AVT
Internet-Draft
Intended status: Standards Track
Expires: February 26, 2010

B. VerSteeg
A. Begen
Cisco Systems
T. VanCaenegem
Alcatel-Lucent
Z. Vax
Microsoft Corporation

August 25, 2009

Unicast-Based Rapid Acquisition of Multicast RTP Sessions
draft-ietf-avt-rapid-acquisition-for-rtp-02

Status of this Memo

Copyright Notice

Abstract

When an RTP receiver joins a primary multicast session, it may need to acquire and parse certain Reference Information before it can process any data sent in the multicast session. Depending on the join time, length of the Reference Information repetition interval, size of the Reference Information as well as the application and transport properties, the time lag before an RTP receiver can usefully consume the multicast data, which we refer to as the Acquisition Delay, varies and may be large. This is an undesirable phenomenon for receivers that frequently switch among different multicast sessions, such as video broadcasts.

In this document, we describe a method using the existing RTP and RTCP protocol machinery that reduces the acquisition delay. In this method, an auxiliary unicast RTP session carrying the Reference Information to the receiver precedes/accompanies the primary multicast stream. This unicast RTP flow may be transmitted at a faster than natural rate to further accelerate the acquisition. The motivating use case for this capability is multicast applications that carry real-time compressed audio and video. However, the proposed method can also be used in other types of multicast applications where the acquisition delay is long enough to be a problem.

Table of Contents

1. Introduction

Most multicast flows carry a stream of inter-related data. Certain information must first be acquired by the receivers to start processing any data sent in the multicast session. This document refers to this information as Reference Information. The Reference Information is conventionally sent periodically in the multicast session and usually consists of items such as a description of the schema for the rest of the data, references to which data to process, encryption information including keys, as well as any other information required to process the data in the primary multicast stream.

Real-time multicast applications require the receivers to buffer data. The receiver may have to buffer data to smooth out the network jitter, to allow loss-repair methods such as Forward Error Correction and retransmission to recover the missing packets, and to satisfy the data processing requirements of the application layer.

When a receiver joins a multicast session, it has no control over what point in the flow is currently being transmitted. Sometimes the receiver may join the session right before the Reference Information is sent in the session. In this case, the required waiting time is usually minimal. Other times, the receiver may join the session right after the Reference Information has been transmitted. In this case, the receiver has to wait for the Reference Information to appear again in the flow before it can start processing any multicast data. In some other cases, the Reference Information is not contiguous in the flow but dispersed over a large period, which forces the receiver to wait for all of the Reference Information to arrive before starting to process the rest of the data.

The net effect of waiting for the Reference Information and waiting for various buffers to fill up is that the receivers may experience significantly large delays in data processing. In this document, we refer to the difference between the time an RTP receiver joins the multicast session and the time the RTP receiver acquires all the necessary Reference Information as the Acquisition Delay. The acquisition delay may not be the same for different receivers; it usually varies depending on the join time, length of the Reference Information repetition interval, size of the Reference Information as well as the application and transport properties.

The varying nature of the acquisition delay adversely affects the receivers that frequently switch among multicast sessions. In this specification, we address this problem for RTP-based multicast applications and describe a method that uses the fundamental tools offered by the existing RTP and RTCP protocols [RFC3550]. In this

method, either the multicast source (or the distribution source in a single-source multicast (SSM) session) retains the Reference Information for a period after its transmission, or an intermediary network element (that we refer to as Retransmission Server) joins the multicast session and continuously caches the Reference Information as it is sent in the session and acts as a feedback target (See [I-D.ietf-avt-rtcpssm]) for the session. When an RTP receiver wishes to join the same multicast session, instead of simply issuing a Source Filtering Group Management Protocol (SFGMP) Join message, it sends a request to the feedback target for the session asking for the Reference Information. The Retransmission Server starts a unicast retransmission RTP session and sends the Reference Information to the RTP receiver over that session. If there is spare bandwidth, the Retransmission Server may also burst the Reference Information at a faster than its natural rate. As soon as the receiver acquires the Reference Information, it can join the multicast session and start processing the multicast data. This method potentially reduces the acquisition delay. We refer to this method as Unicast-based Rapid Acquisition of Multicast RTP Sessions. A simplified network diagram showing this method through an intermediary network element is depicted in Figure 1.

```
                                 +----------------------+
                            +--->|     Intermediary     |
                            | ...|    Network Element   |
                            | :  |(Retransmission Server)|
                            | :  +----------------------+
                            | :
                            | v
       +-----------+     +----------+          +----------+
       | Multicast |     |  Router  |--------->| Joining  |
       |  Source   |---->|          |.........>|   RTP    |
       +-----------+     +----------+          | Receiver |
                              |                +----------+
                              |
                              |                +----------+
                              +--------------->| Existing |
                                               |    RTP   |
                                               | Receiver |
                                               +----------+

       ...> Unicast RTP Flow
       ---> Multicast RTP Flow
```

Figure 1: Rapid acquisition through an intermediary network element

A primary design goal in this solution is to use the existing tools

in the RTP/RTCP protocol family. This improves the versatility of the existing implementations, and promotes faster deployment and better interoperability. To this effect, we use the unicast retransmission support of RTP [RFC4588] and the capabilities of RTCP to handle the signaling needed to accomplish the acquisition. The packet(s) carrying the Reference Information are sent by the Retransmission Server in the auxiliary unicast RTP session for rapid acquisition. These are constructed as retransmission packets that would have been sent in a unicast RTP session to recover the missing packets at an RTP receiver that has never received any packet. In fact, a single RTP session SHOULD be used for both rapid acquisition and retransmission-based loss repair. This session can be used to simultaneously provide the unicast burst for the rapid acquisition and the repair packets requested by the RTP receivers when they detect lost burst packets or lost RTP packets in the primary multicast stream. The conventional RTCP feedback (NACK) message that requests the retransmission of the missing packets [RFC4585] indicates their sequence numbers. However, upon joining a new session the RTP receiver has never received a packet, and thus, does not know the sequence numbers. Instead, the RTP receiver sends a newly defined RTCP feedback message to request the Reference Information needed to rapidly get on the track with the primary multicast session. It is also worth noting that in order to issue the initial RTCP message to the feedback target, the SSRC of the session to be joined must be known prior to any packet reception, and hence, needs to be signaled out-of-band (or otherwise communicated to the RTP receiver in advance of the initiation of the rapid acquisition operation). In a Session Description Protocol (SDP) description, the SSRC MUST be signaled through the 'ssrc' attribute [I-D.ietf-avt-rtcpssm].

In the rest of this specification, we have the following outline: In Section 4, we describe the delay components in generic multicast applications. Section 5 presents an overview of the protocol design considerations for rapid acquisition. We provide the protocol details of the rapid acquisition method in Section 6 and Section 7. Section 8 and Section 9 discuss the SDP signaling issues with examples and NAT-related issues, respectively.

Note that Section 3 provides a list of the definitions frequently used in this document.

2. Requirements Notation

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

3. Definitions

This document uses the following acronyms and definitions frequently:

Primary Multicast Session: The multicast RTP session to which RTP receivers can join at a random point in time.

Primary Multicast Stream: The RTP stream carried in the primary multicast 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. However, the rapid acquisition method introduced in this document does not depend on a specific version of either of these group management protocols. In the remainder of this document, SFGMP will refer to any group management protocol that has Join and Leave functionalities.

Feedback Target: (Unicast RTCP) Feedback target as defined in [I-D.ietf-avt-rtcpssm].

Retransmission Packet: An RTP packet that is formatted as defined in [RFC4588].

Reference Information: The set of certain media content and metadata information that is sufficient for an RTP receiver to start usefully consuming a media stream. The meaning, format and size of this information are specific to the application and are out of scope of this document.

(Unicast) Burst (Stream): A unicast stream of RTP retransmission packets that enable an RTP receiver to rapidly acquire the Reference Information. The burst stream is typically transmitted at an accelerated rate.

