

Internet Engineering Task Force (IETF)
Request for Comments: 5760
Category: Standards Track
ISSN: 2070-1721

J. Ott
Aalto University
J. Chesterfield
University of Cambridge
E. Schooler
Intel
February 2010

RTP Control Protocol (RTCP) Extensions for Single-Source Multicast Sessions with Unicast Feedback

Abstract

This document specifies an extension to the Real-time Transport Control Protocol (RTCP) to use unicast feedback to a multicast sender. The proposed extension is useful for single-source multicast sessions such as Source-Specific Multicast (SSM) communication where the traditional model of many-to-many group communication is either not available or not desired. In addition, it can be applied to any group that might benefit from a sender-controlled summarized reporting mechanism.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5760>.

Copyright Notice

Copyright (c) 2010 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
2. Conventions and Acronyms	4
3. Definitions	5
4. Basic Operation	6
5. Packet Types	10
6. Simple Feedback Model	11
7. Distribution Source Feedback Summary Model	13
8. Mixer/Translator Issues	36
9. Transmission Interval Calculation	37
10. SDP Extensions	39
11. Security Considerations	41
12. Backwards Compatibility	51
13. IANA Considerations	51
14. References	53
Appendix A. Examples for Sender-Side Configurations	57
Appendix B. Distribution Report Processing at the Receiver	60

1. Introduction

The Real-time Transport Protocol (RTP) [1] provides a real-time transport mechanism suitable for unicast or multicast communication between multimedia applications. Typical uses of RTP are for real-time or near real-time group communication of audio and video data streams. An important component of the RTP protocol is the control channel, defined as the RTP Control Protocol (RTCP). RTCP involves the periodic transmission of control packets between group members, enabling group size estimation and the distribution and calculation of session-specific information such as packet loss and round-trip time to other hosts. An additional advantage of providing a control channel for a session is that a third-party session monitor can listen to the traffic to establish network conditions and to diagnose faults based on receiver locations.

RTP was designed to operate in either a unicast or multicast mode. In multicast mode, it assumes an Any Source Multicast (ASM) group model, where both one-to-many and many-to-many communication are supported via a common group address in the range 224.0.0.0 through 239.255.255.255. To enable Internet-wide multicast communication, intra-domain routing protocols (those that operate only within a single administrative domain, e.g., the Distance Vector Multicast Routing Protocol (DVMRP) [16] and Protocol Independent Multicast (PIM) [17][18]) are used in combination with inter-domain routing protocols (those that operate across administrative domain borders, e.g., Multicast BGP (MBGP) [19] and the Multicast Source Discovery Protocol (MSDP) [20]). Such routing protocols enable a host to join a single multicast group address and send data to or receive data from all members in the group with no prior knowledge of the membership. However, there is a great deal of complexity involved at the routing level to support such a multicast service in the network.

Many-to-many communication is not always available or desired by all group applications. For example, with Source-Specific Multicast (SSM) [8][9] and satellite communication, the multicast distribution channel only supports source-to-receiver traffic. In other cases, such as large ASM groups with a single active data source and many passive receivers, it is sub-optimal to create a full routing-level mesh of multicast sources just for the distribution of RTCP control packets. Thus, an alternative solution is preferable.

Although a one-to-many multicast topology may simplify routing and may be a closer approximation to the requirements of certain RTP applications, unidirectional communication makes it impossible for receivers in the group to share RTCP feedback information with other group members. In this document, we specify a solution to that problem. We introduce unicast feedback as a new method to distribute

RTCP control information amongst all session members. This method is designed to operate under any group communication model, ASM or SSM. The RTP data stream protocol itself is unaltered.

Scenarios under which the unicast feedback method can provide benefit include but are not limited to:

a) SSM groups with a single sender.

The proposed extensions allow SSM groups that do not have many-to-many communication capability to receive RTP data streams and to continue to participate in the RTP control protocol (RTCP) by using multicast in the source-to-receiver direction and unicast to send receiver feedback to the source on the standard RTCP port.

b) One-to-many broadcast networks.

Unicast feedback may also be beneficial to one-to-many broadcast networks, such as a satellite network with a terrestrial low-bandwidth return channel or a broadband cable link. Unlike the SSM network, these networks may have the ability for a receiver to multicast return data to the group. However, a unicast feedback mechanism may be preferable for routing simplicity.

c) ASM with a single sender.

A unicast feedback approach can be used by an ASM application with a single sender to reduce the load on the multicast routing infrastructure that does not scale as efficiently as unicast routing does. Because this is no more efficient than a standard multicast group RTP communication scenario, it is not expected to replace the traditional mechanism.

The modifications proposed in this document are intended to supplement the existing RTCP feedback mechanisms described in Section 6 of [1].

2. Conventions and Acronyms

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 RFC 2119 [13].

The following acronyms are used throughout this document:

ASM Any Source Multicast
SSM Source-Specific Multicast

3. Definitions

Distribution Source:

In an SSM context, only one entity distributes RTP data and redistributes RTCP information to all receivers. This entity is called the Distribution Source. It is also responsible for forwarding RTCP feedback to the Media Senders and thus creates a virtual multicast environment in which RTP and RTCP can be applied.

Note that heterogeneous networks consisting of ASM multiple-sender groups, unicast-only clients, and/or SSM single-sender receiver groups MAY be connected via translators or mixers to create a single-source group (see Section 8 for details).

Media Sender:

A Media Sender is an entity that originates RTP packets for a particular media session. In RFC 3550, a Media Sender is simply called a source. However, as the RTCP SSM system architecture includes a Distribution Source, to avoid confusion, in this document a media source is commonly referred to as a Media Sender. There may often be a single Media Sender that is co-located with the Distribution Source. But although there MUST be only one Distribution Source, there MAY be multiple Media Senders on whose behalf the Distribution Source forwards RTP and RTCP packets.

RTP and RTCP Channels:

The data distributed from the source to the receivers is referred to as the RTP channel and the control information as the RTCP channel. With standard RTP/RTCP, these channels typically share the same multicast address but are differentiated via port numbers as specified in [1]. In an SSM context, the RTP channel is multicast from the Distribution Source to the receivers. In contrast, the RTCP or feedback channel is actually the collection of unicast channels between the receivers and the Distribution Source via the Feedback Target(s). Thus, bidirectional communication is accomplished by using SSM in the direction from Distribution Source to the receivers and using the unicast feedback channel in the direction from receivers to Distribution Source. As discussed in the next section, the nature of the channels between the Distribution Source and the Media Sender(s) may vary.

(Unicast RTCP) Feedback Target:

The Feedback Target is a logical function to which RTCP unicast feedback traffic is addressed. The functions of the Feedback Target and the Distribution Source MAY be co-located or integrated in the same entity. In this case, for a session defined as having

a Distribution Source A, on ports n for the RTP channel and k for the RTCP channel, the unicast RTCP Feedback Target is identified by an IP address of Distribution Source A on port k, unless otherwise stated in the session description. See Section 10 for details on how the address is specified. The Feedback Target MAY also be implemented in one or more entities different from the Distribution Source, and different RTP receivers MAY use different Feedback Target instances, e.g., for aggregation purposes. In this case, the Feedback Target instance(s) MUST convey the feedback received from the RTP receivers to the Distribution Source using the RTCP mechanisms specified in this document. If disjoint, the Feedback Target instances MAY be organized in arbitrary topological structures: in parallel, hierarchical, or chained. But the Feedback Target instance(s) and Distribution Source MUST share, e.g., through configuration, enough information to be able to provide coherent RTCP information to the RTP receivers based upon the RTCP feedback collected by the Feedback Target instance(s) -- as would be done if both functions were part of the same entity.

In order for unicast feedback to work, each receiver MUST direct its RTCP reports to a single specific Feedback Target instance.

SSRC:

Synchronization source as defined in [1]. This 32-bit value uniquely identifies each member in a session.

Report blocks:

Report block is the standard terminology for an RTCP reception report. RTCP [1] encourages the stacking of multiple report blocks in Sender Report (SR) and Receiver Report (RR) packets. As a result, a variable-size feedback packet may be created by one source that reports on multiple other sources in the group. The summarized reporting scheme builds upon this model through the inclusion of multiple summary report blocks in one packet. However, stacking of reports from multiple receivers is not permitted in the Simple Feedback Model (see Section 6).

4. Basic Operation

As indicated by the definitions of the preceding section, one or more Media Senders send RTP packets to the Distribution Source. The Distribution Source relays the RTP packets to the receivers using a source-specific multicast arrangement. In the reverse direction, the receivers transmit RTCP packets via unicast to one or more instances of the Feedback Target. The Feedback Target sends either the original RTCP reports (the Simple Feedback Model) or summaries of these reports (the Summary Feedback Model) to the Distribution

Source. The Distribution Source in turn relays the RTCP reports and/or summaries to the Media Senders. The Distribution Source also transmits the RTCP Sender Reports and Receiver Reports or summaries back to the receivers, using source-specific multicast.

When the Feedback Target(s) are co-located (or integrated) with the Distribution Source, redistribution of original or summarized RTCP reports is trivial. When the Feedback Target(s) are physically and/or topologically distinct from the Distribution Source, each Feedback Target either relays the RTCP packets to the Distribution Source or summarizes the reports and forwards an RTCP summary report to the Distribution Source. Coordination between multiple Feedback Targets is beyond the scope of this specification.

The Distribution Source **MUST** be able to communicate with all group members in order for either mechanism to work. The general architecture is displayed below in Figure 1. There may be a single Media Sender or multiple Media Senders (Media Sender i , $1 \leq i \leq M$) on whose behalf the Distribution Source disseminates RTP and RTCP packets. The base case, which is expected to be the most common case, is that the Distribution Source is co-located with a particular Media Sender. A basic assumption is that communication is multicast (either SSM or ASM) in the direction of the Distribution Source to the receivers ($R(j)$, $1 \leq j \leq N$) and unicast in the direction of the receivers to the Distribution Source.

Communication between Media Sender(s) and the Distribution Source may be performed in numerous ways:

- i. Unicast only: The Media Sender(s) **MAY** send RTP and RTCP via unicast to the Distribution Source and receive RTCP via unicast.
- ii. Any Source Multicast (ASM): The Media Sender(s) and the Distribution Source **MAY** be in the same ASM group, and RTP and RTCP packets are exchanged via multicast.
- iii. Source-Specific Multicast (SSM): The Media Sender(s) and the Distribution Source **MAY** be in an SSM group with the source being the Distribution Source. RTP and RTCP packets from the Media Senders are sent via unicast to the Distribution Source, while RTCP packets from the Distribution Source are sent via multicast to the Media Senders.

Note that this SSM group **MAY** be identical to the SSM group used for RTP/RTCP delivery from the Distribution Source to the receivers or it **MAY** be a different one.

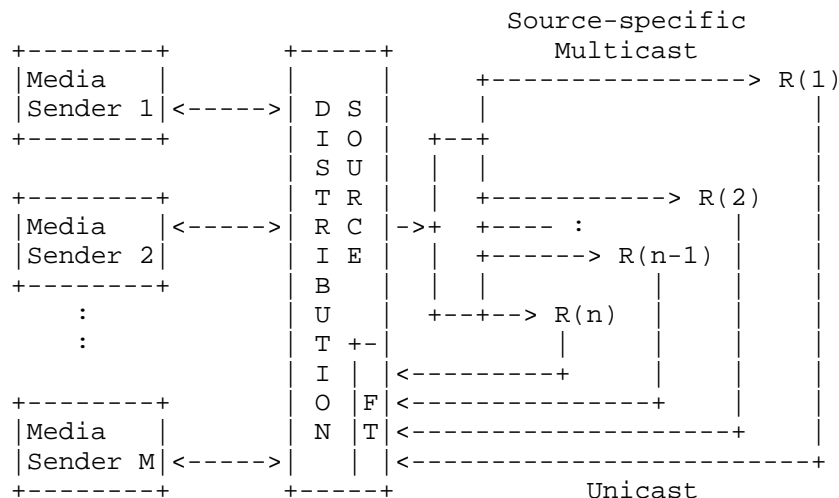
Note that Figure 1 below shows a logical diagram and, therefore, no details about the above options for the communication between Media Sender(s), Distribution Source, Feedback Target(s), and receivers are provided.

Configuration information needs to be supplied so that (among other reasons):

- o Media Sender(s) know the transport address of the Distribution Source or the transport address of the (ASM or SSM) multicast group used for the contribution link;
- o the Distribution Source knows either the unicast transport address(es) or the (ASM or SSM) multicast transport address(es) to reach the Media Sender(s);
- o receivers know the addresses of their respectively responsible Feedback Targets; and
- o the Feedback Targets know the transport address of the Distribution Source.

The precise setup and configuration of the Media Senders and their interaction with the Distribution Source is beyond the scope of this document (appropriate Session Description Protocol (SDP) descriptions MAY be used for this purpose), which only specifies how the various components interact within an RTP session. Informative examples for different configurations of the Media Sources and the Distribution Source are given in Appendix A.

Future specifications may be defined to address these aspects.



FT = Feedback Target

Transport from the Feedback Target to the Distribution Source is via unicast or multicast RTCP if they are not co-located.

