

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: August 13, 2015

M. Montagud, Ed.  
F. Boronat  
UPV  
H. Stokking  
TNO  
February 9, 2015

**Early Event-Driven (EED) RTCP Feedback for Rapid IDMS**  
**draft-montagud-avtcore-eed-rtcp-idms-00**

Abstract

[RFC 7272](#) defines two RTCP messages to enable Inter-Destination Media Synchronization (IDMS). However, if such RTCP messages are exchanged using the Regular RTCP reporting rules specified in [RFC 3550](#), unacceptable delays can be originated when exchanging the synchronization information conveyed into RTCP packets. Accordingly, this document proposes Early Event-Driven (EED) RTCP reporting rules and messages to enable higher interactivity, dynamism, flexibility and accuracy when using RTP/RTCP for IDMS. Such IDMS extensions are targeted at providing faster reaction on dynamic situations, such as out-of-sync situations and channel change delays, as well as a finer granularity for synchronizing media-related events, while still adhering to the RTCP bandwidth bounds specified in [RFC 3550](#). Besides, a new RTP/AVPF transport layer feedback message is proposed to allow the request of re-IDMS setting instructions (e.g., in case of RTCP packet loss) as well as rapid accommodation of latecomers in on-going sessions. Finally, this document also discusses various situations in which the reporting of Reduced-Size RTCP packets ([RFC 3556](#)) can be applicable and beneficial for IDMS.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 13, 2015.

Copyright Notice

Copyright (c) 2015 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 not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

Table of Contents

- [1.](#) Introduction . . . . . [3](#)
- [1.1.](#) RTP/RTCP for IDMS ([RFC 7272](#)) . . . . . [4](#)
- [2.](#) Terminology . . . . . [5](#)
- [3.](#) Background: RTCP Reporting Rules . . . . . [7](#)
- [3.1.](#) Regular RTCP Feedback ([RFC 3550](#)) . . . . . [7](#)
- [3.2.](#) Early RTCP Feedback ([RFC 4585](#)) . . . . . [9](#)
- [3.3.](#) Rapid Synchronization of RTP Flows ([RFC 6051](#)) . . . . . [10](#)
- [3.4.](#) SDP Modifiers for RTCP Bandwidth ([RFC 3556](#)) . . . . . [11](#)
- [4.](#) Early Event-Driven (EED) RTCP Feedback for IDMS . . . . . [12](#)
- [4.1.](#) Immediate Initial RTCP IDMS Settings . . . . . [12](#)
- [4.2.](#) Dynamic EED Reporting of IDMS Settings . . . . . [13](#)
- [4.3.](#) Rapid (Re-)Synchronization Request . . . . . [16](#)
- [4.4.](#) Reduction of Channel Change Delays . . . . . [18](#)
- [5.](#) Reduced-Size RTCP Reporting for IDMS . . . . . [20](#)
- [6.](#) Security Considerations . . . . . [23](#)
- [7.](#) IANA Considerations . . . . . [23](#)
- [8.](#) Contributors . . . . . [23](#)
- [9.](#) Conclusions . . . . . [23](#)
- [10.](#) References . . . . . [23](#)
- [10.1.](#) Normative References . . . . . [23](#)
- [10.2.](#) Informative References . . . . . [24](#)
- Authors' Addresses . . . . . [24](#)



## **1. Introduction**

RTP (Real-Time Transport Protocol) and RTCP (RTP Control Protocol) protocols provide support for both intra-media and inter-media synchronization [[RFC3550](#)]. Moreover, [[RFC7272](#)] extends the RTP/RTCP capabilities to also enable Inter-Destination Media Synchronization (IDMS).

However, if the proposed RTCP messages for IDMS in [[RFC7272](#)] are exchanged using the Regular RTCP reporting rules specified in [[RFC3550](#)], this can lead to unacceptable performance in some delay-sensitive applications requiring stringent synchronization requirements [[Montagud2012](#)]. This is because the RTCP packets are exchanged in a pre-scheduled and inflexible manner, uniquely based on preserving the allowed traffic bounds specified in [[RFC3550](#)]. There is no support for timely feedback that would allow to repair or to manage dynamic events of interest close to their occurrence. Accordingly, there may be a variable time lag (from few seconds to several minutes in large-scale environments) between detecting an event (e.g., an out-of-sync situation) and being able to send an appropriate RTCP packet to handle it. Moreover, the RTCP reports may even not be received at the target side, since RTCP is sent over UDP, which does not provide a reliable control channel.

This document proposes further extensions to RTP/RTCP to allow for a more strategic and efficient usage of the RTCP channel for IDMS. In particular, the RTCP extensions from [[RFC4585](#)] and [[RFC6051](#)] encourage the specification of novel RTCP reporting rules and messages, which in conjunction we call Early Event-Driven (EED) RTCP Feedback for IDMS, to enable higher interactivity, dynamism, flexibility and accuracy when using RTP/RTCP for IDMS. Such IDMS extensions are backward compatible with the solution specified in [[RFC7272](#)], while still adhering to the RTCP traffic bounds specified in [[RFC3550](#)].

More specifically, the proposed EED RTCP Feedback for IDMS provides the following key benefits: i) earlier correction of out-of-sync situations; ii) higher granularity for synchronizing the presentation of dynamically triggered media-related events (e.g., to ensure that important pieces of media content are simultaneously watched by all the users); iii) ability of dynamically requesting IDMS setting instructions (e.g., in case of RTCP packet loss); iv) dynamic and rapid accommodation of latecomers in on-going sessions; and v) reduction of channel-change (i.e., zapping) delays. Additionally, this document also discusses various situations in which the reporting of Reduced-Size RTCP packets [[RFC3556](#)] can be applicable and beneficial for IDMS.



The proposed RTCP extensions for IDMS in this document are applicable to and can have a potentially high impact on a wide spectrum of scenarios with demanding IDMS characteristics [[Montagud2012](#)], such as Social TV, multi-party conferencing, and synchronous e-learning.

A preliminary evaluation of such proposals can be found in [[Montagud2013](#)].

### **1.1. RTP/RTCP for IDMS ([RFC 7272](#))**

IDMS refers to the playout of media streams at two or more geographically distributed locations in a time-synchronized manner. A large number of applications requiring IDMS can be found in [[Montagud2012](#)]. Some relevant examples are: Social TV, multi-party conferencing, networked loudspeakers and networked multi-player games.