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

4. Elements of Delay in Multicast Applications

In an any-source (ASM) or a single-source (SSM) multicast delivery system, there are three major elements that contribute to the overall acquisition delay when an RTP receiver switches from one multicast session to another one. These are:

o Multicast switching delay

o Reference Information latency

o Buffering delays

Multicast switching delay is the delay that is experienced to leave the current multicast session (if any) and join the new multicast session. In typical systems, the multicast join and leave operations are handled by a group management protocol. For example, the receivers and routers participating in a multicast session may use the Internet Group Management Protocol (IGMP) version 3 [RFC3376] or the Multicast Listener Discovery Protocol (MLD) version 2 [RFC3810]. In either of these protocols, when a receiver wants to join a multicast session, it sends a message to its upstream router and the routing infrastructure sets up the multicast forwarding state to deliver the packets of the multicast session to the new receiver. Depending on the proximity of the upstream router, the current state of the multicast tree, the load on the system and the protocol implementation, the join times vary. Current systems provide join latencies usually less than 200 milliseconds (ms). If the receiver had been participating in another multicast session before joining the new session, it needs to send a Leave message to its upstream router to leave the session. In common multicast routing protocols, the leave times are usually smaller than the join times, however, it is possible that the Leave and Join messages may get lost, in which case the multicast switching delay inevitably increases.

Reference Information latency is the time it takes the receiver to acquire the Reference Information. It is highly dependent on the proximity of the actual time the receiver joined the session to the next time the Reference Information will be sent to the receivers in the session, whether the Reference Information is sent contiguously or not, and the size of the Reference Information. For some multicast flows, there is a little or no interdependency in the data, in which case the Reference Information latency will be nil or negligible. For other multicast flows, there is a high degree of interdependency. One example of interest is the multicast flows that carry compressed audio/video. For these flows, the Reference Information latency may become quite large and be a major contributor to the overall delay. Refer to [I-D.begen-avt-rtp-mpeg2ts-preamble] for details.

The buffering component of the overall acquisition delay is driven by the way the application layer processes the payload. In many multicast applications, an unreliable transport protocol such as UDP [RFC0768] is often used to transmit the data packets, and the reliability, if needed, is usually addressed through other means such

as Forward Error Correction and retransmission [I-D.ietf-rmt-pi-norm-revised]. These loss-repair methods require buffering at the receiver side to function properly. In many applications, it is also often necessary to de-jitter the incoming data packets before feeding them to the application. The de-jittering process also increases the buffering delays. Besides these network-related buffering delays, there are also specific buffering needs that are required by the individual applications. For example, standard video decoders typically require an amount, sometimes a significant amount, of coded video data to be available in the pre-decoding buffers prior to starting to decode the video bitstream.

5. Protocol Design Considerations and Their Effect on Resource Management for Rapid Acquisition

Rapid acquisition is an optimization of a system that must continue to work correctly whether or not the optimization is effective, or even fails due to lost control messages, congestion, or other problems. This is fundamental to the overall design requirements surrounding the protocol definition and to the resource management schemes to be employed together with the protocol (e.g., QoS machinery, server load management, etc). In particular, the system needs to operate within a number of constraints:

- First, a rapid acquisition operation must fail gracefully. The user experience must, except perhaps in pathological circumstances, be not significantly worse for trying and failing to complete rapid acquisition compared to simply joining the multicast session.

- Second, providing the rapid acquisition optimizations must not cause collateral damage to either the multicast session being joined, or other multicast sessions sharing resources with the rapid acquisition operation. In particular, the rapid acquisition operation must avoid self-interference with the multicast session that may be simultaneously being received by other hosts. In addition, it must also avoid interference with other multicast sessions sharing the same network resources. These properties are possible, but are usually difficult to achieve.

One challenge is the existence of multiple bandwidth bottlenecks between the receiver and the server(s) in the network providing the rapid acquisition service. In commercial IPTV deployments, for example, bottlenecks are often present in the aggregation network connecting the IPTV servers to the network edge, the access links (e.g., DSL, DOCSIS) and in the home network of the subscribers. Some of these links may serve only a single subscriber, limiting

congestion impact to the traffic of only that subscriber, but others can be shared links carrying multicast sessions of many subscribers. Also note that the state of these links may be varying over time. The receiver may have knowledge of a portion of this network, or may have partial knowledge of the entire network. The methods employed by the devices to acquire this network state information is out of scope for this document. The receiver should be able to signal the server with the bandwidth that it believes it can handle. The server also needs to be able to rate limit the flow in order to stay within the performance envelope that it knows about. Both the server and receiver need to be able to inform the other of changes in the requested and delivered rates. However, the protocol must be robust in the presence of packet loss, so this signaling must include the appropriate default behaviors.

A second challenge is that for some uses (e.g., high-bitrate video) the unicast burst bitrate is high while the flow duration of the unicast burst is short. This is because the purpose of the unicast burst is to allow the RTP receiver to join the multicast quickly and thereby limit the overall resources consumed by the burst. Such high-bitrate, short-duration flows are not amenable to conventional admission control techniques. For example, end-to-end per-flow signaled admission control techniques such as RSVP have too much latency and control channel overhead to be a good fit for rapid acquisition. Similarly, using a TCP (or TCP-like) approach with a 3-way handshake and slow-start to avoid inducing congestion would defeat the purpose of attempting rapid acquisition in the first place by introducing many RTTs of delay.

These observations lead to certain unavoidable requirements and goals for a rapid acquisition protocol. These are:

o The protocol must be designed to allow a deterministic upper bound on the extra bandwidth used (compared to just joining the multicast session). A reasonable size bound is e*B, where B is the "nominal" bandwidth of the primary multicast stream, and e is an "excess-bandwidth" coefficient The total duration of the unicast burst must have a reasonable bound; long unicast bursts devolve to the bandwidth profile of multi-unicast for the whole system.

o The scheme should minimize (or better eliminate) the overlap of the unicast burst and the primary multicast stream. This minimizes the window during which congestion could be induced on a bottleneck link compared to just carrying the multicast or unicast packets alone.

- The scheme must minimize (or better eliminate) any gap between the unicast burst and the primary multicast stream, which has to be repaired later, or in the absence of repair, will result in loss being experienced by the application.

In addition to the above, there are some other protocol design issues to be considered. First, there is at least one RTT of "slop" in the control loop. In starting a rapid acquisition burst, this manifests as the time between the client requesting the unicast burst and the burst description and/or the first unicast burst packets arriving at the receiver. For managing and terminating the unicast burst, there are two possible approaches for the control loop: The receiver can adapt to the unicast burst as received, converge based on observation and explicitly terminate the unicast burst with a second control loop exchange (which takes a minimum of one RTT, just as starting the unicast burst does). Alternatively, the server generating the unicast burst can pre-compute the burst parameters based on the information in the initial request and tell the receiver the burst duration.

The protocol described in the next section allows either method of controlling the rapid acquisition unicast burst.

6. Rapid Acquisition of Multicast RTP Sessions

We start this section with an overview of the rapid acquisition of multicast sessions (RAMS) method.

6.1. Overview

[I-D.ietf-avt-rtcpssm] specifies an extension to the RTP Control Protocol (RTCP) to use unicast feedback in an SSM session. It defines an architecture that introduces the concept of Distribution Source, which - in an SSM context - distributes the RTP data and redistributes RTCP information to all RTP receivers. This RTCP information is retrieved from the Feedback Target, to which RTCP unicast feedback traffic is sent. The Feedback Target MAY be implemented in one or more entities different from the Distribution Source, and different RTP receivers MAY use different Feedback Targets.

This document builds further on these concepts to reduce the acquisition time when an RTP receiver joins a multicast session at a random point in time by introducing the concept of the Burst Source and new RTCP feedback messages. The Burst Source has a cache where the most recent packets from the primary multicast session are continuously stored. When an RTP receiver wants to receive the

primary multicast stream, prior to joining the SSM session, it may first request a unicast burst from the Burst Source. In this burst, the packets are formatted as RTP retransmission packets [RFC4588] and carry the Reference Information. This information allows the RTP receiver to start usefully consuming the RTP packets sent in the primary multicast session.

Using an accelerated rate (as compared to the rate of the primary multicast stream) for the unicast burst implies that at a certain point in time, the payload transmitted in the unicast burst is going to be the same as the payload multicast in the SSM session, i.e., the unicast burst will catch up with the primary multicast stream. At this point, the RTP receiver no longer needs to receive the unicast burst and can join the primary multicast session. This method is referred to as the Rapid Acquisition of Multicast Sessions (RAMS).

This document proposes extensions to [RFC4585] for an RTP receiver to request a unicast burst as well as for additional control messaging that can be leveraged during the acquisition process.

