

Network Working Group
Internet-Draft

Intended status: Informational
Expires: April 30, 2015

H. Stokking
O. van Deventer
R. van Brandenburg
TNO
F. Boronat
M. Montagud
Universitat Politecnica de
Valencia
October 27, 2014

Inter-Destination Media Synchronizaton for IPTV Environments
draft-stokking-avtcore-idms-for-iptv-00.txt

Abstract

[RFC7272] describes the use of RTCP for the purpose of Inter-Destination Media Synchronization (IDMS) between Synchronization Clients (SCs) and an Media Synchronization Application Server (MSAS). This document extends that work for application in the area of IPTV. First, RTCP can be used according to the single source multicast (SSM) principles from [RFC5760] in the IPTV application area. This document specifies the use of a feedback target for collecting and possibly summarizing IDMS reports. For this, the document defines 2 new sub-report blocks for the use of IDMS according to the SSM principles. Alternatively, the MSAS can be co-located with the Feedback Target, for synchronizing small groups of receivers. Secondly, in an IPTV environment, different viewers may receive the same content, but in non-identical streams. The IDMS solution presented in [RFC7272] will no longer work in such a case. This document provides a solution for this.

Requirements Language

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].

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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	3
1.1. IDMS in an IPTV environment	3
1.1.1. IDMS for Single Source Multicast	3
1.1.2. IDMS for different streams providing the same content	4
2. IDMS report aggregation in SSM session	5
2.1. IDMS Packet Received Sub-Report Block	5
2.2. IDMS Packet Presented Sub-Report Block	7
3. IDMS with separate synchronization server with unicast synchronization settings	7
4. IDMS in case of multiple RTP streams	8
5. IANA Considerations	9
6. Security Considerations	9
7. Acknowledgements	9
8. Normative References	9
Authors' Addresses	10

1. Introduction

The Real-time Transport Protocol (RTP) provides a real-time transport mechanism suitable for unicast and multicast communication between multimedia applications. An important component of the RTP protocol is the control channel, defined as the RTP Control Protocol (RTCP). RTP and RTCP have been extended in numerous RFCs. Two such extensions are the extensions for Single-Source Multicast (SSM) Sessions with Unicast Feedback in [RFC5760] and the use of RTCP for Inter-Destination Media Synchronization (IDMS) in [RFC7272].

This Internet draft provides a number of extensions on the use of RTCP for IDMS in an IPTV environment. The IPTV environment has a number of characteristics that are currently not dealt with [RFC7272]. The introduction discusses the IPTV environment and identifies various gaps in the current IDMS solution. The next sections discuss solution directions for dealing with these gaps. The purpose of this Internet Draft is to build upon [RFC7272] so that the IDMS solution can be applied in an IPTV specific environment.

1.1. IDMS in an IPTV environment

An IPTV environment has specific characteristics, which [RFC7272] does not deal with properly. These characteristics are:

- o Single Source Multicast (SSM, [RFC5760]) setting with a large number of viewers.
- o Different receivers may receive different versions of the same content, i.e. they receive non-identical streams, e.g. different unicast streams, different encoded streams, streams from different media senders.

1.1.1. IDMS for Single Source Multicast

The first characteristic of IPTV is the large-scale Single Source Multicast (SSM) setting. Regular linear television is offered using SSM. Such SSM sessions have a large number of viewers, often in the millions, which requires a highly scalable approach. Applying IDMS to such an SSM session can be done in two ways:

1. Synchronize all receivers. In this case, [RFC7272] does not offer the scalability to synchronize all viewers in such large-scale sessions. In such a case each receiver contains a synchronization client (SC) which communicates with the Media Synchronization Application Server (MSAS). [RFC5760] offers a unicast feedback system using feedback targets (FTs) to collect and possibly aggregate RTCP reports of groups of receivers. IDMS can be performed using this

same feedback system, providing more scalability. Section 2 of this document specifies how to accomplish this.

2. Synchronize independent groups of receivers, depending on the application. Use cases for synchronization, such as social TV or such as multiple receivers in a single physical location, require only a limited number of receivers to be synchronized together. For example, when millions of viewers watch the same television show, only specific groups of users viewing the show together would have to be synchronized, and only within their own group. [RFC5760] describes a system that provides all receivers with the same information. If only a limited subset of the receivers are synchronized together, not all receivers need to receive the same synchronization instructions. Section 3 of this document provides a unicast way of sending synchronization instructions to receivers, which requires the MSAS to be co-located with the Feedback Target.