[RFC7272] extends the RTP/RTCP mechanisms to exchange useful feedback information for IDMS between Synchronization Clients (SC) and a centralized Media Synchronization Application Server (MSAS). First, an RTCP XR block for IDMS, called "IDMS report", is defined to enable the SCs to report on packet reception and/or presentation times for a specific RTP stream. Second, a new RTCP IDMS packet type, called "IDMS Settings packet", is defined to allow the MSAS the transmission of IDMS setting instructions to the distributed SCs, based on the collected IDMS timings from them. This IDMS Settings packet will provide a common target playout point, referred to a specific RTP packet (concretely, to its generation timestamp) to which all the SCs belonging to a specific synchronization group (identified by the SyncGroupId field, specified in [[RFC7272](#)]) must synchronize. Besides, a new Session Description Protocol (SDP) parameter, called "rtcp-idms", is defined to notify about the usage of the above IDMS messages and to manage the group membership, thus allowing the co-existence of various independent groups of users in IDMS-enabled sessions.

Figure 1 shows the architectural solution for IDMS adopted in [[RFC7272](#)], which is based on a sync-maestro scheme [[Montagud2012](#)]. In this particular case, the SCs are co-located with RTP Receivers and the MSAS is co-located with the RTP Sender (although it could be co-located with an SC or with a third-party entity). The MSAS is responsible for collecting the IDMS reports from all the SCs belonging to a specific group, computing the delay differences among them and, if needed, sending IDMS Settings packets to make them to enforce the required adjustments to achieve IDMS. Although other architectural solutions are feasible, such as a distributed or a master/slave scheme [[Montagud2012](#)], this document also focuses on the sync-maestro scheme.



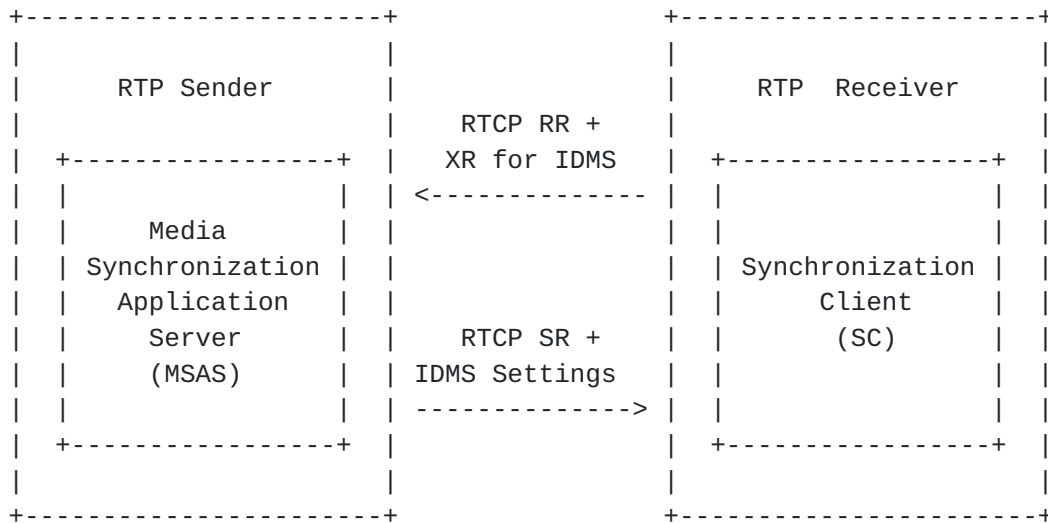


Figure 1: IDMS Architecture

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

The terminology defined in "RTP: A Transport Protocol for Real-Time Applications" [[RFC3550](#)], "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)" [[RFC4585](#)], "Rapid Synchronisation of RTP Flows" [[RFC6051](#)], and in "RTCP for Inter-Destination Media Synchronization" [[RFC7272](#)], apply.

The naming convention for RTCP is sometimes confusing. Below is a list of RTCP related terms and their meaning. Those concepts were defined in [[RFC3550](#)], [[RFC4585](#)], and [[RFC5506](#)], but refreshed here for a better understanding of this document:

- o RTCP packet: There are different types of RTCP packets, each one reporting on a specific metric or event of interest. It consists of a fixed header part followed by structured fields or elements depending on the information conveyed in that RTCP packet type.
- o Regular RTCP mode: Mode of operation in which no preferred transmission of feedback messages is allowed. Instead, RTCP packets are sent following the rules specified in [[RFC3550](#)].
- o Regular RTCP packet: RTCP packet that is sent using Regular RTCP mode.





- o Event: An observation that is (potentially) of interest to the participants of a media session and thus useful to be reported by means of an RTCP feedback message in a timely fashion.
- o Regular RTCP packet: RTCP packet that is sent using Regular RTCP mode.
- o Early RTCP mode: Mode of operation in which a participant is capable of reporting events of interest in a timely fashion (i.e., close to their occurrence).
- o Early RTCP packet: RTCP packet that is transmitted earlier than the scheduled transmission time when following the reporting rules specified in [RFC3550]. In [RFC4585], Early RTCP packets may be sent either as "Immediate Feedback" or as "Early RTCP" modes, depending on the group size. In the context of this document, Early RTCP packets are sent by the MSAS, which is a single centralized entity. Therefore, no dithering interval is needed to avoid feedback implosion, and RTCP packets can be sent in an "Immediate Feedback" mode. Sending an Early RTCP packet is also referred to as sending Early Feedback in this document.
- o Compound RTCP packet: A collection of two or more RTCP packets. In [RFC3550], it is mandatory to send RTCP packets as compound packets including at least an RTCP RR or SR packet, followed by an SDES packet with the CNAME item for the transmitting source identifier. Often the term "compound" is left out, so the interpretation of RTCP packet is therefore dependent on the context.
- o Minimal compound RTCP packet: A compound RTCP packet that contains only mandatory information, such as encryption prefix (if necessary), exactly one RTCP RR or SR packet, exactly one SDES packet with the CNAME item, and possible additional Feedback Messages [RFC4585] that must be sent as an Early RTCP packet.
- o (Full) Compound RTCP packet: A compound RTCP packet that conforms to the requirements on minimal compound RTCP packets and contains any additional number of RTCP packets (additional RRs, further SDES items, etc.). RTCP compound packets are typically sent when using Regular RTCP Feedback [RFC3550], although they may be also sent when using Early RTCP Feedback [RFC4585].
- o Reduced-Size RTCP packet: It only contains one or more RTCP packets, being much smaller than (minimal and full) compound RTCP packets [RFC3550]. Such RTCP packets are sent when using Early RTCP Feedback, but must not be sent when using Regular RTCP Feedback. The full definition is in [Section 4.1 of \[RFC5506\]](#).



Besides, the definition of three key concepts will help understanding (the proposed IDMS extensions in) this document:.

