for the same SSM session, a single unicast retransmission session is established. The RAMS

specification puts a constraint on the RS/FT, stating that an FT (identified by means of IP address and port, and referred to as FT_Ap) must be associated with only a single RTP session. This constraint was put in place, because an RTP_Rx using the RAMS service may not necessarily know the media sender SSRC a priori. However, an RTP_Rx using the retransmission service can of course learn the SSRC by inspecting the SSRC identifier field in the RTP packets, and furthermore the SSRC field needs to be correctly filled out in the RTCP NACK message. Thus, while we do not have this constraint on FTs when only the retransmission service is enabled, the constraint still applies when both RAMS and retransmission services are enabled.

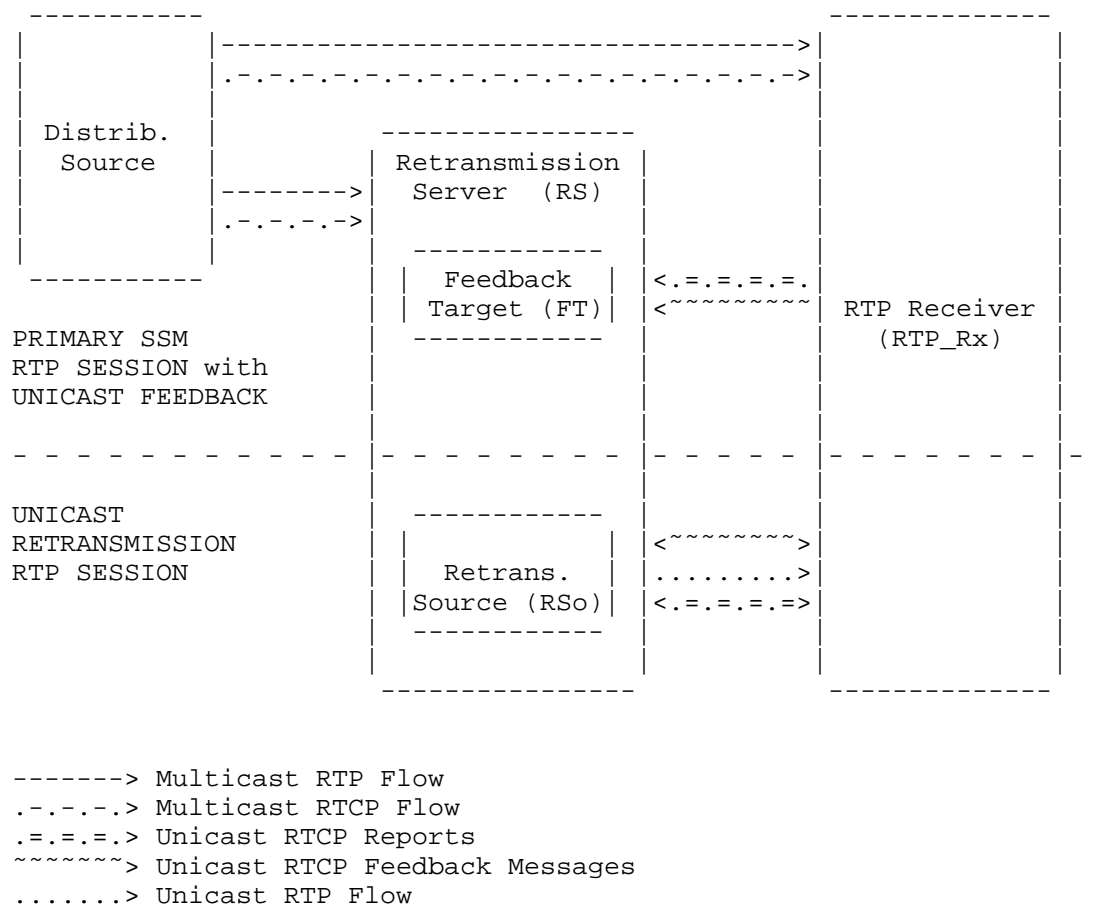


Figure 1: Flow diagram for SSM with retransmission

5. Pitfalls When Combining RTCP Feedback Handling and Retransmission Using RFC 5760 Rules

In this section, we investigate a number of scenarios in terms of SSM topology and behavior of the DS with respect to the forwarding/termination of RTCP feedback messages received from RTP_Rx's. In the first two scenarios, we assume there is only one FT. Note that the DS could be joint with the FT or disjoint from the FT.

