

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 22, 2015

M. Westerlund
B. Burman
Ericsson
R. Even
Huawei Technologies
M. Zanaty
Cisco Systems
July 21, 2014

RTP Header Extension for RTCP Source Description Items
draft-westerlund-avtext-sdes-hdr-ext-02

Abstract

Source Description (SDS) items are normally transported in RTP control protocol (RTCP). In some cases it can be beneficial to speed up the delivery of these items. Mainly when a new source (SSRC) joins an RTP session and the receivers needs this source's relation to other sources and its synchronization context, which are fully or partially identified using SDS items. To enable this optimization, this document specifies a new RTP header extension that can carry any type of SDS items.

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 January 22, 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Definitions	3
2.1. Requirements Language	3
2.2. Terminology	3
3. Motivation	3
4. Specification	5
4.1. SDES Item Header Extension	5
4.1.1. One-Byte Format	5
4.1.2. Two-Byte Format	5
4.2. Usage of the SDES Item Header Extension	6
4.2.1. One or Two Byte Headers	6
4.2.2. MTU and Packet Expansion	6
4.2.3. Transmission Considerations	7
4.2.4. Different Usages	8
4.2.5. SDES Items in RTCP	8
5. IANA Considerations	9
6. Security Considerations	10
7. Acknowledgements	10
8. References	10
8.1. Normative References	10
8.2. Informative References	11
Authors' Addresses	12

1. Introduction

This specification defines an RTP header extension [RFC3550][RFC5285] that can carry RTCP source description (SDES) items. By including selected SDES items in an header extension the determination of relationship and synchronization context for new RTP streams (SSRCs) in an RTP session can be speeded up. Which relationship and what information depends on the SDES items carried. This becomes a complement to using only RTCP for SDES Item delivery.

First, some requirements language is defined. The following section motivates why this header extension is sometimes required or at least provides a significant improvement compared to waiting for regular RTCP packet transmissions of the information. This is followed by a specification of the header extension. Next, a sub-space of the

header-extension URN is defined to be used for existing and future SDES items, and the existing SDES items are registered.

2. Definitions

2.1. 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].

2.2. Terminology

This document uses terminology defined in "A Taxonomy of Grouping Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources" [I-D.ietf-avtext-rtp-grouping-taxonomy] . In particular the following definitions:

Media Source

RTP Stream

Media Encoder

Encoded Stream

Participant

3. Motivation

Source Description (SDES) items are being associated with a particular SSRC and thus RTP stream. The source description items provide various meta data associated with the SSRC. How important it is to have this data no later than when receiving the first RTP packets depends on the item itself. The CNAME item is one item that is commonly needed if not at reception of the first RTP packet for this SSRC, so at least by the time the first media can be played out. If not, the synchronization context cannot be determined and thus any related streams cannot be correctly synchronized. Thus, this is a great example for the need to have this information early when a new RTP stream is received.

The main reason for new SSRCs in an RTP session is that a media sources are added. This either because an end-point is adding a new actual media source, or additional participants in a multi-party session being added to the session. Another reason for a new SSRC can be an SSRC collision that forces the colliding parties to select a new SSRC.

Returning to the case of rapid media synchronization, there exist an RTP header extension for Rapid Synchronization of RTP Flows [RFC6051]. That header extension carries the clock information present in the RTCP sender report (SR) packets. It however assumes that the CNAME binding is known, which can be provided via signaling in some cases, but not all. Thus an RTP header extension for carrying SDES items like CNAME is a powerful combination to enable rapid synchronization in all cases.

The Rapid Synchronization of RTP Flows specification does provide an analysis of the initial synchronization delay for different sessions depending on number of receivers as well as on session bandwidth (Section 2.1 of [RFC6051]). These results are applicable also for other SDES items that have a similar time dependency until the information can be sent using RTCP. Thus the benefit for reduction of initial delay before information is available can be determined for some use cases from these figures.

That document also discusses the case of late joiners, and defines an RTCP Feedback format to request synchronization information, which is another potential use case for SDES items in RTP header extension. It would for example be natural to include CNAME SDES item with the header extension containing the NTP formatted reference clock to ensure synchronization.