Figure 1: System Architecture

The first method proposed to support unicast RTCP feedback, the 'Simple Feedback Model', is a basic reflection mechanism whereby all Receiver RTCP packets are unicast to the Feedback Target, which relays them unmodified to the Distribution Source. Subsequently, these packets are forwarded by the Distribution Source to all receivers on the multicast RTCP channel. The advantage of using this method is that an existing receiver implementation requires little modification in order to use it. Instead of sending reports to a multicast address, a receiver uses a unicast address yet still receives forwarded RTCP traffic on the multicast control channel. This method also has the advantage of being backwards compatible with standard RTP/RTCP implementations. The Simple Feedback Model is specified in Section 6.

The second method, the 'Distribution Source Feedback Summary Model', is a summarized reporting scheme that provides savings in bandwidth by consolidating Receiver Reports at the Distribution Source, optionally with help from the Feedback Target(s), into summary packets that are then distributed to all the receivers. The Distribution Source Feedback Summary Model is specified in Section 7.

The advantage of the latter scheme is apparent for large group sessions where the basic reflection mechanism outlined above generates a large amount of packet forwarding when it replicates all the information to all the receivers. Clearly, this technique requires that all session members understand the new summarized packet format outlined in Section 7.1. Additionally, the summarized scheme provides an optional mechanism to send distribution information or histograms about the feedback data reported by the whole group. Potential uses for the compilation of distribution information are addressed in Section 7.4.

To differentiate between the two reporting methods, a new SDP identifier is created and discussed in Section 10. The reporting method **MUST** be decided prior to the start of the session. A Distribution Source **MUST NOT** change the method during a session.

In a session using SSM, the network **SHOULD** prevent any multicast data from the receiver being distributed further than the first hop router. Additionally, any data heard from a non-unicast-capable receiver by other hosts on the same subnet **SHOULD** be filtered out by the host IP stack so that it does not cause problems with respect to the calculation of the receiver RTCP bandwidth share.

5. Packet Types

The RTCP packet types defined in [1], [26], and [15] are:

Type	Description	Payload number
SR	Sender Report	200
RR	Receiver Report	201
SDES	Source Description	202
BYE	Goodbye	203
APP	Application-Defined	204
RTPFB	Generic RTP feedback	205
PSFB	Payload-specific feedback	206
XR	RTCP Extension	207

This document defines one further RTCP packet format:

Type	Description	Payload number
RSI	Receiver Summary Information	209

Within the Receiver Summary Information packet, there are various types of information that may be reported and encapsulated in optional sub-report blocks:

Name	Long Name	Value
IPv4 Address	IPv4 Feedback Target Address	0
IPv6 Address	IPv6 Feedback Target Address	1
DNS Name	DNS name indicating Feedback Target Address	2
Reserved	Reserved for Assignment by Standards Action	3
Loss	Loss distribution	4
Jitter	Jitter distribution	5
RTT	Round-trip time distribution	6
Cumulative loss	Cumulative loss distribution	7
Collisions	SSRC collision list	8
Reserved	Reserved for Assignment by Standards Action	9
Stats	General statistics	10
RTCP BW	RTCP bandwidth indication	11
Group Info	RTCP group and average packet size	12
-	Unassigned	13 - 255

As with standard RTP/RTCP, the various reports MAY be combined into a single RTCP packet, which SHOULD NOT exceed the path MTU. Packets continue to be sent at a rate that is inversely proportional to the group size in order to scale the amount of traffic generated.

6. Simple Feedback Model

6.1. Packet Formats

The Simple Feedback Model uses the same packet types as traditional RTCP feedback described in [1]. Receivers still generate Receiver Reports with information on the quality of the stream received from the Distribution Source. The Distribution Source still MUST create Sender Reports that include timestamp information for stream synchronization and round-trip time calculation. Both Media Senders and receivers are required to send SDES packets as outlined in [1]. The rules for generating BYE and APP packets as outlined in [1] also apply.

6.2. Distribution Source Behavior

For the Simple Feedback Model, the Distribution Source MUST provide a basic packet-reflection mechanism. It is the default behavior for any Distribution Source and is the minimum requirement for acting as a Distribution Source to a group of receivers using unicast RTCP feedback.

The Distribution Source (unicast Feedback Target) MUST listen for unicast RTCP data sent to the RTCP port. All valid unicast RTCP packets received on this port MUST be forwarded by the Distribution Source to the group on the multicast RTCP channel. The Distribution

Source MUST NOT stack report blocks received from different receivers into one packet for retransmission to the group. Every RTCP packet from each receiver MUST be reflected individually.

If the Media Sender(s) are not part of the SSM group for RTCP packet reflection, the Distribution Source MUST also forward the RTCP packets received from the receivers to the Media Sender(s). If there is more than one Media Sender and these Media Senders do not communicate via ASM with the Distribution Source and each other, the Distribution Source MUST forward each RTCP packet originated by one Media Sender to all other Media Senders.

The Distribution Source MUST forward RTCP packets originating from the Media Sender(s) to the receivers.

The reflected or forwarded RTCP traffic SHOULD NOT be counted as its own traffic in the transmission interval calculation by the Distribution Source. In other words, the Distribution Source SHOULD NOT consider reflected packets as part of its own control data bandwidth allowance. However, reflected packets MUST be processed by the Distribution Source and the average RTCP packet size, RTCP transmission rate, and RTCP statistics MUST be calculated. The algorithm for computing the allowance is explained in Section 9.

6.3. Disjoint Distribution Source and Feedback Target(s)

If the Feedback Target function is disjoint from the Distribution Source, the Feedback Target(s) MUST forward all RTCP packets from the receivers or another Feedback Target -- directly or indirectly -- to the Distribution Source.

6.4. Receiver Behavior

Receivers MUST listen on the RTP channel for data and on the RTCP channel for control. Each receiver MUST calculate its share of the control bandwidth R/n , in accordance with the profile in use, so that a fraction of the RTCP bandwidth, R , allocated to receivers is divided equally between the number of unique receiver SSRCs in the session, n . R may be $rtcp_bw * 0.75$ or $rtcp_bw * 0.5$ (depending on the ratio of senders to receivers as per [1]) or may be set explicitly by means of an SDP attribute [10]. See Section 9 for further information on the calculation of the bandwidth allowance. When a receiver is eligible to transmit, it MUST send a unicast Receiver Report packet to the Feedback Target following the rules defined in Section 9.

When a receiver observes either RTP packets or RTCP Sender Reports from a Media Sender with an SSRC that collides with its own chosen SSRC, it MUST change its own SSRC following the procedures of [1]. The receiver MUST do so immediately after noticing and before sending any (further) RTCP feedback messages.

If a receiver has out-of-band information available about the Media Sender SSRC(s) used in the media session, it MUST NOT use the same SSRC for itself. Even if such out-of-band information is available, a receiver MUST obey the above collision-resolution mechanisms.

Further mechanisms defined in [1] apply for resolving SSRC collisions between receivers.

6.5. Media Sender Behavior

Media Senders listen on a unicast or multicast transport address for RTCP reports sent by the receivers (and forwarded by the Distribution Source) or other Media Senders (forwarded by the Distribution Source if needed). Processing and general operation follows [1].

A Media Sender that observes an SSRC collision with another entity that is not also a Media Sender MAY delay its own collision-resolution actions as per [1], by $5 * 1.5 * T_d$, with T_d being the deterministic calculated reporting interval, for receivers to see whether the conflict still exists. SSRC collisions with other Media Senders MUST be acted upon immediately.

Note: This gives precedence to Media Senders and places the burden of collision resolution on the RTP receivers.

Sender SSRC information MAY be communicated out-of-band, e.g., by means of SDP media descriptions. Therefore, senders SHOULD NOT change their own SSRC aggressively or unnecessarily.

7. Distribution Source Feedback Summary Model

In the Distribution Source Feedback Summary Model, the Distribution Source is required to summarize the information received from all the Receiver Reports generated by the receivers and place the information into summary reports. The Distribution Source Feedback Summary Model introduces a new report block format, the Receiver Summary Information (RSI) report, and a number of optional sub-report block formats, which are enumerated in Section 7.1. As described in Section 7.3, individual instances of the Feedback Target may provide preliminary summarization to reduce the processing load at the Distribution Source.

Sub-reports appended to the RSI report block provide more detailed information on the overall session characteristics reported by all receivers and can also convey important information such as the feedback address and reporting bandwidth. Which sub-reports are mandatory and which ones are optional is defined below.

From an RTP perspective, the Distribution Source is an RTP receiver, generating its own Receiver Reports and sending them to the receiver group and to the Media Senders. In the Distribution Source Feedback Summary Model, an RSI report block MUST be appended to all RRs the Distribution Source generates.

In addition, the Distribution Source MUST forward the RTCP SR reports and SDES packets of Media Senders without alteration. If the Distribution Source is actually a Media Sender, even if it is the only session sender, it MUST generate its own Sender Report (SR) packets for its role as a Media Sender and its Receiver Reports in its role as the Distribution Source.

The Distribution Source MUST use its own SSRC value for transmitting summarization information and MUST perform proper SSRC collision detection and resolution.

The Distribution Source MUST send at least one Receiver Summary Information packet for each reporting interval. The Distribution Source MAY additionally stack sub-report blocks after the RSI packet. If it does so, each sub-report block MUST correspond to the RSI packet and constitutes an enhancement to the basic summary information required by the receivers to calculate their reporting time interval. For this reason, additional sub-report blocks are not required but recommended. The compound RTCP packets containing the RSI packet and the optional corresponding sub-report blocks MUST be formed according to the rules defined in [1] for receiver-issued packets, e.g., they MUST begin with an RR packet, contain at least an SDES packet with a CNAME, and MAY contain further RTCP packets and SDES items.

Every RSI packet MUST contain either a Group and Average Packet Size sub-report or an RTCP Bandwidth sub-report for bandwidth indications to the receivers.

7.1. Packet Formats

All numeric values comprising multiple (usually two or four) octets MUST be encoded in network byte order.

7.1.1. RSI: Receiver Summary Information Packet

The RSI report block has a fixed header size followed by a variable length report:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|V=2|P|reserved | PT=RSI=209 | length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               SSRC                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Summarized SSRC                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               NTP Timestamp (most significant word) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               NTP Timestamp (least significant word) |
+-----+-----+-----+-----+-----+-----+-----+-----+
:                               Sub-report blocks                     :
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The RSI packet includes the following fields:

Length: 16 bits

As defined in [1], the length of the RTCP packet in 32-bit words minus one, including the header and any padding.

SSRC: 32 bits

The SSRC of the Distribution Source.

Summarized SSRC: 32 bits


The SSRC (of the Media Sender) of which this report contains a summary.

Timestamp: 64 bits

Indicates the wallclock time when this report was sent. Wallclock time (absolute date and time) is represented using the timestamp format of the Network Time Protocol (NTP), which is in seconds relative to 0h UTC on 1 January 1900 [1]. The wallclock time MAY (but need not) be NTP-synchronized but it MUST provide a consistent behavior in the advancement of time (similar to NTP). The full-resolution NTP timestamp is used, which is a 64-bit, unsigned, fixed-point number with the integer part in the first 32 bits and the fractional part in the last 32 bits. This format is similar to RTCP Sender Reports (Section 6.4.1 of [1]). The timestamp value is used to enable detection of duplicate packets, reordering, and to provide a chronological profile of the feedback reports.

7.1.2. Sub-Report Block Types

For RSI reports, this document also introduces a sub-report block format specific to the RSI packet. The sub-report blocks are appended to the RSI packet using the following generic format. All sub-report blocks MUST be 32-bit aligned.



SRBT: 8 bits

Sub-Report Block Type. The sub-report block type identifier. The values for the sub-report block types are defined in Section 5.

Length: 8 bits

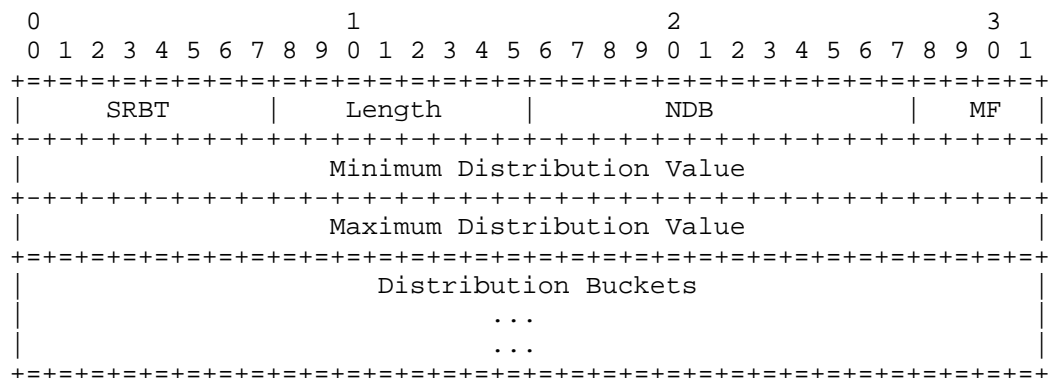
The length of the sub-report in 32-bit words.

SRBT-specific data: <length * 4 - 2> octets

This field may contain type-specific information based upon the SRBT value.

7.1.3. Generic Sub-Report Block Fields

For the sub-report blocks that convey distributions of values (Loss, Jitter, RTT, Cumulative Loss), a flexible 'data bucket'-style report is used. This format divides the data set into variable-size buckets that are interpreted according to the guide fields at the head of the report block.



The SRBT and length fields are calculated as explained in Section 7.1.2.