The choice between options 1 and 2 depends on a number of factors. If only a limited number of receivers use a service that requires IDMS, it is inefficient to synchronize all viewers. Also, playout timing differences between various receivers can be relatively large due to e.g. variable propagation delays. If that is the case, and every receiver is synchronizing to the slowest receiver, a lot of buffering needs to be done. This is not efficient, and also significantly increases channel changing delays. In these cases, it makes sense to use option 2. If on the other hand many viewers use synchronization sensitive services, and playout timing differences are relatively small, it may make sense to synchronize all receivers by using option 1.

1.1.2. IDMS for different streams providing the same content

Different viewers may watch the same content, but use a different media stream in an IPTV environment. An example of this is when one viewer receives an High Definition (HD) stream and another viewer receives an Standard Definition (SD) stream. Another example is when multiple receivers view the same video-on-demand, receiving this using different unicast streams. Services such as social TV, where different viewers remotely view media content together, require synchronization of these different streams. The IDMS solution is based on RTP timestamps. For different streams, these timestamps are not aligned, i.e. there is no relation between the timestamps in one stream and the timestamps in another stream, because of the different random offset of the RTP timestamps, as well as potential clock skew. Because of this, the MSAS cannot determine proper synchronization instructions.

There are two possible solution directions for this problem.

The first is to have the media source output the different streams with the same timing. The NTP timestamp in the RTP packets will then be synchronous, i.e. IDMS can be based on this NTP timestamp analogue to inter-stream synchronization (lip-sync). This does require the MSAS to be informed on the RTP/NTP relationships of the various streams. This information is available in the RTCP SRs of the various streams. If the MSAS is part of the media source, this is implicitly available. If this is not the case, the MSAS should receive these SRs somehow. This document presents various options to achieve this.

The other solution for this problem, is to have the media source determine and signal the relationship between the various RTP timestamps of the various streams. Again, if the MSAS is part of with the media source, then this information is locally available. If the MSAS is a separate entity, the media source can sent this information to the MSAS. Section 4 of this document shows how that is done.

2. IDMS report aggregation in SSM session

[RFC5760] describes how to use Feedback Targets (FTs) or the Distribution Source (DS) for summarizing RTCP packets from receivers, using the Receiver Summary Information (RSI) Packet. This section describes two new sub-report blocks, to be used in those RSI packets.

One sub-report block is for summarizing reported RTP packet received timestamps, the other is for summarizing reported RTP packet presented timestamps.

A feedback target or distribution source MUST only summarize IDMS information from SCs, if they belong to the same synchronization group, i.e. when the reports from the receivers contain the same Media Stream Correlation Identifier [RFC7272]. If at least one of the receivers in a certain synchronization group reports on both packet received timestamps and packet presented timestamps, a feedback target or distribution source SHOULD also include packet presented timestamps. If all receivers report on packet presented timestamps, a feedback target or distribution source MUST include packet presented timestamps. If a feedback target or distribution source summarizes the packet received timestamps, it SHOULD also summarize the packet presented timestamps.

2.1. IDMS Packet Received Sub-Report Block

To summarize the packet received timestamps in the IDMS information from SCs, a feedback target or distribution source can use the following sub-report block. The name of this sub-report block is "IDMS Received", the long name is "IDMS Packet Received NTP

Timestamp".