6.2. Message Flows

Figure 2 shows the main entities involved in rapid acquisition:

- o Multicast Source

- o Feedback Target (FT)

- o Burst/Retransmission Source

- o RTP Receiver (RR)

```
             +-----------------------------------------------+
             |            Retransmission Server              |
             |                    (RS)                       |
             |  +-----------+ +---------------------------+  |
             |  | Feedback  | |    Burst/Retransmission   |  |
             |  |  Target   | |           Source          |  |
             |  +-----------+ +---------------------------+  |
             +-----------------------------------------------+
                                 ^ ^ :
                                 | ' :
                                 | ' :
                                 | ' v
         +-----------+        +----------+            +----------+
         |           |        |          |'''''''''''>|          |
         | Multicast |------->|  Router  |...........>|   RTP    |
         |  Source   |        |          |<~~~~~~~~~~~| Receiver |
         |           |        |          |----------->|   (RR)   |
         +-----------+        +----------+            +----------+


         '''> Unicast RTCP Messages
         ~~~> SFGMP Messages
         ...> Unicast RTP Flow
         ---> Multicast RTP Flow
```

Figure 2: Flow diagram for unicast-based rapid acquisition

The feedback target (FT) is the entity as defined in [I-D.ietf-avt-rtcpssm], to which RR sends its RTCP feedback messages indicating packet loss in the primary multicast stream by means of an RTCP NACK message or indicating RR's desire to rapidly acquire the primary multicast stream by means of an RTCP feedback message defined in this document. While the Burst/Retransmission Source is responsible for responding to these messages and for further RTCP interaction with RR in the case of a rapid acquisition process, it is assumed in the remainder of the document that these two logical entities (FT and Burst/Retransmission Source) are combined in a single physical entity and they share state. In the remainder of the text, the term Retransmission Server will be used whenever appropriate, to refer to the combined functionality of the FT and Burst/Retransmission Source.

However, it must be noted that only FT is involved in the primary multicast session, whereas the Burst/Retransmission Source transmits the unicast burst and retransmission packets both formatted as RTP retransmission packets [RFC4588] in a single separate unicast RTP retransmission session to each RR. In the retransmission session, as in any other RTP session, RS and RR regularly send the periodic

sender and receiver reports, respectively.

Note also that the same method (with the identical message flows) would also apply in a scenario where rapid acquisition is performed by a feedback target co-located with the media source.

The unicast burst is triggered by an RTCP feedback message that is defined in this document, whereas an RTP retransmission is triggered by an RTCP NACK message defined in [RFC4585]. Based on its design, in an RAMS implementation, there may be a gap between the end of the burst and the reception of the primary multicast stream because of the imperfections in the switch-over. If needed, RR may make use of the RTCP NACK message to request a retransmission for the missing packets in the gap.

Figure 3 depicts an example of messaging flow for rapid acquisition. The RTCP feedback messages are explained below. Note that the messages indicated in parentheses may or may not be present during rapid acquisition.

```
     +-----------+   +----------------+   +----------+   +------------+
     | Multicast |   | Retransmission |   |          |   |    RTP     |
     |  Source   |   |     Server     |   |  Router  |   |  Receiver  |
     |           |   |      (RS)      |   |          |   |    (RR)    |
     +-----------+   +----------------+   +----------+   +------------+
         |                   |                 |                |
         |-- RTP Multicast ------------------->|                |
         |                   |                 |                |
         |-- RTP Multicast ->|                 |                |
         |                   |                 |                |
         |                   |<'''''''''''''''''' RTCP RAMS-R ''|
         |                   |                 |                |
         |                   |                 |                |
         |                   |'' (RTCP RAMS-I) ''''''''''''''''>|
         |                   |                 |                |
         |                   |                 |                |
         |                   |.. Unicast RTP Burst ............>|
         |                   |                 |                |
         |                   |<''''''''''''''''''(RTCP RAMS-R)''|
         |                   |                 |                |
         |                   |                 |                |
         |                   |'' (RTCP RAMS-I) ''''''''''''''''>|
         |                   |                 |                |
         |                   |                 |                |
         |                   |                 |<~ SFGMP Join ~~|
         |                   |                 |                |
         |                   |                 |                |
         |-- RTP Multicast ------------------------------------>|
         |                   |                 |                |
         |                   |                 |                |
         |                   |<'''''''''''''''''' RTCP RAMS-T ''|
         |                   |                 |                |
         |                   |                 |                |
         |                   |<''''''''''''''''''' (RTCP NACK)''|
         |                   |                 |                |
         |                   |                 |                |
         |                   |.. (Unicast Retransmissions) ....>|
         |                   |                 |                |
         |                   |                 |                |
         |                   |<''''''''''''''''''' (RTCP BYE) ''|
         |                   |                 |                |
         |                   |                 |                |


     '''> Unicast RTCP Messages
     ~~~> SFGMP Messages
     ...> Unicast RTP Flow
     ---> Multicast RTP Flow
```

Figure 3: Message flows for unicast-based rapid acquisition

This document defines the expected behaviors of RS and RR. It is instructive to have independently operating implementations on RS and RR that request the burst, describe the burst, start the burst, join the multicast session and stop the burst. These implementations send messages to each other, but there MUST be provisions for the cases where the control messages get lost, or re-ordered, or are not being delivered to their destinations.

The following steps describe rapid acquisition in detail:

1. Request: RR sends a rapid acquisition request for the new multicast RTP session to the feedback target address of that session. The request contains the SSRC of RR and the SSRC of the media source. This RTCP feedback message is defined as the RAMS-Request (RAMS-R) message and MAY contain parameters, which may constrain the burst, such as the bandwidth limit. Other parameters may be related to the amount of buffering capacity available at RR, which may be used by RS to prepare a burst that conforms with RR's requirements.

   Before joining the primary multicast session, a new joining RR learns the addresses associated with the new multicast session (addresses for the multicast source, group and retransmission server) by out-of-band means. Also note that since no RTP packets have been received yet for this session, the SSRC must be obtained out-of-band. See Section 8 for details.

2. Response: RS receives the RAMS-R message and decides whether to accept it or not. RS MUST send an (at least one) RAMS-Information (RAMS-I) message to RR. The first RAMS-I message MAY precede the unicast burst or it MAY be sent during the burst. Additional RAMS-I messages MAY be sent during the burst and these RAMS-I messages may or may not be a direct response to an RAMS-R message. The RAMS-I message is sent by the Burst/Retransmission Source logical entity that is part of RS.

   Note that RS learns the IP address information for RR from the RAMS-R message it received. (This description glosses over the NAT details. Refer to Section 9 for a discussion of NAT-related issues.)

   If RS cannot provide a rapid acquisition service, RS rejects the request and informs RR immediately via an RAMS-I message. If RR receives a message indicating that its rapid acquisition request has been denied, it abandons the rapid acquisition attempt and MAY immediately join the multicast session by sending an SFGMP

Join message towards its upstream multicast router for the new primary multicast session.

If RS accepts the request, it sends an RAMS-I message to RR (before commencing the unicast burst or during the unicast burst) that comprises fields that can be used to describe the unicast burst (e.g., the maximum bitrate and the duration of the unicast burst). A particularly important, thus mandatory, field in the RAMS-I message carries the RTP sequence number of the first burst packet.

It is RECOMMENDED to include a sender report with the RAMS-I message in the same compound RTCP packet. This also allows rapid synchronization among multiple RTP flows [I-D.ietf-avt-rapid-rtp-sync].

The unicast burst duration MAY be calculated by RS, and its value MAY be updated by messages received from RR. The join time information (for the new multicast session) SHOULD be populated in at least one of the RAMS-I messages. Note that RS MAY send the RAMS-I message after a significant delay, so RR SHOULD NOT make protocol dependencies on quickly receiving an RAMS-I message.

3. Unicast Burst: If the request is accepted, RS starts sending the unicast burst that comprises one or more RTP retransmission packets (The burst packet(s) are sent by the Burst/Retransmission Source logical entity). In addition, there MAY be optional payload-specific information that RS chooses to send to RR. Such an example is discussed in [I-D.begen-avt-rtp-mpeg2ts-preamble] for transmitting the payload-specific information for MPEG2 Transport Streams.

4. Updated Request: RR MAY send a new RAMS-R message (to the FT entity of RS) with a different value for one or more fields of an earlier RAMS-R message. Upon receiving an updated request, RS MAY use the updated values for sending/shaping the burst, or refine the values and use the refined values for sending/shaping the burst.