5.1. Case of DS Forwarding/Reflecting All RTCP NACKs

The forwarding of RTCP feedback messages is one option allowed for by the Feedback Summary Model of [RFC5760], where reflection is what the DS must do based on the Simple Feedback Model of RFC 5760. In this case, RTP_Rx's applying the [RFC4585] suppression rule will result in the following: when an RTP_Rx X reports a packet loss to the FT by sending an RTCP NACK, this message is distributed, via the FT and, by the DS via the SSM to all other RTP_Rx's. If an RTP_Rx Y, different from RTP_Rx X, detected the same packet loss, but prior to sending out a NACK, this receiver Y gets the NACK message originally sent by X, it will refrain from sending a NACK, following [RFC4585]. The forwarding/reflection of RTCP NACK messages effectively results in feedback suppression, but obviously this is at the expense of limited visibility to the RS on which RTP_Rx suffered packet loss. In general this will result in a decreased quality at the service layer for any RTP_Rx X, because the SSM receiver does not even get an opportunity to report its packet loss by means of an RTCP NACK. Note that for packet losses upstream of the DS and which subsequently impact all RTP_Rx's, the DS does not need to be informed by each individual RTP_Rx about this packet loss. When the DS is capable of recovering the lost packet in due time, it may send an unsolicited retransmission to all RTP_Rx's. The reduced visibility to the RS of packet losses taking place downstream of the DS, could be compensated for by having the RS sending an unsolicited retransmission, whenever an RTCP NACK is received. This will increase the chance on network congestion and will also consume RS processing resources when retransmissions are provided over unicast. The conclusion is that reflection/forwarding of RTCP NACK messages by the DS is not (always) the "right" thing to do.

5.2. Case of DS Terminating All RTCP NACKs

Not forwarding and hence terminating RTCP feedback messages is the alternative option for a DS behaving in [RFC5760] Feedback Summary Model. In this case every RTP_Rx that detects a packet loss event, will also transmit an RTCP NACK (assuming no congestion is sensed by the receiver, see Section 6.4.2). We look at two example scenarios resulting in different consequences: As a first example: two or

multiple RTP_Rx's detect the same packet loss, which was caused by two or multiple un-related packet loss events taking place (quasi-) simultaneously. This could be caused by link transmission errors, such as on xDSL links or in cable networks. In this case, the RS will receive the NACKs from all impacted receivers, and can act accordingly by sending a unicast (solicited) retransmission to the impacted receivers. As a second example, take the case considered previously where a packet loss event took place on the link between a media sender and the DS. Because, even though the transmissions could be dispersed in a time frame ΔT , when a very large number of SSM receivers are impacted by the single packet loss event, a feedback implosion might occur. This may load the network link(s) downstream of the RS and the RS itself, beyond its processing capabilities. Clearly, this is a case where the feedback suppression rule of [RFC4585] would be useful. The conclusion is that termination of RTCP NACK messages by the DS is not (always) the "right" thing to do.

5.3. Case of Multiple FTs

A specific case of SSM with unicast feedback, is where there are multiple FTs, disjoint from the DS. Similar as before, in the considered architectures, each FT is combined with a retransmission source, constituting a retransmission server. Note that the RS are generally not positioned in the direct SSM path between the DS and the RTP_Rx's, i.e., an RS is typically a leaf of the SSM tree. This architecture provides a scalable solution for SSM with a large population of receivers, because it is able to distribute RTCP feedback processing loads across different entities in different parts of the network. It is an architecture that is well suited for IPTV networks, where the DS is the head-end sourcing the SSM that carries video broadcast streams over IP. Note that we again assume that the RTP_Rx's behave compliant to [RFC4585], inclusive conformance to the feedback suppression rule. The previous discussions on feedback suppression (or absence of feedback suppression) and its consequences on retransmissions for the SSM with single FT topology remains applicable and valid for the SSM with multiple (disjoint) FTs topology. A distributed FT architecture brings in an additional aspect that should be addressed, and a dedicated example packet loss event use case visualizes this. The considered topology is a DS with two disjoint FT/RS entities, FT/RS 1 and FT/RS 2, where each FT receives RTCP messages from a separate group of RTP_Rx's. The assumption is that