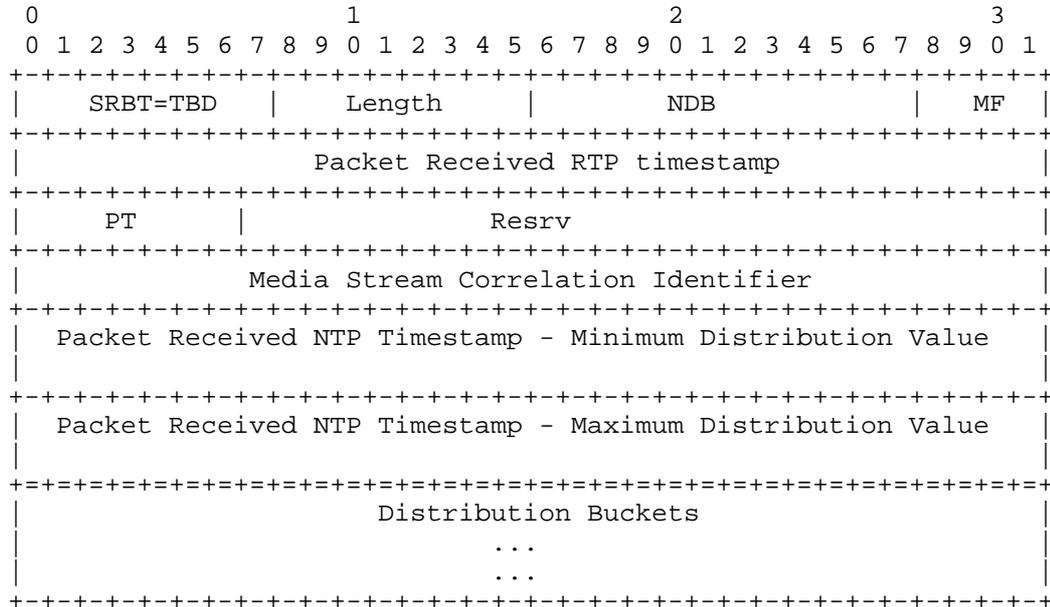


Figure 1: IDMS Packet Received Sub-Report Block

Sub-Report Block Type (SRBT): 8 bits, TBD upon IANA registration.

Length: 8 bits, the length of the sub-report in 32-bit words, as defined in RFC 5760.

Number of Distribution Buckets (NDB): 12 bits, as defined in RFC 5760, except for the calculation of the size of each bucket. Since the header is longer than the sub-report blocks defined in RFC 5760, the size of each bucket can be calculated using the formula $((\text{length} * 4) - 32) * 8 / \text{NDB}$ number of bits.

Multiplicative Factor (MF): 4 bits, as defined in [RFC5760].

Packet Received RTP Timestamp: 32 bits, as defined in [RFC7272].

Payload Type (PT): 7 bits, as defined in [RFC7272].

Reserved Bits (Resrv): 25 bits, as defined in [RFC7272].

Media Stream Correlation Identifier: 32 bits, as defined in [RFC7272].

Packet Received NTP timestamp - Minimum Distribution Value: 64 bits, as defined in [RFC7272].

Packet Received NTP Timestamp - Maximum Distribution Value: 64 bits, as defined in [RFC7272].

Distribution Buckets: each bucket has $((\text{Length} * 4) - 28) * 8 / \text{NDB}$ bits.

The whole sub-report block contains only a single packet received RTP timestamp value. Since various receivers will normally report on different packet received RTP timestamps, a feedback target MUST recalculate all packet received NTP timestamps to match the single packet received RTP timestamp. This will give an overview of the packet received times of all receivers for that specific RTP timestamp. This recalculation is necessary for all reported timestamps: minimum, maximum and those in the distribution buckets.

2.2. IDMS Packet Presented Sub-Report Block

To summarize the packet presented timestamps in the IDMS information from SCs, a feedback target or distribution source can use the following sub-report block. The name of this sub-report block is "IDMS Presented", the long name is "IDMS Packet Presented NTP Timestamp".