   RS MAY send a new RAMS-I message to indicate the changes it made. However, note that RS does not have to send a new RAMS-I, or the new RAMS-I message may get lost. It is also possible that the updated RAMS-R message could have been lost. Thus, RR SHOULD NOT make protocol dependencies on quickly (or ever) receiving a new RAMS-I message, or assume that RS will honor the requested changes.

RR may be in an environment where the available resources are time-varying, which may or may not deserve sending a new updated request. Determining the circumstances where RR should or should not send an updated request and the methods that RR can use to detect and evaluate the time-varying available resources are not specified in this document.

5. Updated Response: RS may send more than one RAMS-I messages, e.g., to update the value of one or more fields in an earlier RAMS-I message and/or to signal RR in real time to join the primary multicast session. RR usually depends on RS to learn the join time, which can be conveyed by the first RAMS-I message, or can be sent/revised in a later RAMS-I message. If RS is not capable of determining the join time in the first RAMS-I message, it MUST send another RAMS-I message (with the join time information) later.

6. Multicast Join Signaling: In principal, RR can join the primary multicast session any time during or after the end of the unicast burst via an SFGMP Join message. However, there may be missing packets if RR joins the primary multicast session too early or too late. For example, if RR starts receiving the primary multicast stream while it is still receiving the unicast burst at a high excess bitrate, this may result in an increased risk of packet loss. Or, if RR joins the primary multicast session some time after the unicast burst is finished, there may be a gap between the burst and multicast data (a number of RTP packets may be missing). In both cases, RR MAY issue retransmissions requests (via RTCP NACK messages) [RFC4585] to fill the gap.

   Yet, there are cases where the remaining available bandwidth may limit the number of retransmissions that can be provided within a certain time period, causing the retransmission data to arrive too late at RR (from an application-layer point of view). To cope with such cases, the RAMS-I message allows RS to signal explicitly when RR should send the SFGMP Join message. Alternatively, RS may pre-compute the burst duration and the time RR should send the SFGMP Join message. This information may be conveyed in the RAMS-I message and can be updated in a subsequent RAMS-I message. While RR MAY use a locally calculated join time, it SHOULD use the information from the most recent RAMS-I message.

7. Multicast Receive: After the join, RR starts receiving the primary multicast stream.

8. Terminate: RS may know when it needs to stop the unicast burst based on the burst parameters, or RR MAY explicitly let RS know

the sequence number of the first RTP packet it received from the multicast session, or RR MAY request RS to terminate the burst immediately.

Regardless of whether or not RS knows when it needs to stop the burst, RR SHALL use the RAMS-Termination (RAMS-T) message at an appropriate time. RR can choose to send the RAMS-T message before or after it starts receiving the multicast data. In the latter case, RR SHALL include the sequence number of the first RTP packet received in the primary multicast session in the RAMS-T message, and RS SHOULD terminate the burst after it sends the unicast burst packet whose Original Sequence Number (OSN) field in the RTP retransmission payload header matches this number minus one.

If RR wants to stop the burst prior to receiving the multicast data, it sends an RAMS-T message without an RTP sequence number.

RR MUST send at least one RAMS-T message (if an RTCP BYE message has not been issued yet as described in Step 9), and the RAMS-T message MUST be addressed to the RTCP port of the retransmission session. Against the possibility of a message loss, RR MAY repeat the RAMS-T message multiple times as long as it follows the RTCP timer rules defined in [RFC4585].

9. Terminate with RTCP BYE: When RR is receiving the burst, if RR becomes no longer interested in the primary multicast stream, RR SHALL issue an RTCP BYE message for the RTP retransmission session and another RTCP BYE message for the primary multicast session.

   Upon receiving an RTCP BYE message, RS MUST terminate the rapid acquisition operation, and cease transmitting any further regular retransmission packets as well as retransmission packets associated with the unicast burst. If support for [RFC5506] has been signaled, the RTCP BYE message MAY be sent in a reduced-size RTCP packet. Otherwise, Section 6.1 of [RFC3550] mandates the RTCP BYE message always to be sent with a sender or receiver report in a compound RTCP packet (If no data has been received, an empty receiver report MUST be included). With the information contained in the receiver report, RS can also figure out how many duplicate RTP packets have been delivered to RR (Note that this will be an upper-bound estimate as one or more packets might have been lost during the burst transmission). The impact of duplicate packets and measures that can be taken to minimize the impact of receiving duplicate packets will be addressed in Section 6.3.

Note that an RTCP BYE message issued for the RTP retransmission session terminates the whole session and ceases transmitting any further packets in that RTP session. Thus, in this case there is no need for sending an (explicit) RAMS-T message, which would only terminate the burst.

Note that for the purpose of gathering detailed information about RR's rapid acquisition experience, [I-D.begen-avt-rapid-sync-rtcp-xr] defines an RTCP Extended Report (XR) Block. This report is designed to be payload-independent, thus, it can be used by any multicast application that supports rapid acquisition. Support for this XR report is, however, optional.

6.3. Shaping the Unicast Burst

This section provides informative guidelines about how RS can shape the transmission of the unicast burst.

A higher bitrate for the unicast burst naturally conveys the Reference Information and media content to RR faster. This way, RR can start consuming the data sooner, which results in a faster acquisition.

A higher rate also represents a better utilization of RS resources. As the burst may continue until it catches up with the primary multicast stream, the higher the bursting rate, the less data RS needs to transmit. However, a higher rate for the burst also increases the chances for congestion-caused packet loss. Thus, as discussed in Section 5, there must be an upper bound on the extra bandwidth used by the burst.

When RS transmits the burst, it SHOULD take into account all available information to prevent any packet loss that may take place during the bursting as a result of buffer overflow on the path between RS and RR and at RR itself. The bursting rate may be determined by taking into account the following data, when available:

a. Information obtained via the RAMS-R message, such as Max RAMS Buffer Fill Requirement and/or Max Receive Bitrate (See Section 7.2).

b. Information obtained via RTCP receiver reports provided by RR in the retransmission session, allowing in-session rate adaptations for the burst. When these receiver reports indicate packet loss, this may indicate a certain congestion state in the path from RS to RR. Heuristics or algorithms that deduce such congestion state and how subsequently the RS should act, are outside the scope of this document.

c. Information obtained via RTCP NACKs provided by RR in the primary multicast session, allowing in-session rate adaptations for the burst. Such RTCP NACKs are transmitted by RR in response to packet loss detection by RR in the burst. NACKs may indicate a certain congestion state on the path from RS to RR. Heuristics or algorithms that deduce such congestion state and how subsequently the RS should act, are outside the scope of this document.

d. There may be other feedback received from RR, e.g., in the form of ECN-CE RTCP feedback messages [I-D.westerlund-avt-ecn-for-rtp] that may influence in-session rate adaptations.

e. Information obtained via updated RAMS-R messages, allowing in-session rate adaptations, if supported by RS.

f. Pre-configured settings for each RR or a set of RRs that indicate the upper-bound bursting rates for which no packet loss will occur as a result of congestion along the path of RS to RR. For example, in managed IPTV networks, where the bottleneck bandwidth along the end-to-end path is known (which is generally the access network link) and where the network between RS and this link is provisioned and dimensioned to carry the burst streams, the bursting rate does not exceed the provisioned value. These settings may also be dynamically adapted using application-aware knowledge.

The initial bursting rate of the unicast burst to RR is determined by parameters directly obtained from RR (a) or by pre-configured settings (f). If such information is not available, RS may choose an appropriate initial bursting rate, and could increase or decrease the rate based on the feedback information (b, c, d or e). However, this may not be an easy task as by the time packet loss is reported back to RS triggering a rate reduction, packet loss may have occurred.

A specific situation occurs near the end of the unicast burst, when RS has almost no more additional data to sustain the relatively higher bursting rate, thus, the upper-bound bursting rate automatically gets limited by the nominal rate of the primary multicast stream. During this time frame, RR will join the primary multicast session because it was instructed to do so via an RAMS-I message or based on some heuristics. This means that both the burst packets and the primary multicast stream packets will be simultaneously received by RR for a period of time.

In this case, when the unicast burst is close to catch up with the primary multicast stream, RS may, for example, keep on sending burst packets but should reduce the rate accordingly by taking the nominal

rate of the primary multicast stream into account. Alternatively, RS may immediately cease transmitting the burst packets, when being close to catch-up. Any gap resulting from an imperfect switch by RR in receiving first the burst packets and then only primary multicast stream packets, can be later repaired by requesting retransmissions of the missing packets from RS. The retransmissions may also be shaped by RS to make sure that they do not cause collateral loss in the primary multicast and retransmission sessions.