- o Inter-Stream Synchronization Delay (or Latency): It refers to the time difference between the instant at which a user joins an on-going multimedia session, probably involving different media (e.g., audio and video, or when using layered and/or multi-description media) carried in separate streams, and the instant at which these correlated streams can be presented to that user in a synchronized manner. It is defined in [[RFC6051](#)].
- o IDMS Delay (or Latency): It refers to the time difference between the instant at which an IDMS-related event occurs (e.g., a user joins an on-going session or an out-of-sync situation is detected) and the instant at which the IDMS Setting instructions to handle/repair such event have been received and processed at the target side/s or SCs.
- o Latecomer: In multi-party multimedia services, users may join and leave the session quite frequently. A user who joins a session in progress is usually called a latecomer (or late joiner) [[RFC6051](#)].

### **3. Background: RTCP Reporting Rules**

In this section, an overview of the RTCP reporting rules is provided. This is important to help understanding the benefits of the RTCP reporting rules and reports proposed in this document.

#### **3.1. Regular RTCP Feedback ([RFC 3550](#))**

During the media session's lifetime, the participants of an RTP Session (i.e., senders and receivers) regularly exchange RTCP reports (typically conveyed into compound RTCP packets) to inform mainly about QoS (Quality of Service) statistics, either in a unicast or multicast way (depending on the specific networked environment). On the one hand, a low frequency of RTCP feedback reporting can lead to faulty behavior due to outdated statistics. On the other hand, excessive reports can be redundant and cause unnecessary control traffic, probably leading to potential congestion situations.

Also, if the RTCP packets were exchanged at a constant rate, the control traffic would grow linearly with the number of participants. Accordingly, a trade-off between up-to-date information and the amount of control traffic must be met. This would allow an application to (automatically) scale over session sizes ranging from few participants to tens of thousands.



The total amount of control traffic added by RTCP should be limited to a small (so that the primary function of media data transport is not impaired) and known (so that each participant can independently calculate its share) fraction of the allocated RTP session bandwidth. A fraction of 5 % is recommended in [[RFC3550](#)]. In such process, media senders are given special consideration to allow a more frequent report exchange of their RTCP statistics, some of which are indeed very relevant for multimedia synchronization. In particular, if the proportion of senders constitute less than one quarter of the session membership, this percentage is further divided into two parts, where 25 % must be dedicated to active senders and the remaining can be consumed by receivers. Otherwise, the RTCP bandwidth is equally shared between senders and receivers.

Based on the above aspects, the RTCP report interval is dynamically and locally computed in each RTP entity, every time an RTCP packet is sent, according to the available session bandwidth, the average size of all received and sent RTCP packets, the number of participants in the session, their role (senders or receivers), as well as the unicast or multicast nature of the session (see [Section 6.2 of \[RFC3550\]](#) for further details).

This (deterministically) calculated RTCP report interval should also have a lower bound to avoid having bursts (cumpling) of RTCP packets when the number of participants is small and the law of large numbers is not helping to smooth out the traffic overhead. This would also help to avoid excessive frequent reports during transient outages like a network partition. The recommended value in [[RFC3550](#)] for that fixed minimum interval is 5 s.

Besides, a delay should be imposed to each participant before sending the first RTCP packet upon joining the session. This allows a quicker convergence of the RTCP report interval to the correct value. This initial delay may be set to half the minimum RTCP report interval (i.e., 2.5 s) in multicast sessions, whilst it may be set to zero in unicast sessions.

In some cases (e.g. if the data rate is high and the application demands more frequent RTCP reports), an implementation may scale the minimum RTCP interval to a smaller value given by 360 divided by the session bandwidth (in kbps). This yields to an interval smaller than 5 s when the session bandwidth becomes greater than 72 kbps. In multicast sessions, only active senders may use that reduced minimum interval, whilst in unicast sessions it also may be used by receivers. In the above cases, however, the minimum interval of 5 s must be still taken into account during the membership accounting procedure to not prematurely time out participants (who can indeed be using it) because of inactivity.



After that, the interval between RTCP packets is varied randomly over the range  $[0.5, 1.5]$  times that RTCP report interval to prevent floods of RTCP reports (i.e., to avoid that all RTCP packets are sent and received almost at the same time, in every report interval).

Additionally, "timer reconsideration" algorithms are introduced to allow for a more rapid adaptation of the RTCP report interval in large-scale sessions, where the membership can largely vary (e.g., many receivers join and leave the session quite frequently). To compensate for the fact that the "timer reconsideration" algorithms converge to a lower value than the intended average RTCP bandwidth, the (randomized) report interval is finally divided by  $e^{-3/2}=1.21828$ .

### **3.2. Early RTCP Feedback ([RFC 4585](#))**

In [[RFC4585](#)], further RTCP reporting mechanisms are specified to enable receivers to provide, statistically, more immediate RTCP feedback to the senders. This Early RTCP Feedback profile, which is known as RTP Audio-Visual Profile with Feedback (RTP/AVPF), allows for short-term adaptation and efficient feedback-based repairing mechanisms to be implemented, while maintaining the RTCP bandwidth constraints and preserving scalability to large groups.

The RTCP report interval specified in [[RFC3550](#)] is denoted as Regular RTCP interval in [[RFC4585](#)]. In addition, [[RFC4585](#)] enables to send RTCP reports earlier than the next scheduled Regular RTCP transmission time if a receiver detects the need to inform about media stream related events (e.g., picture or slice loss) close to their occurrence.

[NOTE] A feedback suppression mechanism is adopted, in which receivers wait for a random dithering interval to avoid feedback implosion (i.e., lots of receivers reporting on the same event).

The reporting rules for Regular RTCP packets in [[RFC4585](#)] are similar than the ones in [[RFC3550](#)]. However, the recommended minimum RTCP report interval of 5 s in [[RFC3550](#)] is dropped in [[RFC4585](#)]. Instead, an optional attribute, called "trr-int", is specified as an offset parameter (in ms) to the computed Regular RTCP Report interval.

Note that providing "trr-int" as an independent variable is intended to restrain from sending too frequent Regular RTCP packets (i.e., saving RTCP bandwidth) while enabling higher flexibility to transmit Early RTCP packets (i.e., using the saved RTCP bandwidth) in response to dynamic events. This could not be achieved by reducing the overall RTCP bandwidth, because the frequency of Early RTCP packets would be affected as well. Values between 4 and 5 s for "trr-int"





are recommended to assure interworking with RTP entities only using Regular RTCP Feedback. However, as "trr-int" is an optional attribute, it may be set to zero (default value) if a specific application would benefit from a higher frequency of Regular RTCP packets. In such a case, the only difference between the RTCP timing rules from [\[RFC3550\]](#) and [\[RFC4585\]](#) for transmitting Regular RTCP packets resides in the minimum value for the report interval, which is dropped in [\[RFC4585\]](#).

In order to preserve the RTCP traffic bounds, only one Early RTCP packet can be transmitted between two consecutive Regular RTCP packets (i.e. receivers cannot send two consecutive Early RTCP packets). After sending an Early RTCP packet, the RTCP reporting engine must schedule the transmission time for the next RTCP packet by skipping the next Regular RTCP interval.