Number of distribution buckets (NDB): 12 bits

The number of distribution buckets of data. The size of each bucket can be calculated using the formula $((\text{length} * 4) - 12) * 8 / \text{NDB number of bits}$. The calculation is based on the length of the whole sub-report block in octets ($\text{length} * 4$) minus the header of 12 octets. Providing 12 bits for the NDB field enables bucket sizes as small as 2 bits for a full-length packet of MTU 1500 bytes. The bucket size in bits must always be divisible by 2 to ensure proper byte alignment. A bucket size of 2 bits is fairly restrictive, however, and it is expected that larger bucket sizes will be more practical for most distributions.

Multiplicative Factor (MF): 4 bits

2^{MF} indicates the multiplicative factor to be applied to each distribution bucket value. Possible values of MF are 0 - 15, creating a range of values from MF = 1, 2, 4 ... 32768. Appendix B gives an example of the use of the multiplicative factor; it is meant to provide more "bits" without having them -- the bucket values get scaled up by the MF.

Length: 8 bits

The length field tells the receiver the full length of the sub-report block in 32-bit words (i.e., $\text{length} * 4$ bytes) and enables the receiver to identify the bucket size. For example, given no MTU restrictions, the data portion of a distribution packet may be only as large as 1008 bytes ($255 * 4 - 12$), providing up to 4032 data buckets of length 2 bits, or 2016 data buckets of length 4 bits, etc.

Minimum distribution value (min): 32 bits

The minimum distribution value, in combination with the maximum distribution value, indicates the range covered by the data bucket values.

Maximum distribution value (max): 32 bits

The maximum distribution value, in combination with the minimum distribution value, indicates the range covered by the data bucket values. The significance and range of the distribution values is defined in the individual subsections for each distribution type (DT).

Distribution buckets: each bucket is $((\text{length} * 4) - 12) * 8 / \text{NDB}$ bits

The size and number of buckets is calculated as outlined above based upon the value of NDB and the length of the packet. The values for distribution buckets are equally distributed; according to the following formula, distribution bucket x (with $0 \leq x < \text{NDB}$) covers the value range:

```
x = 0; [min, min + (max - min) / NDB]
x > 0; [min + (x) * (max - min) / NDB,
        min + (x + 1) * (max - min) / NDB]
```

Interpretation of the minimum, maximum, and distribution values in the sub-report block is sub-report-specific and is addressed individually in the sections below. The size of the sub-report block is variable, as indicated by the packet length field.

Note that, for any bucket-based reporting, if the Distribution Source and the Feedback Target(s) are disjoint entities, out-of-band agreement on the bucket-reporting granularity is recommended to allow the Distribution Source to forward accurate information in these kinds of reports to the receivers.

7.1.4. Loss Sub-Report Block

The Loss sub-report block allows a receiver to determine how its own reception quality relates to the other recipients. A receiver may use this information, e.g., to drop out of a session (instead of sending lots of error repair feedback) if it finds itself isolated at the lower end of the reception quality scale.

The Loss sub-report block indicates the distribution of losses as reported by the receivers to the Distribution Source. Values are expressed as a fixed-point number with the binary point at the left edge of the field similar to the "fraction lost" field in SR and RR packets, as defined in [1]. The Loss sub-report block type (SRBT) is 4.

Valid results for the minimum distribution value field are 0 - 254. Similarly, valid results for the maximum distribution value field are 1 - 255. The minimum distribution value MUST always be less than the maximum.

For examples on processing summarized loss report sub-blocks, see Appendix B.

7.1.5. Jitter Sub-Report Block

A Jitter sub-report block indicates the distribution of the estimated statistical variation of the RTP data packet inter-arrival time reported by the receivers to the Distribution Source. This allows receivers both to place their own observed jitter values in context with the rest of the group and to approximate dynamic parameters for playout buffers. See [1] for details on the calculation of the values and the relevance of the jitter results. Jitter values are measured in timestamp units with the rate used by the Media Sender and expressed as unsigned integers. The minimum distribution value MUST always be less than the maximum. The Jitter sub-report block type (SRBT) is 5.

Since timestamp units are payload-type specific, the relevance of a jitter value is impacted by any change in the payload type during a session. Therefore, a Distribution Source MUST NOT report jitter distribution values for at least 2 reporting intervals after a payload type change occurs.

7.1.6. Round-Trip Time Sub-Report Block

A Round-Trip Time sub-report indicates the distribution of round-trip times from the Distribution Source to the receivers, providing receivers with a global view of the range of values in the group. The Distribution Source is capable of calculating the round-trip time to any other member since it forwards all the SR packets from the Media Sender(s) to the group. If the Distribution Source is not itself a Media Sender, it can calculate the round-trip time from itself to any of the receivers by maintaining a record of the SR sender timestamp and the current time as measured from its own system clock. The Distribution Source consequently calculates the round-trip time from the Receiver Report by identifying the corresponding SR timestamp and subtracting the RR advertised holding time as reported by the receiver from its own time difference measurement, as calculated by the time the RR packet is received and the recorded time the SR was sent.

The Distribution Source has the option of distributing these round-trip time estimations to the whole group, uses of which are described in Section 7.4. The round-trip time is expressed in units of 1/65536 seconds and indicates an absolute value. This is calculated by the Distribution Source, based on the Receiver Report responses declaring the time difference since an original Sender Report and on the holding time at the receiver. The minimum distribution value MUST always be less than the maximum. The Round-Trip Time sub-report block type (SRBT) is 6.

Note that if the Distribution Source and the Feedback Target functions are disjoint, it is only possible for the Distribution Source to determine RTT if all the Feedback Targets forward all RTCP reports from the receivers immediately (i.e., do not perform any preliminary summarization) and include at least the RR packet.

7.1.7. Cumulative Loss Sub-Report Block

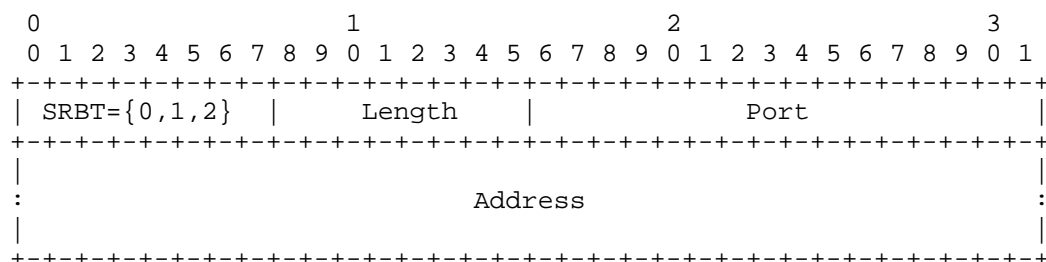
The cumulative loss field in a Receiver Report [1], in contrast to the fraction lost field, is intended to provide some historical perspective on the session performance, i.e., the total number of lost packets since the receiver joined the session. The cumulative loss value provides a longer-term average by summing over a larger sample set (typically the whole session). The Distribution Source MAY record the values as reported by the receiver group and report a distribution of values. Values are expressed as a fixed-point number with the binary point at the left edge of the field, in the same manner as the Loss sub-report block. Since the individual Receiver Reports give the cumulative number of packets lost but not the cumulative number sent, which is required as a denominator to calculate the long-term fraction lost, the Distribution Source MUST maintain a record of the cumulative number lost and extended highest sequence number received, as reported by each receiver at some point in the past. Ideally, the recorded values are from the first report received. Subsequent calculations are then based on ($\text{<the new cumulative loss value>} - \text{<the recorded value>} / (\text{<new extended highest sequence number>} - \text{<recorded sequence number>})$).

Valid results for the minimum distribution value field are 0 - 254. Similarly, valid results for the maximum distribution value field are 1 - 255. The minimum distribution value MUST always be less than the maximum. The Cumulative Loss sub-report block type (SRBT) is 7.

7.1.8. Feedback Target Address Sub-Report Block

The Feedback Target Address block provides a dynamic mechanism for the Distribution Source to signal an alternative unicast RTCP feedback address to the receivers. If a block of this type is included, receivers MUST override any static SDP address information for the session with the Feedback Target address provided in this sub-report block.

If a Feedback Target Address sub-report block is used, it MUST be included in every RTCP packet originated by the Distribution Source to ensure that all receivers understand the information. In this manner, receiver behavior should remain consistent even in the face of packet loss or when there are late session arrivals.



SRBT: 8 bits

The type of sub-report block that corresponds to the type of address is as follows:

- 0: IPv4 address
- 1: IPv6 address
- 2: DNS name

Length: 8 bits

The length of the sub-report block in 32-bit words. For an IPv4 address, this should be 2 (i.e., total length = 4 + 4 = 2 * 4 octets). For an IPv6 address, this should be 5. For a DNS name, the length field indicates the number of 32-bit words making up the string plus 1 byte and any additional padding required to reach the next word boundary.

Port: 2 octets

The port number to which receivers send feedback reports. A port number of 0 is invalid and MUST NOT be used.

Address: 4 octets (IPv4), 16 octets (IPv6), or n octets (DNS name)

The address to which receivers send feedback reports. For IPv4 and IPv6, fixed-length address fields are used. A DNS name is an arbitrary-length string that is padded with null bytes to the next 32-bit boundary. The string MAY contain Internationalizing Domain Names in Applications (IDNA) domain names and MUST be UTF-8 encoded [11].


A Feedback Target Address block for a certain address type (i.e., with a certain SRBT of 0, 1, or 2) MUST NOT occur more than once within a packet. Numerical Feedback Target Address blocks for IPv4 and IPv6 MAY both be present. If so, the resulting transport addresses MUST point to the same logical entity.

If a Feedback Target address block with an SRBT indicating a DNS name is present, there SHOULD NOT be any other numerical Feedback Target Address blocks present.

The Feedback Target Address presents a significant security risk if accepted from unauthenticated RTCP packets. See Sections 11.3 and 11.4 for further discussion.

7.1.9. Collision Sub-Report Block

The Collision SSRC sub-report provides the Distribution Source with a mechanism to report SSRC collisions amongst the group. In the event that a non-unique SSRC is discovered based on the tuple [SSRC, CNAME], the collision is reported in this block and the affected nodes must reselect their respective SSRC identifiers.



Collision SSRC: $n * 32$ bits

This field contains a list of SSRCs that are duplicated within the group. Usually this is handled by the hosts upon detection of the same SSRC; however, since receivers in an SSM session using the Distribution Source Feedback Summary Model no longer have a global view of the session, the collision algorithm is handled by the Distribution Source. SSRCs that collide are listed in the packet. Each Collision SSRC is reported only once for each collision detection. If more Collision SSRCs need to be reported than fit into an MTU, the reporting is done in a round robin fashion so that all Collision SSRCs have been reported once before the second round of reporting starts. On receipt of the packet, receiver(s) MUST detect the collision and select another SSRC, if the collision pertains to them.

The Collision sub-report block type (SRBT) is 8.

Collision detection is only possible at the Distribution Source. If the Distribution Source and Feedback Target functions are disjoint and collision reporting across RTP receiver SSRCs shall be provided, the Feedback Targets(s) MUST forward the RTCP reports from the RTP receivers, including at least the RR and the SDES packets to the Distribution Source.

In system settings in which, by explicit configuration or implementation, RTP receivers are not going to act as Media Senders in a session (e.g., for various types of television broadcasting), SSRC collision detection MAY be omitted for RTP receivers. In system settings in which an RTP receiver MAY become a Media Sender (e.g., any conversational application), SSRC collision detection MUST be provided for RTP receivers.

Note: The purpose of SSRC collision reporting is to ensure unique identification of RTCP entities. This is of particular relevance for Media Senders so that an RTP receiver can properly associate each of the multiple incoming media streams (via the Distribution Source) with the correct sender. Collision resolution for Media Senders is not affected by the Distribution Source's collision reporting because all SR reports are always forwarded among the senders and to all receivers. Collision resolution for RTP receivers is of particular relevance for those receivers capable of becoming a Media Sender. RTP receivers that will, by configuration or implementation choice, not act as a sender in a particular RTP session, do not necessarily need to be aware of collisions as long as the those entities receiving and processing RTCP feedback messages from the receivers are capable of disambiguating the various RTCP receivers (e.g., by CNAME).

Note also that RTP receivers should be able to deal with the changing SSRCs of a Media Sender (like any RTP receiver has to do.) and, if possible, be prepared to continuously render a media stream nevertheless.

7.1.10. General Statistics Sub-Report Block

The General Statistics sub-report block is used if specifying buckets is deemed too complex. With the General Statistics sub-report block, only aggregated values are reported back. The rules for the generation of these values are provided in point b of Section 7.2.1.

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								
SRBT=10										Length										Reserved																			
MFL										HCNL																													
Median inter-arrival jitter																																							

Median fraction lost (MFL): 8 bits

The median fraction lost indicated by Receiver Reports forwarded to this Distribution Source, expressed as a fixed-point number with the binary point at the left edge of the field. A value of all '1's indicates that the MFL value is not provided.

Highest cumulative number of packets lost (HCNL): 24 bits

Highest 'cumulative number of packets lost' value out of the most recent RTCP RR packets from any of the receivers. A value of all '1's indicates that the HCNL value is not provided.

Median inter-arrival jitter: 32 bits

Median 'inter-arrival jitter' value out of the most recent RTCP RR packets from the receiver set. A value of all '1's indicates that this value is not provided.

The General Statistics sub-report block type (SRBT) is 10.