6.4. Failure Cases

All RAMS messages MAY be sent several times against the possibility of message loss as long as RS/RR follows the RTCP timer rules defined in [RFC4585]. In the following, we examine the implications of losing the RAMS-R, RAMS-I or RAMS-T messages.

When RR sends an RAMS-R message to initiate a rapid acquisition but the message gets lost and RS does not receive it, RR will not get an RAMS-I message, nor a unicast burst. In this case, RR MAY resend the request when it is eligible to do so. Or, after a reasonable amount of time, RR MAY time out (based on the previous observed response times) and immediately join the primary multicast session. In this case, RR MUST still send an RAMS-T message.

In the case RR starts receiving a unicast burst but it does not receive a corresponding RAMS-I message within a reasonable amount of time, RR MAY either discard the burst data and stop the burst by sending an RAMS-T message to RS, or decide not to interrupt the unicast burst and be prepared to join the primary multicast session at an appropriate time it determines or indicated in a subsequent RAMS-I message (if available). In either case, RR SHALL send an RAMS-T message to RS at an appropriate time.

In the case the RAMS-T message sent by RR does not reach its destination, RS may continue sending the unicast burst even though RR no longer needs it. In some cases, RS has not pre-computed the burst duration and does not know when to stop the burst. To cover that case, RR MAY repeat the RAMS-T message multiple times as long as it follows the RTCP timer rules defined in [RFC4585]. RS MUST be provisioned to deterministically terminate the burst at some point, even if it never receives an RAMS-T message for an ongoing burst.

If RR becomes no longer interested in receiving the primary multicast stream and the associated unicast burst, RR SHALL issue an RTCP BYE message to RS to terminate the RTP retransmission session. Only after that, RR MAY send a new rapid acquisition request for another primary multicast session.

7. Encoding of the Signaling Protocol in RTCP

This section defines the formats of the RTCP transport-layer feedback messages that are exchanged between the Retransmission Server (RS) and RTP Receiver (RR) during rapid acquisition. These messages are referred to as the RAMS Messages. They are payload-independent and MUST be used by all RTP-based multicast applications that support rapid acquisition regardless of the payload they carry.

Payload-specific feedback messages are not defined in this document, but an extension mechanism is provided where further optional payload-independent and payload-specific information can be included in the exchange.

The common packet format for the RTCP feedback messages is defined in Section 6.1 of [RFC4585]. Each feedback message has a fixed-length field for version, padding, feedback message type (FMT), payload type (PT), length, SSRC of packet sender, SSRC of media source as well as a variable-length field for feedback control information (FCI).

In the RAMS messages, the PT field is set to RTPFB (205) and the FMT field is set to RAMS (6). Individual RAMS messages are identified by a sub-field called Sub Feedback Message Type (SFMT).

Depending on the specific scenario and timeliness/importance of a RAMS message, it may be desirable to send it in a reduced-size RTCP packet [RFC5506]. However, unless support for [RFC5506] has been signaled, compound RTCP packets MUST be used by following [RFC3550] rules.

7.1. Extensions

To improve the functionality of the RAMS method in certain applications, it may be desirable to define new fields in the RAMS Request, Information and Termination messages. Such fields MUST be encoded as TLV elements as described below and sketched in Figure 4:

- o Type: A single-octet identifier that defines the type of the parameter represented in this TLV element.

- o Length: A two-octet field that indicates the length of the TLV element excluding the Type and Length fields in octets. Note that this length does not include any padding that is required for alignment.

- o Value: Variable-size set of octets that contains the specific value for the parameter.

If a TLV element does not fall on a 32-bit boundary, the last word must be padded to the boundary using further bits set to 0.

In an RAMS message any vendor-neutral or private extension MUST be placed after the mandatory fields (if any). The support for extensions is OPTIONAL.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |             Length            |     Value     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Value contd.                         /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Structure of a TLV element

7.1.1. Vendor-Neutral Extensions

If the goal in defining new TLV elements is to extend the functionality in a vendor-neutral manner, they MUST be registered with IANA through the guidelines provided in Section 13.5.

The current document defines several vendor-neutral extensions in the following sections.

7.1.2. Private Extensions

It is desirable to allow vendors to use private extensions in TLV format. For interoperability, such extensions MUST NOT collide with each other.

A certain range of TLV Types is reserved for private extensions (Refer to Section 13.5). IANA management for these extensions is unnecessary and they are the responsibility of individual vendors.

The structure that MUST be used for the private extensions is depicted in Figure 5. Here, the enterprise numbers are used from http://www.iana.org/assignments/enterprise-numbers. This will ensure the uniqueness of the private extensions and avoid any collision.

```
      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |      Type     |             Length            |  Ent. Number  |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |               Ent. Number contd.              |     Value     |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                          Value contd.                         /
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Structure of a private extension

7.2. RAMS Request

The RAMS Request message is identified by SFMT=1.

The FCI field MUST contain only one RAMS Request.

The RAMS Request is used by RR to request rapid acquisition for a new multicast RTP session.

The FCI field has the structure depicted in Figure 6.

Editor's note: We have not finalized whether RAMS-R messages need a sequence number or not.

```
      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |    SFMT=1     |          Optional TLV-encoded Fields          |
     +-+-+-+-+-+-+-+-+                                               |
     :      Optional TLV-encoded Fields (and Padding, if needed)     :
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: FCI field syntax for the RAMS Request message

- o Min RAMS Buffer Fill Requirement (32 bits): Optional TLV element that denotes the minimum milliseconds of data that RR desires to have in its buffer before allowing the data to be consumed by the application.

  RR may have knowledge of its buffering requirements. These requirements may be application and/or device specific. For instance, RR may need to have a certain amount of data in its application buffer to handle transmission jitter and/or to be able to support error-control methods. If RS is told the minimum buffering requirement of the receiver, it may tailor the burst

more precisely, e.g., by choosing an appropriate starting point. The methods used by RR to determine this value are application specific, and thus, out of the scope of this document.

If specified, the amount of backfill that will be provided by the unicast burst SHOULD NOT be smaller than the specified value since it will not be able to build up the desired level of buffer at RR and may cause buffer underruns.

Type: TBD

Length: TBD

- Max RAMS Buffer Fill Requirement (32 bits): Optional TLV element that denotes the maximum milliseconds of data that RR can buffer without losing the burst data due to buffer overflow.

  RR may have knowledge of its buffering requirements. These requirements may be application or device specific. For instance, one particular RR may have more physical memory than another RR, and thus, can buffer more data. If RS knows the buffering ability of the receiver, it may tailor the burst more precisely. The methods used by the receiver to determine this value are application specific, and thus, out of scope.

  If specified, the amount of backfill that will be provided by the unicast burst SHOULD NOT be larger than this value since it may cause buffer overflows at RR.

  Type: TBD

  Length: TBD

- Max Receive Bitrate (32 bits): Optional TLV element that denotes the maximum bitrate (in bits per second) that the RTP receiver can process the unicast burst. This rate should include whatever knowledge the receiver has that would provide an upper bound on the unicast burst bitrate. The limits may include local receiver limits as well as network limits that are known to the receiver.

  If specified, the unicast burst bitrate SHOULD NOT be larger than this value since it may cause congestion and packet loss.

  Type: TBD

  Length: TBD

The semantics of the RAMS-R feedback message is independent of the

payload type.

7.3. RAMS Information

The RAMS Information message is identified by SFMT=2.

The FCI field MUST contain only one RAMS Information.

The RAMS Information is used to describe the unicast burst that will be sent for rapid acquisition. It also includes other useful information for RR as described below. Optional payload-specific information MAY follow RAMS Information.

The FCI field has the structure depicted in Figure 7.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    SFMT=2     |      MSN      |          Response             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|RTP SN of the First Burst Pkt. |  Optional TLV-encoded Fields  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:      Optional TLV-encoded Fields (and Padding, if needed)     :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: FCI field syntax for the RAMS Information message

- o Message Sequence Number (8 bits) : Mandatory field that denotes the sequence number of this RAMS-I message. During rapid acquisition, multiple RAMS-I messages MAY be sent and/or the same RAMS-I message MAY be repeated. The first RAMS-I message SHALL have an MSN value of 0. This value SHALL NOT be changed if the same RAMS-I message is sent to the same RR multiple times for redundancy purposes. If a new information is conveyed in a new RAMS-I message, the MSN value SHALL be incremented by one.