Furthermore, [\[RFC4585\]](#) defines a small number of general-purpose feedback messages as well as a format for codec- and application-specific feedback information for transmission in the RTCP payloads.

### **3.3. Rapid Synchronization of RTP Flows ([RFC 6051](#))**

When using RTP streaming, the inter-stream synchronization delay can greatly increase in certain scenarios, especially in large multicast groups or when Multipoint Conference Units (MCU) are involved in the media delivery process. This increase of delay can be unacceptable and annoying to users, resulting in an overall poor Quality of Experience (QoE).

The aim of [\[RFC6051\]](#) is to minimize the inter-stream synchronization delay when using RTP/RTCP-based streaming. The motivation is that a receiver cannot synchronize playout of the incoming media streams until compound RTCP packets, including a Source Description or SDPS packet (including the media source identification) and a Sender Report or SR (including timing correlation parameters), are received for all the involved RTP senders in a multimedia session. In most implementations, media data will not be played out (watched or listened) until inter-stream synchronization is (initially) achieved. If there is no packet loss, this gives an expected delay equal to the average time for receiving the first RTCP packet from the RTP Session with the longest RTCP report interval. This delay is even more problematic if an RTCP SR packet from one of the involved RTP sessions is lost.

[NOTE] Note that the inter-stream synchronization delay depends on the specific instant at which a user joins the multimedia session or each RTP session (e.g., the user may first receive the RTCP packets from the RTP session with the longest RTCP interval), as well as on



the impact of the randomization processes in all the involved RTP sessions.

In [[RFC6051](#)], three backward compatible extensions to the RTP/RTCP protocols are proposed to reduce the inter-stream synchronization delay. First, the RTCP timing rules are updated to allow Single Source Multicast (SSM) senders [[RFC5760](#)] the immediate transmission of an initial compound RTCP packet upon joining each RTP session in a multimedia session (in parallel with the initial RTP packets). The rationale for not allowing the transmission of immediate RTCP packets to SSM receivers is to avoid feedback implosion in case that many receivers join the session almost simultaneously ("flash crowd" effect). This is clearly not an issue for SSM senders, since there can be at most one sender. Likewise, feedback implosion is a concern for Any Source Multicast (ASM) sessions, so [[RFC6051](#)] does not propose changes to the RTCP timing rules in these kinds of multicast environments. Second, a new RTP/AVPF transport layer feedback message (this type of RTCP messages is defined in [[RFC4585](#)]), called RTCP-SR-REQ, is defined to allow requesting the generation of an Early RTCP SR packet from the media sender. This enables rapid (re-)synchronization in case that an RTCP SR has not been received for a long period (e.g. due to packet loss or in sessions with large RTCP reporting intervals). Likewise, this enables latecomers to achieve inter-stream synchronization as soon as possible upon joining the session. Finally, new RTP header extensions are defined to enable the inclusion of metadata (in particular, NTP-based timestamps) in RTP data packets for in-band synchronization, thus avoiding the need for receiving RTCP SR packets before streams can be synchronized. These RTP header extensions do not eliminate the need for RTCP SR messages, but both mechanisms must be used for the synchronization control process. The use of RTCP SR packets for inter-stream synchronization allows backwards compatibility, but also provides higher robustness in the presence of middle boxes (e.g., RTP translators) that might strip RTP header extensions.

An accurate and rapid inter-stream synchronization is especially relevant when using layered, multi-description and multi-view media encodings. This is because all the individual RTP streams need to be synchronized before starting the decoding processes.

#### **3.4. SDP Modifiers for RTCP Bandwidth ([RFC 3556](#))**

In some applications, it may be appropriate to specify the RTCP bandwidth independently of the allocated RTP session bandwidth. Accordingly, [[RFC3556](#)] defines two additional Session Description Protocol (SDP) attributes to specify modifiers for the RTCP bandwidth for senders and receivers.



On the one hand, using a separate parameter allows rate-adaptive applications to set an RTCP bandwidth consistent with a "typical" data bandwidth that is lower than the maximum bandwidth specified by the session bandwidth parameter. This allows the RTCP bandwidth to be kept under 5% of the session bandwidth when the rate has been adapted downward, e.g. based on the stability of the network conditions. On the other hand, there may be applications that send data at very low rates, but need to communicate quite frequent RTCP information. These applications may need to specify an RTCP bandwidth that is higher than 5% of the data bandwidth.

If any of the SDP attributes for the RTCP bandwidth modifiers are omitted, the default value for that parameter is the one specified in the RTP profile in use for the session. [RFC3556] does not impose limits on the values that may be specified for both RTCP bandwidth modifiers, other than that they must be non-negative. However, the RTP specification and the appropriate RTP profile may specify limits.

#### **4. Early Event-Driven (EED) RTCP Feedback for IDMS**

This section introduces the EED RTCP reporting mechanisms proposed in this document to enable higher flexibility, dynamism, interactivity and accuracy when using RTP/RTCP for IDMS, with a sync-maestro architecture.

##### **4.1. Immediate Initial RTCP IDMS Settings**

[RFC6051] relaxes the timing rules for unicast and layered sessions so that SSM senders are allowed to transmit an initial compound RTCP packet (containing an RTCP SR packet and an RTCP SDES packet with a CNAME item) immediately on joining each RTP session in a multimedia session, in parallel with the initial RTP data packets. Hence, the inter-stream synchronization delay is significantly reduced, provided that the initial RTCP packet is not lost.

The same rationale for reducing the inter-stream synchronization delay in [RFC6051] can also be extrapolated to IDMS. In such a case, it would also be desirable the transmission of a nearly-immediate RTCP IDMS Settings packet by the MSAS upon establishing a multimedia session.

[NOTE]: Note that in this document, the terms (nearly-)immediate, close-to-instant and Early are used as synonymous. This is because the MSAS is a single centralized entity in the media session, and the Early RTCP packets can be sent immediately by this entity without requiring a contention algorithm, as required for receivers in [RFC4585].



If the MSAS is integrated within the RTP Server, it MUST send the IDMS Settings packet in parallel with the initial RTP data packets. If the Sync Manager is co-located within a SC or a third party entity (that also needs to be an RTP receiver for that session), it MUST send the IDMS Settings packet as soon as it receives the initial RTP data packets from the RTP Sender. In either case, as the MSAS is a single centralized RTP entity, it is also allowed to transmit Early RTCP packets [[RFC6051](#)]. This way, the SCs can start consuming the media in a synchronized manner earlier, thus ensuring a reduction of the IDMS Delay.

This mechanism is purely a local change to the MSAS that can be implemented as a configurable option, as stated in [[RFC6051](#)].