- o An RTP packet (with RTP sequence number N) in the primary SSM got dropped in the network upstream of the FT 1 and upstream of the SSM receivers that report to FT 1.

- o The FT 2 and the SSM receivers reporting to FT 2 are not impacted by the packet loss event impacting FT 1. This is because the FT 2 and its reporting SSM receivers do not get the original SSM packets from the DS via the router where the packet loss impacting FT 1 took place.

The FT 1 and its reporting SSM receivers experience a situation that is discussed in a previous section. Because the packet loss event impacts all SSM receivers reporting to FT 1, it is important that those receivers in general suppress sending an RTCP NACK. Hence having the FT 1 forwarding the first received RTCP NACK(s) from an RTP_Rx to the DS which then reflects/forwards the NACK over the original SSM, is the correct thing to do from that point of view. However, the reflection /forwarding of the NACK by the DS means that also the RTP_Rx's reporting to the FT 2 will suppress sending an RTCP NACK for packet N in the SSM. As a consequence, a packet loss event involving an RTP packet with SN N impacting a single RTP_Rx reporting to FT 2 / RS 2 -e.g. due to a link transmission error- will not be reported by this RTP_Rx to RS 2.

This means there is a discrepancy between the network reach of the suppression (covering all SSM receivers) and the actual network domain that was commonly impacted by the packet loss. The RS 2 will in general not know whether there are any SSM receivers (reporting to FT 2) that missed RTP packet with RTP SN N. Note also that an unsolicited retransmission by RS 1 optionally following the loss detection for packet with SN N -can remain confined to the subdomain impacted by the loss, when the FT is co-located with the RS (using either unicast retransmission sessions or a multicast retransmission session sourced by the RS).

SSM with multiple disjoint unicast FTs hence may result in efficient feedback storm suppression across all RTP_Rx's, but this also prevents any RTP_Rx from sending an RTCP NACK for detected packet loss, even when no feedback storm was imminent for the subdomain covered by a particular FT. A solution for maximizing the retransmission service fulfillment may be for the DS to also act as RS and always send retransmissions requested by a particular FT, over a separate retransmission SSM to all RTP_Rx's. However, this unnecessarily loads the network and requires all the SSM RTP receivers to join both a primary SSM and the multicast retransmission session.

5.4. Takeaways

The general conclusion from these three scenarios is that feedback suppression enabled by having the DS forwarding/reflecting RTCP Messages received from the RTP_Rx's or (disjoint) FTs is sometimes

the right thing to do, but sometimes feedback suppression is not appropriate and interferes in a negative way with the packet loss reporting and retransmission service.

6. Rules for SSM with Retransmissions

6.1. Feedback Suppression Behavioral Requirements for DS, FT and RTP_Rx's

The combination of the [RFC4585] suppression rule imposed on RTP_Rx's and the behavior of the DS and the FT(s) -when disjoint from the DS- as specified in [RFC5760], does result in non-ideal situations when deploying an SSM RTP session with Retransmission Server(s), especially for a large group of receivers. In this section we provide some rules for DS and FT on the one hand and RTP_Rx's on the other hand, which impose a deviation respectively from rules defined in [RFC5760] and [RFC4585] , and which results in:

- o All RTP_Rx's offering the opportunity to report detected packet loss by means of RTCP NACK when there is no risk on feedback implosion
- o RTCP NACK suppression (only) when RTCP NACK implosion is imminent
- o A better retransmission efficacy

The rules apply to the transmission and handling of RTCP NACK messages by RTP_Rx's and FT, but similar rules may be applicable for other RTCP feedback messages that may or may not trigger dedicated service actions.