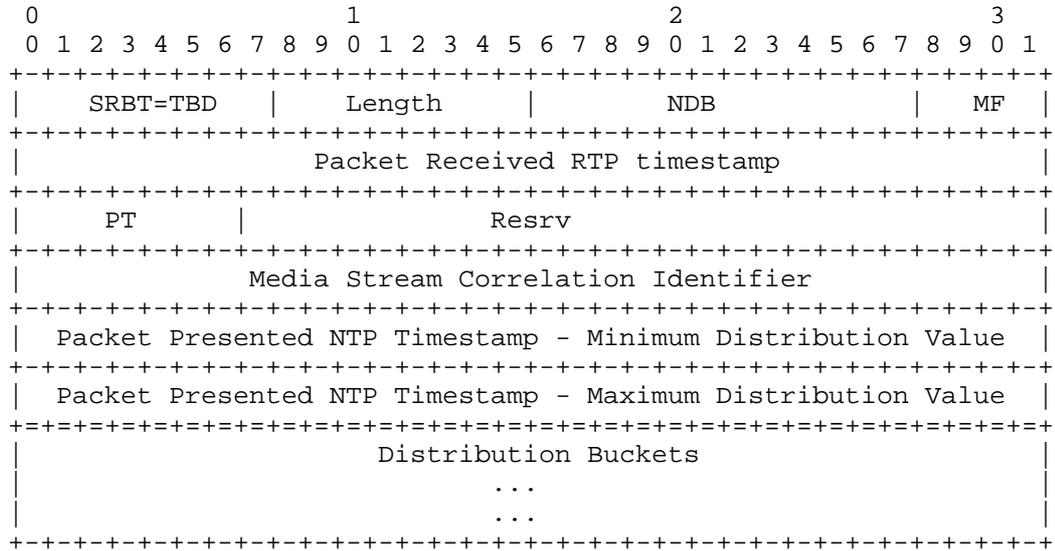


Figure 2: IDMS Packet Presented Sub-Report Block

Sub-Report Block Type (SRBT): 8 bits, TBD upon IANA registration.

Length: 8 bits, the length of the sub-report in 32-bit words, as defined in RFC 5760.

Number of Distribution Buckets (NDB): 12 bits, as defined in RFC 5760, except for the calculation of the size of each bucket. Since the header is longer than of the sub-report blocks defined in RFC 5760, the size of each bucket can be calculated using the formula $((\text{length} * 4) - 28) * 8 / \text{NDB}$ number of bits.

Multiplicative Factor (MF): 4 bits, as defined in [RFC5760].

Packet Received RTP Timestamp: 32 bits, as defined in [RFC7272].

Payload Type (PT): 7 bits, as defined in [RFC7272].

Reserved Bits (Resrv): 25 bits, as defined in [RFC7272].

Media Stream Correlation Identifier: 32 bits, as defined in [RFC7272].

Packet Presented NTP timestamp - Minimum Distribution Value: 32 bits, as defined in [RFC7272].

Packet Presented NTP Timestamp - Maximum Distribution Value: 32 bits,

as defined in [RFC7272].

Distribution Buckets: each bucket has $((\text{Length} * 4) - 28) * 8 / \text{NDB}$ bits.

The whole sub-report block contains only a single packet received RTP timestamp value. Since various receivers will report on different packet presented NTP timestamps, a feedback target MUST recalculate all packet presented NTP timestamps to match the single packet received RTP timestamp. This is true for all reported timestamps: minimum, maximum and those in the distribution buckets.

3. IDMS with separate MSAS with unicast synchronization settings

The second alternative for IDMS in a large-scale SSM context is to synchronize small groups of receivers that need to be synchronized with each other because of the service requirements. Different groups may still receive the same RTP streams, but can be synchronized independent of each other. In that case, each group MUST receive its own synchronization settings instructions in the form of IDMS Settings Packets as defined in [RFC7272]. Normally the Receiver Reports (RRs) or the Received Summary Information (RSI) is sent to all receivers. But, since different groups of receivers may need different synchronization settings instructions, these instructions cannot be multicast. As it happens, multicasting all instructions would lead to a situation where all receivers would receive a multitude of different settings instructions. They would have to find their own instructions based on the MSCSI of their group, which is possible. But, with a large number of groups, this would be highly inefficient. This is why a unicast method is taken here.

To unicast the instructions to the various SCs, the MSAS needs to directly receive the IDMS reports from the various SCs. This means the MSAS MUST be co-located with the feedback target. When supplying SCs with the unicast address to which they should send their reports, different SCs in the same synchronization group MUST be allocated the same feedback target. Also, because the synchronization information is no longer relevant upstream of the MSAS, the feedback target SHOULD terminate these RTCP blocks and not forward them or summarize them.

How the receivers receive the unicast address of the feedback target is out of scope of this draft. [RFC5760] only defines pre-configuration for this. Alternatively, the RTCP-attribute as specified in [RFC3605] can be used on the session level to provide receivers of a shared session with the unicast address of the MSAS, similar to how this is done in [TS183063].

Synchronization settings instructions MUST be sent by the MSAS to the source IP addresses of the received synchronization information, using the same destination port as the received synchronization information.