#### **4.2. Dynamic EED Reporting of IDMS Settings**

During the media session's lifetime, if Regular RTCP Feedback is used, the MSAS may have to wait a nearly-complete RTCP reporting interval to be able to send a new compound RTCP packet (including an IDMS Settings packet) after detecting an out-of-sync situation, which might potentially take several seconds (up to 5 s or even more) [[RFC3550](#)].

This is illustrated in Figure 2. In such a case, if an event (e.g. out-of-sync situation) is detected just after the transmission of an RTCP packet (at instant  $t_{r1}$ ), the next RTCP packet cannot be sent until the next randomized (over the scheduled transmission instant,  $t_{d2}$ ) RTCP transmission time (at instant  $t_{r2}$ ). The figure shows the worst case, in which the randomized RTCP report interval at that moment is near the upper limit, i.e.:

$$(t_{r2} - t_{r1}) = 1.5 * (t_{d2} - t_{r1})$$









synchronization levels (e.g., networked loudspeakers, or networked games) [[Montagud2012](#)].

Accordingly, the MSAS is also allowed to dynamically send Early RTCP IDMS Settings packets once detecting events throughout the duration of the media session. This is illustrated in Figure 4. In such a case, an RTCP IDMS Settings packet is sent just after the detection of an event, despite that this moment is earlier than the next regular RTCP transmission time. Consequently, the IDMS Delay is significantly reduced, mainly due to the fact that the MSAS delay has been minimized (due to the immediate transmission of the IDMS Settings packet).

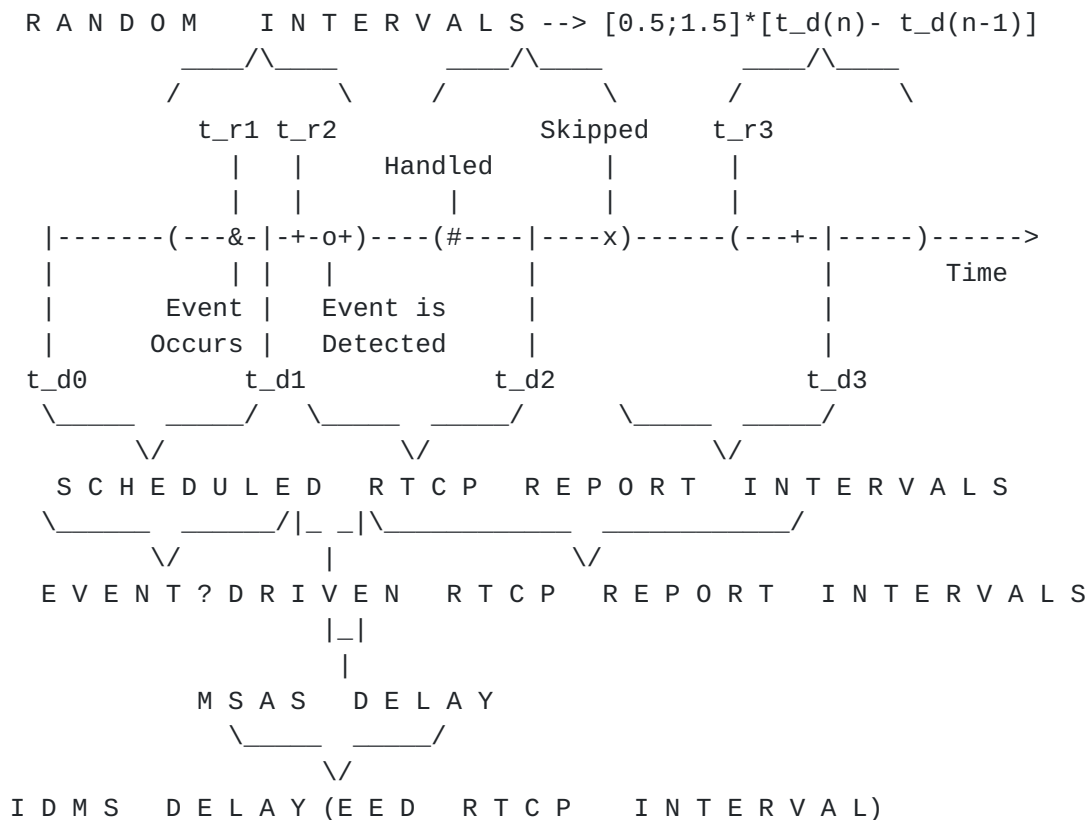


Figure 4: Early Event-Driven (EED) RTCP Timing Rules

Note that if "trr-int" is set to zero, only one Early RTCP packet can be transmitted between two consecutive Regular RTCP packets to preserve the RTCP traffic bounds [[RFC3550](#)]. It means that an Early RTCP packet can only be sent if the previous transmitted RTCP packet was a Regular RTCP packet. Hence, after sending an Early RTCP packet, the RTCP reporting engine MUST schedule the sending time for



the next RTCP packet by delaying (i.e., skipping) one more Regular RTCP report interval, as in [[RFC4585](#)].

In case of high frequency of events, setting an offset value for the Regular RTCP report interval, by means of using the "trr-int" attribute, can help to save RTCP bandwidth (by restraining the transmission of too frequent Regular RTCP packets) while being able to use the (saved) bandwidth when events occur. This is out of scope of this document.

Therefore, the proposed EED RTCP reporting rules are based on the fact that not all the feedback reports from the MSAS are of equal importance. This means that some feedback reports from the MSAS need to be reported in a timely fashion.

The dynamic EED reporting of IDMS Settings packets is also be very useful to provide playout hints for specific events that must be presented to all the involved users in a fine-grained synchronized way with the piece of content they refer to. Those events can be media-related events whose timing can be known in advance (e.g., commercials, start of the match in a sports event...), but the events' timing could be even unknown (e.g., a goal in a football match...) or dynamically triggered by either operators (e.g., a TV quiz show, in-game actions, interesting scenes...) or users (e.g., shared service control, interactive instant messaging...). Therefore, the use of EED RTCP Feedback for IDMS implies an interaction between the application-layer (through which operator or user generated events are triggered) and the transport/control layer (i.e., RTP/RTCP protocols) in order to translate the high-level (i.e., content-based or action-based) events into lower level calls (i.e., transmission of Early RTCP packets), as well as their alignment in terms of timelines. These are not severe issues, since the MSAS will be co-located with the Media Sender in most implementations. However, this differs from the use of Regular RTCP Feedback for IDMS, in which the IDMS adjustments are purely based on packet-level timestamps.