Note that in the context of this document, the FT is part of a Retransmission Server, which is an RTP_Rx on its own. In the remainder of this text, the term 'RTP_Rx' always refers to any RTP receiver reporting to a FT/RS and which does not host a FT/RS functionality, unless mentioned otherwise.

The rules are:

1. A FT/RS that is disjoint from the DS MUST NOT forward nor reflect RTCP NACK messages received from the RTP_Rx's, to the DS.

This behavior for a FT operating in an SSM enabled with retransmission is different from the behavior prescribed by [RFC5760] for a FT with a primary SSM operating in Simple Feedback mode. This rule will prevent feedback suppression at the RTP_Rx's directly triggered by RTCP feedback messages sent by

other RTP_Rx's.

2. A disjoint FT MAY send either an RTCP NACK or a 3rd party loss RTCP feedback message, each time using its own SSRC, to the DS.

Note : The third party loss message format is defined in [I-D.ietf-avtcore-feedback-supression-rtp].

This behavior for a FT operating in an SSM enabled with retransmission is different from the behavior prescribed by [RFC5760] for a FT with a primary SSM operating in Simple Feedback mode.

This rule will enable feedback suppression in the group of RTP_Rx's reporting to this FT (see recommendation 5).

The FT/RS will send an RTCP NACK or 3rd party loss message depending on the location of the packet loss event:

- * The FT/RS SHOULD send an RTCP NACK to the DS -using its own SSRC- when it detected packet loss itself (with the FT/RS acting as an RTP_Rx).
- * The FT/RS SHOULD send an RTCP NACK or a third party loss message to the DS -using its own SSRC- when it did not detect packet loss itself, but there is a risk on feedback implosion.

Feedback implosion risk detection by the FT without detecting directly packet loss, can be based on monitoring the amount of arriving RTCP NACKs from RTP_Rx's reporting the same packet loss, and setting a threshold as function of the known amount of RTP SSM receivers reporting to this FT/RS.

3. A DS MUST forward/reflect any RTCP NACK or RTCP third party loss feedback message received from FT entities on the primary SSM RTP session
4. When capable of detecting packet loss in the RTP streams received from the media senders, the DS MUST send an RTCP NACK using its own SSRC on the primary SSM RTP session and to the media senders
5. An RTP_Rx receiving a third party loss message or RTCP NACK in the primary SSM RTP session MUST conform to the following behavior:

If the SSRC identifier in the 'packet sender SSRC' field in the received RTCP feedback message matches either

- * the SSRC of the (primary SSM) DS (this is the 'media sender SSRC')
- * or the SSRC that is used by the FT/RS entity to which the RTP_Rx reports (see next section)

the RTP_Rx MUST follow the feedback suppression rule defined in [RFC4585], i.e. it abstains from sending an RTCP NACK for the same SN.

If there is no match with these SSRC values, the RTP_Rx SHOULD NOT apply the [RFC4585] feedback suppression rule

This behavior is different from [RFC4585], and prevents feedback suppression when there is no risk on feedback implosion.

If the RTP_Rx does not know the value of the SSRC that is used by the FT/RS entity -in its role towards the DS- to which the receiver reports, it MAY still send its own RTCP feedback message, when having detected the same packet loss, but respecting the [RFC4585] timing rules.

RTP_Rx's that do not know the value of the SSRC that is used by its FT/RS entity in its primary SSM role and abstain from sending an RTCP NACK message when receiving such a message, implement feedback suppression and behave as defined in [RFC4585].