Note that, in case the Distribution Source and the Feedback Target functions are disjoint, it is only possible for the Distribution Source to determine the median of the inter-arrival jitter if all the Feedback Targets forward all RTCP reports from the receivers immediately and include at least the RR packet.

7.1.11. RTCP Bandwidth Indication Sub-Report Block

This sub-report block is used to inform the Media Senders and receivers about either the maximum RTCP bandwidth that is supposed to be used by each Media Sender or the maximum bandwidth allowance to be used by each receiver. The value is applied universally to all Media Senders or all receivers. Each receiver MUST base its RTCP transmission interval calculation on the average size of the RTCP packets sent by itself. Conversely, each Media Sender MUST base its RTCP transmission interval calculation on the average size of the RTCP packets sent by itself.

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								
SRBT=11										Length										S R										Reserved									
RTCP Bandwidth																																							

Sender (S): 1 bit

The contained bandwidth value applies to each Media Sender.

Receivers (R): 1 bit

The contained bandwidth value applies to each RTP receiver.

Reserved: 14 bits

MUST be set to zero upon transmission and ignored upon reception.

RTCP Bandwidth: 32 bits

If the S bit is set to 1, this field indicates the maximum bandwidth allocated to each individual Media Sender. This also informs the receivers about the RTCP report frequency to expect from the senders. This is a fixed-point value with the binary point in between the second and third bytes. The value represents the bandwidth allocation per receiver in kilobits per second, with values in the range $0 \leq BW < 65536$.

If the R bit is set to 1, this field indicates the maximum bandwidth allocated per receiver for sending RTCP data relating to the session. This is a fixed-point value with the binary point in between the second and third bytes. The value represents the bandwidth allocation per receiver in kilobits per second, with values in the range $0 \leq BW < 65536$. Each receiver MUST use this value for its bandwidth allowance.

This report block SHOULD only be used when the Group and Average Packet Size sub-report block is NOT included. The RTCP Bandwidth Indication sub-report block type (SRBT) is 11.

7.1.12. RTCP Group and Average Packet Size Sub-Report Block

This sub-report block is used to inform the Media Senders and receivers about the size of the group (used for calculating feedback bandwidth allocation) and the average packet size of the group. This sub-report MUST always be present, appended to every RSI packet, unless an RTCP Bandwidth Indication sub-report block is included (in which case it MAY, but need not, be present).

Note: The RTCP Bandwidth Indication sub-report block allows the Distribution Source to hide the actual group size from the receivers while still avoiding feedback implosion.

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|  SRBT=12  |  Length  |  Average Packet Size  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Receiver Group Size
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Group size: 32 bits

This field provides the Distribution Source's view of the number of receivers in a session. Note that the number of Media Senders is not explicitly reported since it can be derived by observing the RTCP SR packets forwarded by the Distribution Source. The group size is calculated according to the rules outlined in [1]. When this sub-report block is included, this field **MUST** always be present.

Average RTCP packet size: 16 bits

This field provides the Distribution Source's view of the average RTCP packet size as locally calculated, following the rules defined in [1]. The value is an unsigned integer, counting octets. When this sub-report block is included, this field **MUST** always be present.

The Group and Average Packet Size sub-report block type (SRBT) is 12.

7.2. Distribution Source Behavior

In the Distribution Source Feedback Summary Model, the Distribution Source (the unicast Feedback Target) **MUST** listen for unicast RTCP packets sent to the RTCP port. All RTCP packets received on this port **MUST** be processed by the Distribution Source, as described below. The processing **MUST** take place per Media Sender SSRC for which Receiver Reports are received.

The Distribution Source acts like a regular RTCP receiver. In particular, it receives an RTP stream from each RTP Media Sender(s) and **MUST** calculate the proper reception statistics for these RTP streams. It **MUST** transmit the resulting information as report blocks contained in each RTCP packet it sends (in an RR packet).

Note: This information may help to determine the transmission characteristics of the feed path from the RTP sender to the Distribution Source (if these are separate entities).

The Distribution Source is responsible for accepting RTCP packets from the receivers and for interpreting and storing per-receiver information, as defined in [1]. The importance of providing these

functions is apparent when creating the RSI and sub-report block reports since incorrect information can have serious implications. Section 11 addresses the security risks in detail.

As defined in [1], all RTCP reports from the Distribution Source MUST start with an RR report, followed by any relevant SDES fields. Then the Distribution Source MUST append an RSI header and sub-reports containing any summarization-specific data. In addition, either the Group and Average Packet Size sub-report or the Receiver RTCP Bandwidth sub-report block MUST be appended to the RSI header.

A Distribution Source has the option of masking the group size by including only an RTCP Bandwidth sub-report. If both sub-reports are included, the receiver is expected to give precedence to the information contained in the Receiver RTCP Bandwidth sub-report.

The Receiver RTCP Bandwidth sub-report block provides the Distribution Source with the capability to control the amount of feedback from the receivers, and the bandwidth value MAY be increased or decreased based upon the requirements of the Distribution Source. Regardless of the value selected by the Distribution Source for the Receiver RTCP Bandwidth sub-report block, the Distribution Source MUST continue to forward Sender Reports and RSI packets at the rate allowed by the total RTCP bandwidth allocation. See Section 9 for further details about the frequency of reports.

A Distribution Source MAY start out reporting group size and switch to Receiver RTCP Bandwidth reporting later on and vice versa. If the Distribution Source does so, it SHOULD ensure that the correspondingly indicated values for the Receiver RTCP Bandwidth sub-report roughly match, i.e., are at least the same order of magnitude.

In order to identify SSRC collisions, the Distribution Source is responsible for maintaining a record of each SSRC and the corresponding CNAME within at least one reporting interval by the group, in order to differentiate between clients. It is RECOMMENDED that an updated list of more than one interval be maintained to increase accuracy. This mechanism should prevent the possibility of collisions since the combination of SSRC and CNAME should be unique, if the CNAME is generated correctly. If collisions are not detected, the Distribution Source will have an inaccurate impression of the group size. Since the statistical probability is very low that collisions will both occur and be undetectable, this should not be a significant concern. In particular, the clients would have to randomly select the same SSRC and have the same username + IP address (e.g., using private address space behind a NAT router).

7.2.1. Receiver Report Aggregation

The Distribution Source is responsible for aggregating reception-quality information received in RR packets. In doing so, the Distribution Source **MUST** consider the report blocks received in every RR packet and **MUST NOT** consider the report blocks of an SR packet.

Note: the receivers will get the information contained in the SR packets anyway by packet forwarding, so duplication of this information should be avoided.

For the optional sub-report blocks, the Distribution Source **MAY** decide which are the most significant feedback values to convey and **MAY** choose not to include any. The packet format provides flexibility in the amount of detail conveyed by the data points. There is a tradeoff between the granularity of the data and the accuracy of the data based on the multiplicative factor (MF), the number of buckets, and the min and max values. In order to focus on a particular region of a distribution, the Distribution Source can adjust the minimum and maximum values and either increase the number of buckets, and possibly the factorization, or decrease the number of buckets and provide more accurate values. See Appendix B for detailed examples on how to convey a summary of RTCP Receiver Reports as RSI sub-report block information.

For each value the Distribution Source decides to include in an RSI packet, it **MUST** adhere to the following measurement rules.

- a) If the Distribution Source intends to use a sub-report to convey a distribution of values (Sections 7.1.3 to 7.1.7), it **MUST** keep per-receiver state, i.e., remember the last value V reported by the respective receiver. If a new value is reported by a receiver, the existing value **MUST** be replaced by the new one. The value **MUST** be deleted (along with the entire entry) if the receiver is timed out (as per Section 6.3.5 of [1]) or has sent a BYE packet (as per Section 6.3.7 of [1]).

All the values collected in this way **MUST** be included in the creation of the subsequent Distribution sub-report block.

The results should correspond as closely as possible to the values received during the interval since the last report. The Distribution Source may stack as many sub-report blocks as required in order to convey different distributions. If the distribution size exceeds the largest packet length (1008 bytes data portion), more packets **MAY** be stacked with additional information (but in total **SHOULD NOT** exceed the path MTU).

All stacked sub-reports MUST be internally consistent, i.e., generated from the same session data. Overlapping of distributions is therefore allowed, and variation in values should only occur as a result of data set granularity, with the more accurate bucket sizes taking precedence in the event that values differ. Non-divisible values MUST be rounded up or down to the closest bucket value, and the number of data buckets must always be an even number, padding where necessary with an additional zero bucket value to maintain consistency.

Note: This intentionally provides persistent full coverage of the entire session membership to avoid oscillations due to changing membership samples.

Scheduling the transmission of summarization reports is left to the discretion of the Distribution Source. However, the Distribution Source MUST adhere to the bandwidth limitations for Receiver Reports as defined for the respective AV profile in use.

- b) If the Distribution Source intends to report a short summary using the General Statistics sub-report block format, defined in Section 7.1.10, for EACH of the values included in the report (MFL, HCNL, average inter-arrival jitter), it MUST keep a timer `T_summary` as well as a sufficient set of variables to calculate the summaries for the last three reporting intervals. This MAY be achieved by keeping per-receiver state (i.e., remember the last value `V` reported by the respective receiver) or by maintaining summary variables for each of these intervals.

The summary values are included here to reflect the current group situation. By recording the last three reporting intervals, the Distribution Source incorporates reports from all members while allowing for individual RTCP Receiver Report packet losses. The process of collecting these values also aims to avoid resetting any of the values and then having to send out an RSI report based upon just a few values collected. If data is available for less than three reporting intervals (as will be the case for the first two reports sent), only those values available are considered.

The timer `T_summary` MUST be initialized as $T_summary = 1.5 * T_d$, where T_d is the deterministic reporting interval, and MUST be updated following timer reconsideration whenever the group size or the average RTCP size ("`avg_rtcp_size`") changes. This choice should allow each receiver to report once per interval.

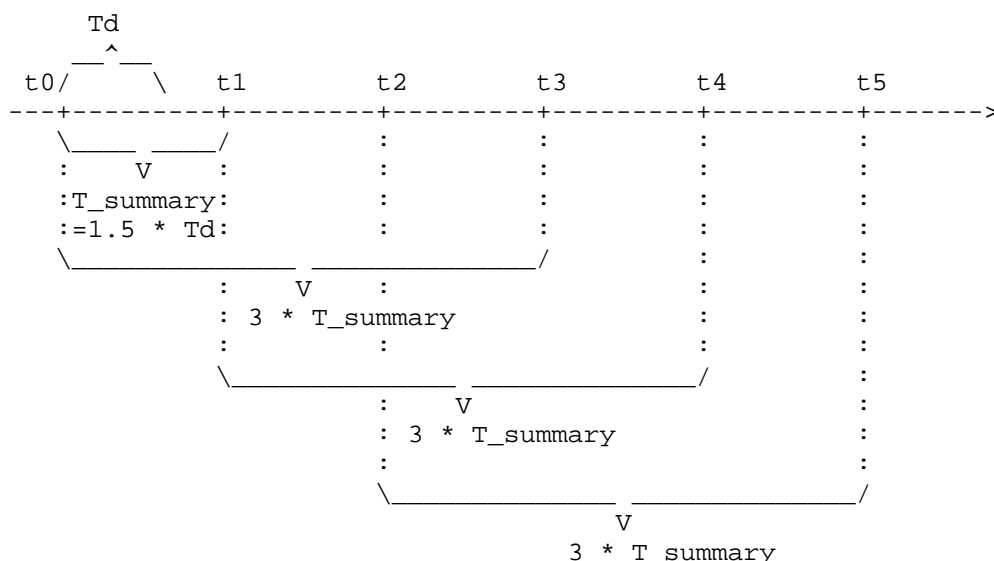


Figure 2: Overview of Summarization Reporting

Figure 2 depicts how the summarization reporting shall be performed. At time t_3 , the RTCP reports collected from t_0 through t_3 are included in the RSI reporting; at time t_4 , those from t_1 through t_4 ; and so on. The RSI summary report sent MUST NOT include any RTCP report from more than three reporting intervals ago, e.g., the one sent at time t_5 , must not include reports received at the Distribution Source prior to t_2 .

7.2.2. Handling Other RTCP Packets from RTP Receivers

When following the Feedback Summary Model, the Distribution Source MAY reflect any other RTCP packets of potential relevance to the receivers (such as APP, RTPFB, PSFB) to the receiver group. Also, it MAY decide not to forward other RTCP packets not needed by the receivers such as BYE, RR, SDDES (because the Distribution Source performs collision resolution), group size estimation, and RR aggregation. The Distribution Source MUST NOT forward RR packets to the receiver group.

If the Distribution Source is able to interpret and aggregate information contained in any RTCP packets other than RR, it MAY include the aggregated information along with the RSI packet in its own compound RTCP packet.

Aggregation MAY be a null operation, i.e., the Distribution Source MAY simply append one or more RTCP packets from receivers to the compound RTCP packet (containing at least RR, SDES, and RSI from the Distribution Source).

Note: This is likely to be useful only for a few cases, e.g., to forward aggregated information from RTPFB Generic NACK packets and thereby maintain the damping property of [15].

Note: This entire processing rule implies that the flow of information contained in non-RR RTCP packets may terminate at the Distribution Source, depending on its capabilities and configuration.

The configuration of the RTCP SSM media session (expressed in SDP) MUST specify explicitly which processing the Distribution Source will apply to which RTCP packets. See Section 10.1 for details.