#### **[4.3.](#) Rapid (Re-)Synchronization Request**

If the initial compound RTCP packet (including an SR, SDES and IDMS Settings packets) is lost, the affected SCs will not be able to synchronize the media playout until the report interval has passed, and the next RTCP packet can be sent. This is undesirable. [[RFC6051](#)] defines a new RTP/AVPF transport layer feedback message to request the generation of an Early RTCP SR, allowing rapid inter-stream (re-)synchronization.



A similar mechanism is proposed in this document to be applied for IDMS purposes. A new RTP/AVPF transport layer feedback message [RFC4585], called RTCP-IDMS-REQ, is introduced to request the rapid generation (and transmission) of an RTCP IDMS Settings packet by the MSAS (see Figure 5).

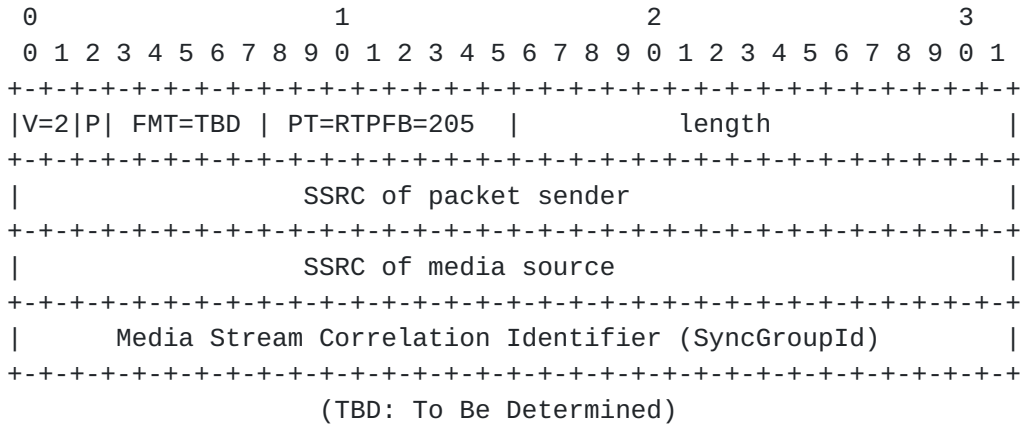


Figure 5: RTCP-IDMS-REQ Feedback Message

The Payload Type (PT) of this type of RTCP message should be 205, as specified in [RFC4585], the Frame Message Type (FMT) MUST be assigned by IANA (Internet Assigned Numbers Authority), and its length must be equal to 3. The SSRC of the packet sender MUST indicate the SC sending the packet, while the SSRC of the media source field MUST indicate the source of the media stream the SC (who sends this RTCP-IDMS-REQ) is unable to synchronize. In contrast to the RTCP-SR-REQ feedback message defined in [RFC6051], in which the Feedback Control Information (FCI) part is kept empty, in the RTCP-IDMS-REQ it must carry the SyncGroupId (specified in [RFC7272]) to which the sender of this message belongs.

Once a new RTCP-IDMS-REQ is received by the MSAS, it MUST generate an Early RTCP IDMS Settings packet as soon as possible, while complying with the Early RTCP feedback rules.

This mechanism can also be employed if a SC has not received RTCP IDMS Settings in a (configurable) long time interval. The RTCP-IDMS-REQ packet MAY be repeated once per RTCP reporting interval if no RTCP IDMS Settings packet is forthcoming. Likewise, the MSAS MAY ignore RTCP-IDMS-REQ packets if its regular schedule for RTCP transmission will allow the SC to synchronize within a reasonable time interval.





Although this mechanism is similar to the one in [[RFC6051](#)] to request rapid RTCP SRs, it is especially necessary since, in most implementations, the IDMS Settings packets will not be regularly sent in each RTCP report interval, as RTCP SRs, but only when the detected asynchrony exceeds an allowable threshold..

When using SSM sessions with unicast feedback, it is possible that the feedback target and the MSAS are not co-located. If a feedback target receives an RTCP-IDMS-REQ feedback message in such a case, the request SHOULD be forwarded to the MSAS. However, the mechanism to be used for forwarding such requests is not defined in this document.

If the feedback target provides a network management interface, it might be useful to provide a log of which SCs send RTCP-IDMS-REQ feedback packets and which do not, since the later will experience slower stream synchronization.

#### **[4.4.](#) Reduction of Channel Change Delays**

The participation and support of latecomers are key issues to enable dynamic IDMS-enabled sessions. One widely applicability of the RTCP-IDMS-REQ messages is the dynamic and rapid accommodation of latecomers.

Once a latecomer joins an IDMS-enabled session, it MUST send an RTCP-IDMS-REQ to the MSAS of that session. Then, the MSAS MUST generate an Early RTCP IDMS Settings packet to rapidly bring the latecomer up-to-date, thus minimizing the IDMS Delay experienced by that latecomer (i.e. the time interval between joining and acquiring IDMS).

Immediately after receiving the RTCP IDMS Settings packet, the latecomer will begin to play out the media stream in a time synchronized way with the other SCs, thus becoming an additional member in the IDMS-enabled session, as shown in Figure 5. This will prevent from both long annoying startup delays and initial playout inconsistencies.