The next section describes how an RTP_Rx knows the value of the SSRC that is used by the FT/RS entity to which the receiver reports

6.2. Informing RTP_Rx's with the SSRC Value of FT/RS

There are different SSRC values that can be associated to the FT/RS:

- o As described in [RFC5760], in summary feedback model, the FT must act as a receiver and choose a unique SSRC for its own reporting towards the Distribution Source. If a FT acts strictly according to the Simple Feedback Model of [RFC5760], it does not need a unique SSRC. However, with the recommendation '2' of the previous section, the FT acts rather as translator when sending its own RTCP NACK or Third Party Loss message to the DS. The FT selects a SSRC value in its role as translator between the RTP_Rx's and the DS.
- o The FT is part of the RS and the RS will typically also act as RTP_Rx, being the most efficient way for obtaining and caching copies of RTP packets that are transmitted down the primary SSM

and providing retransmission from that RS. The RS selects a SSRC value in its role as RTP_Rx as well, when reporting to the DS. Nothing prevents this SSRC value to be the same as the SSRC value used by the FT entity.

- o An RS also takes a role as RTP sender towards each of the RTP_Rx's that established a unicast retransmission session. In this unicast retransmission session, the SSRC used by the RS is the same as the SSRC value used by the DS in the primary SSM RTP session, as per [RFC4588].

Note that the CNAME of the FT/RS associated with all these SSRC identifier values is one and the same.

It is the SSRC value used by the FT acting as receiver/translator in the SSM with unicast feedback topology, that will be present in the 'Packet Sender SSRC' field of the RTCP NACK or Third Party Loss message forwarded by the DS to the RTP_Rx's. There can be several ways by which an RTP_Rx learns this SSRC value:

1. In-band : applies only for the Feedback Summary Model, where by means of a new RSI "FT SSRC list" message, the DS provides a listing of all deployed FT_Ap's with the corresponding SSRC for each of these FTs/RSs. FTs/RSs are identified by means of their IPv4/IPv6 address, optional port, and optional CNAME. The DS can learn the SSRC value chosen by a FT/RS by inspecting RTCP messages received from the FT. The DS sends regularly the RSI "FT SSRC list" message down the primary SSM, to give every RTP_Rx joining the SSM at a random time a chance to learn this SSRC value.
2. Application-specific : a specific extension can be defined for the RAMS-I message that signals the SSRC value associated with the FT in its role as reporter to the DS. This RAMS-I message is transferred in the Burst/Retransmission Session.
3. Out-of-band for either the Feedback Summary or Simple Feedback Model of SSM operation, by advertising the FT's SSRC as a media attribute for the FT in the SSM RTP session description [RFC5576].

The out-of-band signaling mechanism requires the application signaling to know/learn the SSRCS deployed by the FTs prior to signaling this information to the RTP_Rx's.

6.3. Unsolicited Retransmissions

The method of providing an (unsolicited) RTP retransmission packet to a set of RTP_Rx's enables loss recovery for packet loss events impacting a large set of receivers without risk for feedback implosion. Unsolicited retransmissions can be provided over each of the established unicast retransmission sessions or only once in a single multicast retransmission session, sourced either by the DS or by the RS. Providing (unsolicited) retransmissions over a separate multicast retransmission RTP session has the following advantages:

- o The retransmission packet is transmitted by the DS/RS only once, saving both network and server resources.
- o The multicast retransmission packet is less likely to be blocked by any intermediate gateway or a middlebox on the path between the DS/RS and RTP_Rx.
- o An RTP_Rx that joins the multicast retransmission RTP session via SFGMP, implicitly indicates it is willing to accept unsolicited retransmissions.

The rules related to the feedback suppression in Section 6.1 do not impose an RS/DS implementation to provide (support for) unsolicited retransmissions. In this document, we assume that whenever the FT/RS/DS supports solicited retransmission and it actively suppresses feedback by sending a RTCP NACK or 3rd party loss message, it might try to provide one or more subsequent unsolicited retransmissions. However, as it is the case with any RTP retransmission, such unsolicited retransmissions may or may not arrive at their destinations.

The requirements for unsolicited retransmissions are:

1. Unsolicited retransmissions MAY be provided over a separate multicast retransmission that is sourced by the RS that also acts as retransmission server for the unicast retransmission sessions or by the DS that also sources the primary SSM session.
2. An RS SHOULD only send unsolicited retransmissions when it knows with a minimum certainty that the receiver has implemented feedback suppression for this specific packet (unicast retransmission) or a majority of the receivers have implemented feedback suppression (multicast retransmission).