- o Response (16 bits): Mandatory field that denotes the RS response code for this RAMS-I message.

  Editor's note: HTTP/SIP-like response codes will be defined and registered with IANA in a later version.

- o RTP SN of the First Burst Pkt. (16 bits): Mandatory field that specifies the RTP sequence number of the first packet that will be sent as part of the burst. This allows RR to know if one or more packets have been dropped at the beginning of the burst.

- o Earliest Multicast Join Time (32 bits): Optional TLV element that specifies the time difference (i.e., delta time) between the arrival of this RAMS-I message and the earliest time instant when RR could join the primary multicast session in RTP ticks. A zero value in this field means that RR can join the primary multicast session right away.

  Note that if the RAMS request has been accepted, RS SHOULD send this field at least once, so that RR knows when to join the primary multicast session. If the burst request has been rejected as indicated in the Response field, this field MAY be omitted or set to 0. In that case, it is up to RR when or whether to join the primary multicast session.

  Type: TBD

  Length: TBD

- o Burst Duration (32 bits): Optional TLV element that denotes the duration of the burst that RS is planning to send (in RTP ticks). In the absence of additional stimulus, RS will send a burst of this duration. However, the burst duration may be modified by subsequent events, including changes in the primary multicast stream and reception of RAMS-T messages.

  Note that RS MUST terminate the flow in a deterministic timeframe, even if it does not get an RAMS-T or a BYE from RR. It is optional to send this field in an RAMS-I message when the burst request is accepted. If the burst request has been rejected as indicated in the Response field, this field MAY be omitted or set to 0.

  Type: TBD

  Length: TBD

- o Max Burst Bitrate (32 bits): Optional TLV element that denotes the maximum bitrate (in bits per second) that will be used by RS for the unicast burst.

  Type: TBD

  Length: TBD

The semantics of the RAMS-I feedback message is independent of the payload type.

The RAMS-I message MAY be sent multiple times at the start of, prior

to, or during the unicast burst. The subsequent RAMS-I messages MAY signal changes in any of the fields.

7.4. RAMS Termination

The RAMS Termination message is identified by SFMT=3.

The FCI field MUST contain only one RAMS Termination.

The RAMS Termination may be used to assist RS in determining when to stop the burst.

If prior to sending the RAMS-T message RR has already joined the primary multicast session and received at least one RTP packet from the multicast session, RR includes the sequence number of the first RTP packet in the RAMS-T message. With this information, RS can decide when to terminate the unicast burst.

If RR issues the RAMS-T message before it has joined and/or begun receiving RTP packets from the primary multicast session, RR does not specify any sequence number in the RAMS-T message, which indicates RS to stop the burst immediately. However, the RAMS-T message may get lost and RS may not receive this message.

The FCI field has the structure depicted in Figure 8.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    SFMT=3     |          Optional TLV-encoded Fields          |
+-+-+-+-+-+-+-+-+                                               |
:      Optional TLV-encoded Fields (and Padding, if needed)     :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 8: FCI field syntax for the RAMS Termination message

o Extended RTP Seqnum of First Multicast Packet (32 bits): Optional TLV element that specifies the extended RTP sequence number of the of the first multicast packet received by RR. If no RTP packet has been received from the primary multicast session, this field does not exist and tells RS to stop the burst immediately.

   Type: TBD

   Length: TBD

The semantics of the RAMS-T feedback message is independent of the

payload type.

8. SDP Definitions and Examples

8.1. Definitions

The syntax of the 'rtcp-fb' attribute has been defined in [RFC4585]. Here we add the following syntax to the 'rtcp-fb' attribute (the feedback type and optional parameters are all case sensitive):

(In the following ABNF [RFC5234], fmt, SP and CRLF are used as defined in [RFC4566].)

```
rtcp-fb-syntax = "a=rtcp-fb:" rtcp-fb-pt SP rtcp-fb-val CRLF

rtcp-fb-pt         = "*"   ; wildcard: applies to all formats
                   / fmt   ; as defined in SDP spec

rtcp-fb-val        = "nack" SP "ssli"
```

The following parameter is defined in this document for use with 'nack':

o 'ssli' stands for Stream Synchronization Loss Indication and indicates the use of RAMS messages as defined in Section 7.

This document also defines a new SDP attribute ('rams-updates') that indicates whether RS supports updated request messages or not. This attribute is used in a declarative manner. If RS supports updated request messages and this attribute is included in the SDP description, RR MAY send updated requests. RS may or may not be able to accept value changes in every field in an RAMS-R message. However, if the 'rams-updates' attribute is not included in the SDP description, RR SHALL NOT send updated requests (RR MAY repeat its initial request without changes, though).

8.2. Examples

This section provides a declarative SDP [RFC4566] example for enabling rapid acquisition of multicast RTP sessions. The following example uses the SDP grouping semantics [RFC3388], the RTP/AVPF profile [RFC4585], the RTP retransmissions [RFC4588], the RTCP extensions for SSM sessions with unicast feedback [I-D.ietf-avt-rtcpssm] and the source-specific media attributes [I-D.ietf-mmusic-sdp-source-attributes].

In the example shown Figure 9, we have a primary multicast stream and

a unicast retransmission stream. The source stream is multicast from a distribution source (with a source IP address of 192.0.2.2) to the multicast destination address of 233.252.0.2 and port 41000. A Retransmission Server including feedback target functionality (with an address of 192.0.2.1 and port of 41001) is specified with the 'rtcp' attribute. The RTP receiver(s) can report missing packets on the source stream to the feedback target and request retransmissions. In the RAMS context, the parameter 'rtx-time' specifies the time in milliseconds that the Retransmission Server keeps an RTP packet in its buffer available for retransmission (measured from the time the packet was received by the Retransmission Server).

The RTP retransmissions are sent on a unicast session with a destination port of 41002.

Editor's note: This text will be updated in a later version to reflect the capability for RRs to use their desired ports to receive the burst and retransmission packets.

The RTCP port for the unicast session (41003) is specified with the 'rtcp' attribute. In this example, both the conventional retransmission and rapid acquisition support are enabled. This is achieved by the additional "a=rtcp-fb:98 nack ssli" line. Note that this SDP includes the "a=sendonly" line for the media description of the retransmission stream and is for the Retransmission Server (RS). Its counterpart for the RTP Receiver (RR) includes the "a=recvonly" line as shown in Figure 10.

When an RTP receiver requires rapid acquisition for a new multicast session it wants to join, it sends an RAMS-R message to the feedback target of that primary multicast session. This feedback message has to have the SSRC of the primary multicast stream for which rapid acquisition is requested for. However, since this RTP receiver has not received any RTP packets from the primary multicast session yet, the RTP receiver MUST learn the SSRC value from the 'ssrc' attribute of the media description [I-D.ietf-avt-rtcpssm]. In addition to the SSRC value, the 'cname' source attribute MUST also be present in the SDP description [I-D.ietf-mmusic-sdp-source-attributes].

Note that listing the SSRC values for the primary multicast sessions in the SDP file does not create a problem in SSM sessions when an SSRC collision occurs. This is because in SSM sessions, an RTP receiver that observed an SSRC collision with a media source MUST change its own SSRC [I-D.ietf-avt-rtcpssm] by following the rules defined in [RFC3550].