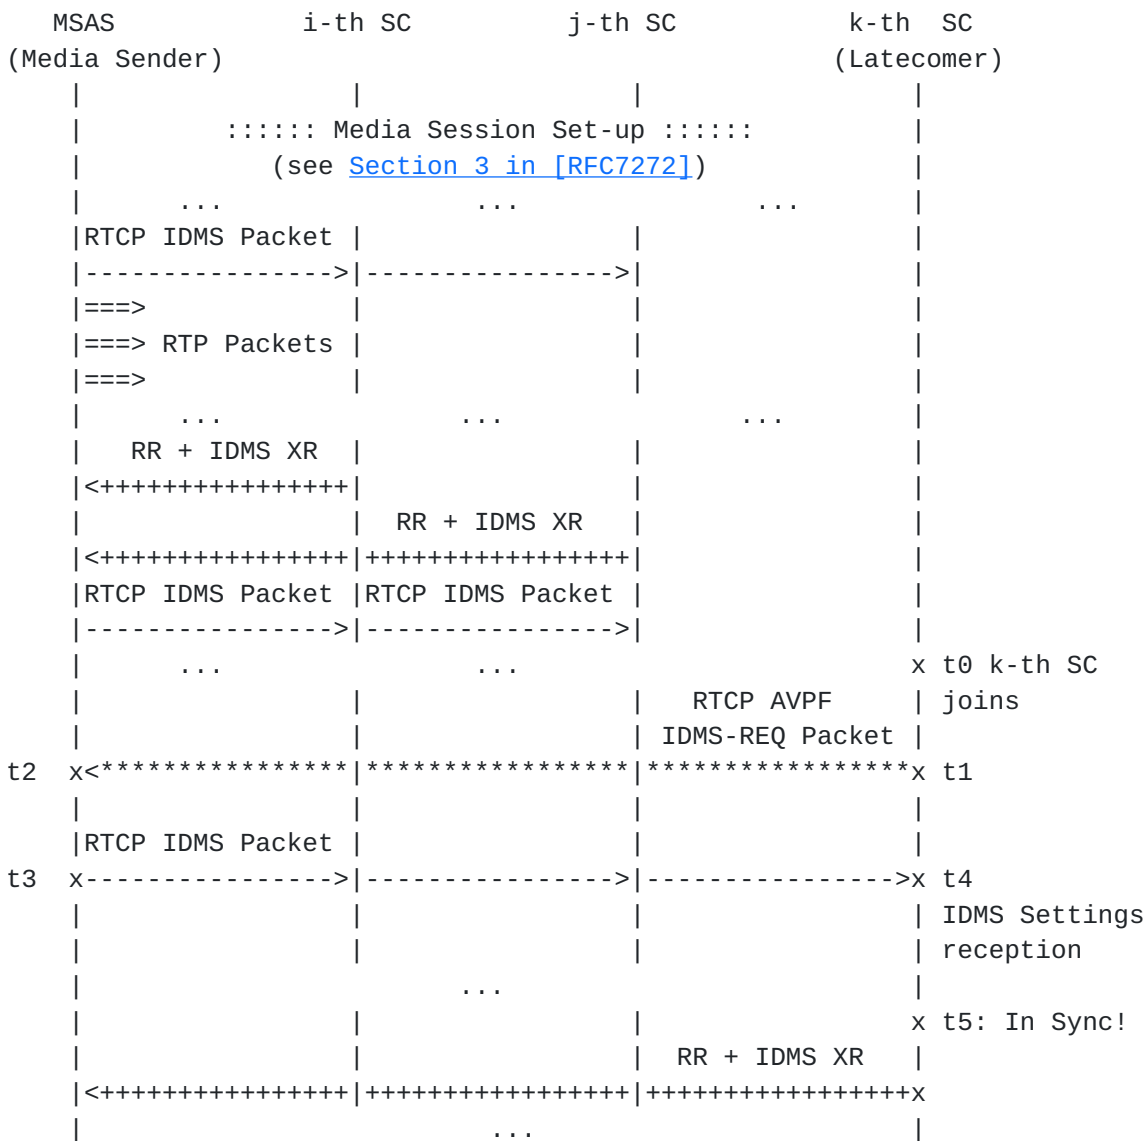


Figure 6: IDMS Operation and Rapid Accommodation of Latecomers

The timing diagram for the RTCP exchange processes is illustrated in Figure 6. It can be seen that, if enabling EED RTCP Feedback for IDMS, the IDMS delay (which is the time interval between joining and acquiring IDMS, represented by (t5 - t0) in Figure 6) for the latecomer (k-th SC) can be significantly reduced, mainly due to the fact that the MSAS delay ((t3 - t2) in Figure 6) can be minimized.

Although the above discussion has been focused for latecomers, it is also applicable for reducing channel change (i.e., zapping) delays in IDMS-enable sessions, which is currently a hot research topic in the IPTV area. Concretely, the relevance of channel change delays and their variability in IDMS-sensitive services is threefold. First, as



for inter-stream synchronization, the required time to receive the IDMS setting instructions must not contribute to further increase the channel change delays. Second, apart from the magnitudes of channel change delays, their variability (i.e., the delay variation for each user) will also impact the IDMS performance. Third, when a group of users are watching IPTV together and they (simultaneously) change (or must change) to another channel, any playout time differences among them will also influence the resulting delay. Therefore, in this context, the use of EED RTCP Feedback for IDMS is very beneficial because: i) it significantly reduces the time needed to receive the initial RTCP IDMS Settings packet; and ii) it enables the compensation of the delay differences when changing channels.

Two additional mechanisms could contribute to further reduce the IDMS Delay. The first one consists of employing priority mechanisms for the transport of RTCP messages, e.g., by adopting a Differentiated Services (DiffServ) policy. This would help to decrease the Round Trip Time (RTT) delays and the loss probability for RTCP packets (out of the scope of this document). The second one is based on the transmission of Early RTCP-IDMS-REQ messages by latecomers upon joining the session. According to [\[RFC6051\]](#), the delay since joining and sending an RTCP-IDMS-REQ message ((t1-t0) in Figure 6) should not be reduced to avoid flooding of requests at specific time instants (e.g., at the time a broadcasted sport event begins). Although we initially adhere to this standard compliant rule in this document, we believe that it should be discussed within the AVTCORE WG if this flash crowd effect is a real limiting issue in real SSM scenarios (e.g., networked quiz shows, gaming, IPTV, etc.). Our initial assumption is that the upstream bandwidth availability by the SCs (which is not used for other purposes) and the aggregation and re-distribution mechanisms by Feedback Targets [\[RFC5760\]](#) do not entail a real constraint for allowing the transmission of Early RTCP-IDMS-REQ by the SCs. Moreover, it is assumed in [\[RFC6051\]](#) that all SCs switch channels simultaneously, but even though using automated procedures (e.g., through notifications via the Electronic Program Guides in IPTV), this would not be a matter of a few seconds, but most probably of minutes.

## **5. Reduced-Size RTCP Reporting for IDMS**

RTCP SR and RR contain relevant QoS statistics usable for monitoring of the media stream transmission and thus enable media adaptation. The regular transmission of these RTCP packets becomes very useful to infer trends between successive reports due to the dynamic nature of the conveyed information. Likewise, the tracking of the session size thanks to that regular reports warrants bounded RTCP traffic bandwidth. Moreover, compound RTCP packets are useful to establish and maintain multimedia synchronization. Due to the above issues,



the regular transmission of compound RTCP packets must be maintained throughout the RTP session's lifetime.

However, [[RFC5506](#)] defines certain changes to the RTCP reporting rules to allow feedback messages to be sent as Reduced-Size RTCP packets under certain conditions when using the RTP/AVPF profile [[RFC4585](#)]. The motivation is that, in some cases, it is more useful to report certain events of interest the more frequently or the sooner as possible to enable short-term adaptation, rather than sending full compound RTCP packets including periodic statistics.

In low bitrate links (e.g. radio access technologies), there are some benefits of reporting Reduced-Size RTCP packets. First, if channels conditions are poor, smaller RTCP packets are much more likely to be successfully transmitted than larger compound RTCP packets, and they will also introduce lower traffic overhead. The last issue is critical, as those messages are likely to occur when channel conditions are poor (e.g., for reporting picture or slice loss). Second, the serialization time when transmitting small size RTCP packets is shorter than when transmitting full (larger) RTCP compound packets. Third, when the bandwidth availability is scarce, smaller RTCP packets will enable more frequent feedback messages.