The same retransmission of a packet originally transmitted in the primary SSM can be sent in both a unicast retransmission session and multicast retransmission session. The RTP_Rx will handle these two

(solicited or unsolicited) packets as duplicate packets.

Editor's note: Should we define a means by which an RTP_Rx can indicate it does not want to receive unsolicited retransmissions in the unicast retransmission session?

6.4. Congestion Control Considerations

6.4.1. Congestion Control in RFC 4585 and in RAMS

The RAMS document indicates that an RTP SSM distribution equipped with a retransmission/burst server can work reliably in both engineered and best-effort networks. However, there is a difference in the core operation of RAMS and the core operation of a Retransmission-enabled SSM in terms of congestion risk and avoidance:

1. In RAMS, retransmission packets, constituting a burst from the RAMS server, are requested by an RTP_Rx prior to the moment the receiver joins the SSM session. The Retransmission server bursts the retransmission packets to the receiver and instructs the receiver when to join the SSM, which will take place near the end of this burst. The congestion control for the bursting process can hence be decoupled from the SSM distribution (and any congestion control taking place on the SSM). RAMS discusses in more detail how to limit the risk on congestion for the RAMS bursts, especially in a best effort environment.
2. When using retransmission service for an SSM, retransmission packets are transmitted to and received by an RTP_Rx, while the RTP_Rx also receives simultaneously the SSM RTP session. When used in best effort networks, packet losses can be caused by congestion or (access) transmission link errors. If packets in the SSM are dropped because of congestion, there is a good chance that a subsequent retransmission will actually exaggerate the congestion situation (this will depend on the relative positioning of the congestion, the RS and the RTP_Rx in the considered SSM topology).

[RFC4588] section 7 states: 'In a best-effort environment, the sender SHOULD NOT send retransmission packets without reducing the packet rate and bitrate of the original stream (for example, by encoding the data at a lower rate).' In the context of SSM, some RTP_Rx's may observe packet loss (caused by congestion), other receivers receiving the same SSM may not. As the DS transmits a single original stream to all SSM receivers, the recommendation from [RFC4588] for the DS to adapt its packet rate whenever an RTP_Rx is reporting packet loss, will impact all SSM receivers. This would make SSM with retransmission service practically infeasible in best

effort networks. It is also important to note that the retransmission server/sender may be different from the DS sending the original stream, especially in large-scale SSM applications. Hence, to a large extent the SSM path from DS to an SSM receiver may be decoupled from the network path between RS and SSM receiver. I.e. it is quite possible that the congestion on the SSM does not impact the network path taken by the retransmission packet(s). For these reasons, a different recommendation is in place in terms of congestion handling when retransmission is applied in SSM with unicast feedback.

6.4.2. Congestion Handling for SSM with Retransmissions

In this section we formulate a set of rules for the RS in order to be cautious that retransmissions will not contribute to enhanced congestion situations on the SSM path between DS and RTP_Rx, but without imposing an explicit role to the RS to assist in (helping to) relieve any (potential) congestion between DS and RTP_Rx. This results in a congestion control mechanism for an RS operating in an SSM deployment that is decoupled from the behavior of the DS which may use its own congestion control.

Note: in this section only congestion on the path from the RS towards the RTP_Rx is considered. Congestion on the upstream path towards the FT/RS has been discussed in the context of feedback implosion risk and avoidance in the previous sections.

Case: Unmanaged Networks

The general requirement is that for unmanaged networks, an RTP retransmission server for an SSM RTP session SHOULD keep track, on a per RTP_Rx basis, of the recent retransmission request pattern for that RTP_Rx. The RS then behaves as follows:

- o Initially and when an RTP_Rx sends a retransmission request for the first time to the FT/RS, the RS SHOULD assume there is no congestion between the DS and the RTP_Rx, unless the RS has information deduced or received not by way of RTCP NACKs through which it can assume the SSM path is likely congested.