4. IDMS in case of multiple RTP streams

IDMS is based on various receivers reporting on the packet received times or packet presentation times. This document describes situations in which the MSAS is not part of the media source. If all receivers receive the exact same RTP streams, e.g. in case of multiple receivers of a single multicast streams, this will work fine. The MSAS can relate the various received IDMS information. Even if different receivers report on different RTP timestamps, the MSAS can calculate the timing differences between clients by extrapolation using the RTP clock frequency derived from the reported payload type.

However, when the MSAS is not co-located with the media source and the receivers receive the same content in different RTP streams, an MSAS cannot perform the necessary calculations for achieving synchronization. To perform these calculations, there has to exist some common timeline in the reports by the various receivers. To determine a common timeline, the MSAS needs some kind of information correlating the RTP timestamps in the various streams. This section provides two alternatives for this.

The first alternative is to use the RTCP Sender Reports that belong to the various RTP streams. In these SRs, the RTP timestamps are linked to the employed wallclock time. This is normally used for intra- and inter-media synchronization. The media sources need to ensure that the same part of the content in different streams corresponds to the same wallclock time (NTP timestamp in the SR). This way the SRs of the various RTP streams can be used to establish a common timeline between those RTP streams. The easiest way to send the SR to the MSAS is by having the receiver append it to its report blocks. Another option is to have the MSAS act as a third party monitor, as described in [RFC3550].

The second alternative is that the media source sends information on the correlation of the various timestamps to the MSAS. This can be done by using the IDMS report block from [RFC7272], using the Synchronization Packet Sender Types (SPST) 3 and 4, as specified in [TS183063] Annex W, and registered with IANA in the IDMS XR Block SPST Registry. These SPSTs are normally used for synchronization in case a transcoder is changing the media stream such that the RTP timestamps also change. In this case, synchronization would be impossible between users receiving the original stream and users

receiving the transcoded version. A transcoder can link an incoming RTP timestamps (SPST=3) to an outgoing RTP timestamp (SPST=4), and thus enable correlating the timelines. These SPSTs 3 and 4 can also be used if one source sends out two version of the same content, linking the timestamps of one stream to those of the other stream.

5. IANA Considerations

This document defines two new RSI Sub-Report Blocks, the "IDMS Received" and the "IDMS Presented". Based on the specification in section 2, these two sub-report blocks are added to the IANA registry for Sub-Report Block Type (SRBT) Values for the RSI Packet, as part of the RTP parameters registration.

6. Security Considerations

The content of this ID does not pose any security risks above or beyond those mentioned in [RFC5760] and [RFC7272].

7. Acknowledgements

8. 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.
- [RFC3605] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.
- [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.
- [RFC7272] Brandenburg, R. van, Stokking, H., Deventer, O. van, Boronat, F., Montagud, M., Gross, K., "Inter-Destination Media Synchronization (IDMS) Using the RTP Control Protocol (RTCP)", RFC 7272, June 2014.
- [TS183063] ETSI, "Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); IMS-based IPTV stage 3 specification", TS 183 063 v3.6.1, August

2014.

Authors' Addresses

Hans Stokking
TNO
Brassersplein 2
Delft, 2292CH
The Netherlands

Phone: 0031-(0)88-86 67 278
Fax:
Email: hans.stokking@tno.nl
URI:

Oskar van Deventer
TNO
Brassersplein 2
Delft, 2292CH
The Netherlands

Phone: 0031-(0)88-86 67 078
Fax:
Email: oskar.vandeventer@tno.nl
URI:

Ray van Brandenburg
TNO
Brassersplein 2
Delft, 2292CH
The Netherlands

Phone: 0031-(0)88-86 63 609
Fax:
Email: ray.vanbrandenburg@tno.nl
URI:

Fernando Boronat
Universitat Politecnica de Valencia
IGIC Institute, Universitat Politecnica de Valencia-Campus de Gandia
(UPV), C/ Paraninfo, 1, Grao de Gandia Valencia, 46730
Spain

Phone: +34 962 849 341
Fax:

Email: fboronat@dcom.upv.es
URI:

Mario Montagud
Universitat Politecnica de Valencia
IGIC Institute, Universitat Politecnica de Valencia-Campus de Gandia
(UPV), C/ Paraninfo, 1, Grao de Gandia Valencia, 46730
Spain

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