Some new SDES items are currently proposed, which can all benefit from timely delivery:

SRCNAME: This is a media source and encoding identifier to enable support for simulcast and improve some scalable encoding usages [I-D.westerlund-avtext-rtcp-sdes-srcname]. This SDES item could be used both for new sources and late joiners.

APPID: This SDES item provides an application specific identifier dynamically assigned to a particular RTP stream. The intention is to provide a receiver with information about the current role of the received RTP stream or its usage in an application [I-D.even-mmusic-application-token]. Thus a particular ID can be reassigned many times during the lifetime of an RTP session. This puts additional timing requirements, not only for new sources and late joiners, but also whenever the Application token is reassigned to another stream.

Based on the above, there appear to be good reasons why an RTP header extension for SDES items is worthwhile to pursue.

4. Specification

This section first specifies the SDES item RTP header extension format, followed by some usage considerations.

4.1. SDES Item Header Extension

The RTP header extension scheme that allows for multiple extensions to be included is defined in "A General Mechanism for RTP Header Extensions" [RFC5285]. That specification defines both short and long item headers. The short headers (One-byte) are restricted to 1 to 16 bytes of data, while the long format (Two-byte) supports a data length of 0 to 255 bytes. Thus that RTP header extension format is capable of supporting any SDES item from a data length perspective.

The ID field, independent of short or long format, identifies both the type of RTP header extension and, in the case of the SDES item header extension, the type of SDES item. The mapping is done in signaling by identifying the header extension and SDES item type using a URN, which is defined in the IANA consideration (Section 5) for all existing SDES items.

4.1.1. One-Byte Format

The one-byte header format for an SDES item extension element consists of the One-Byte header (defined in Section 4.2 of [RFC5285]), which consists of a 4-bit ID followed by a 4-bit length field (len) that identifies how many bytes (len value +1) of data that follows the header. The data part consists of len+1 bytes of UTF-8 text. The type of text is determined by the ID field value and its mapping to the type of SDES item.




Figure 1

4.1.2. Two-Byte Format

The two-byte header format for an SDES item extension element consists of the two-byte header (defined in Section 4.3 of [RFC5285]), which consists of an 8-bit ID followed by an 8-bit length field (len) that identifies how many bytes of data that follows the header. The data part consists of len bytes of UTF-8 text. The type

of text is determined by the ID field value and its mapping to the type of SDES item.




Figure 2

4.2. Usage of the SDES Item Header Extension

This section discusses various usage considerations; which form of header extension to use, the packet expansion, and when to send SDES items in header extension.

4.2.1. One or Two Byte Headers

The RTP header extensions for SDES items MAY use either the one-byte or two-byte header formats, depending on the text value size for the used SDES items. The one-byte header SHOULD be used when all non SDES item header extensions supports the one-byte format and all SDES item text values contain at most 16 bytes. Note that the RTP header extension specification does not allow mixing one-byte and two-byte headers for the same RTP stream (SSRC), so if the value size of any of the SDES items value requires the two-byte header, the all other header extensions MUST also use the two-byte header format.

For example using CNAMEs that are generated according to "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)" [RFC7022], using short term persistent values, and if 96-bit random values prior to base64 encoding are sufficient, then they will fit into the One-Byte header format.

4.2.2. MTU and Packet Expansion

The RTP packet size will clearly increase when they include the header extension. How much depends on which header extensions and their data parts. The SDES items can vary in size. There are also some use-cases which require transmitting multiple SDES items in the same packet to ensure that all relevant data reaches the receiver. An example of that is when you need both the CNAME, a SRCNAME and an appId plus the rapid time synchronization extension from RFC 6051. Such a combination is quite likely to result in at least 16+3+1+8 bytes of data plus the headers, which will be another 8 bytes for one-byte headers, thus in total 36 bytes.

The packet expansion can cause an issue when it cannot be taken into account when producing the RTP payload. Thus an RTP payload that is created to meet a particular IP level Maximum Transmission Unit (MTU), taking the addition of IP/UDP/RTP headers into account but excluding RTP header extensions suddenly exceeds the MTU, resulting in IP fragmentation. IP fragmentation is known to negatively impact the loss rate due to middleboxes unwilling or not capable of dealing with IP fragments.

As this is a real issue, the media encoder and payload packetizer should be flexible and be capable of handling dynamically varying payload size restrictions to counter the packet expansion caused by header extensions. If that is not possible, some reasonable worst case packet expansion should be calculated and used to reduce the RTP payload size of all RTP packets the sender transmits.