A feedback target that receives an RAMS-R feedback message becomes aware that the prediction chain at the RTP receiver side has been

broken or does not exist any more. If the necessary conditions are satisfied (as outlined in Section 7 of [RFC4585]) and available resources exist, RS MAY react to the RAMS-R message by sending any transport-layer and payload-specific feedback message(s) and starting the unicast burst.

```
v=0
o=ali 1122334455 1122334466 IN IP4 rams.example.com
s=Rapid Acquisition Example
t=0 0
a=group:FID 3 4
a=rtcp-unicast:rsi
m=video 41000 RTP/AVPF 98
i=Primary Multicast Stream #2
c=IN IP4 233.252.0.2/255
a=source-filter: incl IN IP4 233.252.0.2 192.0.2.2
a=recvonly
a=rtpmap:98 MP2T/90000
a=rtcp:41001 IN IP4 192.0.2.1
a=rtcp-fb:98 nack
a=rtcp-fb:98 nack ssli
a=ssrc:123321 cname:iptv-ch32@rams.example.com
a=rams-updates
a=mid:3
m=video 41002 RTP/AVPF 99
i=Unicast Retransmission Stream #2 (Ret. and Rapid Acq. Support)
c=IN IP4 192.0.2.1
a=sendonly
a=rtpmap:99 rtx/90000
a=rtcp:41003
a=fmtp:99 apt=98; rtx-time=5000
a=mid:4
```

Figure 9: Example SDP for RS when RAMS support is enabled

```
v=0
o=ali 1122334455 1122334466 IN IP4 rams.example.com
s=Rapid Acquisition Example
t=0 0
a=group:FID 3 4
a=rtcp-unicast:rsi
m=video 41000 RTP/AVPF 98
i=Primary Multicast Stream #2
c=IN IP4 233.252.0.2/255
a=source-filter: incl IN IP4 233.252.0.2 192.0.2.2
a=recvonly
a=rtpmap:98 MP2T/90000
a=rtcp:41001 IN IP4 192.0.2.1
a=rtcp-fb:98 nack
a=rtcp-fb:98 nack ssli
a=ssrc:123321 cname:iptv-ch32@rams.example.com
a=rams-updates
a=mid:3
m=video 41002 RTP/AVPF 99
i=Unicast Retransmission Stream #2 (Ret. and Rapid Acq. Support)
c=IN IP4 192.0.2.1
a=recvonly
a=rtpmap:99 rtx/90000
a=rtcp:41003
a=fmtp:99 apt=98; rtx-time=5000
a=mid:4
```

Figure 10: Example SDP for RR when RAMS support is enabled

The offer/answer model considerations [RFC3264] for the 'rtcp-fb' attribute are provided in Section 4.2 of [RFC4585].

9. NAT Considerations

For a variety of reasons, one or more NAPT devices (hereafter simply called NAT) are expected to exist between RR and RS. NATs have a variety of operating characteristics for UDP traffic [RFC4787]. For a NAT to permit traffic from RS to arrive at RR, the NAT(s) must first either:

a. See UDP traffic sent from RR (which is on the 'inside' of the NAT) to RS (which is on the 'outside' of the NAT). This traffic is sent to the same transport address as the subsequent response traffic, OR;

b. Be configured to forward certain ports (e.g., using HTML configuration, UPnP IGD [UPnP-IGD], DLNA [DLNA]). Details of

this are out of scope of this document.

For both (a) and (b), RR is responsible for maintaining the NAT's state if it wants to receive traffic from the RS on that port. For (a), RR MUST send UDP traffic to keep the NAT binding alive, at least every 30 seconds [RFC4787]. Note that while (a) is more like an automatic/dynamic configuration, (b) is more like a manual/static configuration.

When using (a), RR will need to first learn the UDP port(s) of the NAT's binding(s) from the perspective of RS. This is done by sending a STUN [RFC5389] message from RR to the RTP port of RS which will be used for incoming RTP traffic. If RTP/RTCP multiplexing on a single port [I-D.ietf-avt-rtp-and-rtcp-mux] is not supported by RR, it will need to send a second STUN message to the RTCP port of RS which will be used for incoming RTCP traffic. If RTP/RTCP multiplexing is supported by RR, it only needs to learn one port. RS receives the STUN message(s) and responds to them. RR now knows the UDP ports from the perspective of RS.

Editor's note: The issues related to using ports across multicast and unicast RTP sessions will be discussed in a separate draft and the current document will normatively reference that document. The updated text for this section will be provided in a later version.

10. Known Implementations

10.1. Open Source RTP Receiver Implementation by Cisco

An open source RTP Receiver code that implements the functionalities introduced in this document is available. For documentation, visit the following URL:

http://www.cisco.com/en/US/docs/video/cds/cda/vqe/3_3/user/guide/ch1_over.html

The code is also available at:

ftp://ftpeng.cisco.com/ftp/vqec/

Note that this code is under development and may be based on an earlier version of this document. As we make progress in the draft, the source code will also be updated to reflect the changes.

Some preliminary results based on this code are available in [CCNC09] and [IC2009].

10.2. IPTV Commercial Implementation by Microsoft

Rapid Acquisition of Multicast RTP Sessions is supported as part of the Microsoft Mediaroom Internet Protocol Television (IPTV) and multimedia software platform. This system is in wide commercial deployment. More information can be found at:

http://www.microsoft.com/mediaroom

http://informitv.com/articles/2008/10/13/channelchangetimes/

11. Open Issues

- o Discussion of acquisition for the individual RTP streams vs. the whole RTP session.
- o Updating the NAT section.
- o Completing the TLV types, lengths, etc.
- o Response/status codes for RAMS.

12. Security Considerations

Applications that are using RAMS make heavy use of the unicast feedback mechanism described in [I-D.ietf-avt-rtcpssm] and the payload format defined in [RFC4588]. Thus, these applications are subject to the general security considerations discussed in [I-D.ietf-avt-rtcpssm] and [RFC4588]. In this section, we give an overview of the guidelines and suggestions described in these specifications from a RAMS perspective. We also discuss the security considerations that explicitly apply to RAMS applications.

First of all, much of the session description information is available in the SDP descriptions that are distributed to the media sources, Retransmission Servers and RTP Receivers. Adequate security measures are RECOMMENDED to ensure the integrity and authenticity of the SDP descriptions so that transport addresses of the media sources, Feedback Targets as well as other session-specific information can be authenticated.

Compared to an RTCP NACK message that triggers one or more retransmissions, an RAMS Request (RAMS-R) message may trigger a new burst stream to be sent by the Retransmission Server. Depending on the application-specific requirements and conditions existing at the time of the RAMS-R reception by the Retransmission Server, the

resulting burst stream may contain potentially a large number of retransmission packets. Since these packets are sent at a faster than the nominal rate of the multicast session, RAMS consumes more resources on the Retransmission Server, the RTP Receiver and the network. This particularly makes denial-of-service attacks more intense, and hence, more harmful than attacks that target ordinary retransmission sessions.

Following the suggestions given in [RFC4588], counter-measures SHOULD be taken to prevent tampered or spoofed RTCP packets. Tampered RAMS-R messages may trigger inappropriate burst streams or alter the existing burst streams in an inappropriate way. For example, if the Max Receive Bitrate field is altered by a tampered RAMS-R message, the updated burst may overflow the buffer on the receiver side, or oppositely, may slow down the burst to the point that it is useless. Tampered RAMS Termination (RAMS-T) messages may terminate valid burst streams pre-maturely resulting in gaps in the received RTP packets. RAMS Information (RAMS-I) messages contain fields that are critical for the success of the RAMS operation. Any tampered information in the RAMS-I message may easily cause the RTP Receiver to make wrong decisions. Consequently, the RAMS operation may fail.

While most of the denial-of-service attacks can be prevented by the integrity and authenticity checks enabled by SRTP, an attack can still be started by legitimate endpoints that send several valid RAMS-R messages to a particular Feedback Target in a synchronized fashion and very short amount of time. Since a RAMS operation may temporarily consume a large amount of resources, a series of the RAMS-R messages may temporarily overload the Retransmission Server. In these circumstances, the Retransmission Server may, for example, reject incoming RAMS requests until its resources become available again. One means to ameliorate this threat is to apply a per-endpoint policing mechanism on the incoming RAMS requests. A reasonable policing mechanism should consider application-specific requirements and minimize false negatives.

In addition to the denial-of-service attacks, man-in-the-middle and replay attacks can also be harmful. However, RAMS itself does not bring any new risks or threats other than the ones discussed in [I-D.ietf-avt-rtcpssm].

[RFC4588] RECOMMENDS that the cryptography mechanisms are used for the retransmission payload format to provide protection against known plaintext attacks. As discussed in [RFC4588], the retransmission payload format sets the timestamp field in the RTP header to the media timestamp of the original packet and this does not compromise the confidentiality. Furthermore, if cryptography is used to provide security services on the original stream, then the same services,

with equivalent cryptographic strength, MUST be provided on the retransmission stream per [RFC4588].