In high bitrate environments, the above issues are not real limitations, but using Reduced-Size RTCP packets may also be beneficial to reduce the processing delay and complexity of RTCP packets.

Independently of the link type, additional benefits with sending feedback in small Reduced-Size RTCP packets can be cited. For instance, when using Early RTCP Feedback [[RFC4585](#)], receivers typically need to send frequent event-driven feedback messages. In such cases, if using Reduced-Sized RTCP packets, the risk that the RTCP bandwidth becomes too high during periods of heavy feedback signaling is reduced.

More details about use cases for, benefits of and issues with Reduced-Size RTP reporting can be found in [[RFC5506](#)].

Accordingly, the use of Reduced-Size RTCP packets MAY also be beneficial when enabling the EED RTCP reporting rules for IDMS proposed in this document. [[RFC5506](#)] specifies that in Immediate Feedback mode, as it is the case of RTCP IDMS Settings packets, all feedback messages may be sent as Reduced-Size RTCP packets. However, it is also stated that Reduced-Size RTCP packets shall not be sent until at least one compound RTCP packet has been transmitted.





This implies that if the use of non-compound RTCP packets [[RFC5506](#)] has been previously negotiated between the participants in an IDMS-enabled session, both the RTCP-IDMS-REQ and the RTCP IDMS Settings packets MAY be sent as non-compound RTCP packets, but only if a compound RTCP packet has been already sent by the senders of those feedback messages.

The Reduced-Size RTCP reporting features do not apply to the first IDMS Settings packet, which SHOULD be sent in parallel with the initial RTP data packets ([Section 4.1](#)). In such a case, SCs would also be benefited by the reception of a compound RTCP packet including SR and SDES packets. This is also the case when a latecomer sends an Early RTCP-IDMS-REQ. Here, the media sender and the MSAS would also need an SDES packet from that latecomer for membership accounting.

However, the use of Reduced-Size RTCP reporting MAY be beneficial when an RTCP-IDMS-REQ is sent by an active SC that has not received an RTCP IDMS Settings packet for a long period, requesting re-synchronization setting instructions. Moreover, Reduced-Size RTCP reporting MAY be especially useful for Dynamic EED transmission of IDMS Settings packets ([Section 4.2](#)), e.g. immediately after detecting an out-of-sync situation, or when a specific media event is targeted to be simultaneously presented in all the SCs.

Reduced-Size RTCP packets can become substantially smaller than compound packets. A compound packet is forced to contain both an RR or an SR and an SDES with at least the CNAME item. The RR containing a report block for a single source is 32 bytes, and an SR is 52 bytes. Both may be larger if they contain report blocks for multiple sources. The SDES packet containing a CNAME item will be 10 bytes plus the CNAME string length. Here, it is reasonable that the CNAME string is at least 10 bytes to get a decent collision resistance. If the recommended form of user@host is used, then most strings will be longer than 20 characters. Additionally, for IDMS purposes: i) SCs will regularly send RTCP XR blocks for IDMS, which consist of 10 32-bit words (40 bytes); ii) the MSAS will send RTCP IDMS Settings packets, which consists of 9 32-bit words (36 bytes).

Therefore, Reduced-Size RTCP packets can become at least 70-80 bytes smaller than RTCP compound packets, thus reducing the mean RTCP packet size, decreasing sporadic traffic peaks, reducing the transmission delay, and increasing the overall RTCP feedback frequency.



## **6. Security Considerations**

This document does not impose security considerations, beyond the ones described in [[RFC3550](#)], [[RFC4585](#)], [[RFC6051](#)], and [[RFC7272](#)].

## **7. IANA Considerations**

This document defines a new RTP/AVPF transport layer feedback message [[RFC4585](#)], called RTCP-IDMS-REQ, whose FMT Value needs to be assigned by IANA..

## **8. Contributors**

Authors would like to thank the co-authors of [[RFC7272](#)] for their collaboration on specifying the RTCP extensions for IDMS, which are starting point of this work.

## **9. Conclusions**

This document proposes Early Event-Driven RTCP reporting rules to enable higher interactivity, dynamism and accuracy when using RTP/RTCP for IDMS.

## **10. References**

### **10.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC3556] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", [RFC 3556](#), July 2003.
- [RFC3611] Friedman, T., Caceres, R., and A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", [RFC 3611](#), November 2003.
- [RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", [RFC 4585](#), July 2006.



- [RFC5506] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", [RFC 5506](#), April 2009.
- [RFC5576] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", [RFC 5576](#), June 2009.
- [RFC5760] Ott, J., Chesterfield, J., and E. Schooler, "RTP Control Protocol (RTCP) Extensions for Single-Source Multicast Sessions with Unicast Feedback", [RFC 5760](#), February 2010.
- [RFC6051] Perkins, C. and T. Schierl, "Rapid Synchronisation of RTP Flows", [RFC 6051](#), November 2010.
- [RFC7272] van Brandenburg, R., Stokking, H., van Deventer, O., Boronat, F., Montagud, M., and K. Gross, "Inter-Destination Media Synchronization (IDMS) Using the RTP Control Protocol (RTCP)", [RFC 7272](#), June 2014.

## **10.2. Informative References**

- [Montagud2012]  
Montagud, M., Boronat, F., Stokking, H., and R. van Brandenburg, "Inter-Destination Multimedia Synchronization; Schemes, Use Cases and Standardization", *Multimedia Systems Journal* (Springer), 18(6), pp. 459-482", November 2012, <<http://link.springer.com/article/10.1007%2Fs00530-012-0278-9#page-1>>.
- [Montagud2013]  
Montagud, M., Boronat, F., and H. Stokking, "Early Event-Driven (EED) RTCP Feedback for Rapid IDMS", *The 21st ACM International Conference on Multimedia (ACM MM 2013)*, Barcelona (Spain)", October 2013, <<http://dl.acm.org/citation.cfm?id=2502114>>.
- [RFC5968] Ott, J. and C. Perkins, "Guidelines for Extending the RTP Control Protocol (RTCP)", [RFC 5968](#), September 2010.

Authors' Addresses



Mario Montagud (editor)  
Universitat Politecnica de Valencia  
C/ Paraninf, 1  
Grau de Gandia, Valencia 46730  
Spain

Phone: +34 962 849 341  
Email: mamontor@upv.es

Fernando Boronat  
Universitat Politecnica de Valencia  
C/ Paraninf, 1  
Grau de Gandia, Valencia 46730  
Spain

Phone: +34 962 849 341  
Email: fboronat@dc.com.upv.es

Hans Stokking  
TNO  
Brassersplein 2  
Delft 2612CT  
The Netherlands

Phone: +31-88-866-7000  
Email: hans.stokking@tno.nl