7.2.3. Receiver Report Forwarding

If the Media Sender(s) are not part of the SSM group for RTCP packet reflection, the Distribution Source MUST explicitly inform the Media Senders of the receiver group. To achieve this, the Distribution Source has two options: 1) it forwards the RTCP RR and SDES packets received from the receivers to the Media Sender(s); or 2) if the Media Senders also support the RTCP RSI packet, the Distribution Source sends the RSI packets not just to the receivers but also to the Media Senders.

If the Distribution Source decides to forward RR and SDES packets unchanged, it MAY also forward any other RTCP packets to the senders.

If the Distribution Source decides to forward RSI packets to the senders, the considerations of Section 7.2.2 apply.

7.2.4. Handling Sender Reports

The Distribution Source also receives RTCP (including SR) packets from the RTP Media Senders.

The Distribution Source MUST forward all RTCP packets from the Media Senders to the RTP receivers.

If there is more than one Media Sender and these Media Senders do not communicate via ASM with the Distribution Source and each other, the Distribution Source MUST forward each RTCP packet from any one Media Sender to all other Media Senders.

7.2.5. RTCP Data Rate Calculation

As noted above, the Distribution Source is a receiver from an RTP perspective. The Distribution Sources MUST calculate its deterministic transmission interval T_d as every other receiver; however, it MAY adjust its available data rate depending on the destination transport address and its local operation:

1. For sending its own RTCP reports to the SSM group towards the receivers, the Distribution Source MAY use up to the joint share of all receivers as it is forwarding summaries on behalf of all of them. Thus, the Distribution Source MAY send its reports up to every T_d/R time units, with R being the number of receivers.
2. For sending its own RTCP reports to the Media Senders only (i.e., if the Media Senders are not part of the SSM group), the allocated rate depends on the operation of the Distribution Source:
 - a) If the Distribution Source only sends RSI packets along with its own RTCP RR packets, the same rate calculation applies as for #1 above.
 - b) If the Distribution Source forwards RTCP packets from all other receivers to the Media Senders, then it MUST adhere to the same bandwidth share for its own transmissions as all other receivers and use the calculation as per [1].

7.2.6. Collision Resolution

A Distribution Source observing RTP packets from a Media Sender with an SSRC that collides with its own chosen SSRC MUST change its own SSRC following the procedures of [1] and MUST do so immediately after noticing.

A Distribution Source MAY use out-of-band information about the Media Sender SSRC(s) used in the media session when available to avoid SSRC collisions with Media Senders. Nevertheless, the Distribution Source MUST monitor Sender Report (SR) packets to detect any changes, observe collisions, and then follow the above collision-resolution procedure.

For collision resolution between the Distribution Source and receivers or the Feedback Target(s) (if a separate entity, as described in the next subsection), the Distribution Source and the Feedback Target (if separate) operate similar to ordinary receivers.

7.3. Disjoint Distribution Source and Feedback Target

If the Distribution Source and the Feedback Target are disjoint, the processing of the Distribution Source is limited by the amount of RTCP feedback information made available by the Feedback Target.

The Feedback Target(s) MAY simply forward all RTCP packets incoming from the RTP receivers to the Distribution Source, in which case the Distribution Source will have all the necessary information available to perform all the functions described above.

The Feedback Target(s) MAY also perform aggregation of incoming RTCP packets and send only aggregated information to the Distribution Source. In this case, the Feedback Target(s) MUST use correctly formed RTCP packets to communicate with the Distribution Source and they MUST operate in concert with the Distribution Source so that the Distribution Source and the Feedback Target(s) appear to be operating as a single entity. The Feedback Target(s) MUST report their observed receiver group size to the Distribution Source, either explicitly by means of RSI packets or implicitly by forwarding all RR packets.

Note: For example, for detailed statistics reporting, the Distribution Source and the Feedback Target(s) may need to agree on a common reporting granularity so that the Distribution Source can aggregate the buckets incoming from various Feedback Targets into a coherent report sent to the receivers.

The joint behavior of the Distribution Source and Feedback Target(s) MUST be reflected in the (SDP-based) media session description as per Section 7.2.2.

If the Feedback Target performs summarization functions, it MUST also act as a receiver and choose a unique SSRC for its own reporting towards the Distribution Source. The collision-resolution considerations for receivers apply.

7.4. Receiver Behavior

An RTP receiver MUST process RSI packets and adapt session parameters, such as the RTCP bandwidth, based on the information received. The receiver no longer has a global view of the session and will therefore be unable to receive information from individual receivers aside from itself. However, the information conveyed by the Distribution Source can be extremely detailed, providing the receiver with an accurate view of the session quality overall, without the processing overhead associated with listening to and analyzing all Receiver Reports.

The RTP receiver MUST process the report blocks contained in any RTP SR and RR packets to complete its view of the RTP session.

The SSRC collision list MUST be checked against the SSRC selected by the receiver to ensure there are no collisions as MUST be incoming RTP packets from the Media Senders. A receiver observing RTP packets from a Media Sender with an SSRC that collides with its own chosen SSRC MUST change its own SSRC following the procedures of [1]. The receiver MUST do so immediately after noticing and before sending any (further) RTCP feedback messages.

A Group and Average Packet Size sub-report block is most likely to be appended to the RSI header (either a Group Size sub-report or an RTCP Bandwidth sub-report MUST be included). The group size n allows a receiver to calculate its share of the RTCP bandwidth, r . Given R , the total available RTCP bandwidth share for receivers (in the SSM RTP session) $r = R/(n)$. For the group size calculation, the RTP receiver MUST NOT include the Distribution Source, i.e., the only RTP receiver sending RSI packets.

The receiver RTCP bandwidth field MAY override this value. If the receiver RTCP bandwidth field is present, the receiver MUST use this value as its own RTCP reporting bandwidth r .

If the RTCP bandwidth field was used by the Distribution Source in an RTCP session but this field was not included in the last five RTCP RSI reports, the receiver MUST revert to calculating its bandwidth share based upon the group size information.

If the receiver has not received any RTCP RSI packets from the Distribution Source for a period of five times the sender reporting interval, it MUST cease transmitting RTCP Receiver Reports until the next RTCP RSI packet is received.

The receiver can use the summarized data as desired. This data is most useful in providing the receiver with a more global view of the conditions experienced by other receivers and enables the client to place itself within the distribution and establish the extent to which its reported conditions correspond to the group reports as a whole. Appendix B provides further information and examples of data processing at the receiver.

The receiver SHOULD assume that any sub-report blocks in the same packet correspond to the same data set received by the Distribution Source during the last reporting time interval. This applies to packets with multiple blocks, where each block conveys a different range of values.

A receiver MUST NOT rely on all of the RTCP packets it sends reaching the Media Senders or any other receiver. While RR statistics will be aggregated, BYE packets will be processed, and SSRC collisions will usually be announced, processing and/or forwarding of further RTCP packets is up to the discretion of the Distribution Source and will be performed as specified in the session description.

If a receiver has out-of-band information available about the Media Sender SSRC(s) used in the media session, it MUST NOT use the same SSRC for itself. The receiver MUST be aware that such out-of-band information may be outdated (i.e., that the sender SSRC(s) may have changed) and MUST follow the above collision-resolution procedure if necessary.

A receiver MAY use such Media Sender SSRC information when available but MUST beware of potential changes to the SSRC (which can only be learned from Sender Report (SR) packets).

7.5. Media Sender Behavior

Media Senders listen on a unicast or multicast transport address for RTCP reports sent by the receivers (and forwarded by the Distribution Source) or other Media Senders (optionally forwarded by the Distribution Source).

Unlike in the case of the simple forwarding model, Media Senders MUST be able to process RSI packets from the Distribution Source to determine the group size and their own RTCP bandwidth share. Media Senders MUST also be capable of determining the group size (and their corresponding RTCP bandwidth share) from listening to (forwarded) RTCP RR and SR packets (as mandated in [1]).

As long as they send RTP packets, they MUST also send RTCP SRs, as defined in [1].

A Media Sender that observes an SSRC collision with another entity that is not also a Media Sender MAY delay its own collision-resolution actions, as per [1], by $5 * 1.5 * T_d$, with T_d being the deterministic calculated reporting interval, for receivers to see whether the conflict still exists. SSRC collisions with other Media Senders MUST be acted upon immediately.

Note: This gives precedence to Media Senders and places the burden of collision resolution on RTP receivers.

Sender SSRC information MAY be communicated out-of-band, e.g., by means of SDP media descriptions. Therefore, senders SHOULD NOT change their own SSRC eagerly or unnecessarily.

8. Mixer/Translator Issues

The original RTP specification allows a session to use mixers and translators to help connect heterogeneous networks into one session. There are a number of issues, however, which are raised by the unicast feedback model proposed in this document. The term 'mixer' refers to devices that provide data stream multiplexing where multiple sources are combined into one stream. Conversely, a translator does not multiplex streams but simply acts as a bridge between two distribution mechanisms, e.g., a unicast-to-multicast network translator. Since the issues raised by this document apply equally to either a mixer or translator, the latter are referred to from this point onwards as mixer-translator devices.

A mixer-translator between distribution networks in a session must ensure that all members in the session receive all the relevant traffic to enable the usual operation by the clients. A typical use may be to connect an older implementation of an RTP client with an SSM distribution network, where the client is not capable of unicasting feedback to the source. In this instance, the mixer-translator must join the session on behalf of the client and send and receive traffic from the session to the client. Certain hybrid scenarios may have different requirements.

8.1. Use of a Mixer-Translator

The mixer-translator MUST adhere to the SDP description [5] for the single-source session (Section 11) and use the feedback mechanism indicated. Implementers of receivers SHOULD be aware that when a mixer-translator is present in the session, more than one Media Sender may be active, since the mixer-translator may be forwarding traffic to the SSM receivers either from multiple unicast sources or from an ASM session. Receivers SHOULD still forward unicast RTCP reports in the usual manner to their assigned Feedback Target/Distribution Source, which in this case -- by assumption -- would be the mixer-translator itself. It is RECOMMENDED that the simple packet-reflection mechanism be used under these circumstances, since attempting to coordinate RSI + summarization reporting between more than one source may be complicated unless the mixer-translator is capable of summarization.

8.2. Encryption and Authentication Issues

Encryption and security issues are discussed in detail in Section 11. A mixer-translator MUST be able to follow the same security policy as the client in order to unicast RTCP feedback to the source, and it therefore MUST be able to apply the same authentication and/or encryption policy required for the session. Transparent bridging and

subsequent unicast feedback to the source, where the mixer-translator is not acting as the Distribution Source, is only allowed if the mixer-translator can conduct the same source authentication as required by the receivers. A translator MAY forward unicast packets on behalf of a client but SHOULD NOT translate between multicast-to-unicast flows towards the source without authenticating the source of the feedback address information.

9. Transmission Interval Calculation

The Control Traffic Bandwidth referred to in [1] is an arbitrary amount that is intended to be supplied by a session-management application (e.g., SDR [21]) or decided based upon the bandwidth of a single sender in a session.

The RTCP transmission interval calculation either remains the same as in the original RTP specification [1] or uses the algorithm in [10] when bandwidth modifiers have been specified for the session.

9.1. Receiver RTCP Transmission

If the Distribution Source is operating in Simple Feedback Model (which may be indicated in the corresponding session description for the media session but which the receiver also notices from the absence of RTCP RSI packets), a receiver MUST calculate the number of other members in a session based upon its own SSRC count, derived from the forwarded Sender and Receiver Reports it receives. The receiver MUST calculate the average RTCP packet size from all the RTCP packets it receives.

If the Distribution Source is operating in Distribution Source Feedback Summary Model, the receiver MUST use either the group size field and the average RTCP packet size field or the Receiver Bandwidth field from the respective sub-report blocks appended to the RSI packet.

A receiver uses these values as input to the calculation of the deterministic calculated interval as per [1] and [10].

9.2. Distribution Source RTCP Transmission

If operating in Simple Feedback Model, the Distribution Source MUST calculate the transmission interval for its Receiver Reports and associated RTCP packets, based upon the above control traffic bandwidth, and MUST count itself as RTP receiver. Receiver Reports will be forwarded as they arrive without further consideration. The Distribution Source MAY choose to validate that all or selected receivers roughly adhere to the calculated bandwidth constraints and

MAY choose to drop excess packets for receivers that do not. In all cases, the average RTCP packet size is determined from the forwarded Media Senders' and receivers' RTCP packets and from those originated by the Distribution Source.

If operating in Distribution Source Feedback Summary Model, the Distribution Source does not share the forward RTCP bandwidth with any of the receivers. Therefore, the Distribution Source SHOULD use the full RTCP bandwidth for its Receiver Reports and associated RTCP packets, as well as RTCP RSI packets. In this case, the average RTCP packet size is only determined from the RTCP packets originated by the Distribution Source.

The Distribution Source uses these values as input to the calculation of the deterministic calculated interval as per [1] and [10].

9.3. Media Senders RTCP Transmission

In Simple Feedback Model, the Media Senders obtain all RTCP SRs and RRs as they would in an ASM session, except that the packets are forwarded by the Distribution Source. They MUST perform their RTCP group size estimate and calculation of the deterministic transmission interval as per [1] and [10].

In Distribution Source Feedback Summary Model, the Media Senders obtain all RTCP SRs as they would in an ASM session. They receive either RTCP RR reports as in ASM (in case these packets are forwarded by the Distribution Source) or RSI packets containing summaries. In the former case, they MUST perform their RTCP group size estimate and calculation of the deterministic transmission interval as per [1] and [10]. In the latter case, they MUST combine the information obtained from the Sender Reports and the RSI packets to create a complete view of the group size and the average RTCP packet size and perform the calculation of the deterministic transmission interval, as per [1] and [10], based upon these input values.