13. IANA Considerations

The following contact information shall be used for all registrations in this document:

Ali Begen
abegen@cisco.com

170 West Tasman Drive
San Jose, CA 95134 USA

13.1. Registration of SDP Attributes

This document registers a new attribute name in SDP.

```
SDP Attribute ("att-field"):
Attribute name:     rams-updates
Long form:          Support for Updated RAMS Request Messages
Type of name:       att-field
Type of attribute:  Media level
Subject to charset: No
Purpose:            See this document
Reference:          This document
Values:             None
```

13.2. Registration of SDP Attribute Values

This document registers a new value for the 'nack' attribute to be used with the 'rtcp-fb' attribute in SDP. For more information about 'rtcp-fb', refer to [RFC4585].

```
Value name:     ssli
Long name:      Stream Synchronization Loss Indication
Usable with:    nack
Reference:      This document
```

13.3. Registration of FMT Values

Within the RTPFB range, the following format (FMT) value is registered:

Name:	RAMS
Long name:	Rapid Acquisition of Multicast Sessions
Value:	6
Reference:	This document

13.4. SFMT Values for RAMS Messages Registry

This document creates a new sub-registry for the sub-feedback message type (SFMT) values to be used with the FMT value registered for RAMS messages. The registry is called the SFMT Values for RAMS Messages Registry. This registry is to be managed by the IANA according to the Specification Required policy of [RFC5226].

The length of the SFMT field in the RAMS messages is a single octet, allowing 256 values. The registry is initialized with the following entries:

Value	Name	Reference
1	RAMS Request	This document
2	RAMS Information	This document
3	RAMS Termination	This document

The SFMT values 0 and 255 are reserved for future use.

Any registration for an unassigned SFMT value MUST contain the following information:

- o Contact information of the one doing the registration, including at least name, address, and email.
- o A detailed description of what the new SFMT represents and how it shall be interpreted.

Note that new RAMS functionality should be introduced by using the extension mechanism within the existing RAMS message types not by introducing new message types unless it is absolutely necessary.

13.5. RAMS TLV Space Registry

This document creates a new IANA TLV space registry for the RAMS extensions. The registry is called the RAMS TLV Space Registry. This registry is to be managed by the IANA according to the Specification Required policy of [RFC5226].

The length of the Type field in the TLV elements is a single octet,

allowing 256 values. The registry is initialized with the following entries:

Type	Description	Reference
TBD	TBD	This document
...	...	

The registry entries are TBC.

The TYPE values 0 and 255 are reserved for future use. The TYPE values between (and including) 128 and 254 are reserved for private extensions.

Any registration for an unassigned TYPE value MUST contain the following information:

- o Contact information of the one doing the registration, including at least name, address, and email.
- o A detailed description of what the new TLV element represents and how it shall be interpreted.

14. Acknowledgments

The authors would like to specially thank Peilin Yang for his contributions to this document and detailed reviews.

The authors also thank the following individuals for their contributions, comments and suggestions to this document: Dave Oran, Dan Wing, Tony Faustini, Jeff Goldberg, Muriel Deschanel, Orit Levin, Roni Even, Guy Hirson, Tom Taylor, Xavier Marjou, Ye-Kui Wang, Zixuan Zou, Ingemar Johansson, Haibin Song, Ning Zong, Jonathan Lennox and Sean Sheedy.

15. Change Log

15.1. draft-ietf-avt-rapid-acquisition-for-rtp-02

The following are the major changes compared to version 01:

- o Port mapping discussion has been removed since it will be discussed in a separate draft.

- o Security considerations section has been added.

- o Burst shaping section has been completed.

- o Most of the outstanding open issues have been addressed.

15.2. draft-ietf-avt-rapid-acquisition-for-rtp-01

The following are the major changes compared to version 00:

- o Formal definitions of vendor-neutral and private extensions and their IANA registries have been added.

- o SDP examples were explained in more detail.

- o The sub-FMT field has been introduced in the RAMS messages for message type identification.

- o Some terminology has been fixed.

- o NAT considerations section has been added.

15.3. draft-ietf-avt-rapid-acquisition-for-rtp-00

This is a resubmission of version 03 as a WG item.

15.4. draft-versteeg-avt-rapid-synchronization-for-rtp-03

The following are the major changes compared to version 02:

- o The title and message names have been changed.

- o RTCP message semantics have been added. RAMS protocol has been revised to handle updated requests and responses.

- o Definitions have been revised.

- o RTP/RTCP muxing reference has been added.

15.5. draft-versteeg-avt-rapid-synchronization-for-rtp-02

The following are the major changes compared to version 01:

- o The discussion around MPEG2-TS has been moved to another document.

- o The RAMS-R, RAMS-I and RAMS-T messages have been extensively modified and they have been made mandatory.

o IANA Considerations section has been updated.

o The discussion of RTCP XR report has been moved to another document.

o A new section on protocol design considerations has been added.

15.6. draft-versteeg-avt-rapid-synchronization-for-rtp-01

The following are the major changes compared to version 00:

o The core of the rapid synchronization method is now payload-independent. But, the draft still defines payload-specific messages that are required for enabling rapid synch for the RTP flows carrying MPEG2-TS.

o RTCP APP packets have been removed, new RTCP transport-layer and payload-specific feedback messages have been defined.

o The step for leaving the current multicast session has been removed from Section 6.2.

o A new RTCP XR (Multicast Join) report has been defined.

o IANA Considerations section have been updated.

o Editorial changes to clarify several points.

16. References

16.1. 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.

[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.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

[RFC3388] Camarillo, G., Eriksson, G., Holler, J., and H. Schulzrinne, "Grouping of Media Lines in the Session Description Protocol (SDP)", RFC 3388, December 2002.

[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.

[I-D.ietf-avt-rtcpssm] Schooler, E., Ott, J., and J. Chesterfield, "RTCP Extensions for Single-Source Multicast Sessions with Unicast Feedback", draft-ietf-avt-rtcpssm-18 (work in progress), March 2009.

[I-D.ietf-mmusic-sdp-source-attributes] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", draft-ietf-mmusic-sdp-source-attributes-02 (work in progress), October 2008.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

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

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

16.2. Informative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[I-D.ietf-rmt-pi-norm-revised]
Adamson, B., Bormann, C., London, U., and J. Macker, "NACK-Oriented Reliable Multicast Transport Protocol", draft-ietf-rmt-pi-norm-revised-13 (work in progress), June 2009.

[I-D.begen-avt-rtp-mpeg2ts-preamble]
Begen, A. and E. Friedrich, "RTP Payload Format for MPEG2-TS Preamble", draft-begen-avt-rtp-mpeg2ts-preamble-02 (work in progress), August 2009.

[I-D.begen-avt-rapid-sync-rtcp-xr]
Begen, A. and E. Friedrich, "Multicast Acquisition Report Block Type for RTCP XR", draft-begen-avt-rapid-sync-rtcp-xr-02 (work in progress), August 2009.

[I-D.ietf-avt-rapid-rtp-sync]
Perkins, C. and T. Schierl, "Rapid Synchronisation of RTP Flows", draft-ietf-avt-rapid-rtp-sync-05 (work in progress), July 2009.

[I-D.westerlund-avt-ecn-for-rtp]
Westerlund, M., Johansson, I., and C. Perkins, "Explicit Congestion Notification (ECN) for RTP over UDP", draft-westerlund-avt-ecn-for-rtp-00 (work in progress), July 2009.

[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.

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.

[UPnP-IGD]
Forum, UPnP., "Universal Plug and Play (UPnP) Internet Gateway Device (IGD)", November 2001.

[DLNA] , DLNA., "http://www.dlna.org/home".

[CCNC09] Begen, A., Glazebrook, N., and W. VerSteeg, "A Unified Approach for Repairing Packet Loss and Accelerating Channel Changes in Multicast IPTV (IEEE CCNC)", January 2009.

[IC2009] Begen, A., Glazebrook, N., and W. VerSteeg, "Reducing Channel Change Times in IPTV with Real-Time Transport Protocol (IEEE Internet Computing)", May 2009.

Authors' Addresses

Bill VerSteeg
Cisco Systems
5030 Sugarloaf Parkway
Lawrenceville, GA 30044
USA

Email: billvs@cisco.com

Ali Begen
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Email: abegen@cisco.com

Tom VanCaenegem
Alcatel-Lucent
Copernicuslaan 50
Antwerpen,   2018
Belgium

Email: Tom.Van_Caenegem@alcatel-lucent.be

Zeev Vax
Microsoft Corporation
1065 La Avenida
Mountain View, CA 94043
USA

Email: zeevvax@microsoft.com