4.2.3. Transmission Considerations

The general recommendation is to only send header extensions when needed. This is especially true for SDES items that can be sent in periodic repetitions of RTCP throughout the whole session. Thus, the different usages (Section 4.2.4) have different recommendations. First some general considerations for getting the header extensions delivered to the receiver:

1. The probability for packet loss and burst loss determine how many repetitions of the header extensions will be required to reach a targeted delivery probability, and if burst loss is likely what dispersion would be needed to avoid getting multiple header extensions lost in a single burst.
2. How early the SDES item information is needed, from the first received RTP data or only after some set of packets are received, can guide if the header extension(s) should be in all of the first N packets or be included only once per set of packets, for example once per video frame.
3. The use of RTP level robustness mechanisms, such as RTP retransmission [RFC4588], or Forward Error Correction, e.g., [RFC5109] may treat packets differently from a robustness perspective, and SDES header extensions should be added to packets that get a treatment corresponding to the relative importance of receiving the information.

In summary, the number of header extension transmissions should be tailored to a desired probability of delivery taking the receiver population size into account. For the very basic case, N repetitions of the header extensions should be sufficient, but may not be

optimal. N is selected so that probability of delivery of at least one out of the N reaches the target value when calculating $1-P^N$, where P is the probability of packet loss. For point to point or small receiver populations, it might also be possible to use feedback, such as RTCP, to determine when the information in the header extensions has likely reached all receivers.

4.2.4. Different Usages

4.2.4.1. New SSRC

A new SSRC joins an RTP session. As this SSRC is completely new for everyone, the goal is to ensure that all receivers with high probability receives the information in the header extension. Thus header extension transmission strategies that allow some margins in the delivery probability should be considered.

4.2.4.2. Late Joiner

In a multi-party RTP session where one or a small number of receivers join a session where the majority of receivers already have all necessary information, the use of header extensions to deliver relevant information should be tailored to reach the new receivers. The trigger to send header extensions can for example either be RTCP from new receiver(s) or an explicit request like the Rapid Resynchronization Request defined in [RFC6051].

4.2.4.3. Information Change

In cases when the SDES item text value is changed and the new SDES information is tightly coupled to and thus needs to be synchronized with a related change in the RTP stream, use of a header extension is far superior to RTCP SDES. In this case it is equal or even more important with timely SDES information than in the case of new SSRCS (Section 4.2.4.1). Continued use of the old SDES information can lead to really undesired effects in the application. Application Token [I-D.even-mmusic-application-token] would be one such case. Thus, header extension transmission strategies with high probability of delivery should be chosen.

4.2.5. SDES Items in RTCP

As this RTP header extensions information, i.e. SDES Items can and will be sent also in RTCP it is worth some reflections on this interaction. There also exist the possibility to schedule a non-regular RTCP packet transmission containing important SDES items if one uses a RTP/AVPF based RTP profile. Depending on which mode ones RTCP feedback transmitter is working on extra RTCP packets may be

sent as immediate or early packets, enabling more timely deliver of SDES information.

There is however two aspects that differ between using RTP header extension and any non-regular transmission of RTCP packets. First, as the RTCP packet is a separate packet, there is no direct relation and also no fate sharing between the relevant media data and the SDES information. The order of arrival for the packets will matter. With a header-extension the SDES items can be ensured to arrive if the media data to played out arrives. Secondly, it is difficult to determine if an RTCP packet is actually delivered. This, as the RTCP packets lack both sequence number or a mechanism providing feedback on the RTCP packets themselves.

5. IANA Considerations

This IANA section firstly proposes to:

- o Reserve the SDES item RTP header extension defined in this document for use with current and future SDES items.
- o Register and assign the URN sub-space "urn:ietf:params:rtp-hdext:sdes:" in the RTP Compact Header Extensions registry.

The reason to require registering a URN within that sub-space is that the name represent an RTCP Source Description item, where a specification is strongly recommended. The formal policy is maintained from the main space, i.e. Expert Review.