Such non-RTCP feedback information can be information coming from RTCP RR, other RTCP messages or any out-of-band signaling from the receiver or other network elements.

- o An RS MAY respond to a retransmission request with one or more retransmission packets when the path from the DS to the RTP_Rx originating the request is assumed to be congestion free.

- o When an RS observes an increasing amount of retransmission requests from a single receiver, while having serviced all or most retransmission requests from that receiver, or the server receives information in a way different from RTCP NACKs which predicts or indicates that congestion may likely occur, it SHOULD stop responding to retransmission requests from that receiver. Equally, the RS SHOULD NOT send unicast unsolicited retransmissions to such receiver.
- o When no other information is available, an RS MAY assume absence of congestion for an RTP_Rx when it has not received retransmission requests for a minimum duration of time.
- o An RS sending unsolicited retransmissions over a dedicated retransmission multicast SHOULD only do so when it knows this will not cause or contribute to congestion on the primary SSM path for the large majority of primary SSM receivers.

Detailed algorithms by which a RS can declare an RTP_Rx in a congestion state, or free from congestion, are not elaborated further upon in this document.

Note 1: an RS may correlate information it has on RTP_Rx's to make a judgement on congestion or congestion-free state of the SSM paths to other RTP_Rx's using the same RS.

Note 2: any congestion handling policy is orthogonal to any other policy maintained at the RS having to do with retransmission service access control e.g. using access control lists, IP anti-spoofing measures, authentication, etc.

Note 3: an algorithm similar to the one described above for congestion handling, but to handle a denial of service attack MAY be considered for the RS implementation, next to the other security considerations related to retransmission as discussed in section 12 in [RFC4588].

Equivalent requirements are formulated for an RTP_Rx in an unmanaged network, where retransmission is enabled:

- o An RTP_Rx which senses that the net throughput of the combination of original RTP packets and retransmission packets decreases over time with increasing amount of retransmission requests and subsequent retransmissions, SHOULD cease making retransmission requests.
- o An RTP_Rx may re-start sending out retransmission requests once the original SSM session is received with close-to-zero packet

loss, and where any optional sporadic packet loss can be assumed to be attributed to link transmission errors.

Case: Managed Networks

For managed networks congestion should not happen on the SSM path or otherwise be a rare event. [RFC4588] states : 'If enhanced service is used, it should be made sure that the total bitrate and packet rate do not exceed that of the requested service. It should be further monitored that the requested services are actually delivered.'

Note: The total bitrate and packet rate is the original data (or primary data) and retransmission data.

This means that the network is engineered to make sure that the combined SSM packets and retransmission packets on the network paths to any RTP_Rx will not be impacted by congestion.

An alternative approach is to engineer the network for the congestion-free transmission of the primary SSM data, where-as retransmissions are delivered without dedicated bandwidth engineering provisions or otherwise with limited engineering provisions, in a sense that up to a certain maximum rate and up to certain maximum burst size, the retransmission traffic can be delivered on a congestion-free path.

In the engineered case for SSM where retransmissions are not taken into account for the bandwidth provisioning, the chance on having congestion in the retransmission session for SSM enabled with retransmission-only service is expected to be smaller compared to an SSM enabled with RAMS service. This is because a single RAMS-R request results in a burst of packets in the retransmission session whereas retransmissions requests typically results in the transmission of only one or several retransmission packets, with retransmission requests being sporadic.

6.5. New RSI Message and RAMS-I TLV Extension

6.5.1. RSI Message 'FT SSRC List'

TBD.

6.5.2. RAMS-I TLV Extension 'FT SSRC'

TBD.

7. Security Considerations

Editor's note: A subset of the security aspects discussed in the RAMS specification also applies to this document. TBC.

8. IANA Considerations

TBC.