9.4. Operation in Conjunction with AVPF and SAVPF

If the RTP session is an AVPF session [15] or an SAVPF session [28] (as opposed to a regular AVP [6] session), the receivers MAY modify their RTCP report scheduling, as defined in [15]. Use of AVPF or SAVPF does not affect the Distribution Source's RTCP transmission or forwarding behavior.

It is RECOMMENDED that a Distribution Source and possible separate Feedback Target(s) be configured to forward AVPF/SAVPF-specific RTCP packets in order to not counteract the damping mechanism built into AVPF; optionally, they MAY aggregate the feedback information from

the receivers as per Section 7.2.2. If only generic feedback packets that are understood by the Distribution Source and that can easily be aggregated are in use, the Distribution MAY combine several incoming RTCP feedback packets and forward the aggregate along with its next RTCP RR/RSI packet. In any case, the Distribution Source and Feedback Target(s) SHOULD minimize the extra delay when forwarding feedback information, but the Distribution Source MUST stay within its RTCP bandwidth constraints.

In the event that specific APP packets without a format and summarization mechanism understood by the Feedback Target(s) and/or the Distribution Source are to be used, it is RECOMMENDED that such packets are forwarded with minimal delay. Otherwise, the capability of the receiver to send timely feedback messages is likely to be affected.

10. SDP Extensions

The Session Description Protocol (SDP) [5] is used as a means to describe media sessions in terms of their transport addresses, codecs, and other attributes. Signaling that RTCP feedback will be provided via unicast, as specified in this document, requires another session parameter in the session description. Similarly, other SSM RTCP feedback parameters need to be provided, such as the summarization model at the sender and the target unicast address to which to send feedback information. This section defines the SDP parameters that are needed by the proposed mechanisms in this document (and that have been registered with IANA).

10.1. SSM RTCP Session Identification

A new session-level attribute MUST be used to indicate the use of unicast instead of multicast feedback: "rtcp-unicast".

This attribute uses one parameter to specify the model of operation. An optional set of parameters specifies the behavior for RTCP packet types (and subtypes).

rtcp-unicast:reflection

This attribute MUST be used to indicate the "Simple Feedback" model of operation where packet reflection is used by the Distribution Source (without further processing).

rtcp-unicast:rsi *(SP <processing>:<rtcp-type>)]

This attribute MUST be used to indicate the "Distribution Source Feedback Summary" model of operation. In this model, a list of parameters may be used to explicitly specify how RTCP packets originated by receivers are handled. Options for processing a given RTCP packet type are:

aggr: The Distribution Source has means for aggregating the contents of the RTCP packets and will do so.

forward: The Distribution Source will forward the RTCP packet unchanged.

term: The Distribution Source will terminate the RTCP packet.

The default rules applying if no parameters are specified are as follows:

RR and SDES packets MUST be aggregated and MUST lead to RSI packets being generated. All other RTP packets MUST be terminated at the Distribution Source (or Feedback Target(s)).

The SDP description needs only to specify deviations from the default rules. Aggregation of RR packets and forwarding of SR packets MUST NOT be changed.

The token for the new SDP attribute is "rtcp-unicast" and the formal SDP ABNF syntax for the new attribute value is as follows:

att-value = "reflection"
/ "rsi" *(SP rsi-rule)

rsi-rule = processing ":" rtcp-type

processing = "aggr" / "forward" / "term" / token ;keep it extensible

rtcp-type = 3*3DIGIT ;the RTCP type (192, 193, 202--209)

10.2. SSM Source Specification

In a Source-Specific Multicast RTCP session, the address of the Distribution Source needs to be indicated both for source-specific joins to the multicast group and for addressing unicast RTCP packets on the backchannel from receivers to the Distribution Source.

This is achieved by following the proposal for SDP source filters documented in [4]. According to the specification, only the inclusion model ("a=source-filter:incl") MUST be used for SSM RTCP.

There SHOULD be exactly one "a=source-filter:incl" attribute listing the address of the Distribution Source. The RTCP port MUST be derived from the m= line of the media description.

An alternative Feedback Target Address and port MAY be supplied using the SDP RTCP attribute [7], e.g., a=rtcp:<port> IN IP4 192.0.2.1. This attribute only defines the transport address of the Feedback Target and does not affect the SSM group specification for media stream reception.

Two "source-filter" attributes MAY be present to indicate an IPv4 and an IPv6 representation of the same Distribution Source.

10.3. RTP Source Identification

The SSRC information for the Media Sender(s) MAY be communicated explicitly out of band (i.e., outside the RTP session). One option for doing so is the Session Description Protocol (SDP) [5]. If such an indication is desired, the "ssrc" attribute [12] MUST be used for this purpose. As per [12], the "cname" Source Attribute MUST be present. Further Source Attributes MAY be specified as needed.

If used in an SDP session description of an RTCP-SSM session, the ssrc MUST contain the SSRC intended to be used by the respective Media Sender and the cname MUST equal the CNAME for the Media Sender. If present, the role SHOULD indicate the function of the RTP entity identified by this line; presently, only the "media-sender" role is defined.

Example:

```
a=ssrc:314159 cname:iptv-sender@example.com
```

In the above example, the Media Sender is identified to use the SSRC identifier 314159 and the CNAME iptv-sender@example.com.

11. Security Considerations

The level of security provided by the current RTP/RTCP model MUST NOT be diminished by the introduction of unicast feedback to the source. This section identifies the security weaknesses introduced by the feedback mechanism, potential threats, and level of protection that MUST be adopted. Any suggestions on increasing the level of security

provided to RTP sessions above the current standard are RECOMMENDED but OPTIONAL. The final section outlines some security frameworks that are suitable to conform to this specification.

11.1. Assumptions

RTP/RTCP is a protocol that carries real-time multimedia traffic, and therefore a main requirement is for any security framework to maintain as low overhead as possible. This includes the overhead of different applications and types of cryptographic operations as well as the overhead to deploy or to create security infrastructure for large groups.

Although the distribution of session parameters (typically encoded as SDP objects) through the Session Announcement Protocol (SAP) [22], email, or the web is beyond the scope of this document, it is RECOMMENDED that the distribution method employs adequate security measures to ensure the integrity and authenticity of the information. Suitable solutions that meet the security requirements outlined here are included at the end of this section.

In practice, the multicast and group distribution mechanism, e.g., the SSM routing tree, is not immune to source IP address spoofing or traffic snooping; however, such concerns are not discussed here. In all the following discussions, security weaknesses are addressed from the transport level or above.

11.2. Security Threats

Attacks on media distribution and the feedback architecture proposed in this document may take a variety of forms. A detailed outline of the types of attacks follows:

a) Denial of Service (DoS)

DoS is a major area of concern. Due to the nature of the communication architecture, a DoS can be generated in a number of ways by using the unicast feedback channel to the attacker's advantage.

b) Packet Forgery

Another potential area for attack is packet forgery. In particular, it is essential to protect the integrity of certain influential packets since compromise could directly affect the transmission characteristics of the whole group.

c) Session Replay

The potential for session recording and subsequent replay is an additional concern. An attacker may not actually need to understand packet content but simply have the ability to record the data stream and, at a later time, replay it to any receivers that are listening.

d) Eavesdropping on a Session

The consequences of an attacker eavesdropping on a session already constitutes a security weakness; in addition, eavesdropping might facilitate other types of attacks and is therefore considered a potential threat. For example, an attacker might be able to use the eavesdropped information to perform an intelligent DoS attack.

11.3. Architectural Contexts

To better understand the requirements of the solution, the threats outlined above are addressed for each of the three communication contexts:

a) Source-to-Receiver Communication

The downstream communication channel, from the source to the receivers, is critical to protect since it controls the behavior of the group; it conveys the bandwidth allocation for each receiver, and hence the rate at which the RTCP traffic is unicast, directly back to the source. All traffic that is distributed over the downstream channel is generated by a single source. Both the RTP data stream and the RTCP control data from the source are included in this context, with the RTCP data generated by the source being indirectly influenced by the group feedback.

The downstream channel is vulnerable to the four types of attack outlined above. The denial of service attack is possible but less of a concern than the other types. The worst case effect of DoS would be the transmission of large volumes of traffic over the distribution channel, with the potential to reach every receiver but only on a one-to-one basis. Consequently, this threat is no more pronounced than the current multicast ASM model. The real danger of denial of service attacks in this context comes indirectly via compromise of the source RTCP traffic. If receivers are provided with an incorrect group size estimate or bandwidth allowance, the return traffic to the source may create a distributed DoS effect on the source. Similarly, an incorrect feedback address -- whether as a result of a malicious attack or

by mistake, e.g., an IP address configuration error (e.g., typing) -- could directly create a denial of service attack on another host.

An additional concern relating to Denial of service attacks would come indirectly through the generation of fake BYE packets, causing the source to adjust the advertised group size. A Distribution Source MUST follow the correct rules for timing out members in a session prior to reporting a change in the group size, which allows the authentic SSRC sufficient time to continue to report and, consequently, cancel the fake BYE report.

The danger of packet forgery in the worst case may be to maliciously instigate a denial of service attack, e.g., if an attacker were capable of spoofing the source address and injecting incorrect packets into the data stream or intercepting the source RTCP traffic and modifying the fields.

The replay of a session would have the effect of recreating the receiver feedback to the source address at a time when the source is not expecting additional traffic from a potentially large group. The consequence of this type of attack may be less effective on its own, but in combination with other attacks might be serious.

Eavesdropping on the session would provide an attacker with information on the characteristics of the source-to-receiver traffic, such as the frequency of RTCP traffic. If RTCP traffic is unencrypted, this might also provide valuable information on characteristics such as group size, Media Source SSRC(s), and transmission characteristics of the receivers back to the source.

b) Receiver-to-Distribution-Source Communication

The second context is the return traffic from the group to the Distribution Source. This traffic should only consist of RTCP packets and should include Receiver Reports, SDP information, BYE reports, extended reports (XR), feedback messages (RTPFB, PSFB) and possibly application-specific packets. The effects of compromise on a single or subset of receivers are not likely to have as great an impact as in context (a); however, much of the responsibility for detecting compromise of the source data stream relies on the receivers.

The effects of compromise of critical Distribution Source control information can be seriously amplified in the present context. A large group of receivers may unwittingly generate a distributed

DoS attack on the Distribution Source in the event that the integrity of the source RTCP channel has been compromised and the compromise is not detected by the individual receivers.

An attacker capable of instigating a packet forgery attack could introduce false RTCP traffic and create fake SSRC identifiers. Such attacks might slow down the overall control channel data rate since an incorrect perception of the group size may be created. Similarly, the creation of fake BYE reports by an attacker would cause some group size instability, but should not be effective as long as the correct timeout rules are applied by the source in removing SSRC entries from its database.

A replay attack on receiver return data to the source would have the same implications as the generation of false SSRC identifiers and RTCP traffic to the source. Therefore, ensuring authenticity and freshness of the data source is important.

Eavesdropping in this context potentially provides an attacker with a great deal of potentially personal information about a large group of receivers available from SDES packets. It would also provide an attacker with information on group traffic-generation characteristics and parameters for calculating individual receiver bandwidth allowances.

c) Receiver-to-Feedback-Target Communication

The third context is the return traffic from the group to the Feedback Target. It suffers from the same threats as the receiver-to-source context, with the difference being that now a large group of receivers may unwittingly generate a distributed DoS (DDoS) attack on the Feedback Target, where it is impossible to discern if the DDoS is deliberate or due merely to a misconfiguration of the Feedback Target Address. While deliberate attacks can be mitigated by properly authenticating messages that communicate the Feedback Target Address (i.e., the SDP session description and the Feedback Target Address sub-report block carried in RTCP), a misconfigured address will originate from an authenticated source and hence cannot be prevented using security mechanisms.

Furthermore, the Feedback Target is unable to communicate its predicament with either the Distribution Source or the session receivers. From the feedback packets received, the Feedback Target cannot tell either which SSM multicast group the feedback belongs to or the Distribution Source, making further analysis and suppression difficult. The Feedback Target may not even support RTCP or listen on the port number in question.

Note that because the DDoS occurs inside of the RTCP session and because the unicast receivers adhere to transmission interval calculations ([1], [10]), the bandwidth misdirected toward the Feedback Target in the misconfigured case will be limited to a percentage of the session bandwidth, i.e., the Control Traffic Bandwidth established for the session.

11.4. Requirements in Each Context

To address these threats, this section presents the security requirements for each context.

- a) The main threat in the source-to-receiver context concerns denial of service attacks through possible packet forgery. The forgery may take the form of interception and modification of packets from the source, or it may simply inject false packets into the distribution channel. To combat these attacks, data integrity and source authenticity **MUST** be applied to source traffic. Since the consequences of eavesdropping do not affect the operation of the protocol, confidentiality is not a requirement in this context. However, without confidentiality, access to personal and group characteristics information would be unrestricted to an external listener. Therefore, confidentiality is **RECOMMENDED**.
- b) The threats in the receiver-to-source context concern the same kinds of attacks, but are considered less important than the downstream traffic compromise. All the security weaknesses are also applicable to the current RTP/RTCP security model, and therefore only recommendations towards protection from compromise are made. Data integrity is **RECOMMENDED** to ensure that interception and modification of an individual receiver's RTCP traffic can be detected. This would protect against the false influence of group control information and the potentially more serious compromise of future services provided by the distribution functionality. In order to ensure security, data integrity and authenticity of receiver traffic is therefore also **RECOMMENDED**. With respect to data confidentiality, the same situation applies as in the first context, and it is **RECOMMENDED** that precautions be taken to protect the privacy of the data.
- c) The threats to the receiver-to-feedback-target context are similar to those in the receiver-to-source context, and thus the recommendations to protect against them are similar.