Secondly, it is proposed that all the current existing SDES items are registered for usage in the RTP Compact Header Extensions registry :

URN	SDES Item	Reference
=====		
urn:ietf:params:rtp-hdext:sdes:cname	CNAME	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:name	NAME	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:email	EMAIL	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:phone	PHONE	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:loc	LOC	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:tool	TOOL	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:note	NOTE	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:priv	PRIV	[RFC3550]
urn:ietf:params:rtp-hdext:sdes:h323-caddr	H323-CADDR	[Vineet_Kumar]
urn:ietf:params:rtp-hdext:sdes:apsi	APSI	[RFC6776]

6. Security Considerations

Source Description items may contain data that are sensitive from a security perspective. There exist SDES items that are or may be sensitive from a user privacy perspective, like CNAME, NAME, EMAIL, PHONE, LOC and H323-CADDR. Others may contain sensitive information like NOTE and PRIV, while others may be sensitive from profiling implementations for vulnerability or other reasons, like TOOL. The CNAME sensitivity can vary depending on how it is generated and what persistence it has. A short term CNAME identifier generated using a random number generator may have minimal security implications, while one of the form user@host has privacy concerns and one generated from a MAC address has long term tracking potentials.

The above security concerns may have to be put in relation to needs of third party monitoring. In RTP sessions where any type of confidentiality protection is enabled, the SDES item header extensions SHOULD also be protected per default. This implies that to provide confidentiality, users of SRTP need to implement encrypted header extensions per [RFC6904]. Commonly, it is expected that the same security level is applied both RTCP packets carrying SDES items, as a RTP header extension containing a SDES item. If the security level is different it is important to consider the security properties as the worst in each aspect for the different configurations.

As the SDES items are used by the RTP based application to establish relationships between RTP streams or between an RTP stream and information about the originating Participant, there SHOULD be strong requirements on integrity and source authentication of the header extensions. If not, an attacker can modify the SDES item value to create erroneous relationship bindings in the receiving application.

7. Acknowledgements

The authors likes to thanks the following individuals for feedback and suggestions; Colin Perkins.

8. References

8.1. 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.

- [RFC5285] Singer, D. and H. Desineni, "A General Mechanism for RTP Header Extensions", RFC 5285, July 2008.
- [RFC6904] Lennox, J., "Encryption of Header Extensions in the Secure Real-time Transport Protocol (SRTP)", RFC 6904, April 2013.

8.2. Informative References

- [I-D.even-mmusic-application-token]
Even, R., Lennox, J., and Q. Wu, "The Session Description Protocol (SDP) Application Token Attribute", draft-even-mmusic-application-token-03 (work in progress), April 2014.
- [I-D.ietf-avtext-rtp-grouping-taxonomy]
Lennox, J., Gross, K., Nandakumar, S., and G. Salgueiro, "A Taxonomy of Grouping Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources", draft-ietf-avtext-rtp-grouping-taxonomy-01 (work in progress), February 2014.
- [I-D.westerlund-avtext-rtcp-sdes-srcname]
Westerlund, M., "RTCP Source Description Item SRCNAME to Label Individual Media Sources", draft-westerlund-avtext-rtcp-sdes-srcname-03 (work in progress), October 2013.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, July 2006.
- [RFC5109] Li, A., "RTP Payload Format for Generic Forward Error Correction", RFC 5109, December 2007.
- [RFC6051] Perkins, C. and T. Schierl, "Rapid Synchronisation of RTP Flows", RFC 6051, November 2010.
- [RFC6776] Clark, A. and Q. Wu, "Measurement Identity and Information Reporting Using a Source Description (SDES) Item and an RTCP Extended Report (XR) Block", RFC 6776, October 2012.
- [RFC7022] Begen, A., Perkins, C., Wing, D., and E. Rescorla, "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)", RFC 7022, September 2013.

Authors' Addresses

Magnus Westerlund
Ericsson
Farogatan 6
SE-164 80 Stockholm
Sweden

Phone: +46 10 714 82 87
Email: magnus.westerlund@ericsson.com

Bo Burman
Ericsson
Kistavagen 25
SE-164 80 Stockholm
Sweden

Phone: +46 10 714 13 11
Email: bo.burman@ericsson.com

Roni Even
Huawei Technologies
Tel Aviv
Israel

Email: roni.even@mail01.huawei.com

Mo Zanaty
Cisco Systems
7100 Kit Creek
RTP, NC 27709
USA

Email: mzanaty@cisco.com