9. RTP Retransmission Service in DVB IPTV

RTP retransmission for linear IPTV streams has been specified in Annex D of [DVB-IPTV]. This ETSI document, produced by the DVB Technical Module on IP-Infrastructure (TM-IPI) specifies the interface of a managed IPTV service client, commonly called the Home Network End Device or HNED in DVB IPTV terms. The DVB IPTV RTP retransmission solution in Annex D references [RFC3550], [RFC4585], [RFC4588] and the [RFC5760]. Its solution is hence very well aligned with the IETF RTP specifications and addresses feedback implosion and feedback suppression in a similar way as described in this document:

- o The DVB Linear Media Broadcast (LMB) services are delivered over an RTP SSM that operates according to the summary feedback model defined in [RFC5760]
- o RTCP NACK messages transmitted by HNEDs towards the FT('s) are not forwarded over the primary RTP SSM.
- o An HNED enables feedback suppression when receiving an RTCP Feedforward message from the network. This RTCP Feedforward (FF) message has the same format as an RTCP NACK message.
- o The RTCP FF message stems from a 'so-called' upstream reporter (having its own SSRC identifier, signaled to the HNEDs in a DVB IPTV specific way) , which represents the Retransmission Server in its role as primary RTP_Rx. Alternatively the RTCP FF message is transmitted by the RS acting as RTP translator when the RS assumes a packet loss event took place that is impacting many HNEDs, but not the RS (as RTP_Rx) itself.
- o This RTCP FF message may be sent in the primary RTP SSM but may also be sent in a dedicated RTP SSM sourced by the Retransmission Server.
- o In order to better control and allow the Retransmission Server to better recognize potential feedback implosion situations, an

optional feature has been defined. A subset of RTP_Rx can act as immediate reporters (which do not follow the dithering transmission timing rules from [RFC4585]), and will send an RTCP NACK message immediately upon packet loss detection. Whether an HNED acts as immediate reporter or not is determined by whether its randomly chosen SSRC identifier value maps to a pre-defined SSRC mask/pattern that is signaled to all HNEDs.

10. Contributors

TBC.

11. Acknowledgments

TBC.

12. References

12.1. 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.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [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.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, July 2006.
- [RFC5760] Ott, J., Chesterfield, J., and E. Schooler, "RTP Control Protocol (RTCP) Extensions for Single-Source Multicast

Sessions with Unicast Feedback", RFC 5760, February 2010.

[RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, August 2006.

[I-D.ietf-avtcore-ports-for-ucast-mcast-rtp]
Begen, A., Wing, D., and T. VanCaenegem, "Port Mapping Between Unicast and Multicast RTP Sessions",
draft-ietf-avtcore-ports-for-ucast-mcast-rtp-02 (work in progress), April 2011.

[I-D.ietf-avt-rapid-acquisition-for-rtp]
Steeg, B., Begen, A., Caenegem, T., and Z. Vax, "Unicast-Based Rapid Acquisition of Multicast RTP Sessions",
draft-ietf-avt-rapid-acquisition-for-rtp-17 (work in progress), November 2010.

[I-D.ietf-avtcore-feedback-supression-rtp]
Wu, W., Xia, F., and R. Even, "RTCP Extension for Third-party Loss Report",
draft-ietf-avtcore-feedback-supression-rtp-03 (work in progress), May 2011.

12.2. Informative References

[RFC5576] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, June 2009.

[I-D.ietf-fecframe-dvb-al-fec]
Begen, A. and T. Stockhammer, "Guidelines for Implementing DVB-IPTV Application-Layer Hybrid FEC Protection",
draft-ietf-fecframe-dvb-al-fec-04 (work in progress),
December 2009.

[DVB-IPTV]
"ETSI TS 102034 Digital Video Broadcasting (DVB);
Transport of MPEG-2 TS Based DVB Services over IP Based
Networks", August 2009.

Authors' Addresses

Tom VanCaenegem
Alcatel-Lucent
Copernicuslaan 50
Antwerpen, 2018
Belgium

Email: Tom.Van_Caenegem@alcatel-lucent.com

Bill VerSteeg
Cisco
5030 Sugarloaf Parkway
Lawrenceville, GA 30044
USA

Email: billvs@cisco.com

Ali Begen
Cisco
181 Bay Street
Toronto, ON M5J 2T3
Canada

Email: abegen@cisco.com