However, there are a couple situations with broader issues to solve, which are beyond the scope of this document.

1. An endpoint experiencing DDoS or the side effects of a misconfigured RTCP session may not even be a participant in the session, i.e., may not be listening on the respective port number and may even support RTCP, so it will be unable to react within RTCP. Determining that there is a problem will be up to network administrators and, possibly, anti-malware software that can perform correlation across receiver nodes.
2. With misconfiguration, unfortunately the normally desirable usage of SRTP and SRTCP becomes undesirable. Because the packet content is encrypted, neither the misconfigured Feedback Target nor the network administrator have the ability to determine the root cause of the traffic.

In the case where the misconfigured Feedback Target happens to be a node participating in the session or is an RTCP-enabled node, the Feedback Target Address block provides a dynamic mechanism for the Distribution Source to signal an alternative unicast RTCP feedback address to the receivers. As this type of packet MUST be included in every RTCP packet originated by the Distribution Source, all receivers would be able to obtain the corrected Feedback Target information. In this manner, receiver behavior should remain consistent even in the face of packet loss or when there are late-session arrivals. The only caveat is that the misconfigured Feedback Target is largely uninvolved in the repair of this situation and thus relies on others for the detection of the problem.

An additional security consideration, which is not a component of this specification but which has a direct influence upon the general security, is the origin of the session-initiation data. This involves the SDP parameters that are communicated to the members prior to the start of the session via channels such as an HTTP server, email, SAP, or other means. It is beyond the scope of this document to place any strict requirements on the external communication of such information; however, suitably secure SDP communication approaches are outlined in Section 11.7.

11.5. Discussion of Trust Models

As identified in the previous sections, source authenticity is a fundamental requirement of the protocol. However, it is important to also clarify the model of trust that would be acceptable to achieve this requirement. There are two fundamental models that apply in this instance:

- a) The shared-key model, where all authorized group members share the same key and can equally encrypt/decrypt the data. This method assumes that an out-of-band method is applied to the distribution of the shared group key, ensuring that every key-holder is individually authorized to receive the key and, in the event of member departures from the group, a re-keying exercise can occur. The advantage of this model is that the costly processing associated with one-way key-authentication techniques is avoided, as well as the need to execute additional cipher operations with alternative key sets on the same data set, e.g., in the event that data confidentiality is also applied. The disadvantage is that, for very large groups where the receiver set becomes effectively untrusted, a shared key does not offer much protection.
- b) The public-key authentication model, using cryptosystems such as RSA-based or PGP authentication, provides a more secure method of source authentication at the expense of generating higher processing overhead. This is typically not recommended for real-time data streams but, in the case of RTCP reports, which are distributed with a minimum interval of 5 seconds, this may be a viable option (the processing overhead might still be too great for small, low-powered devices and should therefore be considered with caution). Wherever possible, however, the use of public key source authentication is preferable to the shared key model identified above.

As concerns requirements for protocol acceptability, either model is acceptable although it is RECOMMENDED that the more secure public-key-based options be applied wherever possible.

11.6. Recommended Security Solutions

This section presents some existing security mechanisms that are RECOMMENDED to suitably address the requirements outlined in Section 11.5. This is only intended as a guideline and it is acknowledged that there are other solutions that would also be suitable to comply with the specification.

11.6.1. Secure Distribution of SDP Parameters

- a) SAP, HTTPS, Email -- Initial distribution of the SDP parameters for the session SHOULD use a secure mechanism such as the SAP authentication framework, which allows an authentication certificate to be attached to the session announcements. Other methods might involve HTTPS or signed email content from a trusted source. However, some more commonly used techniques for distributing session information and starting media streams are the Real-Time Streaming Protocol (RTSP) [25] and SIP [14].

- b) RTSP -- RTSP provides a client- or server-initiated stream control mechanism for real-time multimedia streams. The session parameters are conveyed using SDP syntax and may adopt standard HTTP authentication mechanisms in combination with suitable network (e.g., IPsec)- or transport (e.g., Transport Layer Security (TLS))-level security.
- c) SIP -- A typical use of SIP involving a unicast feedback identifier might be a client wishing to dynamically join a multi-party call on a multicast address using unicast RTCP feedback. The client would be required to authenticate the SDP session descriptor information returned by the SIP server. The recommended method for this, as outlined in the SIP specification [14], is to use an S/MIME message body containing the session parameters signed with an acceptable certificate.

For the purposes of this profile, it is acceptable to use any suitably secure authentication mechanism that establishes the identity and integrity of the information provided to the client.

11.6.2. Suitable Security Solutions for RTP Sessions with Unicast Feedback

- a) SRTP -- SRTP [3] is the recommended Audio/Video Transport (AVT) security framework for RTP sessions. It specifies the general packet formats and cipher operations that are used and provides the flexibility to select different stream ciphers based on preference/requirements. It can provide confidentiality of the RTP and RTCP packets as well as protection against integrity compromise and replay attacks. It provides authentication of the data stream using the shared-key trust model. Any suitable key-distribution mechanism can be used in parallel to the SRTP streams.
- b) IPSEC -- A more general group security profile that might be used is the Group Domain of Interpretation [23], which defines the process of applying IPsec mechanisms to multicast groups. This requires the use of the Encapsulating Security Payload (ESP) in tunnel mode as the framework and it provides the capability to authenticate -- either using a shared key or individually through public-key mechanisms. It should be noted that using IPsec would break the 'transport-independent' condition of RTP and would therefore not be useable for anything other than IP-based communication.
- c) TESLA - Timed Efficient Stream Loss-Tolerant Authentication (TESLA) [24] is a scheme that provides a more flexible approach to data authentication using time-based key disclosure. The

authentication uses one-way, pseudo-random key functions based on key chain hashes that have a short period of authenticity based on the key disclosure intervals from the source. As long as the receiver can ensure that the encrypted packet is received prior to the key disclosure by the source, which requires loose time synchronization between source and receivers, it can prove the authenticity of the packet. The scheme does introduce a delay into the packet distribution/decryption phase due to the key disclosure delay; however, the processing overhead is much lower than other standard public-key mechanisms and therefore may be more suited to small or energy-restricted devices.

11.6.3. Secure Key Distribution Mechanisms

- a) MIKEY -- Multimedia Internet KEYing (MIKEY) [29] is the preferred solution for SRTP sessions providing a shared group-key distribution mechanism and intra-session rekeying facilities. If a partly protected communication channel exists, keys may also be conveyed using SDP as per [27].
- b) GSAKMP -- The Group Secure Association Key Management Protocol (GSAKMP) is the general solution favored for Multicast Secure group-key distribution. It is the recommended key distribution solution for Group Domain of Interpretation (GDOI) [RFC3547] sessions.

11.7. Troubleshooting Misconfiguration

As noted above, the security mechanisms in place will not help in case an authorized source spreads properly authenticated and integrity-protected yet incorrect information about the Feedback Target. In this case, the accidentally communicated Feedback Target will receive RTCP packets from a potentially large group of receivers -- the RTCP rate fortunately limited by the RTCP timing rules.

Yet, the RTCP packets do not provide much context information and, if encrypted, do not provide any context, making it difficult for the entity running (the network with) the Feedback Target to debug and correct this problem, e.g., by tracking down and informing the origin of the misconfiguration.

One suitable approach may be to provide explicit context information in RTCP packets that would allow determining the source. While such an RTCP packet could be defined in this specification, it would be of no use when using SRTP/SRTCP and encryption of RTCP reports. Therefore, and because the extensions in this document may not be the

only case that may face such a problem, it is desirable to find a solution that is applicable to RTP at large. Such mechanisms are for further study in the AVT WG.

12. Backwards Compatibility

The use of unicast feedback to the source should not present any serious backwards compatibility issues. The RTP data streams should remain unaffected, as should the RTCP packets from the Media Sender(s) that continue to enable inter-stream synchronization in the case of multiple streams. The unicast transmission of RTCP data to a source that does not have the ability to redistribute the traffic either by simple reflection or through summaries could have serious security implications, as outlined in Section 11, but would not actually break the operation of RTP. For RTP-compliant receivers that do not understand the unicast mechanisms, the RTCP traffic may still reach the group in the event that an ASM distribution network is used, in which case there may be some duplication of traffic due to the reflection channel, but this should be ignored. It is anticipated, however, that typically the distribution network will not enable the receiver to multicast RTCP traffic, in which case the data will be lost and the RTCP calculations will not include those receivers. It is RECOMMENDED that any session that may involve non-unicast-capable clients should always use the simple packet-reflection mechanism to ensure that the packets received can be understood by all clients.

13. IANA Considerations

The following contact information shall be used for all registrations included here:

Contact: Joerg Ott
mail: jo@acm.org
tel: +358-9-470-22460

Based on the guidelines suggested in [2], a new RTCP packet format has been registered with the RTCP Control Packet type (PT) Registry:

Name: RSI
Long name: Receiver Summary Information
Value: 209
Reference: This document.

This document defines a substructure for RTCP RSI packets. A new sub-registry has been set up for the sub-report block type (SRBT) values for the RSI packet, with the following registrations created initially:

Name: IPv4 Address
Long name: IPv4 Feedback Target Address
Value: 0
Reference: This document.

Name: IPv6 Address
Long name: IPv6 Feedback Target Address
Value: 1
Reference: This document.

Name: DNS Name
Long name: DNS Name indicating Feedback Target Address
Value: 2
Reference: This document.

Name: Loss
Long name: Loss distribution
Value: 4
Reference: This document.

Name: Jitter
Long name: Jitter Distribution
Value: 5
Reference: This document.

Name: RTT
Long name: Round-trip time distribution
Value: 6
Reference: This document.

Name: Cumulative loss
Long name: Cumulative loss distribution
Value: 7
Reference: This document.

Name: Collisions
Long name: SSRC Collision list
Value: 8
Reference: This document.

Name: Stats
Long name: General statistics
Value: 10
Reference: This document.

Name: RTCP BW
Long name: RTCP Bandwidth indication
Value: 11
Reference: This document.

Name: Group Info
Long name: RTCP Group and Average Packet size
Value: 12
Reference: This document.

The value 3 shall be reserved as a further way of specifying a Feedback Target Address. The value 3 MUST only be allocated for a use defined in an IETF Standards Track document.

Further values may be registered on a first come, first served basis. For each new registration, it is mandatory that a permanent, stable, and publicly accessible document exists that specifies the semantics of the registered parameter as well as the syntax and semantics of the associated sub-report block. The general registration procedures of [5] apply.

In the registry for SDP parameters, the attribute named "rtcp-unicast" has been registered as follows:

SDP Attribute ("att-field"):

Attribute Name: rtcp-unicast
Long form: RTCP Unicast feedback address
Type of name: att-field
Type of attribute: Media level only
Subject to charset: No
Purpose: RFC 5760
Reference: RFC 5760
Values: See this document.

14. References

14.1. Normative References

- [1] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [2] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

- [3] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [4] Quinn, B. and R. Finlayson, "Session Description Protocol (SDP) Source Filters", RFC 4570, July 2006.
- [5] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [6] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.
- [7] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.
- [8] Meyer, D., Rockell, R., and G. Shepherd, "Source-Specific Protocol Independent Multicast in 232/8", BCP 120, RFC 4608, August 2006.
- [9] 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.
- [10] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, July 2003.
- [11] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [12] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, June 2009.
- [13] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

14.2. Informative References


- [14] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [15] 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.

- [16] Pusateri, T., "Distance Vector Multicast Routing Protocol", Work in Progress, October 2003.
- [17] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [18] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.
- [19] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [20] Fenner, B., Ed., and D. Meyer, Ed., "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.
- [21] Session Directory Rendez-vous (SDR), developed at University College London by Mark Handley and the Multimedia Research Group, <http://www-mice.cs.ucl.ac.uk/multimedia/software/sdr/>.
- [22] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, October 2000.
- [23] Baugher, M., Weis, B., Hardjono, T., and H. Harney, "The Group Domain of Interpretation", RFC 3547, July 2003.
- [24] Perrig, A., Song, D., Canetti, R., Tygar, J., and B. Briscoe, "Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction", RFC 4082, June 2005.
- [25] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, April 1998.
- [26] Friedman, T., Ed., Caceres, R., Ed., and A. Clark, Ed., "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.
- [27] Andreassen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.
- [28] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, February 2008.

- [29] Arkko, J., Carrara, E., Lindholm, F., Naslund, M., and K. Norrman, "MIKEY: Multimedia Internet KEYing", RFC 3830, August 2004.

This appendix describes a few common setups, focusing on the contribution side, i.e., the communications between the Media Sender(s) and the Distribution Source. In all cases, the same session description may be used for the distribution side as defined in the main part of this document. This is because this specification defines only the media stream setup between the Distribution Source and the receivers.

In the simplest case, the Distribution Source is identical to the Media Sender as depicted in Figure 3. Obviously, no further configuration for the interaction between the Media Sender and the Distribution Source is necessary.



In a slightly more complex scenario, the Media Sender and the Distribution Source are separate entities running on the same or different machines. Hence, the Media Sender needs to deliver the media stream(s) to the Distribution Source. This can be done either via unicasting the RTP stream, via ASM multicast, or via SSM. In this case, the Distribution Source is responsible for forwarding the RTP packets comprising the media stream and the RTCP Sender Reports towards the receivers and conveying feedback from the receivers, as well as from itself, to the Media Sender.

This scenario is depicted in Figure 4. The communication setup between the Media Sender and the Distribution Source may be statically configured or SDP may be used in conjunction with some signaling protocol to convey the session parameters. Note that it is a local configuration matter of the Distribution Source how to associate a session between the Media Sender and itself (the Contribution session) with the corresponding session between itself and the receivers (the Distribution session).

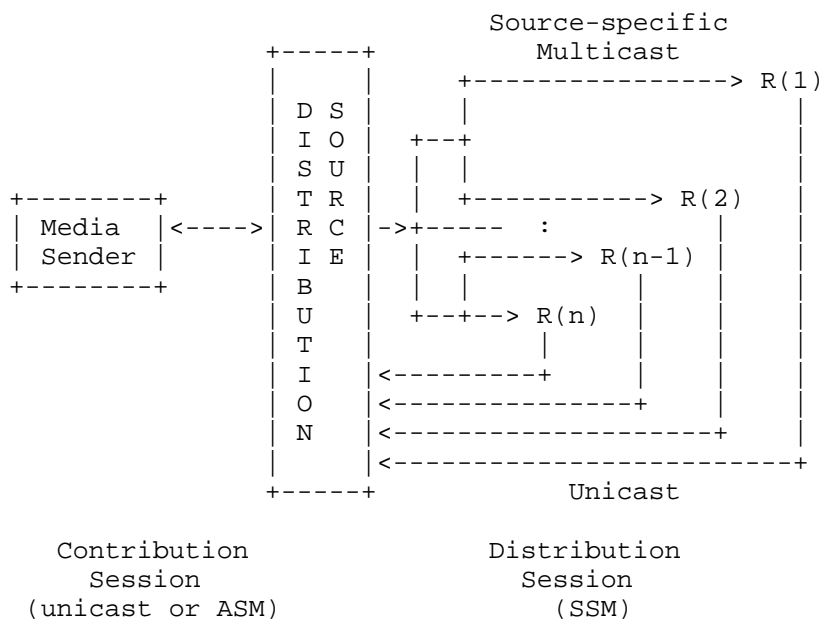


Figure 4: One Media Sender Separate from Distribution Source

A.3. Three Media Senders, Unicast Contribution

Similar considerations apply if three Media Senders transmit to an SSM multicast group via the Distribution Source and individually send their media stream RTP packets via unicast to the Distribution Source.

In this case, the responsibilities of the Distribution Source are a superset to the previous case; the Distribution Source also needs to relay media traffic from each Media Sender to the receivers and to forward (aggregated) feedback from the receivers to the Media Senders. In addition, the Distribution Source relays RTCP packets (SRs) from each Media Sender to the other two.

The configuration of the Media Senders is identical to the previous case. It is just the Distribution Source that must be aware that there are multiple senders and then perform the necessary relaying. The transport address for the RTP session at the Distribution Source may be identical for all Media Senders (which may make correlation easier) or different addresses may be used.

This setup is depicted in Figure 5.

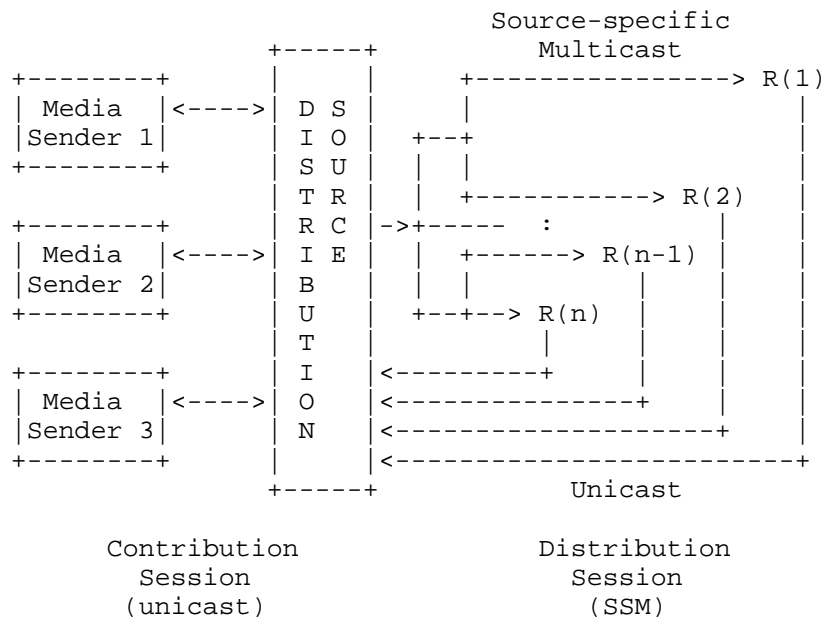



Figure 5: Three Media Senders, Unicast Contribution

A.4. Three Media Senders, ASM Contribution Group

In this final example, the individual unicast contribution sessions between the Media Senders and the Distribution Source are replaced by a single ASM contribution group (i.e., a single common RTP session). Consequently, all Media Senders receive each other's traffic by means of IP-layer multicast and the Distribution Source no longer needs to perform explicit forwarding between the Media Senders. Of course, the Distribution Source still forwards feedback information received from the receivers (optionally as summaries) to the ASM contribution group.

Figure 6 shows this scenario.



Appendix B. Distribution Report Processing at the Receiver

B.1. Algorithm

Example processing of Loss Distribution Values

X values represent the loss percentage.

Y values represent the number of receivers.

Number of x values is the NDB value

```
xrange = Max Distribution Value(MaDV) - Min Distribution Value(MnDV)
```

```
First data point = MnDV,first ydata
```

```
then
```

```
For each ydata => xdata += (MnDV + (xrange / NDB))
```

B.2. Pseudo-Code

```
Packet Variables -> factor,NDB,MnDVL,MaDV
```

```
Code variables -> xrange, ydata[NDB],x,y
```

```
xrange = MaDV - MnDV
```

```
x = MnDV;
```

```
for(i=0; i<NDB; i++) {  
    y = (ydata[i] * factor);  
    /*OUTPUT x and y values*/  
    x += (xrange / NDB);  
}
```

B.3. Application Uses and Scenarios

Providing a distribution function in a feedback message has a number of uses for different types of applications. Although this appendix enumerates potential uses for the distribution scheme, it is anticipated that future applications might benefit from it in ways not addressed in this document. Due to the flexible nature of the summarization format, future extensions may easily be added. Some of the scenarios addressed in this section envisage potential uses beyond a simple SSM architecture, for example, single-source group topologies where every receiver may in fact also be capable of becoming the source. Another example may be multiple SSM topologies, which, when combined, make up a larger distribution tree.

A distribution of values is useful as input into any algorithm, multicast or otherwise, that could be optimized or tuned as a result of having access to the feedback values for all group members. Following is a list of example areas that might benefit from distribution information:

- The parameterization of a multicast Forward Error Correction (FEC) algorithm. Given an accurate estimate of the distribution of reported losses, a source or other distribution agent that does not have a global view would be able to tune the degree of redundancy built into the FEC stream. The distribution might help to identify whether the majority of the group is experiencing high levels of loss, or whether in fact the high loss reports are only from a small subset of the group. Similarly, this data might enable a

receiver to make a more informed decision about whether it should leave a group that includes a very high percentage of the worst-case reporters.

- The organization of a multicast data stream into useful layers for layered coding schemes. The distribution of packet losses and delay would help to identify what percentage of members experience various loss and delay levels, and thus how the data stream bandwidth might be partitioned to suit the group conditions. This would require the same algorithm to be used by both senders and receivers in order to derive the same results.
- The establishment of a suitable feedback threshold. An application might be interested to generate feedback values when above (or below) a particular threshold. However, determining an appropriate threshold may be difficult when the range and distribution of feedback values is not known a priori. In a very large group, knowing the distribution of feedback values would allow a reasonable threshold value to be established, and in turn would have the potential to prevent message implosion if many group members share the same feedback value. A typical application might include a sensor network that gauges temperature or some other natural phenomenon. Another example is a network of mobile devices interested in tracking signal power to assist with hand-off to a different distribution device when power becomes too low.
- The tuning of Suppression algorithms. Having access to the distribution of round-trip times, bandwidth, and network loss would allow optimization of wake-up timers and proper adjustment of the Suppression interval bounds. In addition, biased wake-up functions could be created not only to favor the early response from more capable group members but also to smooth out responses from subsequent respondents and to avoid bursty response traffic.
- Leader election among a group of processes based on the maximum or minimum of some attribute value. Knowledge of the distribution of values would allow a group of processes to select a leader process or processes to act on behalf of the group. Leader election can promote scalability when group sizes become extremely large.

B.4. Distribution Sub-Report Creation at the Source

The following example demonstrates two different ways to convey loss data using the generic format of a Loss sub-report block (Section 7.1.4). The same techniques could also be applied to representing other distribution types.

- 1) The first method attempts to represent the data in as few bytes as possible.
- 2) The second method conveys all values without providing any savings in bandwidth.

Data Set

X values indicate loss percentage reported; Y values indicate the number of receivers reporting that loss percentage.

X - 0	1	2	3	4	5	6	7	8	9
Y - 1000	800	6	1800	2600	3120	2300	1100	200	103

X - 10	11	12	13	14	15	16	17	18	19
Y - 74	21	30	65	60	80	6	7	4	5

X - 20	21	22	23	24	25	26	27	28	29
Y - 2	10	870	2300	1162	270	234	211	196	205

X - 30	31	32	33	34	35	36	37	38	39
Y - 163	174	103	94	76	52	68	79	42	4

Constant value

Due to the size of the multiplicative factor field being 4 bits, the maximum multiplicative value is 15.

The distribution type field of this packet would be value 1 since it represents loss data.

Example: 1st Method

Description

The minimal method of conveying data, i.e., small amount of bytes used to convey the values.

Algorithm

Attempt to fit the data set into a small sub-report size, selected length of 8 octets

Can we split the range (0 - 39) into 16 4-bit values? The largest bucket value would, in this case, be the bucket for X values 5 - 7.5, the sum of which is 5970. An MF value of 9 will generate a multiplicative factor of 2^9 , or 512 -- which, multiplied by the max bucket value, produces a maximum real value of 7680.

The packet fields will contain the values:

Header distribution Block

Distribution Type:	1
Number of Data Buckets:	16
Multiplicative Factor:	9
Packet Length field:	5 (5 * 4 => 20 bytes)
Minimum Data Value:	0
Maximum Data Value:	39
Data Bucket values:	(each value is 16-bits)

Results, 4-bit buckets:

X - 0 - 2.5		2.5 - 5		5 - 7.5		7.5 - 10	
(Y - 1803		4403		5970		853) ACTUAL
Y - 4		9		12		2	
X - 10 - 12.5		12.5 - 15		15 - 17.5		17.5 - 20	
(Y - 110		140		89.5		12.5)	ACTUAL
Y - 0		0		0		0	
X - 20 - 22.5		22.5 - 25		25 - 27.5		27.5 - 30	
(Y - 447		3897		609.5		506.5)	ACTUAL
Y - 1		8		1		1	
X - 30 - 32.5		32.5 - 35		35 - 37.5		37.5 - 40	
(Y - 388.5		221.5		159.5		85.5)	ACTUAL
Y - 1		0		0		0	

Example: 2nd Method

Description

This demonstrates the most accurate method for representing the data set. This method doesn't attempt to optimise any values.

Algorithm

Identify the highest value and select buckets large enough to convey the exact values, i.e., no multiplicative factor.

The highest value is 3120. This requires 12 bits (closest 2 bit boundary) to represent, therefore it will use 60 bytes to represent the entire distribution. This is within the max packet size; therefore, all data will fit within one sub-report block. The multiplicative value will be 1.

The packet fields will contain the values:

Header Distribution Block	
Distribution Type:	1
Number of Data Buckets:	40
Multiplicative Factor:	0
Packet Length field:	18 (18 * 4 => 72 bytes)
Minimum Loss Value:	0
Maximum Loss Value:	39

Bucket values are the same as the initial data set.

Result

Selecting one of the three methods outlined above might be done by a congestion parameter or by user preference. The overhead associated with processing the packets is likely to differ very little between the techniques. The savings in bandwidth are apparent, however, using 20, 52, and 72 octets respectively. These values would vary more widely for a larger data set with less correlation between results.

Acknowledgements

The authors would like to thank Magnus Westerlund, Dave Oran, Tom Taylor, and Colin Perkins for detailed reviews and valuable comments.

Authors' Addresses

Joerg Ott
Aalto University
School of Science and Technology
Department of Communications and Networking
PO Box 13000
FIN-00076 Aalto

EMail: jo@acm.org

Julian Chesterfield
University of Cambridge
Computer Laboratory,
15 JJ Thompson Avenue
Cambridge
CB3 0FD
UK

EMail: julianchesterfield@cantab.net

Eve Schooler
Intel Research / CTL
MS RNB6-61
2200 Mission College Blvd.
Santa Clara, CA, USA 95054-1537

EMail: eve_schooler@acm.org

